

Acronis

acronis.com

Cyber Protection

24.01

Table of contents

- About Secure Zone 19**
 - Why use Secure Zone? 19
 - Limitations 19
 - How creating Secure Zone transforms the disk 19
 - How to create Secure Zone 20
 - How to delete Secure Zone 21
- Getting started with Cyber Protection 23**
 - Activating the account 23
 - Password requirements 23
 - Two-factor authentication 23
 - Privacy settings 25
 - Accessing the Cyber Protection service 26
 - Software requirements 27
 - Supported web browsers 27
 - Supported operating systems and environments 27
 - Supported Microsoft SQL Server versions 33
 - Supported Microsoft Exchange Server versions 33
 - Supported Microsoft SharePoint versions 34
 - Supported Oracle Database versions 34
 - Supported SAP HANA versions 34
 - Supported MySQL versions 34
 - Supported MariaDB versions 34
 - Supported virtualization platforms 35
 - Compatibility with encryption software 47
 - Compatibility with Dell EMC Data Domain storages 48
 - Supported protection features by operating system 49
 - Supported operating systems and versions 49
 - Supported file systems 58
- Installing and deploying Cyber Protection agents 61**
 - Preparation 61
 - Step 1 61
 - Step 2 61
 - Step 3 61
 - Step 4 61
 - Step 5 62

Step 6	63
Which agent do I need?	63
Agent-based and agentless backup	67
Which backup type do I need?	67
System requirements for agents	68
Linux packages	70
Are the required packages already installed?	71
Installing the packages from the repository	72
Installing the packages manually	73
Configuring proxy server settings	74
Installing protection agents	78
Downloading protection agents	78
Installing protection agents in Windows	78
Installing protection agents in Linux	80
Installing protection agents in macOS	83
Granting the required system permissions to the Connect Agent	84
Changing the logon account on Windows machines	85
Dynamic installation and uninstallation of components	87
Unattended installation or uninstallation	87
Unattended installation or uninstallation in Windows	87
Examples	88
Example	89
Examples	90
Examples	97
Example	98
Examples	99
Unattended installation or uninstallation in Linux	104
Unattended installation and uninstallation in macOS	110
Registering and unregistering workloads manually	119
Passwords with special characters or blank spaces	123
Changing the registration of a workload	124
Autodiscovery of machines	124
Prerequisites	125
How autodiscovery works	125
How remote installation of agents works	127
Autodiscovery and manual discovery	127
Managing discovered machines	132

Troubleshooting	133
Deploying Agent for VMware (Virtual Appliance)	134
Before you start	134
Deploying the OVF template	135
Configuring the virtual appliance	135
Deploying Agent for Scale Computing HC3 (Virtual Appliance)	138
Before you start	138
Deploying the QCOW2 template	139
Configuring the virtual appliance	140
Agent for Scale Computing HC3 – required roles	142
Deploying Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance)	143
Before you start	143
Configuring networks in Virtuozzo Hybrid Infrastructure	144
Configuring user accounts in Virtuozzo Hybrid Infrastructure	145
Deploying the QCOW2 template	147
Configuring the virtual appliance	148
Deploying Agent for oVirt (Virtual Appliance)	151
Before you start	151
Deploying the OVA template	152
Configuring the virtual appliance	153
Agent for oVirt – required roles and ports	156
Deploying Agent for Synology	157
Before you start	157
Downloading the setup program	158
Installing Agent for Synology	158
Updating Agent for Synology	163
Deploying agents through Group Policy	165
Prerequisites	165
Generating a registration token	166
Creating the transform file and extracting the installation packages	168
Setting up the Group Policy object	169
SSH connections to a virtual appliance	170
Starting the Secure Shell daemon	170
Setting the root password on a virtual appliance	170
Accessing a virtual appliance via an SSH client	171
Updating agents	171
Updating agents manually	172

Updating agents automatically	174
Updating agents on BitLocker-protected workloads	176
Preventing unauthorized uninstallation or modification of agents	176
Uninstalling agents	177
Protection settings	179
Automatic updates for components	179
Updating the Cyber Protection definitions by schedule	180
Updating the Cyber Protection definitions on-demand	180
Cache storage	180
Changing the service quota of machines	181
Cyber Protection services installed in your environment	182
Services installed in Windows	182
Services installed in macOS	183
Saving an agent log file	183
Site-to-site Open VPN - Additional information	183
License management for on-premises management servers	190
Defining how and what to protect	191
The Management tab	191
Plan statuses	191
Protection plans	192
Backup plans for cloud applications	192
Backup scanning plans	192
Off-host data processing	193
VM heartbeat	201
Screenshot validation	201
Intermediate snapshots	208
Protection plans and modules	208
Creating a protection plan	209
Actions with protection plans	210
Resolving plan conflicts	214
Default protection plans	215
Individual protection plans for hosting control panel integrations	221
#CyberFit Score for machines	221
How it works	222
Running a #CyberFit Score scan	226
Cyber Scripting	228
Prerequisites	228

Limitations	228
Supported platforms	228
Scripts	229
Script repository	234
Scripting plans	235
Script quick run	244
User roles and Cyber Scripting rights	245
Protection of collaboration and communication applications	247
Understanding your current level of protection	248
Monitoring	248
The Overview dashboard	248
The Activities dashboard	249
The Alerts dashboard	250
Alert types	251
Alert widgets	269
Cyber Protection	270
Protection status	270
Endpoint Detection and Response (EDR) widgets	271
#CyberFit Score by machine	275
Disk health monitoring	276
Data protection map	280
Vulnerability assessment widgets	281
Patch installation widgets	282
Backup scanning details	284
Recently affected	284
Cloud applications	285
Software inventory widgets	286
Hardware inventory widgets	287
Remote sessions widget	288
Smart protection	288
The Activities tab	295
Cyber Protect Monitor	296
Configuring proxy server settings in Cyber Protect Monitor	297
Reports	297
Actions with reports	299
Reported data according to widget type	301
Managing workloads in the Cyber Protect console	303

The Cyber Protect console	303
What's new in the Cyber Protect console	304
Using the Cyber Protect console as a partner administrator	304
Workloads	308
Adding workloads to the Cyber Protect console	310
Removing workloads from the Cyber Protect console	315
Device groups	318
Built-in groups and custom groups	319
Static groups and dynamic groups	319
Cloud-to-cloud groups and non-cloud-to-cloud groups	320
Creating a static group	321
Adding workloads to a static group	322
Creating a dynamic group	322
Editing a dynamic group	339
Deleting a group	340
Applying a plan to a group	340
Revoking a plan from a group	341
Working with the Device control module	342
Using device control	344
Access settings	351
Device types allowlist	356
USB devices allowlist	357
Excluding processes from access control	361
Device control alerts	363
Wiping data from a managed workload	366
Managing the isolation of workloads	367
Isolating a workload from the network	368
Managing network exclusions	369
Viewing workloads managed by RMM integrations	370
CyberApp workloads	371
Aggregated workloads	371
Working with CyberApp workloads	371
Working with aggregated workloads	372
Linking workloads to specific users	373
Find the last logged in user	374
Managing the backup and recovery of workloads and files	375
Backup	375

Protection plan cheat sheet	377
Selecting data to back up	379
Selecting entire machine	379
Selecting disks or volumes	380
Selecting files or folders	383
Selecting system state	385
Selecting ESXi configuration	386
Continuous data protection (CDP)	386
How it works	387
Supported data sources	388
Supported destinations	389
Configuring a CDP backup	389
Selecting a destination	390
Advanced storage option	391
About Secure Zone	392
Backup schedule	395
Backup schemes	395
Backup types	397
Running a backup on a schedule	397
Running a backup manually	410
Retention rules	411
Important tips	412
Retention rules according to the backup scheme	412
Configuring retention rules	415
Replication	416
Usage examples	416
Supported locations	416
Encryption	417
Configuring encryption in the protection plan	418
Configuring encryption as a machine property	418
Notarization	420
How to use notarization	421
How it works	421
Default backup options	421
Backup options	422
Availability of the backup options	422
Alerts	424

Backup consolidation	425
Backup file name	425
Backup format	430
Backup validation	431
Changed block tracking (CBT)	432
Cluster backup mode	432
Compression level	433
Error handling	434
Fast incremental/differential backup	435
File filters (Inclusions/Exclusions)	435
File-level backup snapshot	437
Forensic data	438
Log truncation	446
LVM snapshotting	447
Mount points	447
Multi-volume snapshot	448
One-click recovery	448
Performance and backup window	453
Physical Data Shipping	457
Pre/Post commands	458
Pre/Post data capture commands	460
Scheduling	463
Sector-by-sector backup	463
Splitting	464
Task failure handling	464
Task start conditions	465
Volume Shadow Copy Service (VSS)	465
Volume Shadow Copy Service (VSS) for virtual machines	467
Weekly backup	469
Windows event log	469
Recovery	469
Recovery cheat sheet	469
Safe recovery	471
Recovering a machine	472
Prepare drivers	481
Check access to the drivers in bootable environment	482
Automatic driver search	482

Mass storage drivers to install anyway	482
Recovering files	484
Recovering system state	490
Recovering ESXi configuration	490
Recovery options	491
Operations with backups	499
The Backup storage tab	499
Mounting volumes from a backup	501
Validating backups	502
Exporting backups	503
Deleting backups	504
Understanding the detection of bottlenecks	506
Backing up workloads to public clouds	510
Defining a backup location in Microsoft Azure	510
Viewing and updating Microsoft Azure backup locations	512
Managing public cloud account access	513
Protecting Microsoft applications	517
Protecting Microsoft SQL Server and Microsoft Exchange Server	517
Protecting Microsoft SharePoint	517
Protecting a domain controller	517
Recovering applications	517
Prerequisites	518
Database backup	520
Application-aware backup	526
Mailbox backup	528
Recovering SQL databases	530
Recovering Exchange databases	538
Recovering Exchange mailboxes and mailbox items	540
Changing the SQL Server or Exchange Server access credentials	546
Protecting mobile devices	547
Supported mobile devices	547
What you can back up	547
What you need to know	547
Where to get the Cyber Protect app	548
How to start backing up your data	548
How to recover data to a mobile device	549
How to review data via the Cyber Protect console	549

Protecting Hosted Exchange data	550
What items can be backed up?	550
What items can be recovered?	550
Selecting Exchange Online mailboxes	551
Recovering mailboxes and mailbox items	551
Protecting Microsoft 365 data	553
Why back up Microsoft 365 data?	553
Cloud agent and local agent	554
Required user rights	556
Limitations	557
Microsoft 365 seats licensing report	558
Logging	558
Using the locally installed Agent for Office 365	558
Using the cloud Agent for Microsoft 365	562
Protecting Google Workspace data	593
What does Google Workspace protection mean?	593
Required user rights	594
About the backup schedule	594
Limitations	595
Logging	595
Adding a Google Workspace organization	595
Creating a personal Google Cloud project	596
Discovering Google Workspace resources	599
Setting the frequency of Google Workspace backups	600
Protecting Gmail data	601
Protecting Google Drive files	604
Protecting Shared drive files	608
Notarization	612
Search in cloud-to-cloud backups	613
Full-text search	614
Search indexes	614
Checking the size of a search index	615
Updating, rebuilding, or deleting indexes	615
Enabling enhanced search in encrypted backups	616
Enabling or disabling enhanced search in existing plans	617
Disabling full-text search for Gmail backups	617
Protecting Oracle Database	618

Protecting SAP HANA	618
Protecting MySQL and MariaDB data	618
Configuring an application-aware backup	619
Recovering data from an application-aware backup	620
Protecting websites and hosting servers	624
Protecting websites	624
Protecting web hosting servers	627
Special operations with virtual machines	628
Running a virtual machine from a backup (Instant Restore)	628
Working in VMware vSphere	632
Backing up clustered Hyper-V machines	650
Limiting the total number of simultaneously backed-up virtual machines	650
Machine migration	652
Microsoft Azure and Amazon EC2 virtual machines	655
Creating bootable media to recover operating systems	656
Custom or ready-made bootable media?	656
Linux-based or WinPE/WinRE-based bootable media?	657
Creating physical bootable media	657
Bootable Media Builder	658
Recovery from the cloud storage	662
Recovery from a network share	662
Files of a script	663
Structure of autostart.json	664
Top-level object	664
Variable object	664
Control type	665
Connecting to a machine booted from bootable media	672
Local operations with bootable media	672
Remote operations with bootable media	674
Startup Recovery Manager	676
Implementing disaster recovery	679
About Cyber Disaster Recovery Cloud	679
The key functionality	679
Software requirements	680
Supported operating systems	680
Supported virtualization platforms	680
Limitations	681

Cyber Disaster Recovery Cloud trial version	682
Limitations when using Geo-redundant Cloud Storage	682
Disaster Recovery compatibility with encryption software	682
Compute points	683
Setting up the disaster recovery functionality	684
Create a disaster recovery protection plan	684
Editing the Recovery server default parameters	686
Cloud network infrastructure	687
Setting up connectivity	687
Networking concepts	688
Initial connectivity configuration	698
Prerequisites	700
Network management	706
Prerequisites	721
Setting up recovery servers	722
Creating a recovery server	722
How failover works	725
How failback works	733
Prerequisites	735
Prerequisites	740
Working with encrypted backups	743
Operations with Microsoft Azure virtual machines	744
Setting up primary servers	744
Creating a primary server	744
Operations with a primary server	747
Managing the cloud servers	747
Firewall rules for cloud servers	748
Setting firewall rules for cloud servers	749
Checking the cloud firewall activities	751
Backing up the cloud servers	752
Orchestration (runbooks)	752
Why use runbooks?	753
Creating a runbook	753
Operations with runbooks	756
Configuring your antivirus and antimalware protection	758
Supported platforms	758
Supported features per platform	759

Antivirus and antimalware protection	761
Antimalware features	762
Scanning types	762
Antivirus and antimalware protection settings	763
Active Protection in the Cyber Backup Standard edition	778
Active Protection settings in Cyber Backup Standard	779
URL filtering	785
How it works	786
URL filtering configuration workflow	788
URL filtering settings	788
Description	794
Microsoft Defender Antivirus and Microsoft Security Essentials	794
Schedule scan	795
Default actions	795
Real-time protection	796
Advanced	796
Exclusions	797
Firewall management	797
Quarantine	798
How do files get into the quarantine folder?	798
Managing quarantined files	799
Quarantine location on machines	799
Self-service custom folder on-demand	799
Corporate whitelist	800
Automatic adding to the whitelist	800
Manual adding to the whitelist	800
Adding quarantined files to the whitelist	801
Whitelist settings	801
Viewing details about items in the whitelist	801
Antimalware scan of backups	801
Limitations	802
Working with Advanced protection features	804
Advanced Data Loss Prevention	806
Creating the data flow policy and policy rules	806
Enabling Advanced Data Loss Prevention in protection plans	815
Automated detection of destination	818
Sensitive data definitions	818

Data Loss Prevention events	824
Advanced Data Loss Prevention widgets on the Overview dashboard	825
Custom sensitivity categories	826
Organization map	828
Known issues and limitations	831
Endpoint Detection and Response (EDR)	831
Why you need Endpoint Detection and Response (EDR)	832
Enabling Endpoint Detection and Response (EDR) functionality	834
How to use Endpoint Detection and Response (EDR)	836
Viewing which incidents are currently not mitigated	839
Understanding the scope and impact of incidents	840
How to navigate attack stages	848
Enabling monitoring mode for Endpoint Detection and Response (EDR)	882
How to test if Endpoint Detection and Response (EDR) is working correctly	883
Assessing vulnerabilities and managing patches	885
Vulnerability assessment	885
Supported Microsoft and third-party products	885
Supported Apple and third-party products	887
Supported Linux products	888
Vulnerability assessment settings	888
Vulnerability assessment for Windows machines	890
Vulnerability assessment for Linux machines	891
Vulnerability assessment for macOS devices	891
Managing found vulnerabilities	892
Patch management	893
The patch management workflow	894
Patch management settings in the protection plan	894
Viewing the list of available patches	899
Automatic patch approval	901
Approving patches manually	906
Installing patches on demand	906
Managing your software and hardware inventory	908
Software inventory	908
Enabling the software inventory scanning	908
Running a software inventory scan manually	909
Browsing the software inventory	909
Viewing the software inventory of a single device	911

Hardware inventory	912
Enabling the hardware inventory scanning	912
Running a hardware inventory scan manually	913
Browsing the hardware inventory	913
Viewing the hardware of a single device	916
Connecting to workloads for remote desktop or remote assistance	918
Supported remote desktop and assistance features	919
Supported platforms	922
Remote connection protocols	923
NEAR	923
RDP	924
Apple Screen Sharing	924
Remote sound redirection	924
Connections to remote workloads for remote desktop or remote assistance	925
Remote management plans	926
Creating a remote management plan	926
Adding a workload to a remote management plan	934
Removing workloads from a remote management plan	934
Additional operations with existing remote management plans	935
Compatibility issues with remote management plans	937
Resolving compatibility issues with remote management plans	938
Workload credentials	939
Adding credentials	939
Assigning credentials to a workload	940
Deleting credentials	940
Unassigning credentials from a workload	940
Working with managed workloads	941
Configuring RDP settings	941
Connecting to managed workloads for remote desktop or remote assistance	942
Connecting to a managed workload via a web client	944
Transferring files	945
Performing control actions on managed workloads	946
Monitoring workloads via screenshot transmission	947
Observing multiple managed workloads simultaneously	948
Working with unmanaged workloads	949
Connecting to unmanaged workloads via Acronis Quick Assist	950
Connecting to unmanaged workloads via IP address	950

Transferring files via Acronis Quick Assist	951
Using the toolbar in the Viewer window	952
Recording and playing remote connection sessions	954
Configuring the Connect Client settings	955
The remote desktop notifiers	956
Monitoring the health and performance of workloads	958
Monitoring plans	958
Monitoring types	958
Anomaly-based monitoring	958
Supported platforms for monitoring	959
Configurable monitors	959
Settings of the Disk space monitor	963
Settings of the CPU temperature monitor	966
Settings of the GPU temperature monitor	967
Settings of the Hardware changes monitor	969
Settings of the CPU usage monitor	969
Settings of the Memory usage monitor	971
Settings of the Disk transfer rate monitor	973
Settings of the Network usage monitor	975
Settings of the CPU usage by process monitor	978
Settings of the Memory usage by process monitor	978
Settings of the Disk transfer rate by process monitor	979
Settings of the Network usage by process monitor	980
Settings of the Windows service status monitor	982
Settings of the Process status monitor	982
Settings of the Installed software monitor	983
Settings of the Last system restart monitor	983
Settings of the Windows event log monitor	984
Settings of the Files and folders size monitor	985
Settings of the Windows Update status monitor	986
Settings of the Firewall status monitor	986
Settings of the Failed logins monitor	986
Settings of the Antimalware software status monitor	987
Settings of the AutoRun feature status monitor	988
Settings of the Custom monitor	989
Monitoring plans	990
Creating a monitoring plan	990

Adding workloads to monitoring plans	992
Revoking monitoring plans	993
Configuring automatic response actions	993
Additional operations with monitoring plans	995
Compatibility issues with monitoring plans	998
Resolving compatibility issues with monitoring plans	998
Resetting the machine learning models	999
Monitoring alerts	1000
Configuring monitoring alerts	1000
Monitoring alert variables	1001
Manual response actions	1003
Viewing the monitoring alerts for a workload	1006
Viewing the alert log of monitoring alerts	1006
Configuring email notification policies	1007
Viewing monitor data	1008
Monitor widgets	1009
Additional Cyber Protection tools	1011
Enhanced security mode	1011
Limitations	1011
Unsupported features	1011
Setting the encryption password	1011
Changing the encryption password	1012
Recovering backups for tenants in the Enhanced security mode	1012
Immutable storage	1013
Immutable storage modes	1013
Supported storages and agents	1013
Enabling immutable storage	1014
Disabling immutable storage	1014
Accessing deleted backups in immutable storage	1015
Geo-redundant storage	1015
Enabling and disabling geo-redundant storage	1015
Geo-replication status	1016
Limitations	1016
Glossary	1018
Index	1022

About Secure Zone

Secure Zone is a secure partition on a disk of the backed-up machine. It can store backups of disks or files of this machine.

Should the disk experience a physical failure, the backups located in the Secure Zone may be lost. That's why Secure Zone should not be the only location where a backup is stored. In enterprise environments, Secure Zone can be thought of as an intermediate location used for backup when an ordinary location is temporarily unavailable or connected through a slow or busy channel.

Why use Secure Zone?

Secure Zone:

- Enables recovery of a disk to the same disk where the disk's backup resides.
- Offers a cost-effective and handy method for protecting data from software malfunction, virus attack, human error.
- Eliminates the need for a separate media or network connection to back up or recover the data. This is especially useful for roaming users.
- Can serve as a primary destination when using replication of backups.

Limitations

- Secure Zone cannot be organized on a Mac.
- Secure Zone is a partition on a basic disk. It cannot be organized on a dynamic disk or created as a logical volume (managed by LVM).
- Secure Zone is formatted with the FAT32 file system. Because FAT32 has a 4-GB file size limit, larger backups are split when saved to Secure Zone. This does not affect the recovery procedure and speed.

How creating Secure Zone transforms the disk

- Secure Zone is always created at the end of the hard disk.
- If there is no or not enough unallocated space at the end of the disk, but there is unallocated space between volumes, the volumes will be moved to add more unallocated space to the end of the disk.
- When all unallocated space is collected but it is still not enough, the software will take free space from the volumes you select, proportionally reducing the volumes' size.
- However, there should be free space on a volume, so that the operating system and applications can operate; for example, create temporary files. The software will not decrease a volume where free space is or becomes less than 25 percent of the total volume size. Only when all volumes on the disk have 25 percent or less free space, will the software continue decreasing the volumes proportionally.

As is apparent from the above, specifying the maximum possible Secure Zone size is not advisable. You will end up with no free space on any volume, which might cause the operating system or applications to work unstably and even fail to start.

Important

Moving or resizing the volume from which the system is booted requires a reboot.

How to create Secure Zone

1. Select the machine that you want to create Secure Zone on.
2. Click **Details > Create Secure Zone** .
3. Under **Secure Zone disk**, click **Select**, and then select a hard disk (if several) on which to create the zone.

The software calculates the maximum possible size of Secure Zone.

4. Enter the Secure Zone size or drag the slider to select any size between the minimum and the maximum ones.

The minimum size is approximately 50 MB, depending on the geometry of the hard disk. The maximum size is equal to the disk's unallocated space plus the total free space on all of the disk's volumes.

5. If all unallocated space is not enough for the size you specified, the software will take free space from the existing volumes. By default, all volumes are selected. If you want to exclude some volumes, click **Select volumes**. Otherwise, skip this step.

✕ Create Secure Zone

Secure Zone disk

Disk 1, 60.0 GB

Maximum possible size of Secure Zone: 35.9 GB

Secure Zone size:

20 GB

There is not enough unallocated space. Free space will be taken from all volumes where it is present.

Select volumes

Password protection

Off

6. [Optional] Enable the **Password protection** switch and specify a password.
The password will be required to access the backups located in Secure Zone. Backing up to Secure Zone does not require a password, unless the backup is performed under bootable media.
7. Click **Create**.
The software displays the expected partition layout. Click **OK**.
8. Wait while the software creates Secure Zone.

You can now choose Secure Zone in **Where to back up** when creating a protection plan.

How to delete Secure Zone

1. Select a machine with Secure Zone.
2. Click **Details**.
3. Click the gear icon next to **Secure Zone**, and then click **Delete**.
4. [Optional] Specify the volumes to which the space freed from the zone will be added. By default, all volumes are selected.
The space will be distributed equally among the selected volumes. If you do not select any volumes, the freed space will become unallocated.
Resizing the volume from which the system is booted requires a reboot.
5. Click **Delete**.

As a result, Secure Zone will be deleted along with all backups stored in it.

Getting started with Cyber Protection

Activating the account

When an administrator creates an account for you, an email message is sent to your email address. The message contains the following information:

- **Your login.** This is the user name that you use to log in. Your login is also shown on the account activation page.
- **Activate account** button. Click the button and set the password for your account. Ensure that your password is at least nine characters long. For more information about the password, refer to "Password requirements" (p. 23).

If your administrator has enabled two-factor authentication, you will be prompted to set it up for your account. For more information about it, refer to "Two-factor authentication" (p. 23).

Password requirements

The password for a user account must be at least 9 characters long. Passwords are also checked for complexity, and fall into one of the following categories:

- Weak
- Medium
- Strong

You cannot save a weak password, even though it might contain 9 characters or more. Passwords that repeat the user name, the login, the user email, or the name of the tenant to which a user account belongs are always considered weak. Most common passwords are also considered weak.

To strengthen a password, add more characters to it. Using different types of characters, such as digits, uppercase and lowercase letters, and special characters, is not mandatory but it results in stronger passwords that are also shorter.

Two-factor authentication

Two-factor authentication (2FA) provides extra protection from unauthorized access to your account. When 2FA is set up, you are required to enter your password (the first factor) and a one-time code (the second factor) to log in to the Cyber Protect console. The one-time code is generated by a special application that must be installed on your mobile phone or another device that belongs to you. Even if someone discovers your login and password, they will not be able to log in to your account without having access to your second-factor device.

To set up two-factor authentication for your account

You must set up 2FA for your account if the administrator has enabled it for your organization. If the administrator enables 2FA while you are logged in to the Cyber Protect console, you will have to set it up when your current session expires.

Prerequisites

- Two-factor authentication is enabled for your organization by an administrator.

To set up two-factor authentication for your account

1. Install an authenticator app on your mobile device.

Examples of authenticator apps:

- Twilio Authy
- Microsoft Authenticator
- Google Authenticator

2. Scan the QR code using your authenticator app, and then enter the 6-digit code displayed on the authenticator app in the **Set up two-factor authentication** window.

3. Click **Next**.

The instructions on how to restore your access to your account if you lose your 2FA device or uninstall the authenticator app are displayed.

4. Save or print the PDF file.

Note

Ensure that you save the PDF file in a safe place or print it for further reference. This is the best way to restore your access.

5. Return to the Cyber Protect console login page and enter the generated code.

A one-time code is valid for 30 seconds. If you wait longer than 30 seconds, use the next generated code.

Next time you log in, you can select the **Trust this browser...** check box. In this case, the code will not be required for subsequent logins by using this browser on this machine.

Note

We recommend that you leave this check box clear. Otherwise, you will lose the access to 2FA for your account.

To restore two-factor authentication on a new device (2FA)

If you have access to the previously set-up mobile authentication app

1. Install an authenticator app on your new device.
2. Use the PDF file that you saved when you configured 2FA on your device. This file contains the 32-digit code that you must enter in the authenticator app to link the authenticator app to your Acronis account again.

Important

If the code is not working, ensure that the time in the authenticator mobile app is synced with your device.

If you did not save the PDF file during the setup:

- a. Click **Reset 2FA**, and then enter the one-time password shown in the mobile authenticator app.
- b. Follow the on-screen instructions.

If you do not have access to the previously set-up mobile authenticator app

1. Take a new mobile device.
2. Use the stored PDF file to link a new device (default name of the file is `cyberprotect-2fa-backupcode.pdf`).
3. Restore access to your account from backup. Ensure that backups are supported by your mobile app.
4. Open the app under the same account from another mobile device if it is supported by the app.

Privacy settings

Privacy settings help you indicate whether or not you give consent for the collection, use and disclosure of your personal information.

Depending on the country in which you are using Cyber Protect Cloud and the Cyber Protect Cloud data center that provides services to you, on the initial launch of Cyber Protect Cloud you may be asked to confirm whether you agree to use Google Analytics in Cyber Protect Cloud.

Google Analytics helps us better understand user behavior and improve user experience in Cyber Protect Cloud by collecting pseudonymized data.

If you enabled or refused to enable Google Analytics on the initial launch of Cyber Protect Cloud, you can change your decision at any time later.

To enable or disable Google Analytics

1. In the Cyber Protect console, click **Manage account**.
2. Click the account icon in the upper-right corner.
3. Select **My privacy settings**. The **My privacy settings** window is displayed.
4. In the **Google Analytics data collection** section, click one of the following buttons:
 - **On** to enable Google Analytics
 - **Off** to disable Google Analytics

In the **How to delete cookies** section, you can control and manage cookies directly in your browser.

Note

If you do not see Google Analytics section, it means that Google Analytics is not used in your country.

In the **In-product onboarding and interactive help** section, shown initially during trial period, you can stop or keep receiving the information about the improvements and new features in the program in the future. This feature is enabled by default, but you can disable it by switching the toggle to **Off**.


Accessing the Cyber Protection service

After you activate your account, you can access the Cyber Protection service by logging in to the Cyber Protect console or via the management portal.

To log in to the Cyber Protect console

1. Go to the Cyber Protection service login page.
2. Type your login, and then click **Next**.
3. Type your password, and then click **Next**.
4. [If you use more than one Cyber Protect Cloud service] Click **Cyber Protection**.

Users who only have access the Cyber Protection service, log in directly to the Cyber Protect console.

If **Cyber Protection** is not the only service you have access to, you can switch between the services by using the  icon in the upper-right corner. Administrators can also use this icon for switching to the management portal.

The timeout period for the Cyber Protect console is 24 hours for active sessions and 1 hour for idle sessions.

You can change the language of the web interface by clicking the account icon in the upper-right corner.

To access the Cyber Protect console via the management portal

1. In the management portal, go to **Monitoring > Usage**.
2. Under **Cyber Protect**, select **Protection**, and then click **Manage service**.
Alternatively, under **Clients**, select a customer, and then click **Manage service**.

As a result, you are redirected to the Cyber Protect console.

Important

If the customer is in **Self-service** management mode, you cannot manage services for him. Only the customer administrators can change the customer mode to **Managed by service provider**, and then manage the services.

To reset your password

1. Go to the Cyber Protection service login page.
2. Type your login, and then click **Next**.
3. Click **Forgot password?**
4. Confirm that you want further instructions by clicking **Send**.
5. Follow the instructions in the email that you have received.
6. Set up your new password.

Software requirements

Supported web browsers

The Cyber Protect console supports the following web browsers:

- Google Chrome 29 or later
- Mozilla Firefox 23 or later
- Opera 16 or later
- Microsoft Edge 25 or later
- Safari 8 or later running in the macOS and iOS operating systems

In other web browsers (including Safari browsers running in other operating systems), the user interface might be displayed incorrectly or some functions may be unavailable.

Supported operating systems and environments

Agent for Windows

This agent includes a component for Antivirus & Antimalware protection and URL Filtering. See "Supported protection features by operating system" (p. 49) for details about supported functionality by operating system.

- Windows XP Professional SP1 (x64), SP2 (x64), SP3 (x86)
- Windows Server 2003 SP1/2003 R2 and later – Standard and Enterprise editions (x86, x64)
- Windows Small Business Server 2003/2003 R2
- Windows Server 2008, Windows Server 2008 SP2* – Standard, Enterprise, Datacenter, Foundation, and Web editions (x86, x64)
- Windows Small Business Server 2008, Windows Small Business Server 2008 SP2*
- Windows 7 – all editions

Note

To use Cyber Protection with Windows 7, you must install the following updates from Microsoft before installing the protection agent:

- [Windows 7 Extended Security Updates \(ESU\)](#)
- [KB4474419](#)
- [KB4490628](#)

For more information on the required updates, refer to [this knowledge base article](#).

- Windows Server 2008 R2* – Standard, Enterprise, Datacenter, Foundation, and Web editions
- Windows Home Server 2011*
- Windows MultiPoint Server 2010*/2011*/2012
- Windows Small Business Server 2011* – all editions

- Windows 8/8.1 – all editions (x86, x64), except for the Windows RT editions
- Windows Server 2012/2012 R2 – all editions
- Windows Storage Server 2003/2008/2008 R2/2012/2012 R2/2016
- Windows 10 – Home, Pro, Education, Enterprise, IoT Enterprise and LTSC (formerly LTSB) editions
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows 11 – all editions
- Windows Server 2022 – all installation options, except for Nano Server

Note

* To use Cyber Protection with this version of Windows, you must install the SHA2 code signing support update from Microsoft ([KB4474419](#)) before installing the protection agent.

For information on issues related to the SHA2 code signing support update, refer to [this knowledge base article](#).

Agent for SQL, Agent for Active Directory, Agent for Exchange (for database backup and application-aware backup)

Each of these agents can be installed on a machine running any operating system listed above and a supported version of the respective application.

Agent for Data Loss Prevention

Device control

- Microsoft Windows 7 Service Pack 1 and later
- Microsoft Windows Server 2008 R2 and later
- macOS 10.15 (Catalina)
- macOS 11.2.3 (Big Sur)
- macOS 12 (Monterey)
- macOS 13 (Ventura)

Note

Agent for Data Loss Prevention for macOS supports only x64 processors. Apple silicon ARM-based processors are not supported.

Data loss prevention

- Microsoft Windows 7 Service Pack 1 and later
- Microsoft Windows Server 2008 R2 and later

Note

Agent for Data Loss Prevention might be installed on unsupported macOS systems because it is an integral part of Agent for Mac. In this case, the Cyber Protect console will indicate that Agent for Data Loss Prevention is installed on the computer, but the device control and data loss prevention functionality will not work. Device control functionality will only work on macOS systems that are supported by Agent for Data Loss Prevention.

Agent for Advanced Data Loss Prevention

- Microsoft Windows 7 Service Pack 1 and later
- Microsoft Windows Server 2008 R2 and later

Agent for File Sync & Share

For the list of supported operating systems, refer to the [Cyber Files Cloud user guide](#).

Agent for Exchange (for mailbox backup)

- Windows Server 2008 – Standard, Enterprise, Datacenter, Foundation, and Web editions (x86, x64)
- Windows Small Business Server 2008
- Windows 7 – all editions
- Windows Server 2008 R2 – Standard, Enterprise, Datacenter, Foundation, and Web editions
- Windows MultiPoint Server 2010/2011/2012
- Windows Small Business Server 2011 – all editions
- Windows 8/8.1 – all editions (x86, x64), except for the Windows RT editions
- Windows Server 2012/2012 R2 – all editions
- Windows Storage Server 2008/2008 R2/2012/2012 R2
- Windows 10 – Home, Pro, Education, and Enterprise editions
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows 11 – all editions
- Windows Server 2022 – all installation options, except for Nano Server

Agent for Microsoft 365

- Windows Server 2008 – Standard, Enterprise, Datacenter, Foundation, and Web editions (x64 only)
- Windows Small Business Server 2008
- Windows Server 2008 R2 – Standard, Enterprise, Datacenter, Foundation, and Web editions
- Windows Home Server 2011
- Windows Small Business Server 2011 – all editions

- Windows 8/8.1 – all editions (x64 only), except for the Windows RT editions
- Windows Server 2012/2012 R2 – all editions
- Windows Storage Server 2008/2008 R2/2012/2012 R2/2016 (x64 only)
- Windows 10 – Home, Pro, Education, and Enterprise editions (x64 only)
- Windows Server 2016 – all installation options (x64 only), except for Nano Server
- Windows Server 2019 – all installation options (x64 only), except for Nano Server
- Windows 11 – all editions
- Windows Server 2022 – all installation options, except for Nano Server

Agent for Oracle

- Windows Server 2008R2 – Standard, Enterprise, Datacenter, and Web editions (x86, x64)
- Windows Server 2012R2 – Standard, Enterprise, Datacenter, and Web editions (x86, x64)
- Linux – any kernel and distribution supported by Agent for Linux (listed below)

Agent for MySQL/MariaDB

- Linux – any kernel and distribution supported by Agent for Linux (listed below)

Agent for Linux

This agent includes a component for Antivirus & Antimalware protection and URL Filtering. See "Supported protection features by operating system" (p. 49) for details about supported functionality by operating system.

The following Linux distributions and kernel versions have been specifically tested. However, even if your Linux distribution or kernel version is not listed below, it may still work correctly in all required scenarios, due to the specifics of the Linux operating systems.

If you encounter issues while using Cyber Protection with your combination of Linux distribution and kernel version, contact the Support team for further investigation.

Linux with kernel from 2.6.9 to 5.19 and glibc 2.3.4 or later, including the following x86 and x86_64 distributions:

- Red Hat Enterprise Linux 4.x, 5.x, 6.x, 7.x, 8.x*, 9.0*, 9.1*, 9.2*, 9.3*
- Ubuntu 9.10, 10.04, 10.10, 11.04, 11.10, 12.04, 12.10, 13.04, 13.10, 14.04, 14.10, 15.04, 15.10, 16.04, 16.10, 17.04, 17.10, 18.04, 18.10, 19.04, 19.10, 20.04, 20.10, 21.04, 21.10, 22.04, 22.10, 23.04
- Fedora 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 37, 38
- SUSE Linux Enterprise Server 10, 11, 12, 15

Important

Configurations with Btrfs are not supported for SUSE Linux Enterprise Server 12 and SUSE Linux Enterprise Server 15.

- Debian 4.x, 5.x, 6.x, 7.0, 7.2, 7.4, 7.5, 7.6, 7.7, 8.0, 8.1, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7, 8.8, 8.11, 9.0, 9.1, 9.2, 9.3, 9.4, 9.5, 9.6, 9.7, 9.8, 10, 11
- CentOS 8.x*
- CentOS Stream 8*, 9*
- Oracle Linux 5.x, 6.x, 7.x, 8.x*, 9.0*, 9.1*, 9.2* – both Unbreakable Enterprise Kernel and Red Hat Compatible Kernel

Note

Installing the protection agent on Oracle Linux 8.6 and later, on which Secure Boot is enabled, requires manual signing of kernel modules. For more information on how to sign a kernel module, refer to [this knowledge base article](#).

- CloudLinux 5.x, 6.x, 7.x, 8.x*
- ClearOS 5.x, 6.x, 7.x
- AlmaLinux 8.x*, 9.0*, 9.1*, 9.2*
- Rocky Linux 8.x*, 9.0*, 9.1*, 9.2*, 9.3*
- ALT Linux 7.0

* Starting from version 8.4, supported only with kernels from 4.18 to 5.19

Agent for Mac

This agent includes a component for Antivirus & Antimalware protection and URL Filtering. See "Supported protection features by operating system" (p. 49) for details about supported functionality by operating system.

Both x64 and ARM architecture (used in Apple silicon processors such as Apple M1 and M2) are supported.

Note

You cannot recover disk-level backups of Intel-based Macs to Macs that use Apple silicon processors, and vice-versa. You can recover files and folders.

- macOS High Sierra 10.13
- macOS Mojave 10.14
- macOS Catalina 10.15
- macOS Big Sur 11
- macOS Monterey 12
- macOS Ventura 13
- macOS Sonoma 14

Important

Starting from version C23.07, Cyber Protect Cloud does not support the following operating systems: OS X Yosemite 10.10, OS X El Capitan 10.11, and macOS Sierra 10.12.

We strongly recommend that you upgrade your operating system to a supported version in order to ensure compatibility and be able to use the full functionality of Cyber Protect Cloud.

Agent for VMware (Virtual Appliance)

This agent is delivered as a virtual appliance for running on an ESXi host.

VMware ESXi 4.1, 5.0, 5.1, 5.5, 6.0, 6.5, 6.7, 7.0, 8.0

Agent for VMware (Windows)

This agent is delivered as a Windows application for running in any operating system listed above for Agent for Windows with the following exceptions:

- 32-bit operating systems are not supported.
- Windows XP, Windows Server 2003/2003 R2, and Windows Small Business Server 2003/2003 R2 are not supported.

Agent for Hyper-V

- Windows Server 2008 (x64 only) with Hyper-V role, including Server Core installation mode
- Windows Server 2008 R2 with Hyper-V role, including Server Core installation mode
- Microsoft Hyper-V Server 2008/2008 R2
- Windows Server 2012/2012 R2 with Hyper-V role, including Server Core installation mode
- Microsoft Hyper-V Server 2012/2012 R2
- Windows 8, 8.1 (x64 only) with Hyper-V
- Windows 10 – Pro, Education, and Enterprise editions with Hyper-V
- Windows Server 2016 with Hyper-V role – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2016
- Windows Server 2019 with Hyper-V role – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2019
- Windows Server 2022 – all installation options, except for Nano Server

Agent for Virtuozzo

- Virtuozzo 6.0.10, 6.0.11, 6.0.12, 7.0.13, 7.0.14
- Virtuozzo Hybrid Server 7.5

Agent for Virtuozzo Hybrid Infrastructure

Virtuozzo Hybrid Infrastructure 3.5, 4.0, 4.5, 4.6, 4.7, 5.0, 5.1, 5.2, 5.3

Agent for Scale Computing HC3

Scale Computing Hypercore 8.8, 8.9, 9.0, 9.1

Agent for oVirt

Red Hat Virtualization 4.2, 4.3, 4.4, 4.5

Agent for Synology

DiskStation Manager 6.2.x, 7.x

Agent for Synology only supports NAS devices with x86_64 processors. You cannot install the agent on devices with ARM processors.

Cyber Protect Monitor

- Windows 7 and later
- Windows Server 2008 R2 and later
- All macOS versions that are supported by Agent for Mac

Supported Microsoft SQL Server versions

- Microsoft SQL Server 2022
- Microsoft SQL Server 2019
- Microsoft SQL Server 2017
- Microsoft SQL Server 2016
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2008
- Microsoft SQL Server 2005

The SQL Server Express editions of the above SQL server versions are supported as well.

Supported Microsoft Exchange Server versions

- Microsoft Exchange Server 2019 – all editions.
- Microsoft Exchange Server 2016 – all editions.
- Microsoft Exchange Server 2013 – all editions, Cumulative Update 1 (CU1) and later.
- Microsoft Exchange Server 2010 – all editions, all service packs. Mailbox backup and granular recovery from database backups are supported starting with Service Pack 1 (SP1).
- Microsoft Exchange Server 2007 – all editions, all service packs. Mailbox backup and granular recovery from database backups are not supported.

Supported Microsoft SharePoint versions

Cyber Protection supports the following Microsoft SharePoint versions:

- Microsoft SharePoint 2013
- Microsoft SharePoint Server 2010 SP1
- Microsoft SharePoint Foundation 2010 SP1
- Microsoft Office SharePoint Server 2007 SP2*
- Microsoft Windows SharePoint Services 3.0 SP2*

*In order to use SharePoint Explorer with these versions, you need a SharePoint recovery farm to attach the databases to.

The backups or databases from which you extract data must originate from the same SharePoint version as the one where SharePoint Explorer is installed.

Supported Oracle Database versions

- Oracle Database version 11g, all editions
- Oracle Database version 12c, all editions
- Oracle Database version 19c, all editions
- Oracle Database version 21c, all editions

Only single-instance configurations are supported.

Supported SAP HANA versions

HANA 2.0 SPS 03 installed in RHEL 7.6 running on a physical machine or VMware ESXi virtual machine.

Because SAP HANA does not support recovery of multitenant database containers by using storage snapshots, this solution supports SAP HANA containers with only one tenant database.

Supported MySQL versions

- 5.5.x – Community Server, Enterprise, Standard, and Classic editions
- 5.6.x – Community Server, Enterprise, Standard, and Classic editions
- 5.7.x – Community Server, Enterprise, Standard, and Classic editions
- 8.0.x – Community Server, Enterprise, Standard, and Classic editions

Supported MariaDB versions

- 10.0.x
- 10.1.x
- 10.2.x
- 10.3.x

- 10.4.x
- 10.5.x
- 10.6.x
- 10.7.x

Supported virtualization platforms

The following table summarizes how various virtualization platforms are supported.

For more information about the differences between the agent-based and agentless backup, refer to "Agent-based and agentless backup" (p. 67).

Note

The following hypervisor vendors and versions supported via the **Agent-based backup (Backup from inside a guest OS)** method have been specifically tested. However, even if you run a hypervisor from a vendor or hypervisor with a version that is not listed below, the **Agent-based backup (Backup from inside a guest OS)** method may still work correctly in all required scenarios.

If you encounter issues while using Cyber Protection with your combination of hypervisor vendor and version, contact the Support team for further investigation.

VMware

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
VMware vSphere versions: 4.1, 5.0, 5.1, 5.5, 6.0, 6.5, 6.7, 7.0, 8.0 VMware vSphere editions: VMware vSphere Essentials* VMware vSphere Essentials Plus* VMware vSphere Standard* VMware vSphere Advanced VMware vSphere Enterprise VMware vSphere Enterprise Plus	Supported Devices > Add > Virtualization hosts > VMware ESXi > Agent for installation in Windows or Devices > Add > Virtualization hosts > VMware ESXi > Virtual appliance (OVF)	Supported Devices > Add > Workstations or Servers > Windows or Linux
VMware vSphere Hypervisor (Free ESXi)**	Not supported	Supported Devices > Add > Workstations or Servers > Windows or Linux

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
VMware Server (VMware Virtual server) VMware Workstation VMware ACE VMware Player	Not supported	Supported Devices > Add > Workstations or Servers > Windows or Linux

* In these editions, the HotAdd transport for virtual disks is supported on vSphere 5.0 and later. On version 4.1, backups may run slower.

** Backup at a hypervisor level is not supported for vSphere Hypervisor because this product restricts access to Remote Command Line Interface (RCLI) to read-only mode. The agent works during the vSphere Hypervisor evaluation period while no serial key is entered. Once you enter a serial key, the agent stops functioning.

Note

Acronis officially supports any update within the supported major vSphere version.

For example, vSphere 8.0 support includes support for any update within this version, unless stated otherwise. For example, vSphere 8.0 Update 1 is also supported along with originally released vSphere 8.0.

Support for specific VMware vSphere version means that vSAN of the corresponding version is also supported. For example, support for vSphere 8.0 means that vSAN 8.0 is also supported.

Limitations

- **Fault tolerant machines**

Agent for VMware backs up a fault tolerant machine only if fault tolerance was enabled in VMware vSphere 6.0 and later. If you upgraded from an earlier vSphere version, it is enough to disable and enable fault tolerance for each machine. If you are using an earlier vSphere version, install an agent in the guest operating system.

- **Independent disks and RDM**

Agent for VMware does not back up Raw Device Mapping (RDM) disks in physical compatibility mode or independent disks. The agent skips these disks and adds warnings to the log. You can avoid the warnings by excluding independent disks and RDMs in physical compatibility mode from the protection plan. If you want to back up these disks or data on these disks, install an agent in the guest operating system.

- **In-guest iSCSI connection**

Agent for VMware does not back up LUN volumes connected by an iSCSI initiator that works within the guest operating system. Because the ESXi hypervisor is not aware of such volumes, the volumes are not included in hypervisor-level snapshots and are omitted from a backup without a

warning. If you want to back up these volumes or data on these volumes, install an agent in the guest operating system.

- **Encrypted virtual machines** (introduced in VMware vSphere 6.5)
 - Encrypted virtual machines are backed up in an unencrypted state. If encryption is critical to you, enable encryption of backups [when creating a protection plan](#).
 - Recovered virtual machines are always unencrypted. You can manually enable encryption after the recovery is complete.
 - If you back up encrypted virtual machines, we recommend that you also encrypt the virtual machine where Agent for VMware is running. Otherwise, operations with encrypted machines may be slower than expected. Apply the **VM Encryption Policy** to the agent's machine by using vSphere Web Client.
 - Encrypted virtual machines will be backed up via LAN, even if you configure the SAN transport mode for the agent. The agent will fall back on the NBD transport because VMware does not support SAN transport for backing up encrypted virtual disks.
- **Secure Boot**
 - VMware virtual machines: (introduced in VMware vSphere 6.5) **Secure Boot** is disabled after a virtual machine is recovered as a new virtual machine. You can manually enable this option after the recovery is complete. This limitation applies to VMware.
 - Hyper-V virtual machines: For all GEN2 VMs, Secure Boot is disabled after the virtual machine is recovered to both new virtual machine or an existing virtual machine.
- **ESXi configuration backup** is not supported for VMware vSphere 7.0.

- **Linux machines containing logical volumes (LVM)**

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.
For more information about the migrations scenarios, see "Machine migration" (p. 652).
- Running a virtual machine from a backup.

Microsoft

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
<p>Windows Server 2008 (x64) with Hyper-V</p> <p>Windows Server 2008 R2 with Hyper-V</p> <p>Microsoft Hyper-V Server 2008/2008 R2</p> <p>Windows Server 2012/2012 R2 with Hyper-V</p> <p>Microsoft Hyper-V Server 2012/2012 R2</p> <p>Windows 8, 8.1 (x64) with Hyper-V</p> <p>Windows 10 with Hyper-V</p> <p>Windows Server 2016 with Hyper-V – all installation options, except for Nano Server</p> <p>Microsoft Hyper-V Server 2016</p> <p>Windows Server 2019 with Hyper-V – all installation options, except for Nano Server</p> <p>Microsoft Hyper-V Server 2019</p> <p>Windows Server 2022 with Hyper-V – all installation options, except for Nano Server</p>	<p>Supported</p> <p>Devices > Add > Virtualization hosts > Hyper-V</p>	<p>Supported</p> <p>Devices > Add > Workstations or Servers > Windows or Linux</p>
<p>Microsoft Virtual PC 2004, 2007</p> <p>Windows Virtual PC</p>	<p>Not supported</p>	<p>Supported</p> <p>Devices > Add > Workstations or Servers > Windows or Linux</p>
<p>Microsoft Virtual Server 2005</p>	<p>Not supported</p>	<p>Supported</p> <p>Devices > Add > Workstations or Servers > Windows or Linux</p>

Note

Hyper-V virtual machines running on a hyper-converged cluster with Storage Spaces Direct (S2D) are supported. Storage Spaces Direct is also supported as a backup storage.

Limitations

- **Pass-through disks**

Agent for Hyper-V does not back up pass-through disks. During backup, the agent skips these disks and adds warnings to the log. You can avoid the warnings by excluding pass-through disks from the protection plan. If you want to back up these disks or data on these disks, install an agent in the guest operating system.

- **Hyper-V guest clustering**

Agent for Hyper-V does not support backup of Hyper-V virtual machines that are nodes of a Windows Server Failover Cluster. A VSS snapshot at the host level can even temporarily disconnect the external quorum disk from the cluster. If you want to back up these machines, install agents in the guest operating systems.

- **In-guest iSCSI connection**

Agent for Hyper-V does not back up LUN volumes connected by an iSCSI initiator that works within the guest operating system. Because the Hyper-V hypervisor is not aware of such volumes, the volumes are not included in hypervisor-level snapshots and are omitted from a backup without a warning. If you want to back up these volumes or data on these volumes, install an agent in the guest operating system.

- **Secure Boot**

For all GEN2 VMs, Secure Boot is disabled after the virtual machine is recovered to both new virtual machine or an existing virtual machine.

- **Linux machines containing logical volumes (LVM)**

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.
For more information about the migrations scenarios, see "Machine migration" (p. 652).
- Running a virtual machine from a backup.

- **VHD/VHDX file names with ampersand symbols**

On Hyper-V hosts running Windows Server 2016 or later, you cannot back up legacy virtual machines (version 5.0) originally created with Hyper-V 2012 R2 or older, if the names of their VHD/VHDX files contain the ampersand symbol (&).

To be able to back up such machines, in Hyper-V Manager, detach the corresponding virtual disk from the virtual machine, edit the VHD/VHDX file name by removing the ampersand symbol, and then attach the disk back to the virtual machine.

- **Dependency on the Microsoft WMI subsystem**

Agentless backups of Hyper-V virtual machines depend on the Microsoft WMI subsystem, and in particular on the `Msvm_VirtualSystemManagementService` class. If the WMI queries fail, the backups will also fail. For more information about the `Msvm_VirtualSystemManagementService` class, see the [Microsoft documentation](#).

Scale Computing

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Scale Computing Hypercore 8.8, 8.9, 9.0, 9.1, 9.2, 9.3	Supported Devices > Add > Virtualization hosts > Scale Computing HC3	Supported Devices > Add > Workstations or Servers > Windows or Linux

Limitations

Linux machines containing logical volumes (LVM)

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.

For more information about the migrations scenarios, see "Machine migration" (p. 652).

- Running a virtual machine from a backup.

Citrix

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Citrix XenServer/Citrix Hypervisor 4.1.5, 5.5, 5.6, 6.0, 6.1, 6.2, 6.5, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 8.0, 8.1, 8.2	Not supported	Supported only for fully virtualized (aka HVM) guests. Paravirtualized (aka PV) guests are not supported. Devices > Add > Virtualization hosts > Citrix XenServer > Windows or Linux

Red Hat and Linux

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Red Hat Enterprise Virtualization (RHEV) 2.2, 3.0, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6 Red Hat Virtualization (RHV) 4.0, 4.1	Not supported	Supported Devices > Add > Workstations or Servers > Windows or Linux
Red Hat Virtualization (managed by oVirt) 4.2, 4.3, 4.4, 4.5	Supported Devices > Add > Virtualization hosts > Red Hat Virtualization (oVirt)	Supported Devices > Add > Workstations or Servers > Windows or Linux
Kernel-based Virtual Machines (KVM)	Not supported	Supported Devices > Add > KVM > Windows or Linux
Kernel-based Virtual Machines (KVM) managed by oVirt 4.3 running on Red Hat Enterprise Linux 7.6, 7.7 or CentOS 7.6, 7.7	Supported Devices > Add > Virtualization hosts > Red Hat Virtualization (oVirt)	Supported Devices > Add > Workstations or Servers > Windows or Linux
Kernel-based Virtual Machines (KVM) managed by oVirt 4.4	Supported Devices > Add >	Supported Devices > Add > Workstations

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
running on Red Hat Enterprise Linux 8.x or CentOS Stream 8.x	Virtualization hosts > Red Hat Virtualization (oVirt)	or Servers > Windows or Linux
Kernel-based Virtual Machines (KVM) managed by oVirt 4.5 running on Red Hat Enterprise Linux 8.x or CentOS Stream 8.x	Supported Devices > Add > Virtualization hosts > Red Hat Virtualization (oVirt)	Supported Devices > Add > Workstations or Servers > Windows or Linux

Limitations

Linux machines containing logical volumes (LVM)

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.
For more information about the migrations scenarios, see "Machine migration" (p. 652).
- Running a virtual machine from a backup.

Parallels

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Parallels Workstation	Not supported	Supported Devices > Add > Workstations or Servers > Windows or Linux

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Parallels Server 4 Bare Metal	Not supported	Supported Devices > Add > Workstations or Servers > Windows or Linux

Oracle

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Oracle Virtualization Manager (based on oVirt)* 4.3	Supported Devices > Add > Virtualization hosts > Red Hat Virtualization (oVirt)	Supported Devices > Add > Workstations or Servers > Windows or Linux
Oracle VM Server 3.0, 3.3, 3.4	Not supported	Supported only for fully virtualized (aka HVM) guests. Paravirtualized (aka PV) guests are not supported. Devices > Add > Virtualization hosts > Oracle > Windows or Linux
Oracle VM VirtualBox 4.x	Not supported	Supported Devices > Add > Virtualization hosts > Oracle > Windows or Linux

*Oracle Virtualization Manager is supported by [Agent for oVirt](#).

Limitations

Linux machines containing logical volumes (LVM)

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.

- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.

For more information about the migrations scenarios, see "Machine migration" (p. 652).

- Running a virtual machine from a backup.

Nutanix

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Nutanix Acropolis Hypervisor (AHV) 20160925.x through 20180425.x	Not supported	Supported Devices > Add > Virtualization hosts > Nutanix AHV > Windows or Linux

Virtuozzo

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Virtuozzo 6.0.10, 6.0.11, 6.0.12	Supported Devices > Add > Virtualization hosts > Virtuozzo	Supported for virtual machines only. Containers are not supported. Devices > Add > Workstations or Servers > Windows or Linux
Virtuozzo 7.0.13, 7.0.14	Supported for ploop containers only. Virtual machines are not supported. Devices > Add > Virtualization hosts > Virtuozzo	Supported for virtual machines only. Containers are not supported. Devices > Add > Workstations or Servers > Windows or Linux
Virtuozzo Hybrid Server 7.5	Supported	Supported for virtual machines

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
	Devices > Add > Virtualization hosts > Virtuozzo	only. Containers are not supported. Devices > Add > Workstations or Servers > Windows or Linux

Limitations

Linux machines containing logical volumes (LVM)

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.
For more information about the migrations scenarios, see "Machine migration" (p. 652).
- Running a virtual machine from a backup.

Vituzo Hybrid Infrastructure

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Virtuozzo Hybrid Infrastructure 3.5, 4.5, 4.6, 4.7, 5.0, 5.1, 5.2, 5.3	Supported Devices > Add > Virtualization hosts > Virtuozzo Hybrid infrastructure	Supported Devices > Add > Workstations or Servers > Windows or Linux

Limitations

- **Agentless backup of VMs with disks on an external iSCSI storage**

You cannot back up VMs from Virtuozzo Hybrid Infrastructure, if VM disks are placed on external iSCSI volumes (attached to the VHI cluster).

- **Linux machines containing logical volumes (LVM)**

The following operations are not supported for Linux machines with LVM that you back up in the agentless mode:

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- The file filters (Inclusions/Exclusions) are not applicable. Any configured inclusions or exclusions will be ignored. For more information about the file filters, see "File filters (Inclusions/Exclusions)" (p. 435).

The following operations are not supported for Linux machines with LVM that you back up in the agent-based mode (that is, by Agent for Linux installed on the backed-up machine):

- Performing a machine migration by recovering its backup as a virtual machine (for example, by using Agent for VMware, Agent for Hyper-V, Agent for oVirt, Agent for Virtuozzo, Agent for Virtuozzo Hybrid Infrastructure, or Agent for Scale Computing for P2V, V2P, or V2V migration). To recover data from such a backup, use a bootable media.
For more information about the migrations scenarios, see "Machine migration" (p. 652).
- Running a virtual machine from a backup.

Amazon

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Amazon EC2 instances	Not supported	Supported Devices > Add > Workstations or Servers > Windows or Linux

Microsoft Azure

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
Azure virtual machines	Not supported	Supported Devices > Add > Workstations

Platform	Agentless backup (Backup at the hypervisor level)	Agent-based backup (Backup from inside a guest OS)
		or Servers > Windows or Linux

Compatibility with encryption software

There are no limitations on backing up and recovering data that is encrypted by *file-level* encryption software.

Disk-level encryption software encrypts data on the fly. This is why data contained in the backup is not encrypted. Disk-level encryption software often modifies system areas: boot records, or partition tables, or file system tables. These factors affect disk-level backup and recovery, the ability of the recovered system to boot and access to Secure Zone.

You can back up the data encrypted by the following disk-level encryption software:

- Microsoft BitLocker Drive Encryption
- McAfee Endpoint Encryption
- PGP Whole Disk Encryption

To ensure reliable disk-level recovery, follow the common rules and software-specific recommendations.

Common installation rule

We strongly recommend that you install the encryption software before you install the protection agents.

The way of using Secure Zone

Secure Zone must not be encrypted with disk-level encryption. This is the only way to use Secure Zone:

1. Install the encryption software; then, install the agent.
2. Create Secure Zone.
3. Exclude Secure Zone when encrypting the disk or its volumes.

Common backup rule

You can do a disk-level backup in the operating system.

Software-specific recovery procedures

Microsoft BitLocker Drive Encryption

To recover a system that was encrypted by BitLocker:

1. Boot from the bootable media.
2. Recover the system. The recovered data will be unencrypted.
3. Reboot the recovered system.
4. Turn on BitLocker.

If you only need to recover one partition of a multi-partitioned disk, do so under the operating system. Recovery under bootable media may make the recovered partition undetectable for Windows.

McAfee Endpoint Encryption and PGP Whole Disk Encryption

You can recover an encrypted system partition by using bootable media only.

If the recovered system fails to boot, rebuild Master Boot Record as described in the following Microsoft knowledge base article: <https://support.microsoft.com/kb/2622803>

Compatibility with Dell EMC Data Domain storages

You can use Dell EMC Data Domain devices as backup storage.

With this storage, we recommend that you use a backup scheme that regularly creates full backups, for example **Always full**. To learn more about the available backup schemes, see "Backup schemes" (p. 395).

Retention lock (Governance mode) is supported. If retention lock is enabled, you need to add the AR_RETENTION_LOCK_SUPPORT environment variable to the machine with the protection agent that uses this storage as a backup destination.

Note

Dell EMC Data Domain storages with enabled retention lock are not supported by Agent for Mac.

To add the AR_RETENTION_LOCK_SUPPORT environment variable

In Windows

1. Log in as administrator to the machine with the protection agent.
2. In **Control Panel**, go to **System and Security > System > Advanced system settings**.
3. On the **Advanced tab**, click **Environment Variables**.
4. In the **System variables** panel, click **New**.
5. In the **New System Variable** window, add the new variable as follows:
 - Variable name: AR_RETENTION_LOCK_SUPPORT
 - Variable value: 1
6. Click **OK**.
7. In the **Environment Variables** window, click **OK**.
8. Restart the machine.

In Linux

1. Log in as administrator to the machine with the protection agent.
2. Go to the /sbin directory, and then open the acronis_mms file for editing.
3. Above the line `export LD_LIBRARY_PATH`, add the following line:

```
export AR_RETENTION_LOCK_SUPPORT=1
```

4. Save the acronis_mms file.
5. Restart the machine.

In a virtual appliance

1. Log in as administrator to the virtual appliance.
2. Go to the /bin directory, and then open the autostart file for editing.
3. Under the line `export LD_LIBRARY_PATH`, add the following line:

```
export AR_RETENTION_LOCK_SUPPORT=1
```

4. Save the autostart file.
5. Restart the virtual appliance machine.

Supported protection features by operating system

This topic contains information about the protection features of Cyber Protect Cloud. It does not list the backup and recovery features.

The protection features are only supported on machines on which a protection agent is installed. They are not available for virtual machines that are backed up in the agentless mode, for example, by Agent for Hyper-V, Agent for VMware, Agent for Virtuozzo Hybrid Infrastructure, Agent for Scale Computing, or Agent for oVirt.

Some features might require additional licensing, depending on the applied licensing model.

Supported operating systems and versions

Windows

Unless stated otherwise for a specific feature set, the following Windows versions are supported:

- Windows 7 Service Pack 1 and later
- Windows Server 2008 R2 Service Pack 1 and later

Note

For Windows 7, you must install the following updates from Microsoft before installing the protection agent.

- [Windows 7 Extended Security Updates \(ESU\)](#)
- [KB4474419](#)
- [KB4490628](#)

For more information on the required updates, refer to [this knowledge base article](#).

Linux

Supported Linux distributions and their versions depend on the feature sets, and are shown at the bottom of each table.

macOS

Supported macOS versions depend on the feature sets, and are shown at the bottom of each table.

Feature set	Windows	Linux	macOS
Default protection plans			
Remote Workers	Yes	No	No
Office Workers (third-party antivirus)	Yes	No	No
Office Workers (Cyber Protect antivirus)	Yes	No	No
Cyber Protect Essentials (only for Cyber Protect Essentials edition)	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Forensic backup			
Collecting memory dump	Yes	No	No
Snapshot of running processes	Yes	No	No
Notarization of local image forensic backup	Yes	No	No
Notarization of cloud image forensic backup	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Features	Windows	Linux	macOS
Continuous data protection (CDP)			
CDP for files and folders	Yes	No	No
CDP for changed files via application tracking	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Autodiscovery and remote installation			
Network-based discovery	Yes	No	No
Active Directory-based discovery	Yes	No	No
Template-based discovery (importing machines from a file)	Yes	No	No
Manual adding of devices	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Active Protection			
Process Injects detection	Yes	No	No
Automatic recovery of affected files from the local cache	Yes	Yes	Yes
Self-defense for Acronis backup files	Yes	No	No
Self-defense for Acronis software	Yes	No	Yes (Only Active Protection and antimalware components)
Trusted/blocked process management	Yes	No	Yes
Processes/folders exclusions	Yes	Yes	Yes
Ransomware detection based on a process behavior (AI-based)	Yes	Yes	Yes
Cryptomining process detection based on process behavior	Yes	No	No

Feature set	Windows	Linux	macOS
Active Protection			
External drives protection (HDD, flash drives, SD cards)	Yes	No	Yes
Network folder protection	Yes	Yes	Yes
Server-side protection	Yes	No	No
Zoom, Cisco Webex, Citrix Workspace, and Microsoft Teams protection	Yes	No	No
For more information about the supported operating systems and their versions, see "Supported platforms" (p. 758).			

Feature set	Windows	Linux	macOS
Antivirus and Antimalware protection			
Fully-integrated Active Protection functionality	Yes	No	No
Real-time antimalware protection	Yes	Yes, with the Advanced Antimalware pack	Yes, with the Advanced Antimalware pack
Advanced real-time antimalware protection with local signature-based detection	Yes	Yes	Yes
Static analysis for portable executable files	Yes	No	Yes*
On-demand antimalware scanning	Yes	Yes**	Yes
Network folder protection	Yes	Yes	No
Server-side protection	Yes	No	No
Scan of archive files	Yes	No	Yes
Scan of removable drives	Yes	No	Yes
Scan of new and changed files only	Yes	No	Yes
File/folder exclusions	Yes	Yes	Yes***
Processes exclusions	Yes	No	Yes
Behavioral analysis engine	Yes	No	Yes

Feature set	Windows	Linux	macOS
Antivirus and Antimalware protection			
Exploit prevention	Yes	No	No
Quarantine	Yes	Yes	Yes
Quarantine auto clean-up	Yes	No	Yes
URL filtering (http/https)	Yes	No	No
Corporate-wide whitelist	Yes	No	Yes
Firewall management****	Yes	No	No
Microsoft Defender Antivirus management*****	Yes	No	No
Microsoft Security Essentials management	Yes	No	No
Registering and managing Antivirus and Antimalware protection via Windows Security Center	Yes	No	No
For more information about the supported operating systems and their versions, see "Supported platforms" (p. 758).			

* Static analysis for portable executable files is supported only for scheduled scans on macOS.

** Start conditions are not supported for on-demand scanning on Linux.

*** File/folder exclusions are only supported for the case when you specify files and folders that will not be scanned by real-time protection or scheduled scans on macOS.

**** Firewall management is supported on Windows 8 and later. Windows Server is not supported.

***** Microsoft Defender Antivirus management is supported on Windows 8.1 and later.

Feature set	Windows	Linux	macOS
Vulnerability assessment			
Vulnerability assessment of operating system and its native applications	Yes	Yes*****	Yes
Vulnerability assessment for 3rd-party applications	Yes	No	Yes
For more information about the supported operating systems and their versions, refer to "Supported Microsoft and third-party products" (p. 885), "Supported Linux products" (p. 888), and "Supported Apple and third-party products" (p. 887).			

***** The vulnerability assessment depends on the availability of official security advisories for specific distribution, for example <https://lists.centos.org/pipermail/centos-announce>, <https://lists.centos.org/pipermail/centos-cr-announce>, and others.

Feature set	Windows	Linux	macOS
Patch management			
Patch auto-approval	Yes	No	No
Patch auto-installation	Yes	No	No
Patch testing	Yes	No	No
Manual patch installation	Yes	No	No
Patch scheduling	Yes	No	No
Fail-safe patching: backup of machine before installing patches as part of protection plan	Yes	No	No
Cancellation of a machine reboot if a backup is running	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Features	Windows	Linux	macOS
Data protection map			
Adjustable definition of important files	Yes	No	No
Scanning machines to find unprotected files	Yes	No	No
Unprotected locations overview	Yes	No	No
Ability to start the protection action from the Data protection map widget (Protect all files action)	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Disk health			
AI-based HDD and SSD health control	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Features	Windows	Linux	macOS
Smart protection plans based on Acronis Cyber Protection Operations Center (CPOC) alerts			
Threat feed	Yes	No	No
Remediation wizard	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Backup scanning			
Antimalware scan of image backups as part of backup plan	Yes	No	No
Scanning of image backups for malware in cloud	Yes	No	No
Malware scan of encrypted backups	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Safe recovery			
Antimalware scanning with Antivirus and Antimalware protection during the recovery process	Yes	No	No
Safe recovery for encrypted backups	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Remote desktop connection			
Connection via NEAR	Yes	Yes	Yes
Connection via RDP	Yes	No	No
Connection via Apple Screen Sharing	No	No	Yes
Connection via web client	Yes	No	No
Connection via Quick Assist	Yes	Yes	Yes
Remote assistance	Yes	Yes	Yes
File transfer	Yes	Yes	Yes

Feature set	Windows	Linux	macOS
Remote desktop connection			
Screenshot transmission	Yes	Yes	Yes
For more information about the supported operating systems and their versions, see "Supported platforms" (p. 922).			

Feature set	Windows	Linux	macOS
#CyberFit Score			
#CyberFit Score status	Yes	No	No
#CyberFit Score standalone tool	Yes	No	No
#CyberFit Score recommendations	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Data loss prevention			
Device control	Yes	No	Supported on Macs with Intel processors running macOS 10.15 and later or macOS 11.2.3 or later. Not supported on ARM-based Apple silicon processors, such as Apple M1 / M2.
Advanced Data Loss Prevention	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49).			

Feature set	Windows	Linux	macOS
Management options			
Upsell scenarios to promote Cyber Protect editions	Yes	Yes	Yes

Feature set	Windows	Linux	macOS
Management options			
Web-based centralized and remote management console	Yes	Yes	Yes
Supported operating systems and versions: Platform independent.			

Feature set	Windows	Linux	macOS
Protection options			
Remote wipe	Yes	No	No
Supported for Windows 10 and later.			

Feature set	Windows	Linux	macOS
Cyber Protect Monitor			
Cyber Protect app	Yes	No	Yes
Protection status for Zoom	Yes	No	No
Protection status for Cisco Webex	Yes	No	No
Protection status for Citrix Workspace	Yes	No	No
Protection status for Microsoft Teams	Yes	No	No
See the supported Windows versions in "Supported operating systems and versions" (p. 49). On macOS, Cyber Protect Monitor is supported for all versions on which you can install Agent for Mac. For more information, see "Agent for Mac" (p. 31).			

Feature set	Windows	Linux	macOS
Software inventory			
Software inventory scanning	Yes	No	Yes
Software inventory monitoring	Yes	No	Yes
See the supported Windows versions in "Supported operating systems and versions" (p. 49). On macOS, Software inventory is supported for versions 10.13.x – 13.x.			

Feature set	Windows	Linux	macOS
Hardware inventory			
Hardware inventory scanning	Yes	No	Yes
Hardware inventory monitoring	Yes	No	Yes
See the supported Windows versions in "Supported operating systems and versions" (p. 49). On macOS, Hardware inventory is supported for versions 10.13.x – 13.x.			

Supported file systems

A protection agent can back up any file system that is accessible from the operating system where the agent is installed. For example, Agent for Windows can back up and recover an ext4 file system if the corresponding driver is installed in Windows.

The following table summarizes the file systems that can be backed up and recovered (bootable media supports only recovery). The limitations apply to both the agents and bootable media.

File system	Supported by			Limitations
	Agents	Bootable media for Windows and Linux	Bootable media for Mac	
FAT16/32	All agents	+	+	No limitations
NTFS	All agents	+	+	
ext2/ext3/ext4	All agents	+	-	
HFS+	Agent for Mac	-	+	
APFS	Agent for Mac	-	+	<ul style="list-style-type: none"> Supported starting with macOS High Sierra 10.13 Disk configuration should be re-created manually when recovering to a non-original machine or bare metal.
JFS	Agent for Linux	+	-	<ul style="list-style-type: none"> File filters (Inclusions/Exclusions) are not supported Fast incremental/

File system	Supported by			Limitations
	Agents	Bootable media for Windows and Linux	Bootable media for Mac	
ReiserFS3	Agent for Linux	+	-	differential backup cannot be enabled
ReiserFS4	Agent for Linux	+	-	<ul style="list-style-type: none"> File filters (Inclusions/Exclusions) are not supported Fast incremental/differential backup cannot be enabled Volumes cannot be resized during a recovery
ReFS	All agents	+	+	
XFS	All agents	+	+	<ul style="list-style-type: none"> File filters (Inclusions/Exclusions) are not supported Fast incremental/differential backup cannot be enabled Volumes cannot be resized during a recovery The fast-incremental backup mode is not supported for the XFS file system. Incremental and differential backups of XFS volumes to the cloud may be significantly slower than comparable ext4 backups that use the fast-incremental mode.
Linux swap	Agent for Linux	+	-	No limitations
exFAT	All agents	+ Bootable media cannot be used for recovery if	+	<ul style="list-style-type: none"> Only disk/volume backup is supported File filters (Inclusions/Exclusions) are not supported

File system	Supported by			Limitations
	Agents	Bootable media for Windows and Linux	Bootable media for Mac	
		the backup is stored on exFAT		<ul style="list-style-type: none"> Individual files cannot be recovered from a backup

The software automatically switches to the sector-by-sector mode when backing up drives with unrecognized or unsupported file systems (for example, Btrfs). A sector-by-sector backup is possible for any file system that:

- is block-based
- spans a single disk
- has a standard MBR/GPT partitioning scheme

If the file system does not meet these requirements, the backup fails.

Data Deduplication

In Windows Server 2012 and later, you can enable the Data Deduplication feature for an NTFS volume. Data Deduplication reduces the used space on the volume by storing duplicate fragments of the volume's files only once.

You can back up and recover a data deduplication-enabled volume at a disk level, without limitations. File-level backup is supported, except when using Acronis VSS Provider. To recover files from a disk backup, either [run a virtual machine](#) from your backup, or [mount the backup](#) on a machine running Windows Server 2012 or later, and then copy the files from the mounted volume.

The Data Deduplication feature of Windows Server is unrelated to the Acronis Backup Deduplication feature.

Installing and deploying Cyber Protection agents

Preparation

Step 1

Choose an agent, depending on what you are going to back up. For more information on the possible choices, refer to [Which agent do I need?](#)

Step 2

Ensure that there is enough free space on your hard drive to install an agent. For detailed information about the required space, refer to "System requirements for agents" (p. 68).

Step 3

Download the setup program. To find the download links, click **All devices > Add**.

The **Add devices** page provides web installers for each agent that is installed in Windows. A web installer is a small executable file that downloads the main setup program from the Internet and saves it as a temporary file. This file is deleted immediately after the installation.

If you want to store the setup programs locally, download a package containing all agents for installation in Windows by using the link at the bottom of the **Add devices** page. Both 32-bit and 64-bit packages are available. These packages enable you to customize the list of components to install. These packages also enable unattended installation, for example, via Group Policy. This advanced scenario is described in "Deploying agents through Group Policy" (p. 165).

To download the setup program for Agent for Microsoft 365, click the account icon in the top-right corner, and then click **Downloads > Agent for Microsoft 365**.

Installation in Linux and macOS is performed from ordinary setup programs.

All setup programs require an Internet connection to register the machine in the Cyber Protection service. If there is no Internet connection, the installation will fail.

Step 4

Cyber Protect features require Microsoft Visual C++ 2017 Redistributable. Please ensure that it is already installed on your machine or install it before installing the agent. After the installation of Microsoft Visual C++, a restart may be required. You can find the Microsoft Visual C++ Redistributable package here <https://support.microsoft.com/help/2999226/update-for-universal-c-runtime-in-windows>.

Step 5

Verify that your firewalls and other components of your network security system (such as a proxy server) allow outbound connections through the following TCP ports.

- Ports **443** and **8443**
These ports are used for accessing the Cyber Protect console, registering the agents, downloading the certificates, user authorization, and downloading files from the cloud storage.
- Ports in the range **7770 – 7800**
The agents use these ports to communicate with the management server.
- Ports **44445** and **55556**
The agents use these ports for data transfer during backup and recovery.

If a proxy server is enabled in your network, refer to "Configuring proxy server settings" (p. 74) to understand whether you need to configure these settings on each machine that runs a protection agent.

The minimum Internet connection speed required for managing an agent from the cloud is 1 Mbit/s (not to be confused with the data transfer rate acceptable for backing up to the cloud). Consider this if you use a low-bandwidth connection technology such as ADSL.

TCP ports required for backup and replication of VMware virtual machines

- Port **443**
Agent for VMware (both Windows and Virtual Appliance) connects to this port on the ESXi host/vCenter server to perform VM management operations, such as create, update, and delete VMs on vSphere during backup, recovery, and VM replication operations.
- Port **902**
Agent for VMware (both Windows and Virtual Appliance) connects to this port on the ESXi host to establish NFC connections to read/write data on VM disks during backup, recovery, and VM replication operations.
- Port **3333**
If the Agent for VMware (Virtual Appliance) is running on the ESXi host/cluster that is the target for VM replication, VM replication traffic does not go directly to the ESXi host on port **902**. Instead, the traffic goes from the source Agent for VMware to TCP port **3333** on the Agent for VMware (Virtual Appliance) located on the target ESXi host/cluster.
The source Agent for VMware that reads data from the original VM disks can be anywhere else and can be of any type: Virtual Appliance or Windows.
The service that is responsible for accepting VM replication data on the target Agent for VMware (Virtual Appliance) is called "Replica disk server." This service is responsible for the WAN optimization techniques, such as traffic compression and deduplication during VM replication, including replica seeding (see [Seeding an initial replica](#)). When no Agent for VMware (Virtual Appliance) is running on the target ESXi host, this service is not available, and therefore the replica seeding scenario is not supported.

Ports required by the Downloader component

The Downloader component is responsible for delivering updates to a computer and distributing them to other Downloader instances. It can run in agent mode which turns its computer into Downloader agent. The Downloader agent downloads updates from the internet and serves as the source of updates distribution to other computers. The Downloader requires the following ports to operate.

- TCP and UDP (incoming) port **6888**
Used by the BitTorrent protocol for torrent peer-to-peer updates.
- UDP port **6771**
Used as the local peer discovery port. Also takes part in peer-to-peer updates.
- TCP port **18018**
Used for communication between updaters working in different modes: Updater and UpdaterAgent.
- TCP port **18019**
Local port, used for communication between the Updater and the protection agent.

Step 6

On the machine where you plan to install the protection agent, verify that the following local ports are not in use by other processes.

- 127.0.0.1:**9999**
- 127.0.0.1:**43234**
- 127.0.0.1:**9850**

Note

You do not have to open them in the firewall.

Changing the ports used by the protection agent

Some of the ports required by the protection agent might be in use by other applications in your environment. To avoid conflicts, you can change the default ports used by the protection agent by modifying the following files.

- In Linux: `/opt/Acronis/etc/aakore.yaml`
- In Windows: `\ProgramData\Acronis\Agent\etc\aakore.yaml`

Which agent do I need?

Selecting an agent depends on what you are going to back up. The table below summarizes the information, to help you decide.

In Windows, Agent for Exchange, Agent for SQL, Agent for Active Directory, and Agent for Oracle require that Agent for Windows is also installed. Thus, if you install, for example, Agent for SQL, you also will be able to back up the entire machine where the agent is installed.

We recommend that you also install Agent for Windows when you install Agent for VMware (Windows) and Agent for Hyper-V.

In Linux, Agent for Oracle, Agent for MySQL/MariaDB, and Agent for Virtuozzo require that Agent for Linux (64-bit) is also installed. These agents are bundled into the Agent for Linux (64-bit) setup file.

What are you going to back up?	Which agent to install?	Where to install it?
Physical machines		
Physical machines running Windows	Agent for Windows	On the machine that will be backed up.
Physical machines running Linux	Agent for Linux	
Physical machines running macOS	Agent for Mac	
Databases		
SQL databases	Agent for SQL	On the machine running Microsoft SQL Server.
MySQL databases	Agent for MySQL/MariaDB (Bundled into the Agent for Linux (64-bit) setup file)	On the machine running MySQL Server.
MariaDB databases	Agent for MySQL/MariaDB (Bundled into the Agent for Linux (64-bit) setup file)	On the machine running MariaDB Server.
Exchange databases	Agent for Exchange	On the machine running the Mailbox role of Microsoft Exchange Server.*
Oracle databases	Agent for Oracle (In Linux, bundled into the Agent for Linux (64-bit) setup)	On the machine running Oracle Database.

	file)	
Cloud-to-cloud workloads		
Microsoft 365 mailboxes (Cloud agent or local agent)	Cloud agent (No installation required)	This functionality is available with a cloud agent that is deployed in the data center. For more information, see "Using the cloud Agent for Microsoft 365" (p. 562).
	Agent for Office 365	On a Windows machine that is connected to the Internet. For more information, see "Using the locally installed Agent for Office 365" (p. 558).
Microsoft 365 OneDrive files and SharePoint Online sites	Cloud agent (No installation required)	This functionality is available with a cloud agent that is deployed in the data center. For more information, see "Using the cloud Agent for Microsoft 365" (p. 562).
Google Workspace Gmail mailboxes, Google Drive files, and Shared drive files	Cloud agent (No installation required)	This functionality is available with a cloud agent that is deployed in the data center. For more information, see "Protecting Google Workspace data" (p. 593).
Active Directory		
Machines running Active Directory Domain Services	Agent for Active Directory	On the domain controller.

Virtual machines		
VMware ESXi virtual machines	Agent for VMware (Windows)	On a Windows machine that has network access to vCenter Server and to the virtual machine storage.**
	Agent for VMware (Virtual Appliance)	On the ESXi host.
Hyper-V virtual machines	Agent for Hyper-V	On the Hyper-V host.
Scale Computing HC3 virtual machines	Agent for Scale Computing HC3 (Virtual Appliance)	On the Scale Computing HC3 host.
Red Hat Virtualization virtual machines (managed by oVirt)	Agent for oVirt (Virtual Appliance)	On the Red Hat Virtualization host.
Virtuozzo virtual machines and containers***	Agent for Virtuozzo (Bundled into the Agent for Linux (64-bit) setup file)	On the Virtuozzo host.
Virtuozzo Hybrid Infrastructure virtual machines	Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance)	On the Virtuozzo Hybrid Infrastructure host.
Virtual machines hosted on Amazon EC2	The same as for physical machines****	On the machine that will be backed up.
Virtual machines hosted on Windows Azure		
Citrix XenServer virtual machines		
Red Hat Virtualization (RHV/RHEV), not managed by oVirt		
Kernel-based Virtual Machines (KVM), not managed by oVirt		
Oracle virtual machines, not managed by oVirt		
Nutanix AHV virtual machines		
Red Hat Virtualization (RHV/RHEV), managed by oVirt		
Kernel-based Virtual Machines (KVM), managed by oVirt		
Oracle virtual machines, managed by oVirt		

Mobile devices		
Mobile devices running Android	Mobile app for Android	On the mobile device that will be backed up.
Mobile devices running iOS	Mobile app for iOS	

*During the installation, Agent for Exchange checks for enough free space on the machine where it will run. Free space equal to 15 percent of the biggest Exchange database is temporarily needed during a granular recovery.

**If your ESXi uses a SAN attached storage, install the agent on a machine connected to the same SAN. The agent will back up the virtual machines directly from the storage rather than via the ESXi host and LAN. For detailed instructions, see "Agent for VMware - LAN-free backup" (p. 638).

***For Virtuozzo 7, only ploop containers are supported. Virtual machines are not supported.

****A virtual machine is considered virtual if it is backed up by an external agent. If an agent is installed in the guest system, the backup and recovery operations are the same as with a physical machine. Nevertheless, if Cyber Protection can identify a virtual machine by using the CPUID instruction, a virtual machine service quota is assigned to it. If you use direct passthrough or another option that masks the CPU manufacturer ID, only service quotas for physical machines can be assigned.

Agent-based and agentless backup

Agent-based backup requires that a protection agent is installed on each protected machine. Agent-based backup is supported on all physical and virtual machines. For more information about which agent you need and where to install it, refer to "Which agent do I need?" (p. 63)

Agentless backup is supported by some virtualization platforms and it is not available for physical machines. Agentless backup requires only one protection agent, which is installed on a dedicated machine in the virtual environment. This agent backs up all other virtual machines in this environment. For more information about the supported backup types per virtualization platform, refer to "Supported virtualization platforms" (p. 35).

For some virtualization platforms, virtual appliances are available. A virtual appliance (VA) is a ready-made virtual machine that contains a protection agent. The virtual appliances are available in hypervisor-specific formats, such as .ovf, .ova, or .qcow.

Which backup type do I need?

We recommend the agent-based backup in the following cases:

- You need additional protection functionality, such as antivirus and antimalware, patch management, or remote desktop connection. For more information about these features refer to "Supported protection features by operating system" (p. 49).

- You need to separate the virtual machines on the tenant level, for example, because you want to provide the users in this tenant with access only to their own backups.
- You need file-level backups that you recover to the guest operating systems.

We recommend the agentless backup in the following cases:

- You need only backup, without any additional protection features.
- You are looking for simplified management—you can back up multiple virtual machines by installing and configuring only one agent.
- You need to minimize resource usage—one dedicated agent uses less CPU and RAM than multiple agents installed on each virtual machine in your environment.
- You use specific backup setups, such as LAN-free backup. For more about this functionality, refer to "Agent for VMware - LAN-free backup" (p. 638).
- You do not want to install and configure agents for different operating systems—the dedicated agent backs up the virtual machines on the hypervisor level, regardless of guest operating systems.

System requirements for agents

Agent	Disk space required for installation
Agent for Windows	1.2 GB
Agent for Linux	2 GB
Agent for Mac	1 GB
Agent for SQL and Agent for Windows	1.2 GB
Agent for Exchange and Agent for Windows	1.3 GB
Agent for Data Loss Prevention	500 MB
Agent for Microsoft 365	500 MB
Agent for Active Directory and Agent for Windows	2 GB
Agent for VMware and Agent for Windows	1.5 GB
Agent for Hyper-V and Agent for Windows	1.5 GB
Agent for Virtuozzo and Agent for Linux	1 GB
Agent for Virtuozzo Hybrid Infrastructure	700 MB
Agent for Oracle and Agent for Windows	2.2 GB

Agent for Oracle and Agent for Linux	2 GB
Agent for MySQL/MariaDB and Agent for Linux	2 GB

Backup operations, including deleting backups, require about 1 GB of RAM per 1 TB of backup size. The memory consumption may vary, depending on the amount and type of data being processed by the agents.

Note

The RAM usage might increase when backing up to extra large backup sets (4 TB and more).

On x64 systems, operations with bootable media and disk recovery with restart require at least 2 GB of memory.

On workloads with modern processors, such as 11th Gen Intel Core or AMD Ryzen 7, that support CET technology, some features of the Agent for Data Loss Prevention are disabled to avoid conflicts. The following table lists the availability of Device Control and Advanced DLP features on systems with such CPUs.

Features	Device Control	Advanced DLP
Local channels		
Removable storage	n/a	Yes
Encrypted removable storage	Yes	n/a
Printers	n/a	No
Redirected mapped drives	n/a	Yes
Redirected clipboard	n/a	No
Network communications		
SMTP emails	n/a	Yes
Microsoft Outlook (MAPI)	n/a	Yes
IBM Notes	n/a	No
Webmails	n/a	Yes
Instant messaging (ICQ)	n/a	No
Instant messaging (Viber)	n/a	No
Instant messaging (IRC, Jabber, Skype, Viber)	n/a	Yes
File sharing services	n/a	Yes
Social networks	n/a	Yes

Local network file sharing (SMB)	n/a	Yes
Web access (HTTP/HTTPS)	n/a	Yes
File transfers (FTP/FTPS)	n/a	Yes
Data transfer allowlisting		
Allowlist for device types	n/a	Yes
Allowlist for network communications	n/a	Yes
Allowlist for remote hosts	n/a	Yes
Allowlist for applications	n/a	Yes
Peripheral devices		
Removable storage	Yes	Yes
Encrypted removable storage	Yes	Yes
Printers	No	No
MTP-connected mobile devices	No	No
Bluetooth adapters	Yes	Yes
Optical drives	Yes	Yes
Floppy drives	Yes	Yes
Windows clipboard	No	No
Screenshot capture	No	No
Redirected mapped drives	Yes	Yes
Redirected clipboard	No	No
Cyber Protect Agent self-protection		
Protection from regular end users	Yes	Yes
Protection from local system administrators	Yes	Yes

Linux packages

To add the necessary modules to the Linux kernel, the setup program needs the following Linux packages:

- The package with kernel headers or sources. The package version must match the kernel version.
- The GNU Compiler Collection (GCC) compiler system. The GCC version must be the one with which the kernel was compiled.

- The Make tool.
- The Perl interpreter.
- The `libelf-dev`, `libelf-devel`, or `elfutils-libelf-devel` libraries for building kernels starting with 4.15 and configured with `CONFIG_UNWINDER_ORC=y`. For some distributions, such as Fedora 28, they need to be installed separately from kernel headers.

The names of these packages vary depending on your Linux distribution.

In Red Hat Enterprise Linux, CentOS, and Fedora, the packages normally will be installed by the setup program. In other distributions, you need to install the packages if they are not installed or do not have the required versions.

Are the required packages already installed?

To check whether the packages are already installed, perform these steps:

1. Run the following command to find out the kernel version and the required GCC version:

```
cat /proc/version
```

This command returns lines similar to the following: `Linux version 2.6.35.6` and `gcc version 4.5.1`

2. Run the following command to check whether the Make tool and the GCC compiler are installed:

```
make -v
gcc -v
```

For **gcc**, ensure that the version returned by the command is the same as in the `gcc version` in step 1. For **make**, just ensure that the command runs.

3. Check whether the appropriate version of the packages for building kernel modules is installed:

- In Red Hat Enterprise Linux, CentOS, and Fedora, run the following command:

```
yum list installed | grep kernel-devel
```

- In Ubuntu, run the following commands:

```
dpkg --get-selections | grep linux-headers
dpkg --get-selections | grep linux-image
```

In either case, ensure that the package versions are the same as in `Linux version` in step 1.

4. Run the following command to check whether the Perl interpreter is installed:

```
perl --version
```

If you see the information about the Perl version, the interpreter is installed.

5. In Red Hat Enterprise Linux, CentOS, and Fedora, run the following command to check whether `elfutils-libelf-devel` is installed:

```
yum list installed | grep elfutils-libelf-devel
```

If you see the information about the library version, the library is installed.

Installing the packages from the repository

The following table lists how to install the required packages in various Linux distributions.

Linux distribution	Package names	How to install
Red Hat Enterprise Linux	kernel-devel gcc make elfutils-libelf-devel	The setup program will download and install the packages automatically by using your Red Hat subscription.
	perl	Run the following command: <pre>yum install perl</pre>
CentOS Fedora	kernel-devel gcc make elfutils-libelf-devel	The setup program will download and install the packages automatically.
	perl	Run the following command: <pre>yum install perl</pre>
Ubuntu Debian	linux-headers linux-image gcc make perl	Run the following commands: <pre>sudo apt-get update sudo apt-get install linux-headers-\$(uname -r) sudo apt-get install linux-image-\$(uname -r) sudo apt-get install gcc-<package version> sudo apt-get install make sudo apt-get install perl</pre>
SUSE Linux OpenSUSE	kernel-source gcc make perl	<pre>sudo zypper install kernel-source sudo zypper install gcc sudo zypper install make sudo zypper install perl</pre>

The packages will be downloaded from the distribution's repository and installed.

For other Linux distributions, please refer to the distribution's documentation regarding the exact names of the required packages and the ways to install them.

Installing the packages manually

You may need to install the packages **manually** if:

- The machine does not have an active Red Hat subscription or Internet connection.
- The setup program cannot find the **kernel-devel** or **gcc** version corresponding to the kernel version. If the available **kernel-devel** is more recent than your kernel, you need to either update the kernel or install the matching **kernel-devel** version manually.
- You have the required packages on the local network and do not want to spend time for automatic search and downloading.

Obtain the packages from your local network or a trusted third-party website, and install them as follows:

- In Red Hat Enterprise Linux, CentOS, or Fedora, run the following command as the root user:

```
rpm -ivh PACKAGE_FILE1 PACKAGE_FILE2 PACKAGE_FILE3
```

- In Ubuntu, run the following command:

```
sudo dpkg -i PACKAGE_FILE1 PACKAGE_FILE2 PACKAGE_FILE3
```

Example: Installing the packages manually in Fedora 14

Follow these steps to install the required packages in Fedora 14 on a 32-bit machine:

1. Run the following command to determine the kernel version and the required GCC version:

```
cat /proc/version
```

The output of this command includes the following:

```
Linux version 2.6.35.6-45.fc14.i686  
gcc version 4.5.1
```

2. Obtain the **kernel-devel** and **gcc** packages that correspond to this kernel version:

```
kernel-devel-2.6.35.6-45.fc14.i686.rpm  
gcc-4.5.1-4.fc14.i686.rpm
```

3. Obtain the **make** package for Fedora 14:

```
make-3.82-3.fc14.i686
```

4. Install the packages by running the following commands as the root user:

```
rpm -ivh kernel-devel-2.6.35.6-45.fc14.i686.rpm
rpm -ivh gcc-4.5.1.fc14.i686.rpm
rpm -ivh make-3.82-3.fc14.i686
```

You can specify all these packages in a single `rpm` command. Installing any of these packages may require installing additional packages to resolve dependencies.

Configuring proxy server settings

The protection agents can transfer data through an HTTP/HTTPS proxy server. The server must work through an HTTP tunnel without scanning or interfering with the HTTP traffic. Man-in-the-middle proxies are not supported.

Because the agent registers itself in the cloud during the installation, you must configure the proxy server settings during the installation of the agent or in advance.

For Windows

If a proxy server is configured in **Control panel > Internet Options > Connections**, the setup program reads the proxy server settings from the registry and uses them automatically.

Use this procedure if you want to perform the following tasks.

- Configure the proxy settings before the installation of the agent.
- Update the proxy settings after the installation of the agent.

To configure the proxy settings during the installation of the agent, see "Installing protection agents in Windows" (p. 78).

Note

This procedure is valid only when the `http-proxy.yaml` file does not exist on the machine. If the `http-proxy.yaml` file exists on the machine, you must update the proxy settings in the file, as it overrides the settings in the `aakore.yaml` file.

The `%programdata%\Acronis\Agent\var\aakore\http-proxy.yaml` file is created when you configure the proxy server settings by using Cyber Protection Monitor. For more information, see "Configuring proxy server settings in Cyber Protect Monitor" (p. 297).

To open the `http-proxy.yaml` file, you must be member of the Administrators group in Windows.

To configure the proxy settings

1. Create a new text document and open it in a text editor, such as Notepad.
2. Copy and paste the following lines into the file.

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\Global\HttpProxy]
"Enabled"=dword:00000001
"Host"="proxy.company.com"
```

```
"Port"=dword:000001bb
>Login="proxy_login"
>Password="proxy_password"
```

3. Replace `proxy.company.com` with your proxy server host name/IP address, and `000001bb` with the hexadecimal value of the port number. For example, `000001bb` is port 443.
4. If your proxy server requires authentication, replace `proxy_login` and `proxy_password` with the proxy server credentials. Otherwise, delete these lines from the file.
5. Save the document as `proxy.reg`.
6. Run the file as an administrator.
7. Confirm that you want to edit the Windows registry.
8. If the agent is not installed on this workload yet, install it now. If the agent is already installed on the workload, continue to the next step.
9. Open the `%programdata%\Acronis\Agent\etc\aaakore.yaml` file in a text editor.
To open this file, you must be member of the Administrators group in Windows.
10. Locate the **env** section or create it, and then add the following lines.

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

11. Replace `proxy_login` and `proxy_password` with the proxy server credentials, and `proxy_address:port` with the address and port number of the proxy server.
12. In the **Start** menu, click **Run**, type: **cmd**, and then click **OK**.
13. Restart the `aaakore` service by running the following commands.

```
net stop aaakore
net start aaakore
```

14. Restart the agent by running the following commands.

```
net stop mms
net start mms
```

For macOS

Use this procedure if you want to perform the following tasks.

- Configure the proxy settings before the installation of the agent.
- Update the proxy settings after the installation of the agent.

To configure the proxy settings during the installation of the agent, see "Installing protection agents in macOS" (p. 83).

To configure the proxy settings

1. Create the /Library/Application Support/Acronis/Registry/Global.config file and open it in a text editor, such as Text Edit.
2. Copy and paste the following lines into the file.

```
<?xml version="1.0" ?>
<registry name="Global">
  <key name="HttpProxy">
    <value name="Enabled" type="Tdwor">"1"</value>
    <value name="Host" type="TString">"proxy.company.com"</value>
    <value name="Port" type="Tdwor">"443"</value>
    <value name="Login" type="TString">"proxy_login"</value>
    <value name="Password" type="TString">"proxy_password"</value>
  </key>
</registry>
```

3. Replace proxy.company.com with your proxy server host name/IP address, and 443 with the decimal value of the port number.
4. If your proxy server requires authentication, replace proxy_login and proxy_password with the proxy server credentials. Otherwise, delete these lines from the file.
5. Save the file.
6. If the agent is not installed on this workload yet, install it now. If the agent is already installed on the workload, continue to the next step.
7. Open the /Library/Application Support/Acronis/Agent/etc/aakore.yaml file in a text editor.
8. Locate the **env** section or create it and then add the following lines.

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

9. Replace proxy_login and proxy_password with the proxy server credentials, and proxy_address:port with the address and port number of the proxy server.
10. Go to **Applications > Utilities > Terminal**.
11. Restart the aakore service by running the following commands.

```
sudo launchctl stop aakore
sudo launchctl start aakore
```

12. Restart the agent by running the following commands.

```
sudo launchctl stop acronis_mms
sudo launchctl start acronis_mms
```

For Linux

Run the installation file with the --http-proxy-host=ADDRESS --http-proxy-port=PORT --http-proxy-login=LOGIN --http-proxy-password=PASSWORD parameters. Use the following procedure to update the proxy settings after the installation of the protection agent.

To configure the proxy settings

1. Open the `/etc/Acronis/Global.config` file in a text editor.
2. Do one of the following:
 - If the proxy settings were specified during the agent installation, locate the following section.

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdwor" >"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdwor" >"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```

- If the proxy settings were not specified during the agent installation, copy the following lines and paste them into the file between the `<registry name="Global">...</registry>` tags.

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdwor" >"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdwor" >"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```

3. Replace `ADDRESS` with the new proxy server host name/IP address, and `PORT` with the decimal value of the port number.
4. If your proxy server requires authentication, replace `LOGIN` and `PASSWORD` with the proxy server credentials. Otherwise, delete these lines from the file.
5. Save the file.
6. Open file `/opt/acronis/etc/aakore.yaml` in a text editor.
7. Locate the **env** section or create it and add the following lines:

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

8. Replace `proxy_login` and `proxy_password` with the proxy server credentials, and `proxy_address:port` with the address and port number of the proxy server.
 9. Restart the `aakore` service by running the following command.
10. Restart the agent by executing the running command in any directory.

```
sudo service aakore restart
```

```
sudo service acronis_mms restart
```

For bootable media

When working under bootable media, you might need to access the cloud storage via a proxy server. To configure the proxy server settings, click **Tools > Proxy server**, and then configure the proxy server host name/IP address, port, and credentials.

Installing protection agents

You can install agents on machines running any of the operating systems listed in "[Supported operating systems and environments](#)". The operating systems that support the Cyber Protect features are listed in "[Supported Cyber Protect features by operating system](#)".

Downloading protection agents

Before you install an agent, you must download its installation file from the Cyber Protect console.

To download an agent while adding a workload to protect

1. In the Cyber Protect console, navigate to **Devices > All devices**.
2. In the upper right, click **Add device**.
3. In the **Add devices** panel, from the **Release channel** drop-down menu, select an agent version.
 - **Previous release** - download the agent version from the previous release.
 - **Current** - download the latest available agent version.
4. Select the agent that corresponds to the operating system of the workload that you are adding. The **Save As** dialog opens.
5. [Only for Macs with Apple silicon (such as Apple M1) processors] Click **Cancel**. In the **Add Mac** panel that opens, click the **Download ARM installer** link.
6. Select a location to save the agent installation file and click **Save**.

To download an agent for later use

1. In the upper right corner of the Cyber Protect console, click the **User** icon.
2. Click **Downloads**.
3. In the **Downloads** dialog, from the **Release channel** drop-down menu, select an agent version.
 - **Previous release** - download the agent version from the previous release.
 - **Current** - download the latest available agent version.
4. Scroll the list of available installers to locate the agent installer that you need and click the download icon at the end of its row. The **Save As** dialog opens.
5. Select a location to save the agent installation file and click **Save**.

Installing protection agents in Windows

Prerequisites

Download the agent that you need on the workload that you plan to protect. See "Downloading protection agents" (p. 78).

To install Agent for Windows

1. Ensure that the machine is connected to the Internet.
2. Log on as an administrator and start the installer.
3. [Optional] Click **Customize installation settings** and make the appropriate changes if you want:
 - To change the components to install (for example, to disable the installation of Cyber Protection Monitor or the Command-Line Tool, or to install the Agent for Antimalware protection and URL filtering).

Note

On Windows machines, the antimalware protection and URL filtering features require the installation of Agent for Antimalware protection and URL filtering. It will be installed automatically for protected workloads if the **Antivirus & Antimalware protection** or the **URL filtering** module is enabled in their protection plans.

- To change the method of registering the workload in the Cyber Protection service. You can switch from **Use service console** (default) to **Use credentials** or **Use registration token**.
 - To change the installation path.
 - To change the user account under which the agent service will run. For details, refer to "Changing the logon account on Windows machines" (p. 85).
 - To verify or change the proxy server host name/IP address, port, and credentials. If a proxy server is enabled in Windows, it is detected and used automatically.
4. Click **Install**.
 5. [Only when installing Agent for VMware] Specify the address and access credentials for the vCenter Server or the stand-alone ESXi host on which you want to back up and recover virtual machines, and then click **Done**.

We recommend that you use a dedicated account for accessing vCenter Server or the ESXi host, instead of using an existing account with the Administrator role. To learn more about the required privileges for the dedicated account, refer to "Agent for VMware – necessary privileges" (p. 647).

6. [Only when installing on a domain controller] Specify the user account under which the agent service will run, and then click **Done**. For security reasons, the setup program does not automatically create new accounts on a domain controller.

Note

The user account that you specify must be granted the Log on as a service right. This account must have already been used on the domain controller, in order for its profile folder to be created on that machine.

For more information about installing the agent on a read-only domain controller, see [this knowledge base article](#).

7. If you kept the default registration method **Use service console** in step 3, wait until the registration screen appears, and then proceed to the next step. Otherwise, no more actions are required.
8. Do one of the following:

- If you log in under a company administrator account, register workloads for your company:
 - a. Click **Register workload**.
 - b. In the opened browser window, sign in to the Cyber Protect console and review the registration details.
 - c. In the **Register for account** list, select the user account under which you want to register the workload.
 - d. Click **Check code**, and then click **Confirm registration**.
- If you log in under a partner administrator account, register workloads for your customers:
 - a. Click **Register workload**.
 - b. In the opened browser window, sign in to the Cyber Protect console and review the registration details.
 - c. In the **Register for account** list, select the user account of your customer under which you want to register the workload.
 - d. Click **Check code**, and then click **Confirm registration**.
- Click **Show registration info**. The setup program shows the registration link and the registration code. If you cannot complete the workload registration on the current machine, copy the registration link and code, and then follow the registration steps on a different machine. In this case, you will need to enter the registration code in the registration form. The registration code is valid for one hour.

Alternatively, you can access the registration form by clicking **All devices > Add**, scrolling down to **Registration via code**, and then clicking **Register**.

Note

Do not quit the setup program until you confirm the registration. To initiate the registration again, you will have to restart the setup program and repeat the installation procedure.

As a result, the workload will be assigned to the account that was used to log in to the Cyber Protect console.

- Register the workload manually by using the command line. For more information on how to do this, refer to "Registering and unregistering workloads manually" (p. 119).
9. [If the agent is registered under an account whose tenant is in the Enhanced security mode] Set the encryption password.

Installing protection agents in Linux

Preparation

- Download the agent that you need on the machine that you plan to protect. See "Downloading protection agents" (p. 78).
- Ensure that the necessary [Linux packages](#) are installed on the machine.
- When installing the agent in SUSE Linux, ensure that you use `su -` instead of `sudo`. Otherwise, the following error occurs when you try to register the agent via the Cyber Protect console: Failed to launch the web browser. No display available.

Some Linux distributions, such as SUSE, do not pass the DISPLAY variable when using sudo, and the installer cannot open the browser in the graphical user interface (GUI).

Installation

To install Agent for Linux, you need at least 2 GB of free disk space.

To install Agent for Linux

1. Ensure that the machine is connected to the Internet.
2. As the root user, navigate to directory with the installation file, make the file executable, and then run it.

```
chmod +x <installation file name>
```

```
./<installation file name>
```

If a proxy server is enabled in your network, when running the installation file, specify the server host name/IP address and port in the following format: `--http-proxy-host=ADDRESS --http-proxy-port=PORT --http-proxy-login=LOGIN --http-proxy-password=PASSWORD`.

If you want to change the default method of registering the machine in the Cyber Protection service, run the installation file with one of the following parameters:

- `--register-with-credentials` – to ask for a user name and password during the installation
- `--token=STRING` – to use a registration token
- `--skip-registration` – to skip the registration

3. Select the check boxes for the agents that you want to install. The following agents are available:
 - Agent for Linux
 - Agent for Virtuozzo
 - Agent for Oracle
 - Agent for MySQL/MariaDB

Agent for Virtuozzo, Agent for Oracle, and Agent for MySQL/MariaDB require that Agent for Linux (64-bit) is also installed.

4. If you kept the default registration method in step 2, proceed to the next step. Otherwise, enter the user name and password for the Cyber Protection service, or wait until the machine will be registered by using the token.
5. Do one of the following:
 - If you log in under a company administrator account, register workloads for your company:
 - a. Click **Register workload**.
 - b. In the opened browser window, sign in to the Cyber Protect console and review the registration details.
 - c. In the **Register for account** list, select the user account under which you want to register the workload.
 - d. Click **Check code**, and then click **Confirm registration**.

- If you log in under a partner administrator account, register workloads for your customers:
 - a. Click **Register workload**.
 - b. In the opened browser window, sign in to the Cyber Protect console and review the registration details.
 - c. In the **Register for account** list, select the user account of your customer under which you want to register the workload.
 - d. Click **Check code**, and then click **Confirm registration**.
- Click **Show registration info**. The setup program shows the registration link and the registration code. If you cannot complete the workload registration on the current machine, copy the registration link and code, and then follow the registration steps on a different machine. In this case, you will need to enter the registration code in the registration form. The registration code is valid for one hour.

Alternatively, you can access the registration form by clicking **All devices > Add**, scrolling down to **Registration via code**, and then clicking **Register**.

Note

Do not quit the setup program until you confirm the registration. To initiate the registration again, you will have to restart the setup program and repeat the installation procedure.

As a result, the workload will be assigned to the account that was used to log in to the Cyber Protect console.

- Register the workload manually by using the command line. For more information on how to do this, refer to "Registering and unregistering workloads manually" (p. 119).
6. [If the agent is registered under an account whose tenant is in the Enhanced security mode] Set the encryption password.
 7. If the UEFI Secure Boot is enabled on the machine, you are informed that you need to restart the system after the installation. Be sure to remember what password (the one of the root user or "acronis") should be used.

Note

The installation generates a new key that is used for signing the kernel modules. You must enroll this new key to the Machine Owner Key (MOK) list by restarting the machine. Without enrolling the new key, your agent will not be operational. If you enable the UEFI Secure Boot after the agent is installed, you need to reinstall the agent.

8. After the installation completes, do one of the following:
 - Click **Restart**, if you were prompted to restart the system in the previous step.

During the system restart, opt for MOK (Machine Owner Key) management, choose **Enroll MOK**, and then enroll the key by using the password recommended in the previous step.
 - Otherwise, click **Exit**.

Troubleshooting information is provided in the file:
`/usr/lib/Acronis/BackupAndRecovery/HOWTO.INSTALL`

Installing protection agents in macOS

Prerequisites

Download the agent that you need on the workload that you plan to protect. See "Downloading protection agents" (p. 78).

To install Agent for Mac (x64 or ARM64)

1. Ensure that the machine is connected to the Internet.
 2. Double-click the installation file (.dmg).
 3. Wait while the operating system mounts the installation disk image.
 4. Double-click **Install**.
 5. If a proxy server is enabled in your network, click **Protection Agent** in the menu bar, click **Proxy server settings**, and then specify the proxy server host name/IP address, port, and credentials.
 6. If prompted, provide administrator credentials.
 7. Click **Continue**.
 8. Wait until the registration screen appears.
 9. Do one of the following:
 - If you log in under a company administrator account, register workloads for your company:
 - a. Click **Register workload**.
 - b. In the opened browser window, sign in to the Cyber Protect console and review the registration details.
 - c. In the **Register for account** list, select the user account under which you want to register the workload.
 - d. Click **Check code**, and then click **Confirm registration**.
 - If you log in under a partner administrator account, register workloads for your customers:
 - a. Click **Register workload**.
 - b. In the opened browser window, sign in to the Cyber Protect console and review the registration details.
 - c. In the **Register for account** list, select the user account of your customer under which you want to register the workload.
 - d. Click **Check code**, and then click **Confirm registration**.
 - Click **Show registration info**. The setup program shows the registration link and the registration code. If you cannot complete the workload registration on the current machine, copy the registration link and code, and then follow the registration steps on a different machine. In this case, you will need to enter the registration code in the registration form. The registration code is valid for one hour.
- Alternatively, you can access the registration form by clicking **All devices** > **Add**, scrolling down to **Registration via code**, and then clicking **Register**.

Note

Do not quit the setup program until you confirm the registration. To initiate the registration again, you will have to restart the setup program and repeat the installation procedure.

As a result, the workload will be assigned to the account that was used to log in to the Cyber Protect console.

- Register the workload manually by using the command line. For more information on how to do this, refer to "Registering and unregistering workloads manually" (p. 119).
10. [If the agent is registered under an account whose tenant is in the Enhanced security mode] Set the encryption password.
 11. If your macOS version is Mojave 10.14.x or later, grant full disk access to the protection agent to enable backup operations.
For instructions, see [Grant the 'Full Disk Access' permission to the Cyber Protection agent \(64657\)](#).
 12. To use the remote desktop functionality, grant the required system permissions to the Connect Agent. For more information, see "Granting the required system permissions to the Connect Agent" (p. 84).

Granting the required system permissions to the Connect Agent

To enable all features from the remote desktop functionality on macOS workloads, in addition to the full disk access permission, you must grant the following permissions to the Connect Agent:

- Screen Recording - enables screen recording of the macOS workload via NEAR. Until this permission is granted, all remote control connections will be denied.
- Accessibility - enables remote connections in control mode via NEAR
- Microphone - enables sound redirection from the remote macOS workload to the local workload via NEAR. To enable the sound redirection feature, a sound capture driver must be installed on the workload. For more information, see "Remote sound redirection" (p. 924).
- Automation - enables the empty Recycle bin action

After you start the agent on the macOS workload, it will check if the agent has these rights and will ask you to grant the permissions, if needed.

To grant the Screen Recording permission

1. In the **Grant required system permissions** for Cyber Protect Agent dialog, click **Set up system permissions**.
2. In the **System permissions** dialog, click **Request Screen Recording** permission.
3. Click **Open System Preferences**.
4. Select **Connect Agent**.

If the agent does not have the permission when you try to access the workload remotely, it will show the Screen Recording permission request dialog. Only the local user may answer the dialog.

To grant the Accessibility permission

1. In the **Grant required system permissions** for Cyber Protect Agent dialog, click **Set up system permissions**.
2. In the **System permissions** dialog, click **Request Accessibility permission**.
3. Click **Open System Preferences**.
4. Click the lock icon in the bottom-left corner of the window so that it changes to an unlocked one. The system will ask you for an administrator password to make changes.
5. Select **Connect Agent**.

To grant the Microphone permission

1. In the **Grant required system permissions** for the Connect Agent dialog, click **Set up system permissions**.
2. In the **System permissions** dialog, click **Request Microphone permission**.
3. Click **OK**.

Note

You must also install a sound capture driver on the macOS workload to let the agent utilize the given permission and redirect the sound of the workload. For more information, see "Remote sound redirection" (p. 924).

To grant the Automation permission

1. In the **Grant required system permissions** for the Connect Agent dialog, click **Set up system permissions**.
2. In the **System permissions** dialog, click **Request Automation permission**.

Changing the logon account on Windows machines

On the **Select components** screen, define the account under which the services will run by specifying **Logon account for the agent service**. You can select one of the following:

- **Use Service User Accounts** (default for the agent service)
Service User Accounts are Windows system accounts that are used to run services. The advantage of this setting is that the domain security policies do not affect these accounts' user rights. By default, the agent runs under the **Local System** account.
- **Create a new account**
The account name will be Agent User for the agent.
- **Use the following account**
If you install the agent on a domain controller, the system prompts you to specify existing accounts (or the same account) for the agent. For security reasons, the system does not automatically create new accounts on a domain controller.
The user account that you specify when the setup program runs on a domain controller must be granted the Log on as a service right. This account must have already been used on the domain controller, in order for its profile folder to be created on that machine.

For more information about installing the agent on a read-only domain controller, see [this knowledge base article](#).

If you chose the **Create a new account** or **Use the following account** option, ensure that the domain security policies do not affect the related accounts' rights. If an account is deprived of the user rights assigned during the installation, the component may work incorrectly or not work.

Privileges required for the logon account

A protection agent is run as a Managed Machine Service (MMS) on a Windows machine. The account under which the agent will run must have specific rights for the agent to work correctly. Thus, the MMS user should be assigned the following privileges:

1. Included in the **Backup Operators** and **Administrators** groups. On a Domain Controller, the user must be included in the group **Domain Admins**.
2. Granted the **Full Control** permission on the folder %PROGRAMDATA%\Acronis (in Windows XP and Server 2003, %ALLUSERSPROFILE%\Application Data\Acronis) and on its subfolders.
3. Granted the **Full Control** permission on certain registry keys in the following key: HKEY_LOCAL_MACHINE\SOFTWARE\Acronis.
4. Assigned the following user rights:
 - Log on as a service
 - Adjust memory quotas for a process
 - Replace a process level token
 - Modify firmware environment values

How to assign the user rights

Follow the instructions below to assign the user rights (this example uses the **Log on as service** user right, the steps are the same for other user rights):

1. Log on to the computer by using an account with administrative privileges.
2. Open **Administrative Tools** from **Control Panel** (or click Win+R, type **control admintools**, and press Enter) and open **Local Security Policy**.
3. Expand **Local Policies** and click on **User Rights Assignment**.
4. In the right pane, right-click **Log on as a service** and select **Properties**.
5. Click on the **Add User or Group...** button to add a new user.
6. In the **Select Users, Computers, Service Accounts, or Groups** window, find the user you wish to enter and click **OK**.
7. Click **OK** in the **Log on as a service Properties** to save the changes.

Important

Ensure that the user which you have added to the **Log on as service** user right is not listed in the **Deny log on as a service** policy in **Local Security Policy**.

Note that we recommend that you do not change logon accounts manually after the installation is completed.

Dynamic installation and uninstallation of components

For Windows workloads protected by agent version 15.0.26986 (released in May 2021) or later, the following components are installed dynamically—that is, only when required by a protection plan:

- Agent for Antimalware protection and URL filtering—required for the operation of the antimalware protection and URL filtering features.
- Agent for Data Loss Prevention—required for the operation of the device control features.

By default, these components are not installed. The respective component is automatically installed if a workload becomes protected by a plan in which any of the following modules is enabled:

- Antivirus & Antimalware protection
- URL filtering
- Device control

Similarly, if no protection plan requires antimalware protection, URL filtering, or device control features anymore, the respective component is automatically uninstalled.

Dynamic installation or uninstallation of components takes up to 10 minutes after you change the protection plan. However, if any of the following operations are running, dynamic installation or uninstallation will start after this operation finishes:

- Backup
- Recovery
- Backup replication
- Virtual machine replication
- Testing a replica
- Running a virtual machine from backup (including finalization)
- Disaster recovery failover
- Disaster recovery failback
- Running a script (for Cyber Scripting functionality)
- Patch installation
- ESXi configuration backup

Unattended installation or uninstallation

Unattended installation or uninstallation in Windows

In Windows, you can perform unattended installation or uninstallation in the following ways:

- By using the EXE file of the setup program and specifying the installation parameters on the command line.

- By using an MSI file that you extract from the setup program, and specifying the installation parameters in one of the following ways:
 - In an MST file
 - Directly on the command line

Unattended installation and uninstallation with an EXE file

For this type of unattended installation, download the setup program, and then start it from the command line with the required installation parameters. To see the parameters that you can use, see "Parameters for unattended installation (EXE)" (p. 90).

You do not need to extract installation packages, MSI, and MST files in advance.

Installing and uninstalling agents and components (EXE)

To perform unattended installation with an EXE file, run the setup program and specify the installation parameters on the command line.

To download the setup program, in the Cyber Protect console, click the account icon in the top-right corner, and then click **Downloads**. The download link is also available in the **Add devices** pane.

To install agents and components

1. Start the command-line interface as administrator, and then navigate to the EXE file of the setup program.
2. To start the setup program and specify the installation parameters, run the following command:

```
<file path>/<EXE file><PARAMETER 1>=<value 1> ... <PARAMETER N>=<value n>
```

Use spaces to separate the parameters, and commas without spaces to separate the values for a parameter. For example:

```
C:\Users\Administrator\Downloads\AgentForWindows_web.exe --add-
components=agentForWindows,agentForSql,commandLine --install-dir="C:\Program
Files\BackupClient" --reg-address=https://eu2-cloud.company.com --reg-token=34F6-
8C39-4A5C --quiet
```

To check the available parameters and their values, see "Parameters for unattended installation (EXE)" (p. 90).

Examples

- Installing Agent for Windows, Agent for Antimalware and URL filtering, Command-Line Tool, and Cyber Protect Monitor. Registering the workload in the Cyber Protection service by using a user name and password.

```
C:\Users\Administrator\Downloads\AgentForWindows_web.exe --add-
components=agentForWindows,agentForAmp,commandLine,trayMonitor --install-
```

```
dir="C:\Program Files\BackupClient" --agent-account=system --reg-  
address=https://cloud.company.com --reg-login=johndoe --reg-password=johnspassword
```

- Installing Agent for Windows, Command-Line Tool, and Cyber Protect Monitor. Creating a new logon account for the agent service in Windows. Registering the workload in the Cyber Protection service by using a token.

```
C:\Users\Administrator\Downloads\AgentForWindows_web.exe --add-  
components=agentForWindows,commandLine,trayMonitor --install-dir="C:\Program  
Files\BackupClient" --agent-account=new --reg-address=https://eu2-cloud.company.com -  
-reg-token=34F6-8C39-4A5C
```

- Installing Agent for Windows, Command-Line Tool, Agent for Oracle and Cyber Protect Monitor. Registering the machine in the Cyber Protection service by using a user name and password.

```
C:\Users\Administrator\Downloads\AgentForWindows_web.exe --add-  
components=agentForWindows,commandLine,agentForOracle,trayMonitor --install-  
dir="C:\Program Files\BackupClient" --language=en --agent-account=system --reg-  
address=https://cloud.company.com --reg-login=johndoe --reg-password=johnspassword
```

- Installing Agent for Windows, Command-Line Tool, and Cyber Protect Monitor. Setting the user interface language to German. Registering the machine in the Cyber Protection service by using a token. Setting an HTTP proxy.

```
C:\Users\Administrator\Downloads\AgentForWindows_web.exe --add-  
components=agentForWindows,commandLine,agentForOracle,trayMonitor --install-  
dir="C:\Program Files\BackupClient"--language=de --agent-account=system --reg-  
address=https://eu2-cloud.company.com --reg-token=34F6-8C39-4A5C --http-proxy-  
address=https://my-proxy.company.com:80 --http-proxy-login=tomsmith --http-proxy-  
password=tomspassword
```

To remove an installed component

1. Start the command-line interface as administrator, and then navigate to %ProgramFiles%\BackupClient\RemoteInstall.
2. Run the following command:

```
web_installer.exe --remove-components=<value 1>,<value 2> --quiet
```

To check the available parameters and their values, see "Parameters for unattended installation (EXE)" (p. 90).

Example

- Uninstalling the Cyber Protect Monitor.

```
C:\Program Files\BackupClient\RemoteInstall\web_installer.exe --remove-  
components=trayMonitor --quiet
```

To uninstall an agent

1. Start the command-line interface as administrator, and then navigate to %Program Files%\Common Files\Acronis\BackupAndRecovery.
2. Run the following command:

```
Uninstaller.exe --quiet --delete-all-settings
```

To check the available parameters and their values, see "Parameters for unattended installation (EXE)" (p. 90).

Examples

- Uninstalling Agent for Windows and all its components. Deleting all logs, tasks, and configuration settings.

```
C:\Program Files\Common Files\Acronis\BackupAndRecovery\Uninstaller.exe --quiet --delete-all-settings
```

- Uninstalling a password-protected Agent for Windows and all its components. Deleting all logs, tasks, and configuration settings.

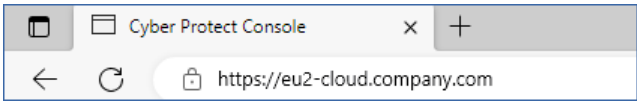
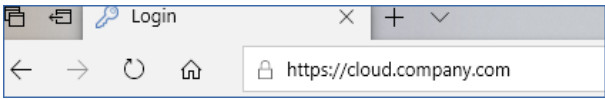
```
C:\Program Files\Common Files\Acronis\BackupAndRecovery\Uninstaller.exe --anti-tamper-password=<password> --quiet --delete-all-settings
```

Parameters for unattended installation (EXE)

The following table summarizes the parameters for unattended installation with an EXE file.

Parameters	Description
General parameters	
--add-components=<component1,component2,...,componentN>	<p>The components to be installed. See the full list of available components in "Components for unattended installation (EXE)" (p. 95).</p> <p>When you specify multiple components, separate them with commas. Do not add spaces before or after the comma.</p> <p>If you specify components that are already installed, these components will be repaired or updated, depending on version of the setup program and the version of the installed components.</p> <p>If you do not specify this parameter, a default set of components will be installed, depending on the machine on which you perform the installation. For example, Agent for SQL is only installed on machines that run MS SQL Server.</p>

Parameters	Description
--install-dir=<path>	<p>The folder in which the selected components will be installed. If the specified folder does not exist, it will be created.</p> <p>If you do not specify this parameter, a default folder is used: C:\Program Files\BackupClient.</p>
--log-dir=<path>	<p>The folder in which the installation logs will be saved.</p> <p>If you do not specify this parameter, a default folder is used: %ProgramData%\Acronis\InstallationLogs.</p>
--language=<code>	<p>The product language.</p> <p>The following values are available: en, bn, bg, cs, da, de, es, fr, ko, id, it, hi, hu, ms, nl, ja, nb, pl, pt, pt_BR, ru, fi, sr, sv, th, tr, vi, zh, zh_TW.</p> <p>If you do not specify this parameter, and the system language of the machine on which you perform the installation is listed above, the system language is used. In all other cases, the value is set to en.</p>
--quiet	<p>Use this parameter to run the setup program without showing the graphical user interface.</p> <p>Do not use it together with the --register-only parameter.</p>
--help	<p>Use this parameter to see a list of all available parameters that you can use on the command line and their descriptions.</p>
--fss-onboarding-auto-start	<p>Use this parameter together with the --quiet parameter to show the File Sync & Share on-boarding wizard after an unattended installation.</p>
Registration parameters	
--registration={skip by-credentials by-token device-flow}	<p>Use this parameter to choose how to register the agent after the installation.</p> <p>To skip the registration, specify skip. You can register the agent later, by using the --register-only parameter.</p> <p>To register the agent by using credentials, specify by-credentials, and then use the --reg-login and --reg-password parameters. Also, you can use only --reg-login and --reg-password parameters, which makes specifying --registration=by-credentials optional.</p> <p>To register the agent with a registration token, specify</p>

Parameters	Description
	<p>by-token, and then use the --reg-token parameter. Also, you can use only the --reg-token parameter, which makes specifying --registration=by-token optional.</p> <p>To register the agent by using the OAuth 2.0 protocol, specify device-flow. After the installation completes, the registration page opens automatically.</p> <p>When you use --registration=device-flow, specify the exact datacenter address as a value for the --reg-address parameter. This is the URL that you see after you log in to the Cyber Protection service. For example, https://eu2-cloud.company.com.</p>  <p>Do not use --registration=device-flow with the --quiet parameter.</p>
<p>--reg-address=<url></p>	<p>The URL of the Cyber Protection service. You can use this parameter either with the --reg-login and --reg-password parameters, or with the --reg-token parameter.</p> <ul style="list-style-type: none"> When you use it with --reg-login and --reg-password parameters, specify the address that you use to log in to the Cyber Protection service. For example, https://cloud.company.com:  <ul style="list-style-type: none"> When you use it with the --reg-token parameter, specify the exact datacenter address. This is the URL that you see after you log in to the Cyber Protection service. For example, https://eu2-cloud.company.com.  <p>Do not use https://cloud.company.com with the --reg-token parameter.</p>
<p>--reg-login=<login></p> <p>--reg-password=<password></p>	<p>The credentials for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.</p> <p>When you use these parameters, specifying the --registration parameter is optional.</p> <p>Do not use these parameters with the --reg-token parameter.</p>

Parameters	Description
--reg-token=<token>	<p>The registration token.</p> <p>The registration token is a series of 12 characters, separated into three segments by hyphens. For more information about how to generate one, see "Generating a registration token" (p. 166).</p> <p>When you use this parameter, specifying the --registration parameter is optional.</p> <p>Do not use this parameter with the --reg-login and --reg-password parameters.</p>
--register-only	<p>Use this parameter to skip the installation and register the agent by using the OAuth 2.0 protocol (device-flow).</p> <p>After the installation completes, the registration page opens automatically.</p> <p>Do not use --register-only with the --quiet parameter.</p>
Logon account for the agent service	
--agent-account={system new custom} or --agent-account-login=<login> --agent-account-password=<password>	<p>Use this parameter to specify the logon account under which agent service will run. For more information about the logon accounts, see "Changing the logon account on Windows machines" (p. 85).</p> <p>To use the Local System account, specify --agent-account=system or do not use the --agent-account parameter in your command.</p> <p>To make the agent service run under a new logon account, Acronis Agent User, which is created automatically, specify new.</p> <p>To make the agent service run under an existing account, specify the account credentials by using the --agent-account-login and --agent-account-password parameters. In this case, specifying the --agent-account=custom parameter is optional.</p>
vCenter/ESXi parameters	
--esxi-address=<host>	<p>The host name or IP address of vCenter Server or the ESXi host.</p> <p>Use this parameter when you install Agent for VMware.</p>
--esxi-login=<login> --esxi-password=<password>	<p>The access credentials to vCenter Server or the ESXi host.</p> <p>Use these parameters when you install Agent for</p>

Parameters	Description
	VMware.
Proxy parameters	
--http-proxy={none system custom}	<p>Use this parameter to specify the HTTP proxy server that you want to use for backup to and recovery from the cloud storage.</p> <p>If disable the proxy server connections, specify --http-proxy=none.</p> <p>To use a system-wide proxy server, specify --http-proxy=system or do not use the --http-proxy parameter in your command.</p> <p>To use another proxy server, specify the proxy server address and credentials by using the --http-proxy-address, --http-proxy-login, and --http-proxy-password parameters. In this case, specifying --http-proxy=custom parameter is optional.</p>
--http-proxy-address=<host>:<port>	The hostname or IP address, and the port of the custom HTTP proxy server.
--http-proxy-login=<login>	Login for the custom HTTP proxy server.
--http-proxy-password=<password>	Password for the custom HTTP proxy server.
Uninstallation parameters	
--remove-components=<component1,component2,...,componentN>	<p>The components to be uninstalled. See the full list of available components in "Components for unattended installation (EXE)" (p. 95).</p> <p>When you specify multiple components, separate them with commas. Do not add spaces before or after the comma.</p> <hr/> <p>Important By using this parameter, you can uninstall only components. To uninstall the product completely, go to Windows Control Panel > Programs and Features, select the product, and then click Uninstall.</p> <hr/>
--delete-all-settings	Use this optional parameter when you use the --remove-components parameter to delete all product logs, tasks, and configuration settings.
--anti-tamper-password=<password>	The password required for uninstalling a password-protected Agent for Windows or modifying its components.

Components for unattended installation (EXE)

The table below summarizes the components that you can use for unattended installation via an EXE file. Use the value names to specify values for the `--add-components` parameter.

For more information, see "Parameters for unattended installation (EXE)" (p. 90) "Parameters for unattended installation (MSI)" (p. 99)

Value name	Component description
agentForWindows	Agent for Windows
agentForSas	Agent for Files Sync & Share
agentForAd	Agent for Active Directory
agentForAmp	Agent for Antimalware protection and URL filtering
agentForDlp	Agent for Data Loss Prevention
agentForEsx	Agent for VMware (Windows)
agentForExchange	Agent for Exchange
agentForHyperV	Agent for Hyper-V
agentForOffice365	Agent for Office 365
agentForOracle	Agent for Oracle
agentForSql	Agent for SQL
commandLine	Command-Line Tool
mediaBuilder	Bootable Media Builder
trayMonitor	Cyber Protect Monitor
all	This value combines all components.
allAgents	This value combines all agents.

Unattended installation and uninstallation with an MSI file

For this type of unattended installation, use the Windows Installer (the `Msiexec` program). Extract the installation packages and the MSI file in advance, by using the graphical user interface of the setup program.

When you install components with an MSI file, you can use an MST transform file to customize the installation parameters. For more information on how to use the combination of MSI and MST files, see "Installing agents and components (MSI and MST combination)" (p. 96). You can use this installation method in an Active Directory domain to install protection agents by using Windows Group Policy. For more information, see "Deploying agents through Group Policy" (p. 165).

Alternatively, you can specify the installation parameters manually on the command line. In this case, you do not need an MST file. For more information, see "Examples" (p. 97).

Extracting the MSI, MST, and CAB files

Extract the MSI, MST, and CAB files with the installation packages by running the graphical user interface of the setup program.

To extract the MSI, MST, and CAB files

1. Run the graphical user interface of the setup program, and then click **Create .mst and .msi files for unattended installation**.
2. In **What to install**, select the components that you want to install, and then click **Done**.
The installation packages for these components will be extracted from the setup program as CAB files.
3. In **Registration settings**, select **Use credentials** or **Use registration token**. Depending on your choice, specify the credentials or the registration token, and then click **Done**.
For more information on how to generate a registration token, see "Generating a registration token" (p. 166).
4. [Only when installing on a domain controller] In **Logon account for the agent service**, select **Use the following account**. Specify the user account under which the agent service will run, and then click **Done**. For security reasons, the setup program does not automatically create new accounts on a domain controller.

Note

The user account that you specify must be granted the Log on as a service right. This account must have already been used on the domain controller, in order for its profile folder to be created on that machine.

For more information about installing the agent on a read-only domain controller, see [this knowledge base article](#).

5. Review or modify other installation settings that will be added to the MST file, and then click **Proceed**.
6. Select the folder in which the MSI, MST, and CAB files will be extracted, and then click **Generate**.

Installing agents and components (MSI and MST combination)

Use the MST file to customize the installation setting for the MSI file. Use the MSI and MST combination when you install agents on multiple machines through a Windows Group Policy. For more information, see "Deploying agents through Group Policy" (p. 165).

To install components with MSI and MST files

1. Extract the MSI and MST files as described in "Extracting the MSI, MST, and CAB files" (p. 96).
2. On the command-line interface of the machine on which you want to install components, run the following command:

```
msiexec /i <MSI file> TRANSFORMS=<MST file>
```

For example:

```
msiexec /i BackupClient64.msi TRANSFORMS=BackupClient64.msi.mst
```

Installing and uninstalling agents and components (MSI and direct selection)

Run the MSI file, manually select the components to install, and specify their installation parameters on the command line. In this case, you do not need the MST file.

To install agents and components

1. Extract the MSI file and the installation packages (CAB files) as described in "Extracting the MSI, MST, and CAB files" (p. 96).

For this installation method, you only need the MSI and CAB files. You do not need the MST file.

2. In the command-line interface of the machine, run the following command:

```
msiexec /i <MSI file><PARAMETER 1>=<value 1> ... <PARAMETER N>=<value n>
```

Use spaces to separate the parameters, and commas without spaces to separate the values for a parameter. For example:

```
msiexec.exe /i BackupClient64.msi  
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,TrayMonitor  
TARGETDIR="C:\Program Files\BackupClient" REGISTRATION_ADDRESS=https://eu2-  
cloud.company.com REGISTRATION_TOKEN=34F6-8C39-4A5C
```

To check the available parameters and their values, see "Parameters for unattended installation (MSI)" (p. 99).

Examples

- Installing Agent for Windows, Agent for Antimalware and URL filtering, Command-Line Tool, and Cyber Protect Monitor. Registering the workload in the Cyber Protection service by using a user name and password.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn  
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,AmpAgentFeature,CommandLineTool,Tray  
Monitor TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress MMS_USE_  
SYSTEM_ACCOUNT=1 REGISTRATION_ADDRESS=https://cloud.company.com REGISTRATION_  
LOGIN=johndoe REGISTRATION_PASSWORD=johnspassword
```

- Installing Agent for Windows, Command-Line Tool, and Cyber Protect Monitor. Creating a new logon account for the agent service in Windows. Registering the workload in the Cyber Protection service by using a token.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,TrayMonitor
TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress MMS_CREATE_NEW_
ACCOUNT=1 REGISTRATION_ADDRESS=https://eu2-cloud.company.com REGISTRATION_TOKEN=34F6-
8C39-4A5C
```

- Installing Agent for Windows, Command-Line Tool, Agent for Oracle and Cyber Protect Monitor. Registering the machine in the Cyber Protection service by using a user name and encoded in base64 password. You might need to encode your password if it contains special characters or blank spaces. For more information about how to encode a password, see "Passwords with special characters or blank spaces" (p. 123).

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,OracleAgentFeature,T
rayMonitor TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress CURRENT_
LANGUAGE=en MMS_USE_SYSTEM_ACCOUNT=1 REGISTRATION_ADDRESS=https://cloud.company.com
REGISTRATION_LOGIN=johndoe REGISTRATION_PASSWORD_ENCODED=am9obnNwYXNzd29yZA==
```

- Installing Agent for Windows, Command-Line Tool, and Cyber Protect Monitor. Registering the machine in the Cyber Protection service by using a token. Setting an HTTP proxy.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,TrayMonitor
TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress CURRENT_LANGUAGE=en
MMS_USE_SYSTEM_ACCOUNT=1 REGISTRATION_ADDRESS=https://eu2-cloud.company.com
REGISTRATION_TOKEN=34F6-8C39-4A5C HTTP_PROXY_ADDRESS=https://my-proxy.company.com
HTTP_PROXY_PORT=80 HTTP_PROXY_LOGIN=tomsmith HTTP_PROXY_PASSWORD=tomspassword
```

To remove an installed component

1. Extract the MSI file and the installation packages (CAB files) as described in "Extracting the MSI, MST, and CAB files" (p. 96).
For this installation method, you only need the MSI and CAB files. You do not need the MST file.
2. In the command-line interface of the machine, run the following command:

```
msiexec /i <MSI file><REMOVE>=<value 1>,<value 2> REBOOT=ReallySuppress /qn
```

To check the available parameters and their values, see "Parameters for unattended installation (MSI)" (p. 99).

Example

- Removing Cyber Protect monitor.

```
msiexec.exe /i BackupClient64.msi /l*v uninstall_log.txt REMOVE=TrayMonitor
REBOOT=ReallySuppress /qn
```

To uninstall an agent

1. Extract the MSI file and the installation packages (CAB files) as described in "Extracting the MSI, MST, and CAB files" (p. 96).
For this installation method, you only need the MSI and CAB files. You do not need the MST file.
2. In the command-line interface of the machine, run the following command:

```
msiexec /x <MSI file> /l*v uninstall_log.txt DELETE_ALL_SETTINGS=1
REBOOT=ReallySuppress /qn
```

To check the available parameters and their values, see "Parameters for unattended installation (MSI)" (p. 99).

Examples

- Uninstalling Agent for Windows and all its components. Deleting all logs, tasks, and configuration settings.

```
msiexec.exe /x BackupClient64.msi /l*v uninstall_log.txt DELETE_ALL_SETTINGS=1
REBOOT=ReallySuppress /qn
```

- Uninstalling a password-protected Agent for Windows and all its components. Deleting all logs, tasks, and configuration settings.

```
msiexec.exe /x BackupClient64.msi /l*v uninstall_log.txt ANTI_TAMPER_
PASSWORD=<password> DELETE_ALL_SETTINGS=1 REBOOT=ReallySuppress /qn
```


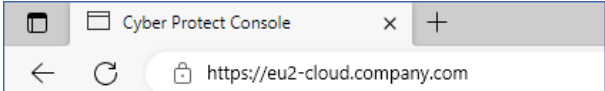
Parameters for unattended installation (MSI)

The following table summarizes the parameters for unattended installation when you use an MSI file.

You can also use additional `msiexec` parameters. For example, use `/qn` to prevent any GUI elements from showing. To learn more about the `msiexec` parameters, see the [Microsoft documentation](#).

Parameters	Description
General parameters	
ADDLOCAL= <component1,component2,...,componentN>	The components to be installed. See the full list of available components in "Components for unattended installation (MSI)" (p. 103). When you specify multiple components, separate them with commas. Do not add spaces before or after the comma.

Parameters	Description
	<p>Note</p> <p>You must extract the installation files for all components that you want to install. For more information about how to extract them, see "Extracting the MSI, MST, and CAB files" (p. 96).</p>
TARGETDIR=<path>	<p>The folder in which the selected components will be installed. If the specified folder does not exist, it will be created.</p> <p>If you do not specify this parameter, a default folder is used: C:\Program Files\BackupClient.</p>
REBOOT=ReallySuppress	Specify this parameter if you want to install components without restarting the machine.
/1*v <log file>	Specify this parameter to save a verbose log. This log is needed if you have to investigate installation issues.
CURRENT_LANGUAGE=<language ID>	<p>The product language.</p> <p>The following values are available: en, bn, bg, cs, da, de, es, fr, ko, id, it, hi, hu, ms, nl, ja, nb, pl, pt, pt_BR, ru, fi, sr, sv, th, tr, vi, zh, zh_TW.</p> <p>If you do not specify this parameter, and the system language of the machine on which you perform the installation is listed above, the system language is used. In all other cases, the value is set to en.</p>
SKIP_SHA2_KB_CHECK={0,1}	<p>Use this parameter to choose whether to check if the SHA2 code signing support update from Microsoft (KB4474419) is installed on the machine. The check only runs on operating systems that require this update. To see if it is required for your operating system, see "Supported operating systems and environments" (p. 27).</p> <p>Use this parameter with value set to 1 to skip the check.</p> <p>If you do not specify the parameter or set its value to 0, and the SHA2 code signing support update is not found on the machine, the installation fails.</p>
FSS_ONBOARDING_AUTO_START={0,1}	<p>Use this parameter with value set to 1 to show the File Sync & Share on-boarding wizard after an unattended installation.</p> <p>If you do not specify this parameter or set its value to 0, the on-boarding wizard will not be shown.</p>

Parameters	Description
Registration parameters	
REGISTRATION_ADDRESS	<p>The URL of the Cyber Protection service. You can use this parameter either with the REGISTRATION_LOGIN and REGISTRATION_PASSWORD parameters, or with REGISTRATION_TOKEN.</p> <ul style="list-style-type: none"> When you use it with REGISTRATION_LOGIN and REGISTRATION_PASSWORD parameters, specify the address that you use to log in to the Cyber Protection service. For example, https://cloud.company.com:  When you use it with the REGISTRATION_TOKEN parameter, specify the exact datacenter address. This is the URL that you see after you log in to the Cyber Protection service. For example, https://eu2-cloud.company.com.  <p>Do not use https://cloud.company.com with the REGISTRATION_TOKEN parameter.</p>
REGISTRATION_LOGIN REGISTRATION_PASSWORD	<p>The credentials for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.</p> <p>Do not use these parameters with the REGISTRATION_TOKEN parameter.</p>
REGISTRATION_PASSWORD_ENCODED	<p>The password for the account under which the agent will be registered in the Cyber Protection service, encoded in base64. For more information on how to encode your password, see "Passwords with special characters or blank spaces" (p. 123).</p>
REGISTRATION_TOKEN	<p>The registration token.</p> <p>The registration token is a series of 12 characters, separated into three segments by hyphens. For more information about how to generate one, see "Generating a registration token" (p. 166).</p> <p>Do not use this parameter with the REGISTRATION_LOGIN and REGISTRATION_PASSWORD parameters.</p>
REGISTRATION_REQUIRED={0,1}	<p>Use this parameter to choose what happens if the</p>

Parameters	Description
	<p>registration fails.</p> <p>If you set the value to 1, the installation also fails. If you set the value to 0 or do not specify the parameter, the installation completes successfully even though the registration fails.</p>
Logon account for the agent service	
MMS_USE_SYSTEM_ACCOUNT={0,1}	<p>Use this parameter with value 1, to make the service run under the Local System logon account.</p> <p>For more information about the logon accounts, see "Changing the logon account on Windows machines" (p. 85).</p>
MMS_CREATE_NEW_ACCOUNT={0,1}	Use this parameter with value 1, to make the agent service run under a new logon account, Acronis Agent User , which is created automatically.
MMS_SERVICE_USERNAME=<user name> MMS_SERVICE_PASSWORD=<password>	Use these parameters to specify an existing logon account under which the agent service will run.
vCenter/ESXi parameters	
SET_ESX_SERVER={0,1}	<p>Use this parameter when you install Agent for VMware.</p> <p>If you set the value to 0, Agent for VMware will not be connected to vCenter Server or an ESXi host.</p> <p>If you set the value to 1, specify the following parameters: ESX_HOST, EXI_USER, ESX_PASSWORD.</p>
ESX_HOST=<host name>	The host name or IP address of vCenter Server or the ESXi host.
ESX_USER=<user name> ESX_PASSWORD=<password>	The access credentials to vCenter Server or the ESXi host.
Proxy parameters	
HTTP_PROXY_ADDRESS=<IP address> HTTP_PROXY_PORT=<port>	<p>Use these parameters to specify the HTTP proxy server that the agent will use.</p> <p>If you do not use a proxy server, do not specify these parameters.</p>
HTTP_PROXY_LOGIN=<login> HTTP_PROXY_PASSWORD=<password>	<p>The credentials for the HTTP proxy server.</p> <p>Use these parameters if the proxy server requires authentication.</p>

Parameters	Description
Uninstallation parameters	
REMOVE={<list of components> ALL}	The components to be uninstalled. When you specify multiple components, separate them with commas. Do not add spaces before or after the comma. To remove all product components, set the value to ALL.
DELETE_ALL_SETTINGS={0, 1}	To delete all product logs, tasks, and configuration settings, set the value to 1. Use this optional parameter when you use the REMOVE parameter.
ANTI_TAMPER_PASSWORD=<password>	The password required for uninstalling a password-protected Agent for Windows or modifying its components.

Components for unattended installation (MSI)

The table below summarizes the components that you can use for unattended installation via an MSI file. Use the value names to specify values for the ADDLOCAL parameter. For more information, see "Parameters for unattended installation (MSI)" (p. 99).

Value name	Component description	Must be installed together with	Bitness
AgentFeature	Core components for agents		32-bit/64-bit
MmsMspComponents	Core components for backup	AgentFeature	32-bit/64-bit
BackupAndRecoveryAgent	Agent for Windows	MmsMspComponents	32-bit/64-bit
AmpAgentFeature	Agent for Antimalware protection and URL filtering	BackupAndRecoveryAgent	32-bit/64-bit
DlpAgentFeature	Agent for Data Loss Prevention	BackupAndRecoveryAgent	32-bit/64-bit
SasAgentFeature	Agent for File	TrayMonitor	32-bit/64-bit

	Sync & Share		bit
ArxAgentFeature	Agent for Exchange	MmsMspComponents	32-bit/64-bit
ArsAgentFeature	Agent for SQL	BackupAndRecoveryAgent	32-bit/64-bit
ARADAgentFeature	Agent for Active Directory	BackupAndRecoveryAgent	32-bit/64-bit
ArxOnlineAgentFeature	Agent for Microsoft 365	MmsMspComponents	32-bit/64-bit
OracleAgentFeature	Agent for Oracle	BackupAndRecoveryAgent	32-bit/64-bit
AcronisESXSupport	Agent for VMware ESX(i) (Windows)	BackupAndRecoveryAgent	64-bit
HyperVAgent	Agent for Hyper-V	BackupAndRecoveryAgent	32-bit/64-bit
CommandLineTool	Command-Line Tool		32-bit/64-bit
TrayMonitor	Cyber Protect Monitor	AgentFeature	32-bit/64-bit
BackupAndRecoveryBootableComponents	Bootable Media Builder		32-bit/64-bit

Unattended installation or uninstallation in Linux

This section describes how to install or uninstall protection agents in the unattended mode on a machine running Linux, by using the command line.

To install an agent

1. Open Terminal.
2. Do one of the following:
 - To start the installation by specifying the parameters on the command line, run the following command:

```
<package name> -a <parameter 1> ... <parameter N>
```

Here, <package name> is the name of the installation package (an .i686 or an .x86_64 file). All available parameters and their values are described in "Unattended installation or uninstallation parameters" (p. 106).

- To start the installation with parameters that are specified in a separate text file, run the following command:

```
<package name> -a --options-file=<path to the file>
```

This approach might be useful if you do not want to enter sensitive information on the command line. In this case, you can specify the configuration settings in a separate text file and ensure that only you can access it. Put each parameter on a new line, followed by the value for that parameter, for example:

```
--rain=https://cloud.company.com  
--login=johndoe  
--password=johnpassword  
--auto
```

or

```
-C  
https://cloud.company.com  
-g  
johndoe  
-w  
johnpassword  
-a  
--language  
en
```

If the same parameter is specified both on the command line and in the text file, the command line value precedes.

3. If UEFI Secure Boot is enabled on the machine, you are informed that you need to restart the system after the installation. Ensure that you remember what password (that of the root user or "acronis") should be used. During the system restart, opt for MOK (Machine Owner Key) management, choose **Enroll MOK**, and then enroll the key by using the recommended password.

If you enable UEFI Secure Boot after the agent installation, repeat the installation, including step 3. Otherwise, backups will fail.

To uninstall an agent

1. Open Terminal.
2. Do one of the following:
 - To uninstall the agent and remove all logs, tasks, and configuration settings, run the following command:

```
/usr/lib/Acronis/BackupAndRecovery/uninstall/uninstall -a
```

- To uninstall the agent but keep its ID (for example, if you plan to install the agent later), run the following command:

```
/usr/lib/Acronis/BackupAndRecovery/uninstall/uninstall -a --no-purge
```

- To uninstall the agent by using the installation file, run the following command:

```
<package name> -a -u
```

Here, <package name> is the name of the installation package (an .i686 or an .x86_64 file). All available parameters and their values are described in "Unattended installation or uninstallation parameters" (p. 106).

Note

Use this command only when the installation package is the same version as the installed agent and if /usr/lib/Acronis/BackupAndRecovery/uninstall/uninstall is corrupted or inaccessible.

Unattended installation or uninstallation parameters

This section describes parameters that are used during unattended installation or uninstallation in Linux.

The minimal configuration for unattended installation includes -a and registration parameters (for example, --login and --password parameters; --rain and --token parameters). You can use more parameters to customize you installation.

Installation parameters

Basic parameters

```
{-i |--id=}<list of components>
```

The components to be installed, separated by commas and without space characters. The following components are available in the .x86_64 installation package:

Component	Component description
BackupAndRecoveryAgent	Agent for Linux
AgentForPCS	Agent for Virtuozzo
OracleAgentFeature	Agent for Oracle
MySQLAgentFeature	Agent for MySQL/MariaDB

Without this parameter, all of the above components will be installed.

Agent for Virtuozzo, Agent for Oracle, and Agent for MySQL/MariaDB require that Agent for Linux is also installed.

The .i686 installation package contains only BackupAndRecoveryAgent.

`{-a|--auto}`

The installation and registration process will complete without any further user interaction. When using this parameter, you must specify the account under which the agent will be registered in the Cyber Protection service, either by using the `--token` parameter, or by using the `--login` and `--password` parameters.

`{-t|--strict}`

If the parameter is specified, any warning that occurs during the installation results in installation failure. Without this parameter, the installation completes successfully even in the case of warnings.

`{-n|--nodeps}`

The absence of required Linux packages will be ignored during the installation.

`{-d|--debug}`

Writes the installation log in the verbose mode.

`--options-file=<location>`

The installation parameters will be read from a text file instead of the command line.

`--language=<language ID>`

The product language. Available values are as follows: en, bg, cs, da, de, es, fr, hu, id, it, ja, ko, ms, nb, nl, pl, pt, pt_BR, ru, fi, sr, sv, tr, zh, zh_TW.

If this parameter is not specified, the product language will be defined by your system language on the condition that it is in the list above. Otherwise, the product language will set to English (en).

Registration parameters

Specify one of the following parameters:

- `{-g|--login=<user name>}` and `{-w|--password=<password>}`

Credentials for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.

- `--token=<token>`


The registration token is a series of 12 characters, separated by hyphens in three segments. You can generate one in the Cyber Protect console, as described in "[Deploying agents through Group Policy](#)".

You cannot use the `--token` parameter along with `--login`, `--password`, and `--register-with-credentials` parameters.


- `{-C|--rain=<service address>}`

The URL of the Cyber Protection service.

You don't need to include this parameter explicitly when you use `--login` and `--password` parameters for registration, because the installer uses the correct address by default – this would be the address that you use **to log in** to the Cyber Protection service. For example:



However, when you use `{-C|--rain=}` with the `--token` parameter, you must specify the exact datacenter address. This is the URL that you see **once you are logged in** to the Cyber Protection service. For example:



- `--register-with-credentials`

If this parameter is specified, the installer's graphical interface will start. To finish the registration, enter the user name and password for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.

- `--skip-registration`

Use this parameter if you need to install the agent but you plan to register it in the Cyber Protection service later. For more information on how to do this, refer to "[Registering machines manually](#)".

Additional parameters

`--http-proxy-host=<IP address>` and `--http-proxy-port=<port>`

The HTTP proxy server that the agent will use for backup and recovery from the cloud, and for connection to the management server. Without these parameters, no proxy server will be used.

`--http-proxy-login=<login>` and `--http-proxy-password=<password>`

The credentials for the HTTP proxy server. Use these parameters if the server requires authentication.

`--tmp-dir=<location>`

Specifies the folder where the temporary files are stored during the installation. The default folder is **`/var/tmp`**.

`{-s|--disable-native-shared}`

Redistributable libraries will be used during the installation, even though they might have already been present on your system.

`--skip-prereq-check`

There will be no check of whether the packages required for compiling the snapapi module are already installed.

`--force-weak-snapapi`

The installer will not compile a snapapi module. Instead, it will use a ready-made module that might not match the Linux kernel exactly. We do not recommend that you use this option.

`--skip-svc-start`

The services will not start automatically after the installation. Most often, this parameter is used with the `--skip-registration` one.

Information parameters

`{-?|--help}`

Shows the description of parameters.

`--usage`

Shows a brief description of the command usage.

`{-v|--version}`

Shows the installation package version.

`--product-info`

Shows the product name and the installation package version.

`--snapapi-list`

Shows the available ready-made snapapi modules.

`--components-list`

Shows the installer components.

Parameters for legacy features

These parameters relate to a legacy component, agent.exe.

`{-e|--ssl=}<path>`

Specifies the path to a custom certificate file for SSL communication.

`{-p|--port=}<port>`

Specifies the port on which agent.exe listens for connections. The default port is 9876.

Uninstallation parameters

`{-u|--uninstall}`

Uninstalls the product.

`--purge`

Uninstalls the product and removes its logs, tasks, and configuration settings. You don't need to specify the `--uninstall` parameter explicitly when you use the `--purge` one.

Examples

- Installing Agent for Linux without registering it.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -i BackupAndRecoveryAgent -a --skip-registration
```

- Installing Agent for Linux, Agent for Virtuozzo, and Agent for Oracle, and registering them by using credentials.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -a --login=johndoe --password=johnpassword
```

- Installing Agent for Oracle and Agent for Linux, and registering them by using a registration token.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -i BackupAndRecoveryAgent,OracleAgentFeature -a --rain=https://eu2-cloud.company.com --token=34F6-8C39-4A5C
```

- Installing Agent for Linux, Agent for Virtuozzo, and Agent for Oracle with configuration settings in a separate text file.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -a --options-file=/home/mydirectory/configuration_file
```

- Uninstalling Agent for Linux, Agent for Virtuozzo, and Agent for Oracle, and removing all their logs, tasks, and configuration settings.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -a --purge
```

Unattended installation and uninstallation in macOS

This section describes how to install, register, and uninstall the protection agent in the unattended mode on a machine running macOS, by using the command line.

Required permissions

Before you initiate an unattended installation on a Mac workload, you must modify the Privacy Preferences Policy Control to allow App access and kernel and system extensions in the macOS of the workload to enable the installation of the Cyber Protection agent. See "Required permissions for unattended installation in macOS" (p. 112).

After you deploy the PPC payload, you can proceed with the procedures below.

To download the installation file (.dmg)

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click **Add**, and then click **Mac**.

To install an agent

1. Open Terminal.
2. Create a temporary directory where you will mount the installation file (.dmg).

```
mkdir <dmg_root>
```

Here, <dmg_root> is a name of your choice.

3. Mount the .dmg file.

```
hdiutil attach <dmg_file> -mountpoint <dmg_root>
```

Here, <dmg_file> is the name of the installation file. For example, **Cyber_Protection_Agent_for_MAC_x64.dmg**.

4. Run the installer.
 - If you use a full installer for Mac, like CyberProtect_AgentForMac_x64.dmg or CyberProtect_AgentForMac_arm64.dmg, run the following command.

```
sudo installer -pkg <dmg_root>/Install.pkg -target LocalSystem
```

Note

If you need to enable auto-onboarding for File Sync & Share, run the following command instead. This option will request the administrator password.

```
open <dmg_root>/Install.app --args --unattended --fss-onboarding-auto-start
```

- If you use an universal installer for Mac, like CyberProtect_AgentForMac_web.dmg, run the following command.

```
sudo <dmg_root>/Install.app/Contents/MacOS/cyber_installer -a
```

5. Detach the installation file (.dmg).

```
hdiutil detach <dmg_root>
```

Example

```
mkdir mydirectory
```

```
hdiutil attach /Users/JohnDoe/Cyber_Protection_Agent_for_MAC_x64.dmg -mountpoint  
mydirectory
```

```
sudo installer -pkg mydirectory/Install.pkg -target LocalSystem
```

```
hdiutil detach mydirectory
```

To uninstall an agent

1. Open Terminal.
2. Do one of the following:
 - To uninstall the agent, run the following command:

```
sudo /Library/Application\ Support/BackupClient/Acronis/Cyber\ Protect\ Agent\ Uninstall.app/Contents/MacOS/AgentUninstall /confirm
```

- To uninstall the agent and remove all logs, tasks and configuration settings, run the following command:

```
sudo /Library/Application\ Support/BackupClient/Acronis/Cyber\ Protect\ Agent\ Uninstall.app/Contents/MacOS/AgentUninstall /confirm /purge
```

Required permissions for unattended installation in macOS

Before you initiate an unattended installation on a Mac workload, you must modify the Privacy Preferences Policy Control to allow App access and kernel and system extensions in the macOS of the workload to enable the installation of the Cyber Protection agent. You can do this by deploying a custom PPPC payload or by configuring the preferences in the graphical user interface of the workload. The following permissions are required.

Requirements for macOS 11 (Big Sur) or later

Tab	Section	Field	Value
-----	---------	-------	-------

Privacy Preferences Policy Control	App Access	Identifier	com.acronis.backup
--	------------	------------	--------------------

		Identifier Type	Bundle ID
--	--	-----------------	-----------

		Code Requirement	identifier "com.acronis.backup" and anchor apple generic and certificate 1[field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf [field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow
	App Access	Identifier	com.acronis.backup.aakore
		Identifier Type	Bundle ID
		Code Requirement	identifier "com.acronis.backup.aakore" and anchor apple generic and certificate 1 [field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf[field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow
	App Access	Identified	com.acronis.backup.activeprotection
		Identifier Type	Bundle ID
		Code Requirement	identifier "com.acronis.backup.activeprotection" and anchor apple generic and certificate 1 [field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf[field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
ACCESS		Allow	

	App Access	Identifier	cyber-protect-service
		Identifier Type	Bundle ID
		Code Requirement	identifier "cyber-protect-service" and anchor apple generic and certificate 1 [field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf[field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow
System Extensions		Allow users to approve system extensions	Enabled
	Allowed Team IDs and System Extensions	Display Name	Acronis Cyber Protection Agent System Extensions
		System Extension Types	Allowed Team Identifiers
		Team Identifier	ZU2TV78AA6

Requirements for macOS versions prior to version 11

Tab	Section	Field	Value
-----	---------	-------	-------

Privacy Preferences Policy Control	App Access	Identifier	com.acronis.backup
--	------------	------------	--------------------

		Identifier Type	Bundle ID
		Code Requirement	identifier "com.acronis.backup" and anchor apple generic and certificate 1[field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf [field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow
	App Access	Identifier	com.acronis.backup.aakore
		Identifier Type	Bundle ID
		Code Requirement	identifier "com.acronis.backup.aakore" and anchor apple generic and certificate 1 [field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf[field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow
	App Access	Identified	com.acronis.backup.activeprotection
		Identifier Type	Bundle ID
		Code Requirement	identifier "com.acronis.backup.activeprotection" and anchor apple generic and certificate 1 [field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf[field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow
	App Access	Identifier	cyber-protect-service
		Identifier Type	Bundle ID
		Code Requirement	identifier "cyber-protect-service" and anchor apple generic and certificate 1[field.1.2.840.113635.100.6.2.6] /* exists */ and certificate leaf [field.1.2.840.113635.100.6.1.13] /* exists */ and certificate leaf[subject.OU] = ZU2TV78AA6
		APP OR SERVICE	SystemPolicyAllFiles
		ACCESS	Allow

Approved Kernel Extensions		Allow users to approve kernel extensions	Enabled
		Allow standard users to approve legacy kernel extensions (macOS 11 or later)	Enabled
	Approved Team IDs and Kernel Extensions	Approved Team ID - Display Name	Acronis Cyber Protection Agent Kernel Extensions
		Team ID	ZU2TV78AA6
		Kernel Extension Bundle IDs	<ul style="list-style-type: none"> com.acronis.systeminterceptors com.acronis.ngscan com.acronis.notifyframework
System Extensions		Allow users to approve system extensions	Enabled
	Allowed Team IDs and System Extensions	Display Name	Acronis Cyber Protection Agent System Extensions
		System Extension Types	Allowed Team Identifiers
		Team Identifier	ZU2TV78AA6

Registering and unregistering workloads manually

Workloads are automatically registered in the Cyber Protection service when you install the protection agent on them. When you uninstall the protection agent, the workloads are automatically unregistered and disappear from the Cyber Protect console.

You can also register a workload manually, by using the command line interface. You might need to use the manual registration, for example, if the automatic registration fails or if you want to move a workload to a new tenant or under a new user account.

To register a workload by using a user name and password

In Windows

At the command line, run the following command:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a <service address> -u <user name> -p <password>
```

For example:

```
"C:\ProgramFiles\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a https://cloud.company.com -u johndoe -p johnspassword
```

In Linux

At the command line, run the following command:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a <service  
address> -u <user name> -p <password>
```

For example:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a  
https://cloud.company.com -u johndoe -p johnspassword
```

In macOS

At the command line, run the following command:

```
sudo "/Library/Application Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent"  
-o register -t cloud -a <service address> -u <user name> -p <password>
```


For example:

```
sudo "/Library/Application Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent"  
-o register -t cloud -a https://cloud.company.com -u johndoe -p johnspassword
```

Note

Use the user name and password for the account under which you want to register the workload. This cannot be a partner administrator account.

The service address is the URL that you use **to log in** to the Cyber Protection service. For example, <https://cloud.company.com>.



Important

If your password contains special characters or blank spaces, refer to "Passwords with special characters or blank spaces" (p. 123).

Important

If you use macOS 10.14 or later, grant full disk access to the protection agent. To do so, go to **Applications >Utilities**, and then run **Cyber Protect Agent Assistant**. Then, follow the instructions in the application window.

To register a workload by using a registration token

In Windows

At the command line, run the following command:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a <service address> --token <registration token>
```

For example:

```
"C:\ProgramFiles\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a https://au1-cloud.company.com --token 3B4C-E967-4FBD
```

In Linux

At the command line, run the following command:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a <service  
address> --token <registration token>
```

For example:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a  
https://eu2-cloud.company.com --token 34F6-8C39-4A5C
```

In macOS

At the command line, run the following command:

```
sudo "/Library/Application Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent"  
-o register -t cloud -a <service address> --token <registration token>
```

For example:

```
sudo "/Library/Application Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent"  
-o register -t cloud -a https://us5-cloud.company.com --token 9DBF-3DA9-4DAB
```

Important

If you use macOS 10.14 or later, grant full disk access to the protection agent. To do so, go to **Applications >Utilities**, and then run **Cyber Protect Agent Assistant**. Then, follow the instructions in the application window.

Virtual appliance

1. In the console of the virtual appliance, press CTRL+SHIFT+F2 to open the command-line interface.
2. At the command prompt, run the following command:

```
register_agent -o register -t cloud -a <service address> --token <registration token>
```


For example:

```
register_agent -o register -t cloud -a https://eu2-cloud.company.com --token 34F6-8C39-4A5C
```

3. To return to the graphical interface of the appliance, press ALT+F1.

Note

When you use a registration token, you must specify the exact data center address. This is the URL that you see **after you log in** to the Cyber Protection service. For example, `https://eu2-cloud.company.com`.



Do not use `https://cloud.company.com` here.

The registration token is a series of 12 characters, separated by hyphens in three segments. For more information on how to generate one, refer to "Generating a registration token" (p. 166).

To unregister a workload

In Windows

At the command line, run the following command:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o unregister
```

For example:

```
"C:\ProgramFiles\BackupClient\RegisterAgentTool\register_agent.exe" -o unregister
```

In Linux

At the command line, run the following command:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o unregister
```

In macOS

At the command line, run the following command:

```
sudo "/Library/Application Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent"  
-o unregister
```

Virtual appliance

1. In the console of the virtual appliance, press CTRL+SHIFT+F2 to open the command-line interface.
2. At the command prompt, run the following command:

```
register_agent -o unregister
```

3. To return to the graphical interface of the appliance, press ALT+F1.

Moving a workload to another tenant

Moving a workload to another tenant is not natively supported. As a workaround, you can unregister the workload, and then register it in another tenant. All applied protection plans will be revoked from that workload, and it will lose access to its backups in the cloud storage of the original tenant.

For more information about how to register a workload in a new tenant or under a new user account, see "Changing the registration of a workload" (p. 124).

Passwords with special characters or blank spaces

If your password contains special characters or blank spaces, enclose it in quotation marks when you type it on the command line.

For example, in Windows, run this command:

Command template:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a <service address> -u <user name> -p "<password">
```

Command example:

```
"C:\ProgramFiles\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a https://cloud.company.com -u johndoe -p "johns password"
```

If this command fails, encode your password into base64 format at <https://www.base64encode.org/>. Then, at the command line, specify the encoded password by using the `-b` or `--base64` parameter.

For example, in Windows, run this command:

Command template:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud  
-a <service address> -u <user name> -b -p <encoded password>
```

Command example:

```
"C:\ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud -a https://cloud.company.com -u johndoe -b -p am9obnNwYXNzd29yZA==
```

Changing the registration of a workload

You can change the current registration of a workload by registering it in a new tenant or under a new user account.

Important

When you change the registration of a workload, all protection plans that are applied to it will be revoked. To continue protecting the workload, apply a new protection plan to it.

If you register the workload in a new tenant, the workload will lose access to the backups in the cloud storage of the original tenant. The backups in non-cloud storages will remain accessible.

You can change the registration of a workload by using the command line or by using the GUI installer. When you use the command line, you do not need to uninstall the agent.

To change the registration of a workload

By using the command line

1. Unregister the protection agent, as described in "To unregister a workload" (p. 122).
2. Register the protection agent in the new tenant or under the new user account, as described in "To register a workload by using a user name and password" (p. 119) or in "To register a workload by using a registration token" (p. 121).

By using the GUI installer

1. Uninstall the protection agent.
2. Install the protection agent, and then register it in the new tenant or under the new user account.

For more information about how to install and register an agent, refer to "Installing protection agents" (p. 78).

Autodiscovery of machines

Using autodiscovery, you can:

- Automate the installation of protection agents and the registration of machines by detecting the machines in your Active Directory domain or local network.
- Install and update protection agents on multiple machines.
- Use synchronization with Active Directory, in order to reduce the efforts for provisioning resources and managing machines in a large Active Directory domain.

Prerequisites

To perform autodiscovery, you need at least one machine with an installed protection agent in your local network or Active directory domain. This agent is used as a discovery agent.

Important

Only agents that are installed on Windows machines can be discovery agents. If there are no discovery agents in your environment, you will not be able to use the **Multiple devices** option in the **Add devices** panel.

Remote installation of agents is supported only for machines running Windows (Windows XP is not supported). For remote installation on a machine running Windows Server 2012 R2, you must have [Windows update KB2999226](#) installed on this machine.

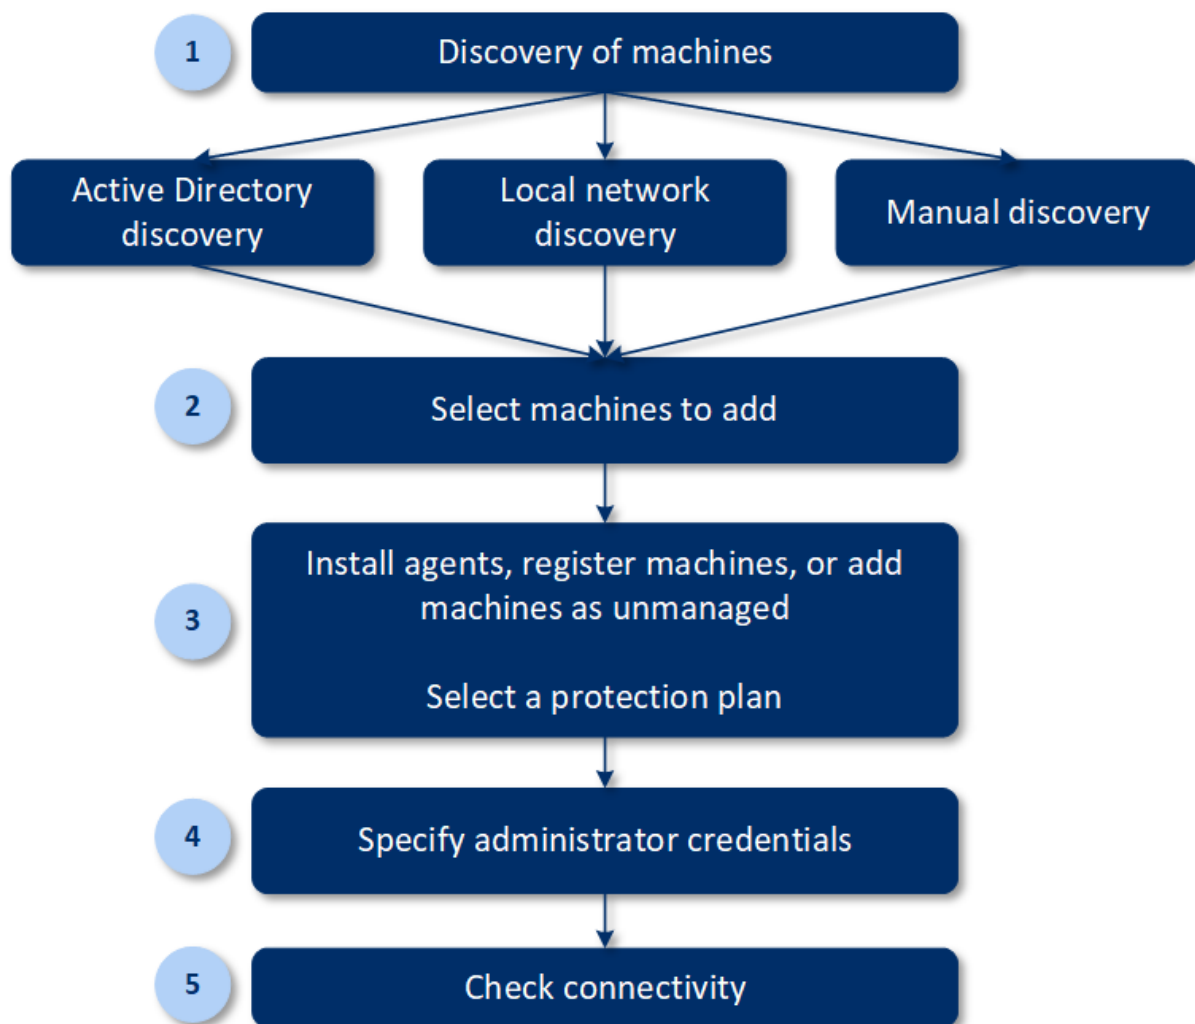
How autodiscovery works

During a local network discovery, the discovery agent collects the following information for each machine in the network, by using NetBIOS discovery, Web Service Discovery (WSD), and the Address Resolution Protocol (ARP) table:

- Name (short/NetBIOS hostname)
- Fully qualified domain name (FQDN)
- Domain/workgroup
- IPv4/IPv6 addresses
- MAC addresses
- Operating system (name/version/family)
- Machine category (workstation/server/domain controller)

During an Active Directory discovery, the discovery agent, in addition to the list above, collects information about the Organizational Unit (OU) of the machines and detailed information about their names and operating systems. However, the IP and MAC addresses are not collected.

The following diagram summarizes the autodiscovery process.



1. Select the discovery method:

- Active Directory discovery
- Local network discovery
- Manual discovery – By using a machine IP address or host name, or by importing a list of machines from a file

The results of an Active directory discovery or a local network discovery exclude machines with installed protection agents.

During a manual discovery, the existing protection agents are updated and re-registered. If you perform autodiscovery by using the same account under which an agent is registered, the agent will only be updated to the latest version. If you perform autodiscovery by using another account, the agent will be updated to the latest version and re-registered under the tenant to which the account belongs.

2. Select the machines that you want to add to your tenant.

3. Select how to add these machines:

- Install a protection agent and additional components on the machines, and register them in the Cyber Protect console.

- Register the machines in the Cyber Protect console (if a protection agent was already installed).
- Add the machines to the Cyber Protect console as **Unmanaged machines**, without installing a protection agent.

You can also apply an existing protection plan to the machines on which you install a protection agent or which you register in the Cyber Protect console.

4. Provide administrator credentials for the selected machines.
5. Verify that you can connect to the machines by using the provided credentials.

The machines that are shown in the Cyber Protect console, fall into the following categories:

- **Discovered** – Machines that are discovered, but a protection agent is not installed on them.
- **Managed** – Machines on which a protection agent is installed.
- **Unprotected** – Machines to which a protection plan is not applied. Unprotected machines include both discovered machines and managed machines with no protection plan applied.
- **Protected** – Machines to which a protection plan is applied.

How remote installation of agents works

1. The discovery agent connects to the target machines by using the host name, IP address, and administrator credentials specified in the discovery wizard, and then uploads the `web_installer.exe` file to these machines.
2. The `web_installer.exe` file runs on the target machines in the unattended mode.
3. The web installer retrieves additional installation packages from the cloud, and then installs them to the target machines via the `msiexec` command.
4. After the installation completes, the components are registered in the cloud.

Note

Remote installation of agents is not supported for Domain Controllers due to the additional permissions required for the agent service to run.

Autodiscovery and manual discovery

Before starting the discovery, ensure that the [prerequisites](#) are met.

Note

Autodiscovery is not supported for adding Domain Controllers due to additional permissions required for the agent service to run.

To discover machines

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click **Add**.
3. In **Multiple devices**, click **Windows-only**. The discovery wizard opens.

4. [If there are units in your organization] Select a unit. Then, in **Discovery agent** you will be able to select the agents associated with the selected unit and its child units.
5. Select the discovery agent that will perform the scan to detect machines.
6. Select the discovery method:
 - **Search Active Directory.** Ensure that the machine with the discovery agent is the Active Directory domain member.
 - **Scan local network.** If the selected discovery agent could not find any machines, select another discovery agent.
 - **Specify manually or import from file.** Manually define the machines to be added or import them from a text file.
7. [If the Active Directory discovery method is selected] Select how to search for machines:
 - **In organizational unit list.** Select the group of machines to be added.
 - **By LDAP dialect query.** Use the [LDAP dialect](#) query to select the machines. **Search base** defines where to search, while **Filter** allows you to specify the criteria for machine selection.
8. [If the Active Directory or local network discovery method is selected] Use a list to select the machines that you want to add.

[If the Manual discovery method is selected] Specify the machine IP addresses or hostnames, or import the machine list from a text file. The file must contain IP addresses/hostnames, one per line. Here is an example of a file:

```
156.85.34.10
156.85.53.32
156.85.53.12
EN-L00000100
EN-L00000101
```

After adding machine addresses manually or importing from a file, the agent tries to ping the added machines and define their availability.

9. Select what actions must be performed after the discovery:
 - **Install agents and register machines.** You can select which components to install on the machines by clicking **Select components**. For more details, refer to "Selecting components for installation" (p. 131).

On the **Select components** screen, define the account under which the services will run by specifying **Logon account for the agent service**. You can select one of the following:

 - **Use Service User Accounts** (default for the agent service)

Service User Accounts are Windows system accounts that are used to run services. The advantage of this setting is that the domain security policies do not affect these accounts' user rights. By default, the agent runs under the **Local System** account.
 - **Create a new account**

The account name will be Agent User for the agent.
 - **Use the following account**

If you install the agent on a domain controller, the system prompts you to specify existing accounts (or the same account) for the agent. For security reasons, the system does not automatically create new accounts on a domain controller.

If you chose the **Create a new account** or **Use the following account** option, ensure that the domain security policies do not affect the related accounts' rights. If an account is deprived of the user rights assigned during the installation, the component may work incorrectly or not work.

- **Register machines with installed agents.** This option is used if the agent is already installed on machines and you need only to register them in Cyber Protection. If no agent is found inside the machines, then they will be added as **Unmanaged** machines.
- **Add as unmanaged machines.** The agent will not be installed on the machines. You will be able to view them in the console and install or register the agent later.

[If the **Install agents and register machines** post-discovery action is selected] **Restart the machine if required** – if the option is enabled, the machine will be restarted as many times as required to complete the installation.

Restart of the machine may be required in one of the following cases:

- Installation of prerequisites is completed and restart is required to continue the installation
- Installation is completed but restart is required as some files are locked during installation
- Installation is completed but restart is required for other previously installed software

[If **Restart the machine if required** is selected] **Do not restart if the user logged in** – if the option is enabled, the machine will not be automatically restarted if the user is logged in to the system. For example, if a user is working while installation requires restart, the system will not be restarted.

If the prerequisites were installed and then the reboot was not done because a user was logged in, then to complete the agent installation you need to reboot the machine and start the installation again.

If the agent was installed but then the reboot was not done, then you need to reboot the machine.

[If there are units in your organization] **User for whom to register the machines** – select the user of your unit or subordinate units for whom the machines will be registered.

If you have selected one of the first two post-discovery actions, then there is also an option to apply the protection plan to the machines. If you have several protection plans, you can select which one to use.

10. Specify the credentials of the user with administrator rights for all of the machines.

Important

Note that remote installation of agent works without any preparations only if you specify the credentials of the built-in administrator account (the first account created when the operating system is installed). If you want to define some custom administrator credentials, then you should do additional manual preparations as described in "Preparing a machine for remote installation" (p. 130).

11. The system checks connectivity to all of the machines. If the connection to some of the machines fails, you can change the credentials for these machines.

When the discovery of machines is initiated, you will find the corresponding task in **Monitoring > Activities > Discovering machines** activity.

Preparing a machine for remote installation

1. For successful installation on a remote machine running Windows 7 or later, the option **Control panel > Folder options > View > Use Sharing Wizard** must be *disabled* on that machine.
2. For successful installation on a remote machine that is *not* a member of an Active Directory domain, User Account Control (UAC) must be *disabled* on that machine. For more information on how to disable it, refer to "[Requirements on User Account Control \(UAC\)](#)" > To disable UAC.
3. By default, the credentials of the built-in administrator account are required for remote installation on any Windows machine. To perform remote installation by using the credentials of another administrator account, User Account Control (UAC) remote restrictions must be *disabled*. For more information on how to disable them, refer to "[Requirements on User Account Control \(UAC\)](#)" > To disable UAC remote restrictions.
4. File and Printer Sharing must be *enabled* on the remote machine. To access this option:
 - On a machine running Windows 2003 Server: go to **Control panel > Windows Firewall > Exceptions > File and Printer Sharing**.
 - On a machine running Windows Server 2008, Windows 7, or later: go to **Control panel > Windows Firewall > Network and Sharing Center > Change advanced sharing settings**.
5. Cyber Protection uses TCP ports 445, 25001, and 43234 for remote installation. Port 445 is automatically opened when you enable File and Printer Sharing. Ports 43234 and 25001 are automatically opened through Windows Firewall. If you use a different firewall, make sure that these three ports are open (added to exceptions) for both incoming and outgoing requests.

After the remote installation is complete, port 25001 is automatically closed through Windows Firewall. Ports 445 and 43234 need to remain open if you want to update the agent remotely in the future. Port 25001 is automatically opened and closed through Windows Firewall during each update. If you use a different firewall, keep all the three ports open.

Requirements on User Account Control (UAC)

On a machine that is running Windows 7 or later and is not a member of an Active Directory domain, centralized management operations (including remote installation) require that UAC and UAC remote restrictions be disabled.

To disable UAC

Do one of the following depending on the operating system:

- **In a Windows operating system prior to Windows 8:**
Go to **Control panel > View by: Small icons > User Accounts > Change User Account Control Settings**, and then move the slider to **Never notify**. Then, restart the machine.
- **In any Windows operating system:**
 1. Open Registry Editor.
 2. Locate the following registry key: **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System**
 3. For the **EnableLUA** value, change the setting to **0**.
 4. Restart the machine.

To disable UAC remote restrictions

1. Open Registry Editor.
2. Locate the following registry key: **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System**
3. For **LocalAccountTokenFilterPolicy** value, change the setting to **1**.
If the **LocalAccountTokenFilterPolicy** value does not exist, create it as DWORD (32-bit). For more information about this value, refer to the Microsoft documentation:
<https://support.microsoft.com/en-us/help/951016/description-of-user-account-control-and-remote-restrictions-in-windows>.

Note

For security reasons, we recommend that after finishing the management operation (for example, remote installation), you revert both settings to their original state: **EnableLUA=1** and **LocalAccountTokenFilterPolicy = 0**

Selecting components for installation

You can find the description of mandatory and additional components in the following table:

Component	Description
Mandatory component	
Agent for Windows	This agent backs up disks, volumes, files and will be installed on Windows machines. It will be always installed, not selectable.
Additional components	
Agent for Data Loss Prevention	This agent enables you to limit the user access to local and redirected peripheral devices, ports, and clipboard on machines under protection plans. It will be installed if selected.
Antimalware and URL filtering	This component enables the Antivirus & Antimalware protection module and URL filtering module in protection plans. Even if you select not to install it, it will be automatically installed later, if any of these modules is enabled in a protection plan for the machine.

Agent for Hyper-V	This agent backs up Hyper-V virtual machines and will be installed on Hyper-V hosts. It will be installed if selected and detected Hyper-V role on a machine.
Agent for SQL	This agent backs up SQL Server databases and will be installed on machines running Microsoft SQL Server. It will be installed if selected and application detected on a machine.
Agent for Exchange	This agent backs up Exchange databases and mailboxes and will be installed on machines running the Mailbox role of Microsoft Exchange Server. I will be installed if selected and application detected on a machine.
Agent for Active Directory	This agent backs up the data of Active Directory Domain Services and will be installed on domain controllers. It will be installed if selected and application detected on a machine.
Agent for VMware (Windows)	This agent backs up VMware virtual machines and will be installed on Windows machines that have network access to vCenter Server. It will be installed if selected.
Agent for Microsoft 365	This agent backs up Microsoft 365 mailboxes to a local destination and will be installed on Windows machines. It will be installed if selected.
Agent for Oracle	This agent backs up Oracle databases and will be installed on machines running Oracle Database. It will be installed if selected.
Cyber Protection Monitor	This component enables a user to monitor execution of running tasks in the notification area and will be installed on Windows machines. It will be installed if selected. Supported on Windows 7 Service Pack 1 and later, and Windows Server 2008 R2 Service Pack 1 and later.

Managing discovered machines

After the discovery process is performed, you can find all of the discovered machines in **Devices > Unmanaged machines**.

This section is divided into subsections by the discovery method used. The full list of machine parameters is shown below (it may vary depending on the discovery method):

Name	Description
Name	The name of the machine. The IP address will be shown if the name of the machine could not be discovered.
IP address	The IP address of the machine.
Discovery type	The discovery method that was used to detect the machine.
Organizational unit	The organizational unit in Active Directory that the machine belongs to. This column is shown if you view the list of machines in Unmanaged machines > Active Directory .

Operating system	The operating system installed in the machine.
-------------------------	--

There is an **Exceptions** section, where you can add the machines that must be skipped during the discovery process. For example, if you do not need the exact machines to be discovered, you can add them to this list.

To add a machine to **Exceptions**, select it in the list and click **Add to exceptions**. To remove a machine from **Exceptions**, go to **Unmanaged machines > Exceptions**, select the machine, and click **Remove from exceptions**.

You can install the protection agent and register a batch of discovered machines in Cyber Protection by selecting them in the list and clicking **Install and register**. The opened wizard also allows you to assign the protection plan to a batch of machines.


After the protection agent is installed on machines, those machines will be shown in the **Devices > Machines with agents** section.

To check your protection status, go to **Monitoring > Overview** and add the **Protection status** widget or the **Discovered machine** widget.

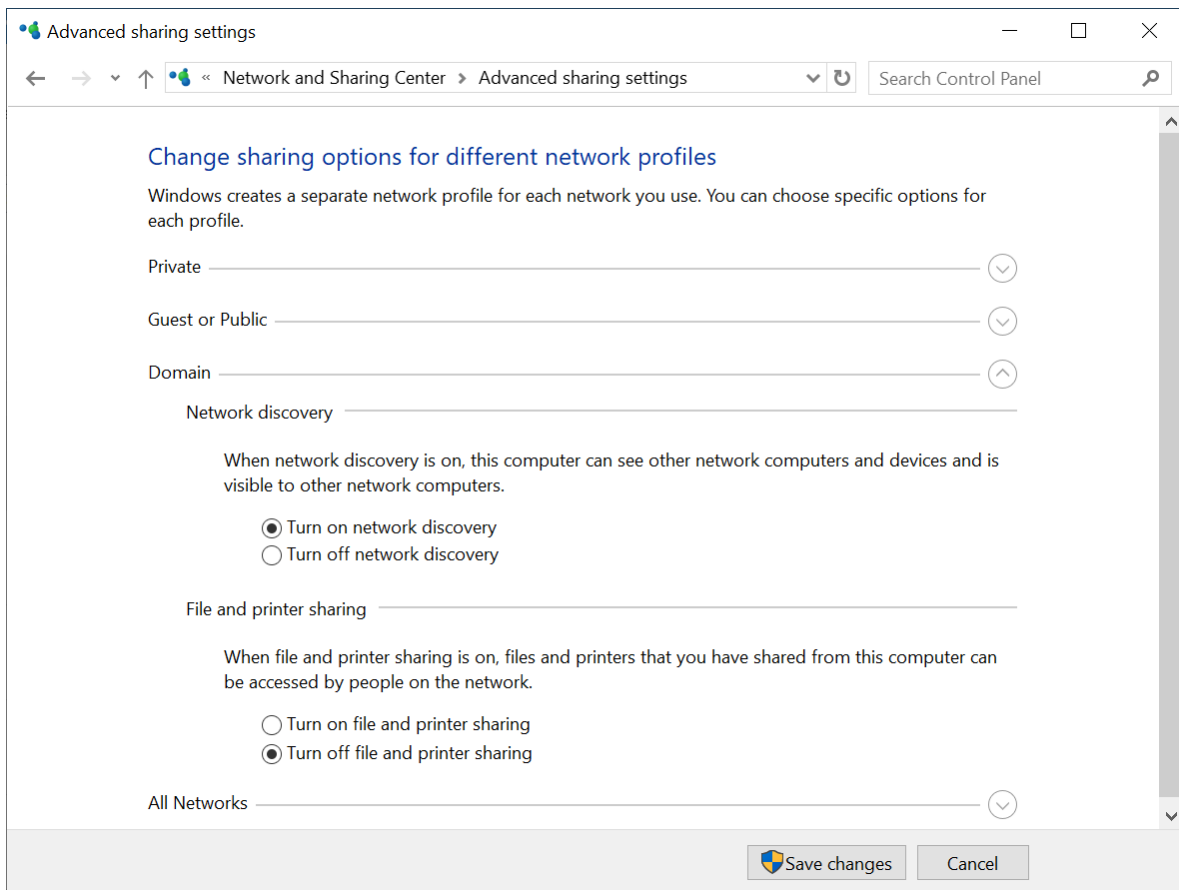
Troubleshooting

If you have any issues with the autodiscovery functionality, try to check the following:

- Check that NetBIOS over TCP/IP is enabled or set to default.



- In the "Control Panel\Network and Sharing Center\Advanced sharing settings" turn on network discovery.



- Check that the Function Discovery Provider Host service is running on the machine that does discovery and on the machines to be discovered.
- Check that the Function Discovery Resource Publication service is running on the machines to be discovered.

Deploying Agent for VMware (Virtual Appliance)

Before you start

System requirements for the agent

By default, the virtual appliance is assigned 4 GB of RAM and 2 vCPUs, which is optimal and sufficient for most operations.

To improve the backup performance and avoid failures related to insufficient RAM memory, we recommend that you increase these resources to 16 GB of RAM and 4 vCPUs in more demanding cases. For example, increase the assigned resources when you expect the backup traffic to exceed 100 MB per second (for example, in 10-Gigabit networks) or if you simultaneously back up multiple virtual machines with large hard drives (500 GB or more).

The appliance's own virtual disks occupy no more than 6 GB. Thick or thin disk format does not matter, it does not affect the appliance performance.

How many agents do I need?

Even though one virtual appliance is able to protect an entire vSphere environment, the best practice is deploying one virtual appliance per vSphere cluster (or per host, if there are no clusters). This makes for faster backups because the appliance can attach the backed-up disks by using the HotAdd transport, and therefore the backup traffic is directed from one local disk to another.

It is normal to use both the virtual appliance and Agent for VMware (Windows) at the same time, as long as they are connected to the same vCenter Server *or* they are connected to different ESXi hosts. Avoid cases when one agent is connected to an ESXi directly and another agent is connected to the vCenter Server which manages this ESXi.

We do not recommend that you use locally attached storage (i.e. storing backups on virtual disks added to the virtual appliance) if you have more than one agent. For more considerations, see "Using a locally attached storage" (p. 641).

Disable automatic DRS for the agent

If the virtual appliance is deployed to a vSphere cluster, be sure to disable automatic vMotion for it. In the cluster DRS settings, enable individual virtual machine automation levels, and then set **Automation level** for the virtual appliance to **Disabled**.

Deploying the OVF template

1. Click **All devices > Add > VMware ESXi > Virtual Appliance (OVF)**.
The .zip archive is downloaded to your machine.
2. Unpack the .zip archive. The folder contains one .ovf file and two .vmdk files.
3. Ensure that these files can be accessed from the machine running vSphere Client.
4. Start vSphere Client and log on to the vCenter Server.
5. Deploy the OVF template.
 - When configuring storage, select the shared datastore, if it exists. Thick or thin disk format does not matter, as it does not affect the appliance performance.
 - When configuring network connections, be sure to select a network that allows an Internet connection, so that the agent can properly register itself in the cloud.

Configuring the virtual appliance

After deploying the virtual appliance, you must configure it so that it can access vCenter Server or the ESXi host and the Cyber Protection service.

To configure the virtual appliance

1. In the vSphere Client, open the console of the virtual appliance.
2. Ensure that the network connection is configured.
The connection is configured automatically via Dynamic Host Configuration Protocol (DHCP).

To change the default configuration, under **Agent options**, in the **eth0** field, click **Change**, and then specify the network settings.

3. Connect the virtual appliance to vCenter Server or the ESXi host.
 - a. Under **Agent options**, in the **vCenter/ESX(i)** field, click **Change**, and then specify the following.
 - [If you use vCenter Server] The vCenter Server name or IP address.
 - [If you do not use vCenter Server] The name or IP address of the ESXi host on which you want to back up and recover virtual machines. For faster backups, deploy the virtual appliance on the same host.
 - The credentials required for the appliance to connect to vCenter Server or the ESXi host. We recommend that you use a dedicated account for accessing vCenter Server or the ESXi host, instead of using an existing account with the Administrator role. To learn more about the required privileges for the dedicated account, refer to "Agent for VMware – necessary privileges" (p. 647).
 - b. Click **Check connection** to verify that the settings are correct.
 - c. Click **OK**.
4. Register the appliance in the Cyber Protection service by using one of the following methods.
 - [Only for tenants without two-factor authentication] Register the appliance in its graphical interface.
 - a. Under **Agent options**, in the **Management Server** field, click **Change**.
 - b. In the **Server name/IP** field, select **Cloud**.

The Cyber Protection service address appears. Do not change this address unless instructed otherwise.
 - c. In the **User name** and **Password** fields, specify the credentials for your account in the Cyber Protection service. The virtual appliance and the virtual machines that the appliance manages are registered under this account.
 - d. Click **OK**.
 - Register the appliance in the command-line interface.

Note


With this method, you need a registration token. For more information about how to generate one, refer to "Generating a registration token" (p. 166).

- a. Press CTRL+SHIFT+F2 to open the command-line interface.
- b. Run the following command:

```
register_agent -o register -t cloud -a <service address> --token <registration token>
```

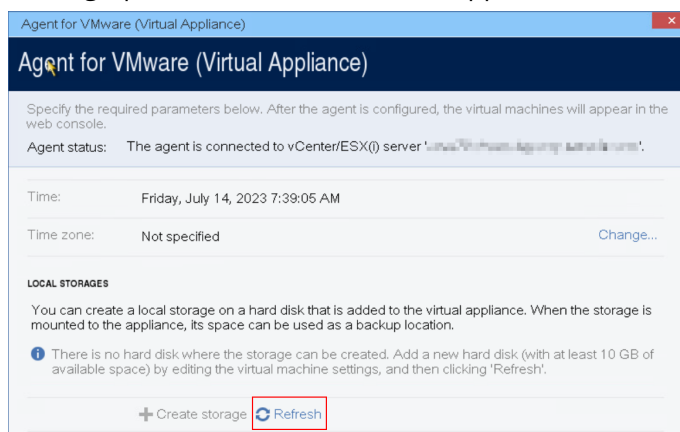
Note

When you use a registration token, you must specify the exact data center address. This is the URL that you see **after you log in** to the Cyber Protect console. For example, `https://eu2-cloud.company.com`.



Do not use `https://cloud.company.com` here.

- c. To return to the graphical interface of the appliance, press ALT+F1.
5. [Optional] Add local storage.
 - a. In the vSphere Client, attach a virtual disk to the virtual appliance. The virtual disk must have at least 10 GB of free space.
 - b. In the graphical user interface of the appliance, click **Refresh**.



The **Create storage** button becomes active.

- c. Click **Create storage**.
 - d. Specify a label for the storage, and then click **OK**.
 - e. Confirm your choice by clicking **Yes**.
6. [If a proxy server is enabled in your network] Configure the proxy server.
 - a. Press CTRL+SHIFT+F2 to open the command-line interface.
 - b. Open the file `/etc/Acronis/Global.config` in a text editor.
 - c. Do one of the following:
 - If the proxy settings were specified during the agent installation, find the following section:

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdword">"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdword">"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```

- Otherwise, copy the above lines and paste them into the file between the <registry name="Global">...</registry> tags.
- d. Replace ADDRESS with the new proxy server host name/IP address, and PORT with the decimal value of the port number.
- e. If your proxy server requires authentication, replace LOGIN and PASSWORD with the proxy server credentials. Otherwise, delete these lines from the file.
- f. Save the file.
- g. Open the file **/opt/acronis/etc/aakore.yaml** in a text editor.
- h. Locate the **env** section or create it and add the following lines:

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

- i. Replace proxy_login and proxy_password with the proxy server credentials, and proxy_address:port with the address and port number of the proxy server.
- j. Run the reboot command.

Note

To be able to update a virtual appliance deployed behind a proxy, edit the appliance `config.yaml` file (`/opt/acronis/etc/va-updater/config.yaml`), by adding the following line to the bottom of that file, and then entering values specific to your environment:

```
httpProxy: http://<proxy_login>:<proxy_password>@<proxy_address>:<port>
```

For example:

```
httpProxy: http://mylogin:mypassword@192.168.2.300:8080
```

Deploying Agent for Scale Computing HC3 (Virtual Appliance)

Before you start

This appliance is a pre-configured virtual machine that you deploy in a Scale Computing HC3 cluster. It contains a protection agent that enables you to administer cyber protection for all virtual machines in the cluster.

System requirements for the agent

By default, the virtual machine with the agent uses 2 vCPUs and 4 GiB of RAM. These settings are sufficient for most operations but you can change them by editing the virtual machine in the Scale Computing HC3 web interface.

To improve the backup performance and avoid failures related to insufficient RAM memory, we recommend that you increase these resources to 4 vCPUs and 8 GiB of RAM in more demanding cases. For example, increase the assigned resources when you expect the backup traffic to exceed 100 MB per second (for example, in 10-Gigabit networks) or if you back up simultaneously multiple virtual machines with large hard drives (500 GB or more).

The size of the appliance virtual disk is about 9 GB.

How many agents do I need?

One agent can protect the entire cluster. However, you can have more than one agent in the cluster if you need to distribute the backup traffic bandwidth load.

If you have more than one agent in a cluster, the virtual machines are automatically evenly distributed between the agents, so that each agent manages a similar number of machines.

Automatic redistribution occurs when the load imbalance among the agents reaches 20 percent. This may happen after you add or remove a machine or an agent. For example, you realize that you need more agents to help with throughput and you deploy an additional virtual appliance to the cluster. The management server will assign the most appropriate machines to the new agent. The old agents' load will reduce. When you remove an agent from the management server, the machines assigned to the agent are redistributed among the remaining agents. However, this will not happen if an agent gets corrupted or is deleted manually from the Scale Computing HC3 cluster. Redistribution will start only after you remove such an agent from the Cyber Protect console.

To check which agent manages a specific machine

1. In the Cyber Protect console, click **Devices**, and then select **Scale Computing**.
2. Click the gear icon in the upper right corner of the table, and under **System**, select the **Agent** check box.
3. Check the name of the agent in the column that appears.

Deploying the QCOW2 template

1. Log in to your Cyber Protection account.
2. Click **Devices > All devices > Add > Scale Computing HC3**.
The .zip archive is downloaded to your machine.
3. Unpack the .zip archive, and then save the .qcow2 file and the .xml file to a folder named **ScaleAppliance**.
4. Upload the **ScaleAppliance** folder to a network share and ensure that the Scale Computing HC3 cluster can access it.
5. Log in to the Scale Computing HC3 cluster as an administrator who has the **VM Create/Edit** role assigned. For more information about the roles required for operations with Scale Computing HC3 virtual machines, refer to "Agent for Scale Computing HC3 – required roles" (p. 142).
6. In the Scale Computing HC3 web interface, import the virtual machine template from the **ScaleAppliance** folder.

- a. Click the **Import HC3 VM** icon.
- b. In the **Import HC3 VM** window, specify the following:
 - A name for the new virtual machine.
 - The network share on which the **ScaleAppliance** folder is located.
 - The user name and password required for accessing this network share.
 - [Optional] A domain tag for the new virtual machine.
 - The path to the **ScaleAppliance** folder on the network share.
- c. Click **Import**.

After the deployment completes, you must configure the virtual appliance. For more information on how to configure it, refer to "Configuring the virtual appliance" (p. 140).

Note

If you need more than one virtual appliance in your cluster, repeat the steps above and deploy additional virtual appliances. Do not clone an existing virtual appliance by using the **Clone VM** option in the Scale Computing HC3 web interface.

Configuring the virtual appliance

After deploying the virtual appliance, you need to configure it so that it can reach both the Scale Computing HC3 cluster that it will protect and the Cyber Protection service.

To configure the virtual appliance

1. Log in to your Scale Computing HC3 account.
2. Select the virtual appliance that you want to configure, and then click the **Console** icon.
3. In the **eth0** field, configure the network interfaces of the appliance.

Ensure that automatically assigned DHCP addresses (if any) are valid within the networks that your virtual machine uses or assign them manually. Depending on the number of networks that the appliance uses, there may be one or more interfaces to configure.
4. In the **Scale Computing** field, click **Change** to specify the Scale Computing HC3 cluster address and credentials for accessing it.
 - a. In the **Server name/IP** field, enter the DNS name or IP address of the cluster.
 - b. In the **User name** and **Password** fields, enter the credentials for the Scale Computing HC3 administrator account.

Ensure that this account has the roles required for operations with Scale Computing HC3 virtual machines. For more information about these roles, refer to "Agent for Scale Computing HC3 – required roles" (p. 142).
 - c. Click **Check connection** to verify that the settings are correct.
 - d. Click **OK**.
5. Register the appliance in the Cyber Protection service by using one of the following methods.
 - [Only for tenants without two-factor authentication] Register the appliance in its graphical interface.

- a. Under **Agent options**, in the **Management Server** field, click **Change**.
 - b. In the **Server name/IP** field, select **Cloud**.
The Cyber Protection service address appears. Do not change this address unless instructed otherwise.
 - c. In the **User name** and **Password** fields, specify the credentials for your account in the Cyber Protection service. The virtual appliance and the virtual machines that the appliance manages are registered under this account.
 - d. Click **OK**.
- Register the appliance in the command-line interface.

Note


With this method, you need a registration token. For more information about how to generate one, refer to "Generating a registration token" (p. 166).

- a. Press CTRL+SHIFT+F2 to open the command-line interface.
- b. Run the following command:

```
register_agent -o register -t cloud -a <service address> --token <registration token>
```

Note

When you use a registration token, you must specify the exact data center address. This is the URL that you see **after you log in** to the Cyber Protect console. For example, <https://eu2-cloud.company.com>.



Do not use <https://cloud.company.com> here.

- c. To return to the graphical interface of the appliance, press ALT+F1.
6. [Optional] In the **Name** field, click **Change** to edit the default name for the virtual appliance, which is **localhost**. This name is shown in the Cyber Protect console.
 7. [Optional] In the **Time** field, click **Change**, and then select the time zone of your location to ensure that the scheduled operations run at the appropriate time.
 8. [If a proxy server is enabled in your network] Configure the proxy server.
 - a. Press CTRL+SHIFT+F2 to open the command-line interface.
 - b. Open the file **/etc/Acronis/Global.config** in a text editor.
 - c. Do one of the following:
 - If the proxy settings were specified during the agent installation, find the following section:

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdwor">"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdwor">"PORT"</value>
```

```
<value name="Login" type="TString">"LOGIN"</value>
<value name="Password" type="TString">"PASSWORD"</value>
</key>
```

- Otherwise, copy the above lines and paste them into the file between the `<registry name="Global">...</registry>` tags.
- d. Replace ADDRESS with the new proxy server host name/IP address, and PORT with the decimal value of the port number.
 - e. If your proxy server requires authentication, replace LOGIN and PASSWORD with the proxy server credentials. Otherwise, delete these lines from the file.
 - f. Save the file.
 - g. Open the file `/opt/acronis/etc/aakore.yaml` in a text editor.
 - h. Locate the **env** section or create it and add the following lines:

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

- i. Replace proxy_login and proxy_password with the proxy server credentials, and proxy_address:port with the address and port number of the proxy server.
- j. Run the reboot command.

Note

To be able to update a virtual appliance deployed behind a proxy, edit the appliance `config.yaml` file (`/opt/acronis/etc/va-updater/config.yaml`), by adding the following line to the bottom of that file, and then entering values specific to your environment:

```
httpProxy: http://<proxy_login>:<proxy_password>@<proxy_address>:<port>
```

For example:

```
httpProxy: http://mylogin:mypassword@192.168.2.300:8080
```

To protect virtual machines in the Scale Computing HC3 cluster

1. Log in to your Cyber Protection account.
2. Navigate to **Devices > Scale Computing HC3** <your cluster> or find your machines in **Devices > All devices**.
3. Select machines and apply a protection plan to them.

Agent for Scale Computing HC3 – required roles

This section describes the roles required for operations with Scale Computing HC3 virtual machines.

Operation	Role
-----------	------

Back up a virtual machine	Backup VM Create/Edit VM Delete
Recover to an existing virtual machine	Backup VM Create/Edit VM Power Control VM Delete Cluster Settings
Recover to a new virtual machine	Backup VM Create/Edit VM Power Control VM Delete Cluster Settings

Deploying Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance)

Before you start

This appliance is a pre-configured virtual machine that you deploy in Virtuozzo Hybrid Infrastructure. It contains a protection agent that enables you to administer cyber protection for all virtual machines in a Virtuozzo Hybrid Infrastructure cluster.

Note

To ensure that backups with enabled **Volume Shadow Copy Service (VSS) for virtual machines** backup option run properly and capture data in application-consistent state, verify that Virtuozzo Guest Tools are installed and up-to-date on the protected virtual machines.

System requirements for the agent

When deploying the virtual appliance, you can choose between different predefined combinations of vCPUs and RAM (flavors). You can also create your own flavors.

2 vCPUs and 4 GB of RAM (medium flavor) are optimal and sufficient for most operations. To improve the backup performance and avoid failures related to insufficient RAM memory, we recommend that you increase these resources to 4 vCPUs and 8 GB of RAM in more demanding cases. For example, increase the assigned resources when you expect the backup traffic to exceed

100 MB per second (for example, in 10-Gigabit networks) or if you back up simultaneously multiple virtual machines with large hard drives (500 GB or more).

How many agents do I need?

One agent can protect the entire cluster. However, you can have more than one agent in the cluster if you need to distribute the backup traffic bandwidth load.

If you have more than one agent in a cluster, the virtual machines are automatically evenly distributed between the agents, so that each agent manages a similar number of machines.

Automatic redistribution occurs when the load imbalance among the agents reaches 20 percent. This may happen after you add or remove a machine or an agent. For example, you realize that you need more agents to help with throughput and you deploy an additional virtual appliance to the cluster. The management server will assign the most appropriate machines to the new agent. The old agents' load will reduce. When you remove an agent from the management server, the machines assigned to the agent are redistributed among the remaining agents. However, this will not happen if an agent gets corrupted or is deleted manually from the Virtuozzo Hybrid Infrastructure node. Redistribution will start only after you remove such an agent from the Cyber Protection web interface.

To check which agent manages a specific machine

1. In the Cyber Protect console, click **Devices**, and then select **Virtuozzo Hybrid Infrastructure**.
2. Click the gear icon in the upper right corner of the table, and under **System**, select the **Agent** check box.
3. Check the name of the agent in the column that appears.

Limitations

- Virtuozzo Hybrid Infrastructure appliance cannot be deployed remotely.
- Application-aware backup of virtual machines is not supported.

Configuring networks in Virtuozzo Hybrid Infrastructure

Before deploying and configuring the virtual appliance, you need to have your networks in Virtuozzo Hybrid Infrastructure configured.

Network requirements for the Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance)

- The virtual appliance requires 2 network adapters.
- The virtual appliance must be connected to Virtuozzo networks with the following network traffic types:
 - Compute API
 - VM Backup

- ABGW Public
- VM Public

For more information about configuring the networks, see [Compute cluster requirements](#) in the Virtuozzo documentation.

Configuring user accounts in Virtuozzo Hybrid Infrastructure

To configure the virtual appliance, you need a Virtuozzo Hybrid Infrastructure user account. This account must have the **Administrator** role in the **Default** domain. For more information about users, refer to [Managing admin panel users](#) in the Virtuozzo Hybrid Infrastructure documentation. Ensure that you granted this account access to all projects in the **Default** domain.

To grant access to all projects in the Default domain

1. Create an environment file for the system administrator. To do this, run the following script in the Virtuozzo Hybrid Infrastructure cluster via the OpenStack Command-Line Interface. For more information on how to connect to this interface, refer to [Connecting to OpenStack command-line interface](#) in the Virtuozzo Hybrid Infrastructure documentation.

```
su - vstoradmin
kolla-ansible post-deploy
exit
```

2. Use the environment file to authorize further OpenStack commands:

```
. /etc/kolla/admin-openrc.sh
```

3. Run the following commands:

```
openstack --insecure user set --project admin --project-domain Default --domain
Default <username>
openstack --insecure role add --domain Default --user <username> --user-domain
Default compute --inherited
```

Here, <username> is the Virtuozzo Hybrid Infrastructure account with the **Administrator** role in the **Default** domain. The virtual appliance will use this account in order to back up and restore the virtual machines in any child project under the **Default** domain.

Example

```
su - vstoradmin
kolla-ansible post-deploy
exit
. /etc/kolla/admin-openrc.sh
openstack --insecure user set --project admin --project-domain Default --domain
Default johndoe
openstack --insecure role add --domain Default --user johndoe --user-domain
Default compute --inherited
```

To manage backups for virtual machines in a domain that is different from the **Default** domain, run the following command as well.

To grant access to all projects in a different domain

```
openstack --insecure role add --domain <domain name> --inherited --user <username> --user-domain Default admin
```

Here, <domain name> is the domain to the projects in which the <username> account will have access.

Example

```
openstack --insecure role add --domain MyNewDomain --inherited --user johndoe --user-domain Default admin
```

After granting access to projects, check what roles are assigned to the account.

To check assigned roles

```
openstack --insecure role assignment list --user <username> --names
```

Here, <username> is the Virtuozzo Hybrid Infrastructure account.

Example

```
openstack --insecure role assignment list --user johndoe --names -c Role -c User -c Project -c Domain
+-----+-----+-----+-----+
| Role      | User           | Project | Domain  |
+-----+-----+-----+-----+
| admin     | johndoe@Default |        | MyNewDomain |
| compute   | johndoe@Default |        | Default    |
| domain_admin | johndoe@Default |        | Default    |
| domain_admin | johndoe@Default |        | Default    |
+-----+-----+-----+-----+
```

In this example, the options -c Role, -c User, -c Project, and -c Domain are used to abridge the command output to fit the page.

To check what effective roles are assigned to the account in all projects, run the following command as well.

To check effective roles in all projects

```
openstack --insecure role assignment list --user <username> --names --effective
```

Here, <username> is the Virtuozzo Hybrid Infrastructure account.

Example

```
openstack --insecure role assignment list --user johndoe --names --effective -c Role -c
User -c Project -c Domain
+-----+-----+-----+-----+
| Role      | User          | Project      | Domain  |
+-----+-----+-----+-----+
| domain_admin | johndoe@Default |             | Default |
| compute     | johndoe@Default | admin@Default |         |
| compute     | johndoe@Default | service@Default |         |
| domain_admin | johndoe@Default | admin@Default |         |
| domain_admin | johndoe@Default | service@Default |         |
| project_user | johndoe@Default | service@Default |         |
| member      | johndoe@Default | service@Default |         |
| reader      | johndoe@Default | service@Default |         |
| project_user | johndoe@Default | admin@Default |         |
| member      | johndoe@Default | admin@Default |         |
| reader      | johndoe@Default | admin@Default |         |
| project_user | johndoe@Default |             | Default |
| member      | johndoe@Default |             | Default |
| reader      | johndoe@Default |             | Default |
+-----+-----+-----+-----+
```

In this example, the options `-c Role`, `-c User`, `-c Project`, and `-c Domain` are used to abridge the command output to fit the page.

Deploying the QCOW2 template

1. Log in to your Cyber Protection account.
2. Click **Devices > All devices > Add > Virtuozzo Hybrid Infrastructure**.
The .zip archive is downloaded to your machine.
3. Unpack the .zip archive. It contains a .qcow2 image file.
4. Log in to your Virtuozzo Hybrid Infrastructure account.
5. Add the .qcow2 image file to the Virtuozzo Hybrid Infrastructure compute cluster as follows:
 - On the **Compute > Virtual machines > Images** tab, click **Add image**.
 - In the **Add image** window, click **Browse**, and then select the .qcow2 file.
 - Specify the image name, select the **Generic Linux OS** type, and then click **Add**.
6. In the **Compute > Virtual machines > Virtual machines** tab, click **Create virtual machine**. A window will open where you need to specify the following parameters:
 - A name for the new virtual machine.
 - In **Deploy from**, choose **Image**.
 - In the **Images** window, select the .qcow2 image file of the appliance, and then click **Done**.
 - In the **Volumes** window, you don't need to add any volumes. The volume that is added automatically for the system disk is sufficient.

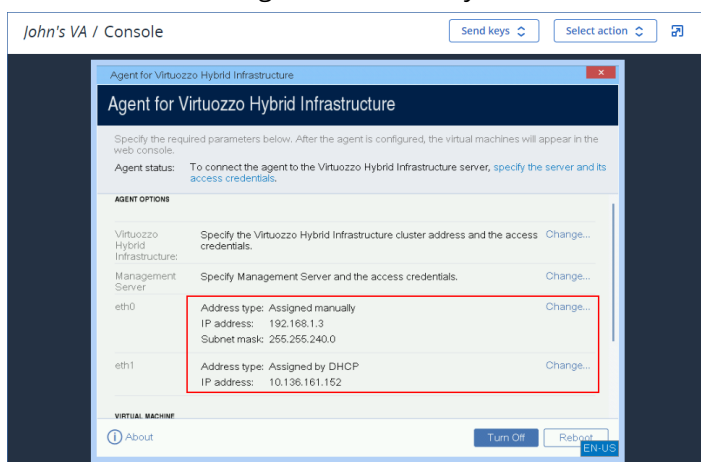
- In the **Flavor** window, choose your desired combination of vCPUs and RAM, and then click **Done**. Usually, 2 vCPUs and 4 GiB of RAM are enough.
 - In the **Network interfaces** window, click **Add**, select the virtual network of type *public*, and then click **Add**. It will appear in the **Network interfaces** list.
If you use a setup with more than one physical network (and thus, with more than one virtual network of type public), repeat this step and select the virtual networks that you need.
7. Click **Done**.
 8. Back in the **Create virtual machine** window, click **Deploy** to create and boot the virtual machine.

Configuring the virtual appliance

After deploying the Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance), you need to configure the virtual appliance so that it can reach both the Virtuozzo Hybrid Infrastructure cluster that it will protect and the Cyber Protection cloud service.

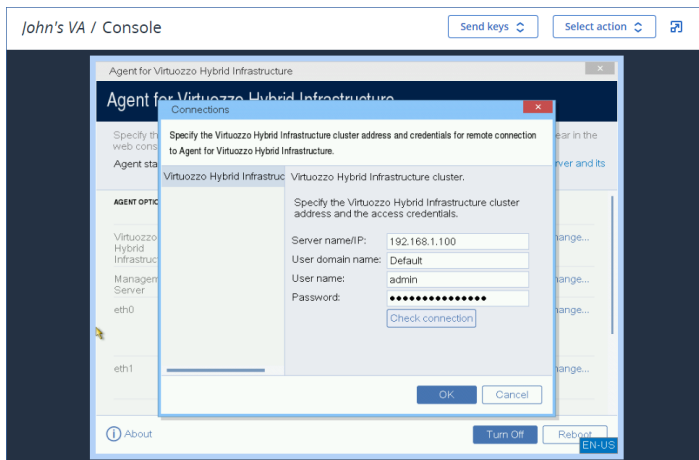
To configure the virtual appliance

1. Log in to your Virtuozzo Hybrid Infrastructure account.
2. On the **Compute > Virtual machines > Virtual Machines** tab, select the virtual machine that you created. Then, click **Console**.
3. Configure the network interfaces of the appliance. There may be one or more interfaces to configure – it depends on the number of virtual networks that the appliance uses. Ensure that automatically assigned DHCP addresses (if any) are valid within the networks that your virtual machine uses or assign them manually.



4. Specify the Virtuozzo cluster address and credentials:
 - DNS name or IP address of the Virtuozzo Hybrid Infrastructure cluster – this is the address of the management node of the cluster. The default port 5000 will be automatically set. If you use a different port, you need to specify it manually.
 - In the **User domain name** field, specify your domain in Virtuozzo Hybrid Infrastructure. For example, **Default**.
The domain name is case-sensitive.

- In the **User name** and **Password** fields, enter the credentials for Virtuozzo Hybrid Infrastructure user account with **Administrator** role in the specified domain. For more information about users, roles, and domains, refer to [Configuring user accounts in Virtuozzo Hybrid Infrastructure](#).



5. Register the appliance in the Cyber Protection service by using one of the following methods.
 - [Only for tenants without two-factor authentication] Register the appliance in its graphical interface.
 - a. Under **Agent options**, in the **Management Server** field, click **Change**.
 - b. In the **Server name/IP** field, select **Cloud**.
The Cyber Protection service address appears. Do not change this address unless instructed otherwise.
 - c. In the **User name** and **Password** fields, specify the credentials for your account in the Cyber Protection service. The virtual appliance and the virtual machines that the appliance manages are registered under this account.
 - d. Click **OK**.
 - Register the appliance in the command-line interface.

Note


With this method, you need a registration token. For more information about how to generate one, refer to "Generating a registration token" (p. 166).

- a. Press CTRL+SHIFT+F2 to open the command-line interface.
- b. Run the following command:

```
register_agent -o register -t cloud -a <service address> --token <registration token>
```

Note

When you use a registration token, you must specify the exact data center address. This is the URL that you see **after you log in** to the Cyber Protect console. For example, <https://eu2-cloud.company.com>.



Do not use <https://cloud.company.com> here.

- c. To return to the graphical interface of the appliance, press ALT+F1.
6. [If a proxy server is enabled in your network] Configure the proxy server.
 - a. Press CTRL+SHIFT+F2 to open the command-line interface.
 - b. Open the file **/etc/Acronis/Global.config** in a text editor.
 - c. Do one of the following:
 - If the proxy settings were specified during the agent installation, find the following section:

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdwor" >"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdwor" >"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```

- Otherwise, copy the above lines and paste them into the file between the `<registry name="Global">...</registry>` tags.
- d. Replace ADDRESS with the new proxy server host name/IP address, and PORT with the decimal value of the port number.
 - e. If your proxy server requires authentication, replace LOGIN and PASSWORD with the proxy server credentials. Otherwise, delete these lines from the file.
 - f. Save the file.
 - g. Open the file **/opt/acronis/etc/aakore.yaml** in a text editor.
 - h. Locate the **env** section or create it and add the following lines:

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

- i. Replace proxy_login and proxy_password with the proxy server credentials, and proxy_address:port with the address and port number of the proxy server.
- j. Run the reboot command.

Note

To be able to update a virtual appliance deployed behind a proxy, edit the appliance `config.yaml` file (`/opt/acronis/etc/va-updater/config.yaml`), by adding the following line to the bottom of that file, and then entering values specific to your environment:

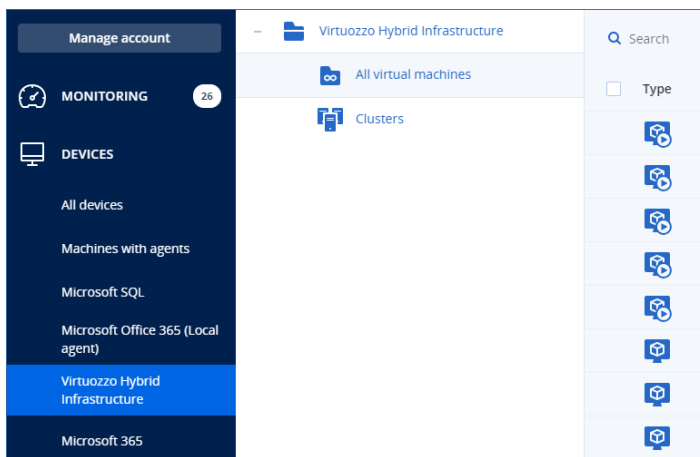
```
httpProxy: http://<proxy_login>:<proxy_password>@<proxy_address>:<port>
```

For example:

```
httpProxy: http://mylogin:mypassword@192.168.2.300:8080
```

To protect the virtual machines in the Virtuozzo Hybrid Infrastructure cluster

1. Log in to your Cyber Protection account.
2. Navigate to **Devices** > **Virtuozzo Hybrid Infrastructure** > <your cluster> > **Default project** > **admin** or find your machines in **Devices** > **All devices**.
3. Select machines and apply a protection plan to them.



Deploying Agent for oVirt (Virtual Appliance)

Before you start

This appliance is a pre-configured virtual machine that you deploy in a Red Hat Virtualization/oVirt data center. The appliance contains a protection agent that enables you to administer cyber protection for all virtual machines in the data center.

System requirements for the agent

By default, the virtual machine with the agent uses 2 vCPUs and 4 GiB of RAM. These settings are sufficient for most operations but you can edit them in Red Hat Virtualization/oVirt Administration Portal.

To improve the backup performance and avoid failures related to insufficient RAM memory, we recommend that you increase these resources to 4 vCPUs and 8 GiB of RAM in more demanding cases. For example, increase the assigned resources when you expect the backup traffic to exceed 100 MB per second (for example, in 10-Gigabit networks) or if you back up simultaneously multiple virtual machines with large hard drives (500 GB or more).

The size of the appliance virtual disk is 8 GiB.

How many agents do I need?

One agent can protect the entire data center. However, you can have more than one agent in the data center if you need to distribute the backup traffic bandwidth load.

If you have more than one agent in the data center, the virtual machines are automatically distributed between the agents, so that each agent manages a similar number of machines.

Automatic redistribution occurs when the load imbalance among the agents reaches 20 percent. This may happen after you add or remove a machine or an agent. For example, you realize that you need more agents to help with throughput and you deploy an additional virtual appliance to the data center. The management server will assign the most appropriate machines to the new agent. The old agents' load will reduce. When you remove an agent, the machines assigned to the agent are redistributed among the remaining agents. However, this will not happen if an agent gets corrupted or is deleted manually from Red Hat Virtualization/oVirt Administration Portal. Redistribution will start only after you remove such an agent from the Cyber Protect console.

To check which agent manages a specific machine

1. In the Cyber Protect console, click **Devices**, and then select **oVirt**.
2. Click the gear icon in the upper right corner of the table, and under **System**, select the **Agent** check box.
3. Check the name of the agent in the column that appears.


Limitations

The following operations are not supported for Red Hat Virtualization/oVirt virtual machines:

- Application-aware backup
- Running a virtual machine from a backup
- Replication of virtual machines
- Changed block tracking

Deploying the OVA template

1. Log in to your Cyber Protection account.
2. Click **Devices > All devices > Add > Red Hat Virtualization (oVirt)**.
The .zip archive is downloaded to your machine.
3. Unpack the .zip archive. It contains one .ova file.

4. Upload the .ova file to a host in the Red Hat Virtualization/oVirt data center that you want to protect.
5. Log in to Red Hat Virtualization/oVirt Administration Portal as an administrator. For more information about the roles required for operations with virtual machines, refer to "Agent for oVirt – required roles and ports" (p. 156).
6. From the navigation menu, select **Compute > Virtual machines**.
7. Click the vertical ellipsis icon  above the main table, and then click **Import**.
8. In the **Import Virtual Machine(s)** window, do the following:
 - a. In **Data center**, select the data center that you want to protect.
 - b. In **Source**, select **Virtual Appliance (OVA)**.
 - c. In **Host**, select the host on which you uploaded the .ova file.
 - d. In **File Path**, specify the path to the directory that contains the .ova file.
 - e. Click **Load**.

The oVirt virtual appliance template from the .ova file appears in the **Virtual Machines on Source** panel.

If the template does not appear in this panel, ensure that you have specified the correct path to the file, the file is not damaged, and the host can be reached.
 - f. In **Virtual Machines on Source**, select the oVirt virtual appliance template, and then click the right arrow.

The template appears in the **Virtual machines to import** panel.
 - g. Click **Next**.
9. In the new window, click the appliance name, and then configure the following settings:
 - On the **Network interfaces** tab, configure the network interfaces.
 - [Optional] On the **General** tab, change the default name of the virtual machine with the agent.

The deployment is now complete. Next, you have to configure the virtual appliance. For more information on how to configure it, refer to "Configuring the virtual appliance" (p. 153).

Note

If you need more than one virtual appliance in your data center, repeat the steps above and deploy additional virtual appliances. Do not clone an existing virtual appliance by using the **Clone VM** option in Red Hat Virtualization/oVirt Administration Portal.

To exclude the virtual appliance from dynamic group backups, you must also exclude it from the list of virtual machines in the Cyber Protect console. To exclude it, in Red Hat Virtualization/oVirt Administration Portal, select the virtual machine with the agent, and then assign the tag `acronis_virtual_appliance` to it.

Configuring the virtual appliance

After deploying the virtual appliance, you need to configure it so that it can reach both the oVirt engine and the Cyber Protection service.

To configure the virtual appliance

1. Log in to Red Hat Virtualization/oVirt Administration Portal.
2. Select the virtual appliance that you want to configure, and then click the **Console** icon.
3. In the **eth0** field, configure the network interfaces of the appliance.
Ensure that automatically assigned DHCP addresses (if any) are valid within the networks that your virtual machine uses or assign them manually. Depending on the number of networks that the appliance uses, there may be one or more interfaces to configure.
4. In the **oVirt** field, click **Change** to specify the oVirt engine address and credentials for accessing it:
 - a. In the **Server name/IP** field, enter the DNS name or IP address of the engine.
 - b. In the **User name** and **Password** fields, enter the administrator credentials for this engine.
Ensure that this administrator account has the roles required for operations with Red Hat Virtualization/oVirt virtual machines. For more information about these roles, refer to "Agent for oVirt – required roles and ports" (p. 156).
If Keycloak is the Single-Sign-On (SSO) provider for the oVirt engine (default in oVirt 4.5.1), use the Keycloak format when specifying the user name. For example, specify the default administrator account as `admin@ovirt@internal.sso` instead of `admin@internal`.
 - c. [Optional] Click **Check connection** to ensure that the provided credentials are correct.
 - d. Click **OK**.
5. Register the appliance in the Cyber Protection service by using one of the following methods.
 - [Only for tenants without two-factor authentication] Register the appliance in its graphical interface.
 - a. Under **Agent options**, in the **Management Server** field, click **Change**.
 - b. In the **Server name/IP** field, select **Cloud**.
The Cyber Protection service address appears. Do not change this address unless instructed otherwise.
 - c. In the **User name** and **Password** fields, specify the credentials for your account in the Cyber Protection service. The virtual appliance and the virtual machines that the appliance manages are registered under this account.
 - d. Click **OK**.
 - Register the appliance in the command-line interface.

Note


With this method, you need a registration token. For more information about how to generate one, refer to "Generating a registration token" (p. 166).

- a. Press CTRL+SHIFT+F2 to open the command-line interface.
- b. Run the following command:

```
register_agent -o register -t cloud -a <service address> --token <registration token>
```

Note

When you use a registration token, you must specify the exact data center address. This is the URL that you see **after you log in** to the Cyber Protect console. For example, <https://eu2-cloud.company.com>.



Do not use <https://cloud.company.com> here.

- c. To return to the graphical interface of the appliance, press ALT+F1.
6. [Optional] In the **Name** field, click **Change** to edit the default name for the virtual appliance, which is **localhost**. This name is shown in the Cyber Protect console.
7. [Optional] In the **Time** field, click **Change**, and then select the time zone of your location to ensure that the scheduled operations run at the appropriate time.
8. [Optional] [If a proxy server is enabled in your network] Configure the proxy server.
 - a. Press CTRL+SHIFT+F2 to open the command-line interface.
 - b. Open the file **/etc/Acronis/Global.config** in a text editor.
 - c. Do one of the following:
 - If the proxy settings were specified during the agent installation, find the following section:

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdwor">"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdwor">"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```

- Otherwise, copy the above lines and paste them into the file between the `<registry name="Global">...</registry>` tags.
- d. Replace ADDRESS with the new proxy server host name/IP address, and PORT with the decimal value of the port number.
 - e. If your proxy server requires authentication, replace LOGIN and PASSWORD with the proxy server credentials. Otherwise, delete these lines from the file.
 - f. Save the file.
 - g. Open the file **/opt/acronis/etc/aakore.yaml** in a text editor.
 - h. Locate the **env** section or create it and add the following lines:

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

- i. Replace proxy_login and proxy_password with the proxy server credentials, and proxy_

address:port with the address and port number of the proxy server.

- j. Run the reboot command.

Note

To be able to update a virtual appliance deployed behind a proxy, edit the appliance `config.yaml` file (`/opt/acronis/etc/va-updater/config.yaml`), by adding the following line to the bottom of that file, and then entering values specific to your environment:

```
httpProxy: http://<proxy_login>:<proxy_password>@<proxy_address>:<port>
```

For example:

```
httpProxy: http://mylogin:mypassword@192.168.2.300:8080
```

To protect virtual machines in the Red Hat Virtualization/oVirt data center

1. Log in to your Cyber Protection account.
2. Navigate to **Devices > oVirt > <your cluster>** or find your machines in **Devices > All devices**.
3. Select machines and apply a protection plan to them.

Agent for oVirt – required roles and ports

Required roles

For its deployment and operation, Agent for oVirt requires an administrator account with the following roles assigned.

oVirt/Red Hat Virtualization 4.2 and 4.3/Oracle Virtualization Manager 4.3

- DiskCreator
- UserVmManager
- TagManager
- UserVmRunTimeManager
- VmCreator

oVirt/Red Hat Virtualization 4.4, 4.5

- SuperUser

Required ports

Agent for oVirt connects to the oVirt engine by using the URL that you specify when you configure the virtual appliance. Usually, the engine URL has the following format: `https://ovirt.company.com`. In this case, the HTTPS protocol and port 443 are used.

Non-default oVirt settings may require another port. You can find the exact port by analyzing the URL format. For example:

oVirt engine URL	Port	Protocol
https://ovirt.company.com/	443	HTTPS
http://ovirt.company.com/	80	HTTP
https://ovirt.company.com:1234/	1234	HTTPS

No additional ports are required for disk Read/Write operations, because the backup is performed in the HotAdd mode.

Deploying Agent for Synology

Before you start

With Agent for Synology, you can back up files and folders from and to Synology NAS devices. The NAS-specific properties and access permissions for shares, folders, and files are preserved.

Agent for Synology runs on the NAS device. Thus, you can use the resources of the device for off-host data processing operations, such as backup replication, validation, and cleanup. To learn more about these operations, refer to "Off-host data processing" (p. 193).

Note

Agent for Synology only supports NAS devices with x86_64 processors. You cannot install the agent on devices with ARM processors.

You can recover a backup to the original or a new location on the NAS device, and to a network folder that is accessible through that device. Backups in the cloud storage can also be recovered to a non-original NAS device on which Agent for Synology is installed.

The table below summarizes the available backup sources and destinations.

What to backup	Items to backup (Backup source)	Where to backup (Backup destination)
Files/folders	Local folder*	Cloud storage
		Local folder*
	Network folder (SMB)**	Network folder (SMB)**
		NFS folder

* Including USB drives that are attached to the NAS device.

Note

Encrypted folders are not supported. These folders are not shown in the Cyber Protection graphical user interface.

** Using external network shares as backup source or backup destination via the SMB protocol is only available for agents running on Synology DiskStation Manager 6.2.3 and later. The data hosted on the Synology NAS itself, including in hosted network shares, can be backed up without limitations.

Limitations

- Backed-up encrypted shares are recovered as non-encrypted.
- Backed-up shares for which the **File compression** option is enabled are recovered with this option disabled.
- You can recover to a Synology NAS device only backups that are created by Agent for Synology.

Downloading the setup program

The setup program for Agent for Synology is available as an SPK file.

Agent for Synology 7.x

To download the setup program

1. In the Cyber Protect console, navigate to **Devices > All devices**.
2. In the upper-right corner, click **Add**.
3. Under **Network attached storage (NAS)**, click **Synology**.

The setup program is downloaded to your machine.

Agent for Synology 6.x

To download the setup program

1. In the Cyber Protect console, navigate to **Devices > All devices**.
2. In the upper-right corner, click **Add**.
3. Under **Network attached storage (NAS)**, click **Synology**.
The setup program for Agent for Synology 7.x is downloaded to your machine.
You can safely stop the download process or ignore the downloaded file.
4. Click **Download Agent for Synology 6.x**.
The setup program for Agent for Synology 6.x is downloaded to your machine.

Installing Agent for Synology

To install Agent for Synology, run the SPK file in Synology DiskStation Manager.

Note

Agent for Synology only supports NAS devices with x86_64 processors. You cannot install the agent on devices with ARM processors.

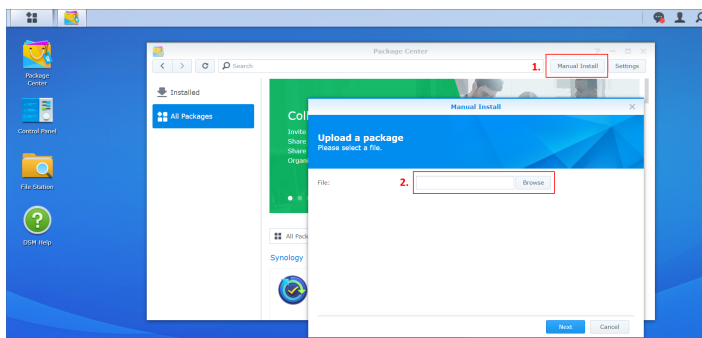
Agent for Synology 7.x

Prerequisites

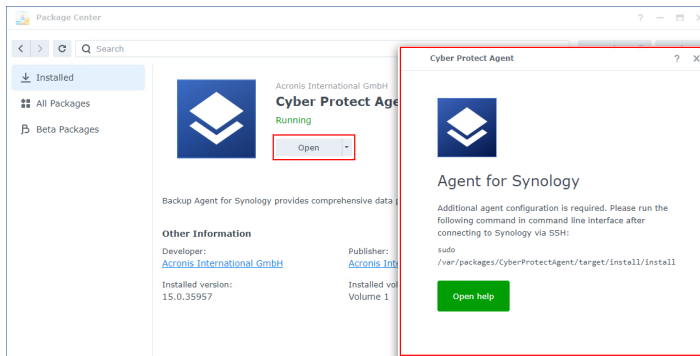
- The NAS device runs DiskStation Manager 7.x.
- You are a member of the **administrators** group on the NAS device.
- There are at least 200 MB of free space on the NAS volume on which you want to install the agent.
- An SSH client is available on your machine. This document uses Putty as an example.

To install Agent for Synology

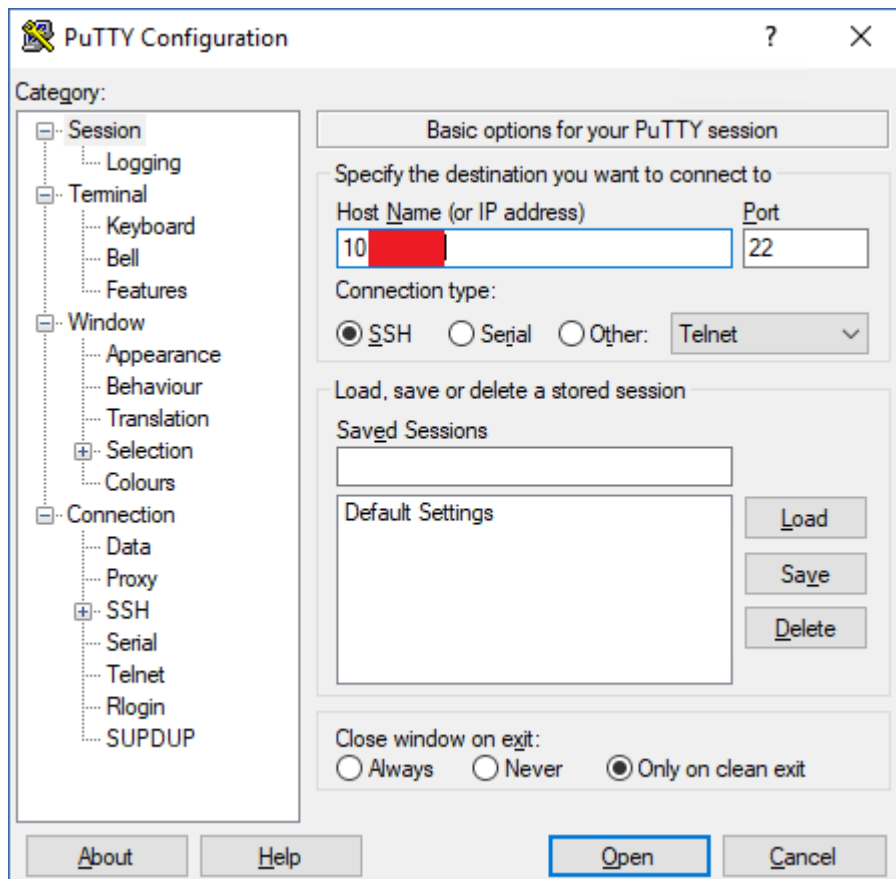
1. Log in to Synology DiskStation Manager.
2. Open **Package Center**.
3. Click **Manual Install**, and then click **Browse**.



4. Select the SPK file that you downloaded from the Cyber Protect console, and then click **Next**.
A warning that you will install a third-party software package is shown. This message is part of the standard installation procedure.
5. To confirm that you want to install the package, click **Agree**.
6. Select the volume on which you want to install the agent, and then click **Next**.
7. Check the settings, and then click **Done**.
8. In Synology DiskStation Manager **Package Center**, open Cyber Protect Agent for Synology, and then verify that you see the following screen.



9. In Synology DiskStation Manager **Control Panel**, go to **Terminal & SNMP**, and then enable the SSH access to the NAS device.
10. Run the `install` script on the NAS device by using an SSH client (in this example, Putty).
 - a. Start Putty, and then specify the IP address or host name of your Synology NAS device.

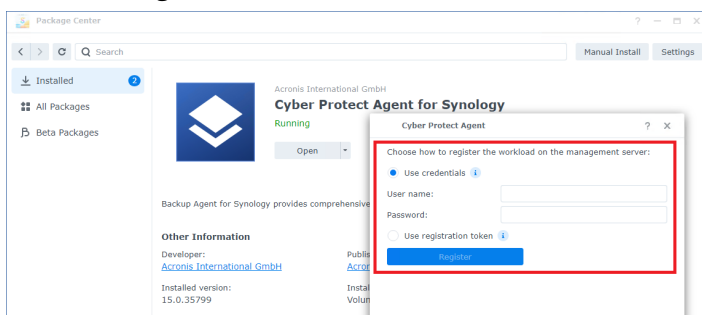


- b. Click **Open**, and then log in as a Synology DSM administrator.
- c. Run the following command.

```
sudo /var/packages/CyberProtectAgent/target/install/install
```

11. In Synology DiskStation Manager **Control Panel**, go to **Terminal & SNMP**, and then disable the SSH access to the NAS device. The SSH access is no longer required.
12. In Synology DiskStation Manager **Package Center**, open Cyber Protect Agent for Synology.

13. Select the registration method.



- [To register the agent by using credentials]
 - In the **User name** and **Password** fields, specify credentials for the account under which the agent will be registered. This account cannot be a partner administrator account.
- [To register the agent by using a registration token]
 - In **Registration address**, specify the exact data center address. The exact data center address is the URL that you see after you log in to the Cyber Protect console. For example, <https://us5-cloud.acronis.com>.

Note

Do not use a URL format without the data center address. For example, do not use <https://cloud.acronis.com>.

- In the **Token** field, specify the registration token.
For more information on how to generate a registration token, see "Generating a registration token" (p. 166).

14. Click **Register**.

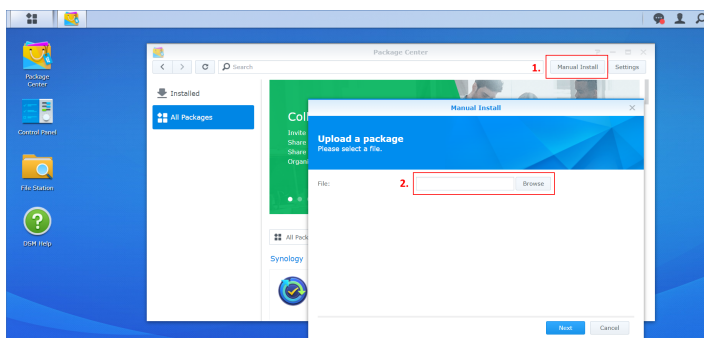
Agent for Synology 6.x

Prerequisites

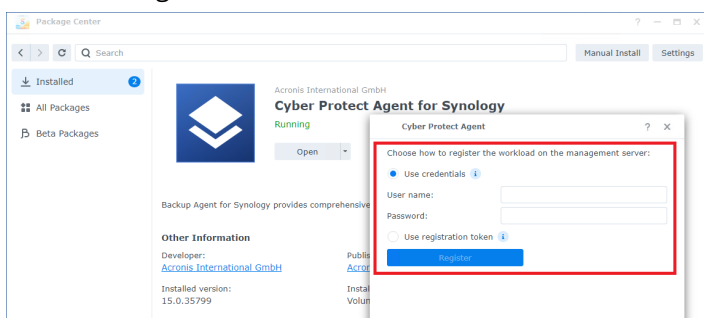
- The NAS device runs DiskStation Manager 6.2.x.
- You are a member of the **administrators** group on the NAS device.
- There are at least 200 MB of free space on the NAS volume on which you want to install the agent.

To install Agent for Synology

1. Log in to Synology DiskStation Manager.
2. Open **Package Center**.
3. Click **Manual Install**, and then click **Browse**.



4. Select the SPK file that you downloaded from the Cyber Protect console, and then click **Next**.
A warning that you will install a package without a digital signature is shown. This message is part of the standard installation procedure.
5. To confirm that you want to install the package, click **Yes**.
6. Select the volume on which you want to install the agent, and then click **Next**.
7. Check the settings, and then click **Apply**.
8. In Synology DiskStation Manager **Package Center**, open Cyber Protect Agent for Synology.
9. Select the registration method.



- [To register the agent by using credentials]
 - In the **User name** and **Password** fields, specify credentials for the account under which the agent will be registered. This account cannot be a partner administrator account.
- [To register the agent by using a registration token]
 - In **Registration address**, specify the exact data center address. The exact data center address is the URL that you see after you log in to the Cyber Protect console. For example, <https://us5-cloud.acronis.com>.

Note

Do not use a URL format without the data center address. For example, do not use <https://cloud.acronis.com>.

- In the **Token** field, specify the registration token.
For more information on how to generate a registration token, see "Generating a registration token" (p. 166).
10. Click **Register**.
When the registration completes, the Synology NAS device appears in the Cyber Protect console, on the **Devices > Network Attached Storage** tab.

To back up the data on the NAS device, apply a protection plan.

Updating Agent for Synology

You can update Agent for Synology 6.x to a newer version of Agent for Synology 6.x. Similarly, you can update Agent for Synology 7.x to a newer version of Agent for Synology 7.x.

To update the agent, run the newer version of the setup program in Synology DiskStation Manager. The original registration of the agent, its settings, and the plans that are applied to the protected workloads will be preserved.

Note

You cannot update the agent from the Cyber Protect console.

Upgrading Agent for Synology 6.x to Agent for Synology 7.x is supported only by uninstalling the older agent and installing the newer agent. In this case, all protection plans are revoked and you must re-apply them manually.

Agent for Synology 7.x

Prerequisites

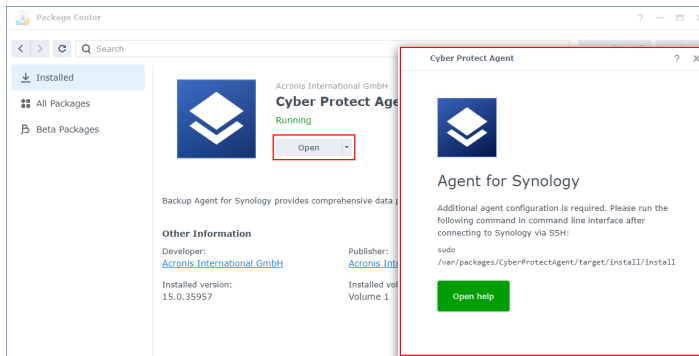
- You are a member of the **administrators** group on the NAS device.
- There are at least 200 MB of free space on the NAS volume on which you want to install the agent.
- An SSH client is available on your machine. This document uses Putty as an example.

To update Agent for Synology

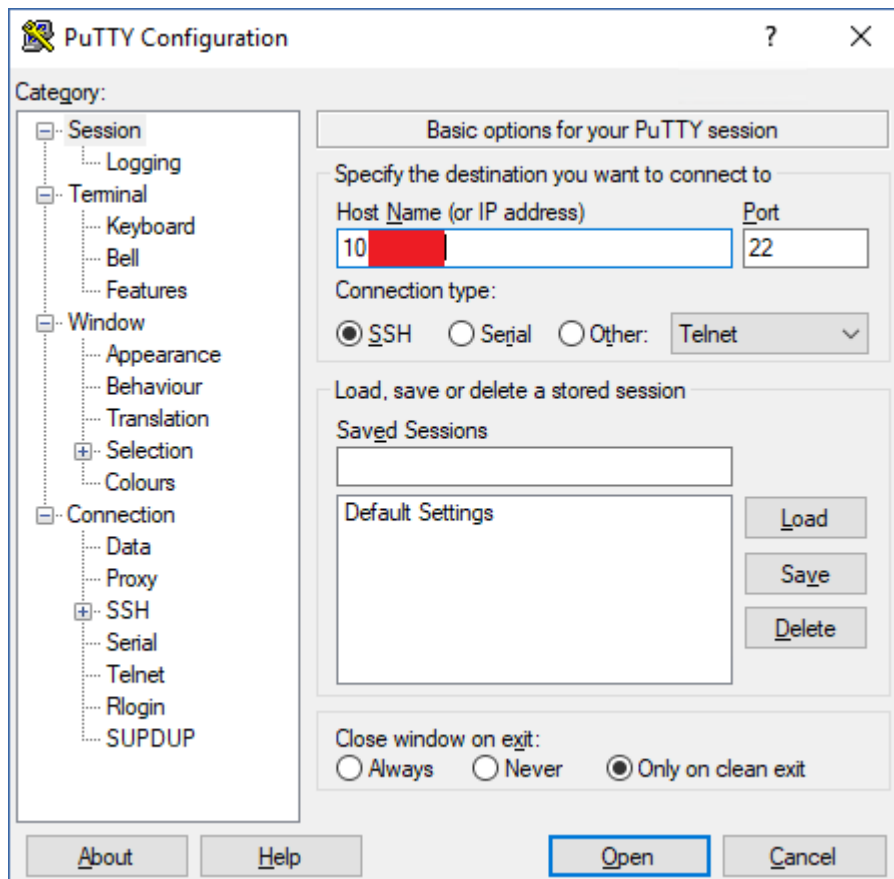
1. In DiskStation Manager, open **Package Center**.
2. Click **Manual Install**, and then click **Browse**.
3. Select the newer SPK file for Agent for Synology 7.x that you downloaded from the Cyber Protect console, and then click **Next**.

A warning that you will install a third-party software package is shown. This message is part of the standard installation procedure.

4. To confirm that you want to install the package, click **Agree**.
5. Check the settings, and then click **Done**.
6. In Synology DiskStation Manager **Package Center**, open Cyber Protect Agent for Synology, and then verify that you see the following screen.



7. In Synology DiskStation Manager **Control Panel**, go to **Terminal & SNMP**, and then enable the SSH access to the NAS device.
8. Run the `install` script on the NAS device by using an SSH client (in this example, Putty).
 - a. Start Putty, and then specify the IP address or host name of your Synology NAS device.



- b. Click **Open**, and then log in as a Synology DSM administrator.
- c. Run the following command.

```
sudo /var/packages/CyberProtectAgent/target/install/install
```

9. In Synology DiskStation Manager **Control Panel**, go to **Terminal & SNMP**, and then disable the SSH access to the NAS device. The SSH access is no longer required.

Agent for Synology 6.x

Prerequisites

- You are a member of the **administrators** group on the NAS device.
- There are at least 200 MB of free space on the NAS volume on which you want to install the agent.

To update Agent for Synology

1. In DiskStation Manager, open **Package Center**.
2. Click **Manual Install**, and then click **Browse**.
3. Select the newer SPK file for Agent for Synology 6.x that you downloaded from the Cyber Protect console, and then click **Next**.
A warning that you will install a package without a digital signature is shown. This message is part of the standard installation procedure.
4. To confirm that you want to install the package, click **Yes**.
5. Check the settings, and then click **Apply**.

Deploying agents through Group Policy

You can centrally install (or deploy) Agent for Windows onto machines that are members of an Active Directory domain, by using Windows Group Policy.

In this section, you will find out how to set up a Group Policy object to deploy agents onto machines in an entire domain or in its organizational unit.

Every time a machine logs on to the domain, the resulting Group Policy object will ensure that the agent is installed and registered.

Prerequisites

- Active Directory domain with a domain controller running Microsoft Windows Server 2003 or later.
- You must be a member of the **Domain Admins** group in this domain.
- You have downloaded the **All agents for Windows** setup program.
To download the setup program, in the Cyber Protect console, click the account icon in the top-right corner, and then click **Downloads**. The download link is also available in the **Add devices** pane.

To deploy agents through Group Policy

1. Generate a registration token as described in "Generating a registration token" (p. 166).
2. Create the .mst file, the .msi file, and the .cab files, as described in "Creating the transform file and extracting the installation packages" (p. 168).
3. Set up the Group Policy object as described in "Setting up the Group Policy object" (p. 169).

Generating a registration token

A registration token passes the identity of a user to the agent setup program, without storing the user credentials for the Cyber Protect console. This enables users to register any number of machines under their account without having to log in.

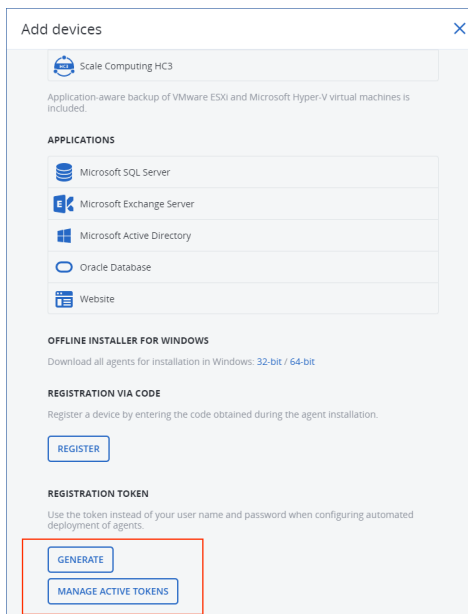
For security reasons, the tokens have limited lifetime, which you can adjust. The default lifetime is 3 days.

Users can generate registration tokens only for their own accounts. Administrators can generate registration tokens for all user accounts in the tenant that they manage.

To generate a registration token

As a user

1. Log in to the Cyber Protect console.
2. Click **Devices > All devices > Add**.
The **Add devices** pane opens on the right.
3. Scroll down to **Registration token**, and then click **Generate**.



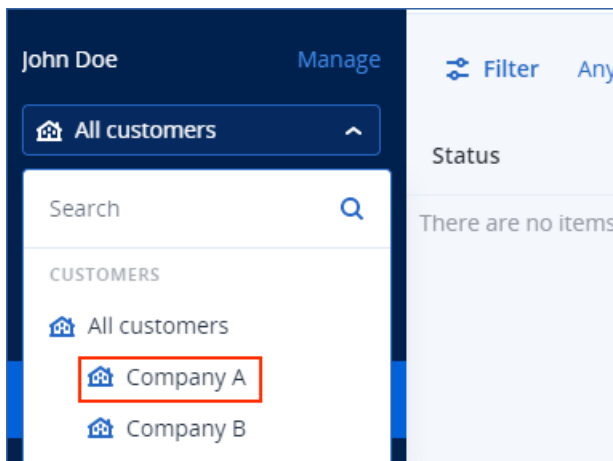
4. Specify the token lifetime.
5. Click **Generate token**.
6. Click **Copy** to copy the token to your device clipboard, or write the token down manually.

As an administrator

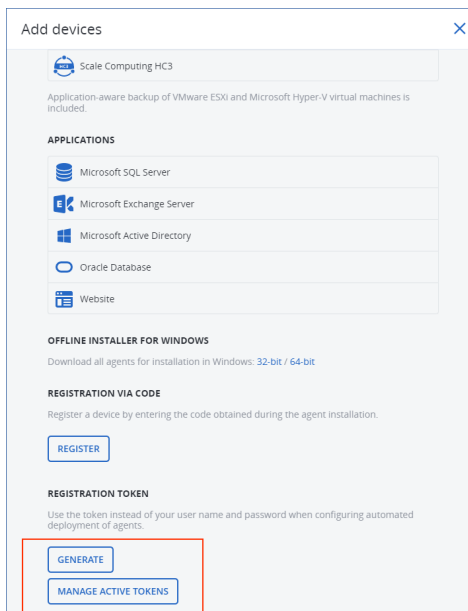
1. Log in to the Cyber Protect console as an administrator.
If you are already signed in to the management portal, you can go to the Cyber Protect console by navigating to **Monitoring > Usage**, and then, under the **Protection** tab, clicking **Manage service**.



[For partner administrators who manage customer tenants] In the Cyber Protect console, select the tenant with the user for whom you want to generate a token. You cannot generate a token on the **All customers** level.



2. Under **Devices**, click **All devices** > **Add**.
The **Add devices** pane opens on the right.
3. Scroll down to **Registration token**, and then click **Generate**.



4. Specify the token lifetime.

5. Select the user for whom you want to generate a token.

Note

When you use the token, workloads will be registered under the user account that you select here.

6. [Optional] To enable the user of the token to apply and revoke a protection plan on the added workloads, select the plan from the drop-down list.
Note that you will need to run a script that will apply or revoke a protection plan on the added workloads. Refer to [this knowledge base article](#) for more details.
7. Click **Generate token**.
8. Click **Copy** to copy the token to your device clipboard, or write the token down manually.

To view or delete registration tokens

1. Log in to the Cyber Protect console.
2. Click **Devices > All devices > Add**.
3. Scroll down to **Registration token**, and then click **Manage active tokens**.
A list with the active tokens that are generated for your tenant opens on the right.

Note

For security reasons, in the **Token** column, only the first two characters of the token value are shown.

4. [To delete a token] Select the token, and then click **Delete**.

Creating the transform file and extracting the installation packages

To deploy protection agents via Windows Group Policy, you need a transform file (.mst), and the installation packages (.msi and .cab files).

Note

The procedure below uses the default registration option, which is registration by token. To learn how to generate a registration token, refer to "Generating a registration token" (p. 166).

To create the .mst file and extract the installation packages (.msi and .cab files)

1. Log in as an administrator on any machine in the Active Directory domain.
2. Create a shared folder that will contain the installation packages. Ensure that domain users can access the shared folder—for example, by leaving the default sharing settings for **Everyone**.
3. Run the agent setup program.
4. Click **Create .mst and .msi files for unattended installation**.
5. In **What to install**, select the components that you want to include in the installation, and then click **Done**.
6. In **Registration settings**, click **Specify**, enter a registration token, and then click **Done**.

You can change the registration method from **Use registration token** (default) to **Use credentials** or **Skip registration**. The **Skip registration** option presumes that you will register the workloads manually later.

7. Review or modify the installation settings, which will be added to the .mst file, and then click **Proceed**.
8. In **Save the files to**, specify the path to the shared folder that you created.
9. Click **Generate**.

As a result, the .mst file, the .msi file, and the .cab files are created and copied to the shared folder that you specified.

Next, set up the Windows Group Policy object. To learn how to do it, refer to "Setting up the Group Policy object" (p. 169).

Setting up the Group Policy object

In this procedure you use the installation packages that you created in "Creating the transform file and extracting the installation packages" (p. 168) to set up a Group Policy object (GPO). The GPO will deploy the agents onto the machines in your domain.

To set up the Group Policy object

1. Log in to the domain controller as a domain administrator.
If the domain has more than one domain controller, log in to any of them as a domain administrator.
2. [If you deploy agents in an organizational unit] Ensure that the organizational unit in which you want to deploy the agents exists in this domain.
3. In the Windows **Start** menu, point to **Administrative Tools**, and then click **Group Policy Management** (or **Active Directory Users and Computers** for Windows Server 2003).
4. [For Windows Server 2008 or later] Right-click the name of the domain or organizational unit, and then click **Create a GPO in this domain, and Link it here**.
5. [For Windows Server 2003] Right-click the name of the domain or organizational unit, and then click **Properties**. In the dialog box, click the **Group Policy** tab, and then click **New**.
6. Name the new Group Policy object **Agent for Windows**.
7. Open the **Agent for Windows** Group Policy object for editing:
 - [In Windows Server 2008 or later] Under **Group Policy Objects**, right-click the Group Policy object, and then click **Edit**.
 - [In Windows Server 2003] Click the Group Policy object, and then click **Edit**.
8. In the Group Policy object editor snap-in, expand **Computer Configuration**.
9. [For Windows Server 2012 or later] Expand **Policies > Software Settings**.
10. [For Windows Server 2003 and Windows Server 2008] Expand **Software Settings**.
11. Right-click **Software installation**, point to **New**, and then click **Package**.
12. Select the agent's .msi installation package in the shared folder that you created, and then click **Open**.

13. In the **Deploy Software** dialog box, click **Advanced**, and then click **OK**.
14. On the **Modifications** tab, click **Add**, and then select the .mst file in the shared folder that you created.
15. Click **OK** to close the **Deploy Software** dialog box.

SSH connections to a virtual appliance

Use a Secure Socket Shell (SSH) connection when you remotely access a virtual appliance, for maintenance purposes.

Starting the Secure Shell daemon

To allow SSH connections to a virtual appliance, start the Secure Shell daemon (sshd) on the appliance.

To start the Secure Shell daemon

1. In the hypervisor software, open the console of the virtual appliance.
2. In the graphical user interface of the appliance, press CTRL+SHIFT+F2 to open the command-line interface.
3. Run the following command:

```
/bin/sshd
```

4. [Only during the first connection to the appliance] Set the password for the root user.
To learn how to set the password, see "Setting the root password on a virtual appliance" (p. 170).

Note

We recommend that you stop the Secure Shell daemon when you do not use the SSH connection.

Setting the root password on a virtual appliance

Before establishing an SSH connection to a virtual appliance for the first time, you must set the root password on the appliance.

To set the root password

1. In the hypervisor software, open the console of the virtual appliance.
2. In the graphical user interface of the appliance, press CTRL+SHIFT+F2 to open the command-line interface.
3. Run the following command:

```
passwd
```

4. Specify a password, and then press Enter.

The password must contain at least nine characters and must have complexity score of three or more. The complexity score is calculated automatically. To reach higher score, use a combination of special symbols, uppercase and lowercase symbols, and digits.

5. Confirm the password, and then press Enter.

Accessing a virtual appliance via an SSH client

Prerequisites

- An SSH client must be available on the remote machine. The procedure below uses the WinSCP client as an example. You can use any SSH client, by adapting the steps accordingly.
- The Secure Shell daemon (sshd) must be started on the virtual appliance. For more information, see "Starting the Secure Shell daemon" (p. 170).

To access a virtual appliance via WinSCP

1. On the remote machine, open WinSCP.
2. Click **Session > New Session**.
3. In **File protocol**, select **SCP**.
4. In **Host name**, specify the IP address of your virtual appliance.
5. In **User name** and **Password**, specify root and the password for the root user.
6. Click **Login**.

A list of all directories on the virtual appliance is shown.

Updating agents

You can update all agents manually either by using the Cyber Protect console or by downloading and running the installation file.

You can configure automatic updates for the following agents:

- Agent for Windows
- Agent for Linux
- Agent for Mac
- Cyber Files Cloud Agent for File Sync & Share

4.2 GB of free space in the following location is required to update an agent automatically, or manually by using the Cyber Protect console:

- For Linux – the root directory
- For Windows – the volume where the agent is installed

5 GB of free space is required to update an agent in macOS – in the root directory.

Note

[For all agents provided in the form of a virtual appliance, including Agent for VMware, Agent for Scale Computing, Agent for Virtuozzo Hybrid Infrastructure, Agent for RHV (oVirt)]

In order to perform automatic or manual update of a virtual appliance located behind a proxy, the proxy server must be configured on each appliance as follows.

In the `/opt/acronis/etc/va-updater/config.yaml` file, add the following line to the bottom of the file and enter the values specific to your environment:

```
httpProxy: http://proxy_login:proxy_password@proxy_address:port
```

Updating agents manually

You can update agents either by using the Cyber Protect console or by downloading and running the installation file.

Virtual appliances with the following versions must be updated only by using the Cyber Protect console:

- Agent for VMware (Virtual Appliance): version 12.5.23094 and later.
- Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance): version 12.5.23094 and later.

Agents with the following versions can also be updated by using the Cyber Protect console:

- Agent for Windows, Agent for VMware (Windows), Agent for Hyper-V: version 12.5.21670 and later.
- Agent for Linux: version 12.5.23094 and later.
- Other agents: version 12.5.23094 and later.

To find the agent version, in the Cyber Protect console, select the machine, and then click **Details**.

To update earlier agent versions of those agents, download and install the newest version manually.

To find the download links, click **All devices > Add**.

Prerequisites

On Windows machines, Cyber Protect features require Microsoft Visual C++ 2017 Redistributable. Ensure that it is already installed on your machine or install it before updating the agent. After the installation, a restart may be required. You can find the Microsoft Visual C++ Redistributable package on the Microsoft website: <https://support.microsoft.com/help/2999226/update-for-universal-c-runtime-in-windows>.

To update an agent by using the Cyber Protect console

1. Click **Settings > Agents**.

The software displays the list of machines. The machines with outdated agent versions are marked with an orange exclamation mark.

2. Select the machines that you want to update the agents on. The machines must be online.
3. Click **Update agent**.

Note

During the update, any backups that are in progress will fail.

To update Agent for VMware (Virtual Appliance) whose version is below 12.5.23094

1. Click **Settings > Agents >** the agent that you want to update **> Details**, and then examine the **Assigned virtual machines** section. You will need to re-enter these settings after the update.
 - a. Make note of the position of the **Automatic assignment** switch.
 - b. To find out what virtual machines are manually assigned to the agent, click the **Assigned:** link. The software displays the list of assigned virtual machines. Make note of the machines that have (M) after the agent name in the **Agent** column.
2. Remove Agent for VMware (Virtual Appliance), as described in "[Uninstalling agents](#)". In step 5, delete the agent from **Settings > Agents**, even though you are planning to install the agent again.
3. Deploy Agent for VMware (Virtual Appliance), as described in "[Deploying the OVF template](#)".
4. Configure Agent for VMware (Virtual Appliance), as described in "[Configuring the virtual appliance](#)".

If you want to reconstruct the locally attached storage, in step 7 do the following:

 - a. Add the disk containing the local storage to the virtual appliance.
 - b. Click **Refresh > Create storage > Mount**.
 - c. The software displays the original **Letter** and **Label** of the disk. Do not change them.
 - d. Click **OK**.
5. Click **Settings > Agents >** the agent that you want to update **> Details**, and then reconstruct the settings that you made note of in step 1. If some virtual machines were manually assigned to the agent, assign them again as described in "[Virtual machine binding](#)".

Once the agent configuration is completed, the protection plans that were applied to the old agent are re-applied automatically to the new agent.
6. The plans with application-aware backup enabled require the guest OS credentials to be re-entered. Edit these plans and re-enter the credentials.
7. The plans that back up ESXi configuration require the "root" password to be re-entered. Edit these plans and re-enter the password.

To update the Cyber Protection definitions on a machine

1. Click **Settings > Agents**.
2. Select the machine on which you want to update the Cyber Protection definitions and click **Update definitions**. The machine must be online.

To assign the Updater role to an agent

1. Click **Settings > Agents**.
2. Select the machine to which you want to assign the [Updater role](#), click **Details**, and then in the **Cyber Protection definitions** section, enable **Use this agent to download and distribute patches and updates**.

Note

An agent with the Updater role can download and distribute patches only for Windows third-party products. For Microsoft products, patch distribution is not supported by the Updater agent.

To clear cached data on an agent

1. Click **Settings > Agents**.
2. Select the machine on which you want to clear the cached data (outdated update files and patch management data) and click **Clear cache**.

Updating agents automatically

To facilitate management of multiple workloads, you can configure automatic updates for Agent for Windows, Agent for Linux, and Agent for Mac. Automatic updates are available for agents version 15.0.26986 (released in May 2021) or later. Older agents must be updated manually to the latest version, first.

Automatic updates are supported on machines running any of the following operating systems:

- Windows XP SP 3 and later
- Red Hat Enterprise Linux 6 and later, CentOS 6 and later
- OS X 10.9 Mavericks and later

The settings for automatic updates are preconfigured on a data center level. A company administrator can customize these settings – for all machines in a company or a unit, or for individual machines. If no custom settings are applied, then the settings from the upper level are used, in this order:

1. Cyber Protection data center
2. Company (customer tenant)
3. Unit
4. Machine

For example, a unit administrator can configure custom auto-update settings for all machines in the unit, which might differ from the setting applied to the machines on the company level. The administrator can also configure different settings for one or more individual machines in the unit, to which neither the unit settings nor the company settings will be applied.

After enabling the automatic updates, you can configure the following options:

- **Update channel**

The update channel defines which version of the agents will be used – the most up-to-date one or the latest version from the previous release.

- **Maintenance window**

The maintenance window defines when updates can be installed. If the maintenance window is disabled, updates can run anytime.

Even within the enabled maintenance window, updates will not be installed while the agent is running any of the following operations:

- Backup
- Recovery
- Backup replication
- Virtual machine replication
- Testing a replica
- Running a virtual machine from backup (including finalization)
- Disaster recovery failover
- Disaster recovery fallback
- Running a script (for Cyber Scripting functionality)
- Patch installation
- ESXi configuration backup

To customize auto-update settings

1. In the Cyber Protect console, go to **Settings > Agents**.
2. Select the scope for the settings:
 - To change the settings for all machines, click **Edit default agent update settings**.
 - To change the settings for specific machines, select the desired machines, and then click **Agent update settings**.
3. Configure the settings according to your needs, and then click **Apply**.

To remove the custom auto-update settings

1. In the Cyber Protect console, go to **Settings > Agents**.
2. Select the scope for the settings:
 - To remove the custom settings for all machines, click **Edit default agent update settings**.
 - To remove the custom settings for specific machines, select the desired machines, and then click **Agent update settings**.
3. Click **Reset to default settings**, and then click **Apply**.

To check the auto-update status

1. In the Cyber Protect console, go to **Settings > Agents**.
2. Click the gear icon in the upper right corner of the table, and then ensure that **Auto-update** check box is selected.
3. Check the status that is shown in the **Auto-update** column.

Updating agents on BitLocker-protected workloads

Agent updates that introduce changes to Startup Recovery Manager interfere with BitLocker on workloads on which both BitLocker and Startup Recovery Manager are enabled. In this case, after a restart, the BitLocker recovery key is required. To mitigate this issue, suspend or disable BitLocker before you update the agent.

Affected agent versions:

- 23.12.36943, released in December 2023

You can also check whether an update introduces changes to Startup Recovery Manager in the release notes of the protection agent.

To update the agent on a workload with BitLocker and Startup Recovery Manager enabled

1. On the workload on which you want to update the agent, suspend or disable BitLocker.
2. Update the agent.
3. Restart the workload.
4. Enable BitLocker.

Preventing unauthorized uninstallation or modification of agents

You can protect Agent for Windows against unauthorized uninstallation or modification, by enabling the **Password protection** setting in a protection plan. This setting is available only when the **Self-protection** setting is enabled.

To enable Password protection

1. In a protection plan, expand the **Antivirus & Antimalware protection** module (**Active Protection** module for Cyber Backup editions).
2. Click **Self-protection** and ensure that the **Self-protection** switch is enabled.
3. Enable the **Password protection** switch.
4. In the window that opens, copy the password that you need to uninstall or modify the components of a protected Agent for Windows.

This password is unique and you will not be able to recover it once you close this window. If you lose or forget this password, you can edit the protection plan and create a new password.

5. Click **Close**.
6. In the **Self-protection** pane, click **Done**.
7. Save the protection plan.

Password protection will be enabled for the machines to which this protection plan is applied. Password protection is only available for Agent for Windows version 15.0.25851 or newer. The machines must be online.

You can apply a protection plan with Password protection enabled to a machine running macOS, but no protection will be provided. You cannot apply such a plan to a machine running Linux.

Also, you cannot apply more than one protection plan with Password protection enabled to the same Windows machine. To learn how to resolve a possible conflict, refer to [Resolving plan conflicts](#).

To change the password in an existing protection plan

1. In the protection plan, expand the **Antivirus & Antimalware protection** module (**Active Protection** module for Cyber Backup edition).
2. Click **Self-protection**.
3. Click **Create new password**.
4. In the window that opens, copy the password that you need to uninstall or modify the components of a protected Agent for Windows.
This password is unique and you will not be able to recover it once you close this window. If you lose or forget this password, you can edit the protection plan and create a new password.
5. Click **Close**.
6. In the **Self-protection** pane, click **Done**.
7. Save the protection plan.

Uninstalling agents

When you uninstall an agent from a workload, the workload is automatically removed from the Cyber Protect console. If the workload is still shown after you uninstall the agent, for example, due to a network problem, manually remove this workload from the console. For more information about how to do it, refer to "Removing workloads from the Cyber Protect console" (p. 315).

Note

Uninstalling an agent does not delete any plans or backups.

To uninstall an agent

Windows

1. Sign in as an administrator to the machine with the agent .
2. In **Control panel**, go to **Programs and Features (Add or Remove Programs** in Windows XP).
3. Right-click **Acronis Cyber Protect**, and then select **Uninstall**.
4. [For password-protected agents] Specify the password that is required to uninstall the agent, and then click **Next**.
5. [Optional] Select the **Remove the logs and configuration settings** check box.

If you are planning to install the agent again, keep this check box cleared. If you select the check box and then install the agent again, this workload might be duplicated in the Cyber Protect console and its old backups might not be associated with it.

6. Click **Uninstall**.

Linux

1. On the machine with the agent, run `/usr/lib/Acronis/BackupAndRecovery/uninstall/uninstall` as the root user.
2. [Optional] Select the **Clean up all product traces (Remove the product's logs, tasks, vaults, and configuration settings)** check box.

If you are planning to install the agent again, keep this check box cleared. If you select the check box and then install the agent again, this workload might be duplicated in the Cyber Protect console and its old backups might not be associated with it.

3. Confirm your decision.

macOS

1. On the machine with the agent, double-click the installation .dmg file.
2. Wait until the operating system mounts the installation disk image.
3. Inside the image, double-click **Uninstall**.
4. If prompted, provide administrator credentials.
5. Confirm your decision.

To uninstall components that are bundled with Agent for Windows

You can uninstall individual components that are bundled with Agent for Windows, such as Cyber Protect Monitor, Agent for Data Loss Prevention, or Bootable Media Builder, without uninstalling Agent for Windows.

1. Sign in as an administrator to the machine with the agent.
2. Run the setup program, and then click **Modify installed components**.
3. Clear the check boxes next to the components that you want to uninstall, and then click **Done**.

To remove Agent for VMware (Virtual appliance)

1. By using the vSphere Client, log in to vCenter Server.
2. [If the virtual appliance is powered on] Right-click the virtual appliance, and then click **Power > Power Off**. Confirm your decision.
3. [If the virtual appliance uses a locally attached storage on a virtual disk and you want to preserve data on that disk] Remove the virtual storage from the virtual appliance.
 - a. Right-click the virtual appliance, and then click **Edit Settings**.
 - b. Select the disk with the storage, and then click **Remove**.
 - c. Under **Removal Options**, click **Remove from virtual machine**.
 - d. Click **OK**.

As a result, the disk remains in the datastore. You can attach the disk to another virtual appliance.

4. Right-click the virtual appliance, and then click **Delete from Disk**. Confirm your decision.
5. [Optional] [If you are not planning to use this appliance again] In the Cyber Protect console, go to **Backup storage > Locations**, and then delete the location corresponding to the locally attached storage.

Protection settings

To configure the general protection settings for Cyber Protection, in the Cyber Protect console, go to **Settings > Protection**.

Automatic updates for components

By default, all agents can connect to the Internet and download updates.

An administrator can minimize the network bandwidth traffic by selecting one or several agents in the environment and assigning the Updater role to them. Thus, the dedicated agents will connect to the Internet and download updates. All other agents will connect to the dedicated updater agents by using peer-to-peer technology, and then download the updates from them.

The agents without the Updater role will connect to the Internet if there is no dedicated updater agent in the environment, or if the connection to a dedicated updater agent cannot be established for about five minutes.

The updater agent distributes updates and patches for Antivirus and Antimalware protection, Vulnerability assessment, and Patch management, but does not include updates of the agent version.

Note

An agent with the Updater role can download and distribute patches only for Windows third-party products. For Microsoft products, patch distribution is not supported by the Updater agent.

Before assigning the Updater role to an agent, ensure that the machine on which the agent runs is powerful enough, and has a stable high-speed Internet connection and enough disk space.

To prepare a machine for the Updater role

1. On agent machine where you plan to enable the Updater role, apply the following firewall rules:
 - Inbound (incoming) "updater_incoming_tcp_ports": allow connection to TCP ports 18018 and 6888 for all firewall profiles (public, private, and domain).
 - Inbound (incoming) "updater_incoming_udp_ports": allow connection to UDP port 6888 for all firewall profiles (public, private, and domain).
2. Restart the Acronis Agent Core Service.
3. Restart the Firewall Service.

If you do not apply these rules and the firewall is enabled, peer agents will download the updates from the Cloud.

To assign the Updater role to a protection agent

1. In the Cyber Protect console, go to **Settings > Agents**.
2. Select the machine with the agent to which you want to assign the Updater role.
3. Click **Details**, and then enable the **Use this agent to download and distribute patches and updates** switch.

The peer-to-peer update works as follows.

1. The agent with the Updater role checks by schedule the index file provided by the service provider to update the core components.
2. The agent with the Updater role starts to download and distribute updates to all agents.

You can assign the Updater role to multiple agents in the environment. Thus, if an agent with the Updater role is offline, other agents with this role can serve as the source for definition updates.

Updating the Cyber Protection definitions by schedule

On the **Schedule** tab, you can set up the schedule for automatic update of the Cyber Protection definitions for each of the following components:

- Antimalware
- Vulnerability assessment
- Patch management

To change the definition updates setting, navigate to **Settings > Protection > Protection definitions update > Schedule**.

Schedule type:

- **Daily** – define on which days of the week to update definitions.
Start at – select at what time to update definitions.
- **Hourly** – define more granular hourly schedule for updates.
Run every – define the periodicity of updates.
From ... To – define a specific time range for the updates.

Updating the Cyber Protection definitions on-demand

To update the Cyber Protection definitions for a particular machine on-demand

1. In the Cyber Protect console, go to **Settings > Agents**.
2. Select the machines on which you want to update the protection definitions, and then click **Update definitions**.

Cache storage

The location of cached data is the following:

- On Windows machines: C:\ProgramData\Acronis\Agent\var\atp-downloader\Cache
- On Linux machines: /opt/acronis/var/atp-downloader/Cache
- On macOS machines: /Library/Application Support/Acronis/Agent/var/atp-downloader/Cache

To change the cache storage setting, navigate to **Settings > Protection > Protection definitions update > Cache Storage**.

In **Outdated update files and patch management data**, specify after what period to remove cached data.

Maximum cache storage size (GB) for agents:

- **Updater role** – define storage size for cache on the machines with the Updater role.
- **Other roles** – define storage size for cache on other machines.

Note

Cyber Protection collects samples of detected malware for additional analysis so that we can improve our software. You can change this setting at any time in the **Protection** tab, by disabling the **Collect and upload malware samples to CPOC** toggle.

Changing the service quota of machines

A service quota is automatically assigned when a protection plan is applied to a machine for the first time.

The most appropriate quota is assigned, depending on the type of the protected machine, its operating system, required level of protection, and the quota availability. If the most appropriate quota is not available in your organization, the second-best quota is assigned. For example, if the most appropriate quota is **Web Hosting Server** but it is not available, the **Server** quota is assigned.

Examples of quota assignment:

- A physical machine that runs a Windows Server or a Linux server operating system (such as Ubuntu Server) is assigned the **Server** quota.
- A physical machine that runs a Windows or a Linux desktop operating system (such as Ubuntu Desktop) is assigned the **Workstation** quota.
- A physical machine that runs Windows 10 with enabled Hyper-V role is assigned the **Workstation** quota.
- A desktop machine that runs on a virtual desktop infrastructure and whose protection agent is installed inside the guest operating system (for example, Agent for Windows), is assigned the **Virtual machine** quota. This type of machine can also use the **Workstation** quota if the **Virtual machine** quota is not available.
- A desktop machine that runs on a virtual desktop infrastructure and which is backed up in the agentless mode (for example, by Agent for VMware or Agent for Hyper-V), is assigned the **Virtual machine** quota.
- A Hyper-V or vSphere server is assigned the **Server** quota.

- A server with cPanel or Plesk is assigned the **Web Hosting Server** quota. It can also use the **Virtual machine** or the **Server** quota, depending on the type of machine on which the web server runs, if the **Web Hosting Server** quota is not available.
- The application-aware backup requires the **Server** quota, even for a workstation.

You can manually change the original assignment later. For example, to apply a more advanced protection plan to the same machine, you might need to upgrade the machine's service quota. If the features required by this protection plan are not supported by the currently assigned service quota, the protection plan will fail.

Alternatively, you can change the service quota if you purchase a more appropriate quota after the original one is assigned. For example, the **Workstation** quota is assigned to a virtual machine. After you purchase a **Virtual machines** quota, you can manually assign this quota to the machine, instead of the original **Workstation** quota.

You can also release the currently assigned service quota, and then assign this quota to another machine.

You can change the service quota of an individual machine or for a group of machines.

To change the service quota of an individual machine

1. In the Cyber Protect console, go to **Devices**.
2. Select the desired machine, and then click **Details**.
3. In the **Service quota** section, click **Change**.
4. In the **Change quota** window, select the desired service quota or **No quota**, and then click **Change**.

To change the service quota for a group of machines

1. In the Cyber Protect console, go to **Devices**.
2. Select more than one machine, and then click **Assign quota**.
3. In the **Change quota** window, select the desired service quota or **No quota**, and then click **Change**.

Cyber Protection services installed in your environment

Cyber Protection installs some or all of the following services, depending on the Cyber Protection options that you use.

Services installed in Windows

Service name	Purpose
Acronis Managed Machine Service	Provides backup, recovery, replication, retention, validation functionality
Acronis Scheduler2 Service	Executes scheduled tasks on certain events

Acronis Active Protection Service	Provides protection against ransomware
Acronis Cyber Protection Service	Provides antimalware protection

Services installed in macOS

Service name and location	Purpose
/Library/LaunchDaemons/com.acronis.aakore.plist	Serves for communication between the agent and management components
/Library/LaunchDaemons/com.acronis.cyber-protect-service.plist	Provides detection of malware
/Library/LaunchDaemons/com.acronis.mms.plist	Provides backup and recovery functionality
/Library/LaunchDaemons/com.acronis.schedule.plist	Executes scheduled tasks

Saving an agent log file

You can save an agent log to a .zip file. If a backup fails for an unknown reason, this file will help the technical support personnel to identify the problem.

By default, the information in the log is optimized for the last three days, but you can change this period.

To collect agent logs

- Do one of the following:
 - Under **Devices**, select the machine from which you want to collect the logs, and then click **Activities**.
 - Under **Settings > Agents**, select the machine from which you want to collect the logs, and then click **Details**.
- [Optional] To change the default period for which system information is included, click the arrow next to the **Collect system information** button, and then select the period.
- Click **Collect system information**.
- If prompted by your web browser, specify where to save the file.

Site-to-site Open VPN - Additional information

When you create a recovery server, you configure its **IP address in production network**, and its **Test IP address**.

After you perform failover (run the virtual machine in the cloud), and log in to the virtual machine to check the IP address of the server, you see the **IP address in production network**.

When you perform test failover, you can reach the test server only by using the **Test IP address**, which is visible only in the configuration of the recovery server.

To reach a test server from your local site, you must use the **Test IP address**.

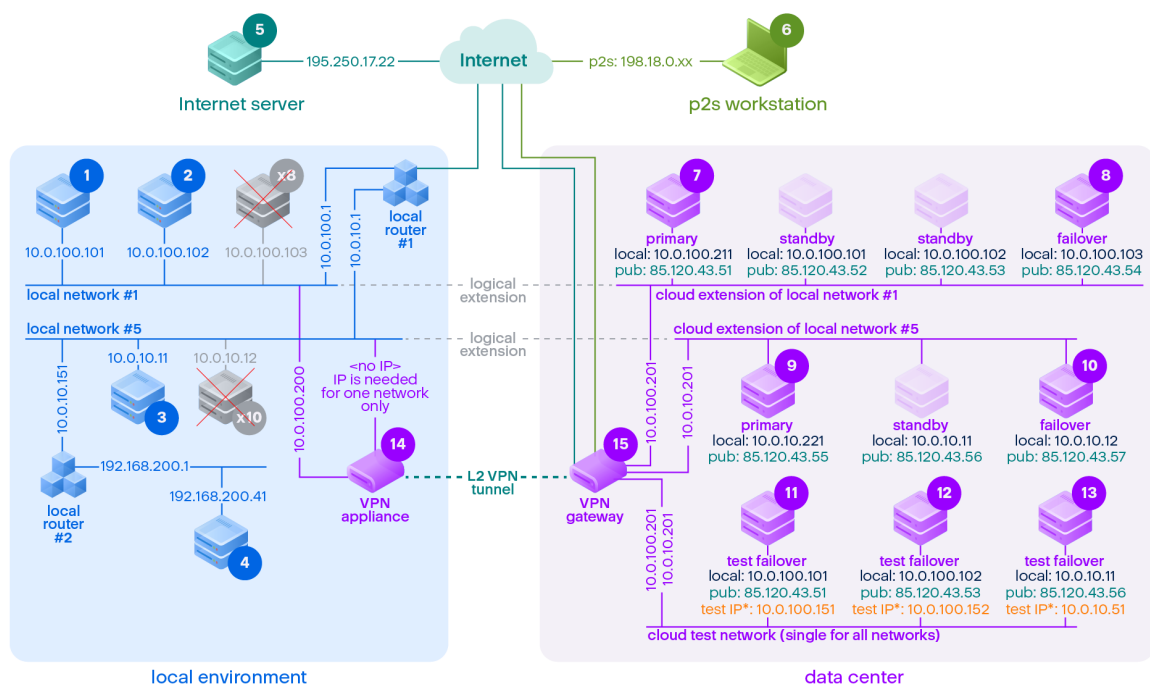
Note

The network configuration of the server always shows the **IP address in production network** (as the test server mirrors how the production server would look). This happens because the test IP address does not belong to the test server, but to the VPN gateway, and is translated to the production IP address using NAT.

The diagram below shows an example of the Site-to-site Open VPN configuration. Some of the servers in the local environment are recovered to the cloud using failover (while the network infrastructure is ok).

- The customer enabled Disaster Recovery by:
 - configuring the VPN appliance (14), and connected it to the dedicated cloud VPN server (15)
 - protecting some of the local servers with Disaster Recovery (1, 2, 3, x8, and x10)

Some servers on the local site (like 4) are connected to networks which are not connected to the VPN appliance. Such servers are not protected with Disaster Recovery.
- Part of the servers (connected to different networks) work in the local site: (1, 2, 3, and 4)
- The protected servers (1, 2, and 3) are being tested with test failover (11, 12, and 13)
- Some servers in the local site are unavailable (x8, x10). After performing failover, they become available in the cloud (8, and 10)
- Some primary servers (7, and 9), connected to different networks, are available in the cloud environment
- (5) is a server in the Internet with a public IP address
- (6) is a workstation connected to the cloud using a Point-to-site VPN connection (p2s)



*The test IP belongs to the VPN gateway and is NATed to the recovery server. The recovery server has the production IP assigned to it.

In this example, the following connection setup is available (for example, "ping") from a server in the **From:** row to a server in the **To:** column.

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
From:		local	local	local	local	internet	primary	primary	failover	primary	failover	test failover	test failover	test failover	VPN appliance	VPN server
1	local		direct	via local router 1	via local router 2	via local router 1 and Internet	no	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: NAT (VPN server) via local router 1 and Internet: pub	via tunnel: NAT (VPN server) via local router 1 and Internet: pub	via local router 1 and Internet: pub	direct	no
2	local	direct		via local router 1	via local router 2	via local router 1 and Internet	no	via tunnel: local via local router 1 and Internet	via tunnel: local via local router 1 and Internet	via tunnel: local via local router 1 and Internet	via tunnel: local via local router 1 and Internet	via tunnel: NAT (VPN server) via local router 1	via tunnel: NAT (VPN server) via local router 1	via local router 1 and Internet: pub	direct	no

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
								t: pub	t: pub	t: pub	t: pub	and Inte rne t: pub	and Inte rne t: pub	via loca l rout er 1 and Inte rne t: pub		
3	local	via loc al rout er 1	via loc al rout er 1		via loc al rout er 2	via local route r 1 and Inter net	n o	via tun nel: loca l via loca l rout er 1 and Inte rne t: pub	via tun nel: loca l via loca l rout er 1 and Inte rne t: pub	via tun nel: loca l via loca l rout er 1 and Inte rne t: pub	via tun nel: loca l via loca l rout er 1 and Inte rne t: pub	via tun nel: NAT (VP N serv er) via loca l rout er 1 and Inte rne t: pub	via tun nel: NAT (VP N serv er) via loca l rout er 1 and Inte rne t: pub	via loca l rout er 1 and tun nel: NAT (VP N serv er) via loca l rout er 1 and Inte rne t: pub	via local rout er	no
4	local	via loc al rout er 2 and ro	via loc al rout er 2 and ro	via loc al rout er 2		via local route r 2, and route r 1, and Inter net	n o	via loca l rout er 2 and tun nel: loca l	via loca l rout er 2 and tun nel: loca l	via loca l rout er 2 and tun nel: loca l	via loca l rout er 2 and tun nel: loca l	via tun nel: NAT (VP N serv er) via	via tun nel: NAT (VP N serv er) via	via tun nel: NAT (VP N serv er) via	via local rout er 2	no

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
		ut er 1	ut er 1					via loca l rout er 2, and loca l rout er 1, and Inte rne t: pub	via loca l rout er 2, and loca l rout er 1, and Inte rne t: pub	via loca l rout er 2, and loca l rout er 1, and Inte rne t: pub	via loca l rout er 2, and loca l rout er 1, and Inte rne t: pub	loca l rout er 2, and rout er 1, and Inte rne t: pub	loca l rout er 2, and rout er 1, and Inte rne t: pub	loca l rout er 2, and rout er 1, and Inte rne t: pub		
5	inter net	no	no	no	no		n / a	via Inte rne t: pub	via Inte rne t: pub	via Inte rne t: pub	via Inte rne t: pub	via Inte rne t: pub	via Inte rne t: pub	via Inte rne t: pub	no	no
6	p2s	no	no	no	no	via Inter net		via p2s VPN (VP N serv er): loca l via Inte rne t: pub	via p2s VPN (VP N serv er): loca l via Inte rne t: pub	via p2s VPN (VP N serv er): loca l via Inte rne t: pub	via p2s VPN (VP N serv er): loca l via Inte rne t: pub	via p2s VPN - NAT (VP N serv er) via Inte rne t: pub	via p2s VPN - NAT (VP N serv er) via Inte rne t: pub	via p2s VPN - NAT (VP N serv er) via Inte rne t: pub	no	no
7	prim ary	via tun nel	via tun nel	via tun nel and loca l ro	via tun nel and loca l ro	via Inter net (via VPN serve r)	no		dire ct in cloud: loca l rout er 1: loca	via tun nel and loca l rout er 1: loca	via tun nel and loca l rout er 1: loca	via VPN serv er: NAT	via VPN serv er: NAT	via tun nel and loca l rout er 1: NAT	no	DHC P and DNS prot ocol s only

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
				ut er 1	ut er 1 an d 2											
8	failover	via tun nel	via tun nel	via tun nel an d loc al ro uter 1	via tun nel an d loc al ro uter 1 an d 2	via Inter net (via VPN serve r)	n o	dire ct in clou d: loc al		via tun nel and loc al ro uter 1: loc al	via tun nel and loc al ro uter 1: loc al	via VPN serv er: NAT	via VPN serv er: NAT	via tun nel and loc al ro uter 1: NAT	no	DHC P and DNS prot ocol s only
9	primary	via tun nel an d loc al ro uter 1	via tun nel an d loc al ro uter 1	via tun nel	via tun nel	via Inter net (via VPN serve r)	n o	via tun nel and loc al ro uter 1: loc al	via tun nel and loc al ro uter 1: loc al		dire ct in clou d: loc al	via tun nel and loc al ro uter 1: NAT	via tun nel and loc al ro uter 1: NAT	via VPN serv er: NAT	no	DHC P and DNS prot ocol s only
10	failover	via tun nel an d loc al ro uter 1	via tun nel an d loc al ro uter 1	via tun nel	via tun nel	via Inter net (via VPN serve r)	n o	via tun nel and loc al ro uter 1: loc al	via tun nel and loc al ro uter 1: loc al	dire ct in clou d: loc al		via tun nel and loc al ro uter 1: NAT	via tun nel and loc al ro uter 1: NAT	via VPN serv er: NAT	no	DHC P and DNS prot ocol s only

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
11	test failover	no	no	no	no	via Internet (via VPN server)	no	no	no	no	no		direct in cloud: local	via VPN server: local (routing)	no	DHCP and DNS protocols only
12	test failover	no	no	no	no	via Internet (via VPN server)	no	no	no	no	no	direct in cloud: local		via VPN server: local (routing)	no	DHCP and DNS protocols only
13	test failover	no	no	no	no	via Internet (via VPN server)	no	no	no	no	no	via VPN server: local (routing)	via VPN server: local (routing)		no	DHCP and DNS protocols only
14	VPN appliance	direct	direct	via local router 1	via local router 2	via Internet (local router 1)	no	no	no	no	no	no	no	no		no
15	VPN server	no	no	no	no	no	no	no	no	no	no	no	no	no	no	

License management for on-premises management servers

For detailed information about how to activate an on-premises management server or how to allocate licenses to it, refer to the [Licensing section in the Cyber Protect user guide](#).

Defining how and what to protect

The Management tab

Note

The availability of this feature depends on the service quotas that are enabled for your account.

All plans that you created are available on the **Management** tab of the Cyber Protect console.

The following sections are available:

- [Protection plans](#)
- [Remote management plans](#)
- [Scripting plans](#)
- [Monitoring plans](#)
- [Script repository](#)
- [Cloud applications backup](#)
- [Backup scanning](#)
- [Backup replication](#)
- [Validation](#)
- [Cleanup](#)
- [Conversion to VM](#)
- [VM replication](#)

Plan statuses

For protection plans and VM replication plans, a status bar shows the following color-coded statuses:

- OK (Green)
- Warning (Orange)
- Error (Dark orange)
- Critical (Red)
- The plan is running (Blue)
- The plan is disabled (Gray)

Click the status bar to see details about the plan statuses on all workloads to which the plan is applied.

Click a specific status to see a list of all workloads with this status.

Protection plans

On the **Management > Protection plans** tab, you can see information about your existing protection plans, perform actions with them, and create new plans.

For more information about the protection plans, refer to "Protection plans and modules" (p. 208).

Backup plans for cloud applications

The **Management > Cloud applications backup** tab shows cloud-to-cloud backup plans. These plans back up applications running in the cloud by means of agents that run in the cloud and use the cloud storage as a backup location.

In this section, you can perform the following operations:

- Create, view, run, stop, edit, and delete a backup plan
- View activities related to each backup plan
- View alerts related to each backup plan

For more information about cloud applications backup, refer to:

- [Protecting Microsoft 365 data](#)
- [Protecting Google Workspace data](#)

Running cloud-to-cloud backups manually

To prevent disrupting the Cyber Protection service, the number of manual cloud-to-cloud backup runs is limited to 10 runs per Microsoft 365 or Google Workspace organization during an hour. After this number has been reached, the number of runs allowed is reset to one per hour, and then an additional run becomes available each hour thereafter (e.g. hour 1, 10 runs; hour 2, 1 run; hour 3, 2 runs) until a total of 10 runs per hour is reached.

Backup plans applied to groups of devices (mailboxes, drives, sites) or containing more than 10 devices cannot be run manually.

Backup scanning plans

To scan backups for malware (including ransomware), create a backup scanning plan.

Important

Backup scanning plans are not supported for all workloads and backup storages. For details, refer to "Limitations" (p. 802).

To create a backup scanning plan

1. In the Cyber Protect console, go to **Management > Backup scanning**.
2. Click **Create plan**.
3. Specify the name of the plan and the following parameters:

- **Scan type:**
 - **Cloud** – this option cannot be changed. An automatically selected cloud agent will perform the backup scan.
- **Backups to scan:**
 - **Locations** – select locations with backup sets that you want to scan.
 - **Backups** – select backup sets that you want to scan.
- **Scan for:**
 - **Malware** – this option cannot be changed. The scan checks the selected backup sets for malware (including ransomware).
- **Encryption** – to scan encrypted backup sets, specify the encryption password. If you select a location or multiple backup sets and the specified password does not match a backup set, an alert is created.
- **Schedule** – this option cannot be changed. In the cloud storage, the scan starts automatically.

4. Click **Create**.

As a result, a backup scanning plan is created and a cloud agent will scan for malware the locations or the backup sets that you specified.

Off-host data processing

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

Replication, validation, and cleanup are usually performed by the protection agent that performs the backup. This puts additional load on the machine on which the agent is running, even after the backup process is complete. To offload the machine, you can create off-host data protection plans – that is, separate plans for replication, validation, cleanup, and conversion to a virtual machine.

With the off-host data protection plans, you can do the following:

- Choose different agents for the backup and off-host data protection operations
- Schedule the off-host data processing operations during off-peak hours to minimize the network bandwidth consumption
- Schedule the off-host data processing operations during non-business hours, if you do not want to install a dedicated agent for off-host data processing

Note

The off-host data processing plans run according to the time settings (including the time zone) of the machine on which the protection agent is installed. For a virtual appliance (for example, Agent for VMware or Agent for Scale Computing HC3), you can configure the time zone in the graphical user interface of the agent.

Backup replication

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

Backup replication is copying a backup to another location. As an off-host data processing operation, it is configured in a backup replication plan.

Backup replication can also be part of a protection plan. For more information about this option, refer to "Replication" (p. 416).

Creating a backup replication plan

To replicate backups as an off-host data processing operation, you create a backup replication plan.

To create a backup replication plan

1. In the Cyber Protect console, click **Management > Backup replication**.
2. Click **Create plan**.
3. In **Agent**, select the agent that will perform the replication.
You can select any agent that has access both to the source location and the replication locations.
4. In **Items to replicate**, select the archives or backup locations to replicate.
To switch between archives and locations, use the **Locations / Backups** switch in the upper-right corner.
If you select multiple encrypted archives, their encryption password must be the same. For archives that use different encryption passwords, create separate plans.
5. In **Destination**, specify the replication location.
6. In **How to replicate**, select which backups (also known as recovery points) to replicate.
The following options are available:
 - **All backups**
 - **Only full backups**
 - **Only the last backup**For more information about these options, refer to "What to replicate" (p. 195).
7. In **Schedule**, configure the replication schedule.
When configuring the schedule of the backup replication plan, ensure that the last replicated backup will still be available in its original location when the backup replication starts. If this backup is not available in the original location, for example, because it was deleted by a retention rule, the whole archive will be replicated as a full backup. This might be very time-consuming and will use additional storage space.
8. In **Retention rules**, specify the retention rules for the target location.

The following options are available:

- **By number of backups**
- **By backup age** (separate settings for monthly, weekly, daily, and hourly backups)
- **By total size of backups**
- **Keep backups indefinitely**

Note

Selecting this option will result in increased storage usage. You must delete the unnecessary backups manually.

9. [If you selected encrypted archives in **Items to replicate**] Enable the **Backup password** switch, and then provide the encryption password.
10. [Optional] To modify the plan options, click the gear icon, and then configure the options as required.
11. Click **Create**.

What to replicate

Note

Some replication operations, such as replicating a whole location or replicating all backups in a backup set, might be very time-consuming.

You can replicate individual backup sets or whole backup locations. When you replicate a backup location, all backup sets in it are replicated.

Backup sets consist of backups (also known as recovery points). You must select which backups to replicate.

The following options are available:

- **All backups**
All backups in the backup set are replicated every time the replication plan runs.
- **Only full backups**
Only the full backups in the backup set are replicated.
- **Only the last backup**
Only the newest backup in the backup set is replicated, regardless of its type (full, differential, or incremental).

Select an option according to your needs and the backup scheme that you use. For example, if you use the **Always incremental (single-file)** backup scheme and you want to replicate only the newest incremental backup, in the backup replication plan, select **Only last backup**.

The following table summarizes which backups will be replicated with different backup schemes.

	Always incremental (single-file)	Always full	Weekly full, Daily incremental	Monthly full, Weekly differential, Daily incremental (GFS)
All backups	All backups in the backup set	All backups in the backup set	All backups in the backup set	All backups in the backup set
Only full backups	Only the first backup, which is full	All backups	One backup every week*	One backup every month*
Only last backup	Only the newest backup in the backup set*	Only the newest backup in the backup set*	Only the newest in the backup set, regardless of its type*	Only the newest in the backup set, regardless of its type*

* When configuring the schedule of the backup replication plan, ensure that the last replicated backup will still be available in its original location when the backup replication starts. If this backup is not available in the original location, for example, because it was deleted by a retention rule, the whole archive will be replicated as a full backup. This might be very time-consuming and will use additional storage space.

Supported locations

The following table summarizes backup locations supported by backup replication plans.

Backup location	Supported as a source	Supported as a target
Cloud storage	+	+
Local folder	+	+
Network folder	+	+
Public cloud	+	+
NFS folder	-	-
Secure Zone	-	-

Validation

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

By validating a backup, you verify that you can recover the data from it.

To validate a backup as an off-host data processing operation, you create a validation plan. For more information about how to create one, refer to "Creating a validation plan" (p. 198).

The following validation methods are available:

- Checksum verification
- Run as virtual machine
 - VM heartbeat
 - Screenshot validation

You can select one or more of these methods. When more than one method is selected, the operations for every validation method run consecutively. For more information about the methods, refer to "VM heartbeat" (p. 201).

You can validate backup sets or backup locations. Validation of a backup location validates all backup sets in it.

Supported locations

The following table shows the supported backup locations and validation methods.

Note

The validation option is not available for public cloud backups due to the prohibitive costs of reading an entire archive from a public cloud.

Backup location	Checksum verification	Run as virtual machine	
		VM heartbeat	Screenshot validation
Cloud storage	+	+	+
Local folder	+	+	+
Network folder	+	+	+
NFS folder	-	-	-
Secure Zone	-	-	-

Validation status

After a successful validation, the backup is marked with a green dot and the label **Validated**.

If the validation fails, the backup is marked with a red dot. The validation fails even when only one of the used validation methods fails. In some cases, this might be the result of a misconfiguration of the validation plan – for example, using the **VM heartbeat** method for virtual machines on a wrong host.

The validation status of a backup is updated with every new validation operation. The status for each validation method is updated separately. That is why the validation of a backup in which one method failed, will be shown as failed until the same validation method succeeds, even if the latest validation operations do not use the failed method and complete successfully.

For more information about how to check the validation status, refer to "Checking the validation status of a backup" (p. 203).

Creating a validation plan

To validate a backup set as an off-host data processing operation, you create a validation plan.

To create a validation plan

1. In the Cyber Protect console, click **Management > Validation**.
2. Click **Create plan**.
The template for a new validation plan opens.
3. [Optional] To modify the plan name, click the default name.
4. In **Agent**, select the agent that will perform the validation, and then click **OK**.
If you want to perform validation by running a virtual machine from a backup, select a machine with Agent for VMware or Agent for Hyper-V. Otherwise, select any machine that has access to the backup location.
5. In **Items to validate**, select the backups sets that you want to validate.
 - a. Select the scope for the plan – individual backup sets or entire locations, by clicking **Locations** or **Backups** in the upper-right corner.
If the selected backups are encrypted, all of them must use the same encryption password. For backups that use different encryption passwords, create separate plans.
 - b. Click **Add**.
 - c. Depending on the scope of the validation plan, select locations or a location and backup sets, and then click **Done**.
 - d. Click **Done**.
6. In **What to validate**, select which backups (also known as recovery points) within the selected backup sets to validate. The following options are available:
 - **All backups**
 - **Only the last backup**
7. In **How to validate**, select the validation method.
You can select one or both of the following options:
 - **Checksum verification**
 - **Run as a virtual machine**For more information about the methods, refer to "VM heartbeat" (p. 201).
8. [If you selected **Checksum verification**] Click **Done**.
9. [If you selected **Run as a virtual machine**]. Configure the settings for this method.

- a. In **Target machine**, select the virtual machine type (ESXi or Hyper-V), the host, and the machine name template, and then click **OK**.
The default name is **[Machine Name]_validate**.
 - b. In **Datastore** (for ESXi) or **Path** (for Hyper-V), select the datastore for the virtual machine.
 - c. Select one or both of the validation methods that **Run as virtual machine** provides:
 - **VM heartbeat**
 - **Screenshot validation**
 - d. [Optional] Click **VM settings** to change the memory size and network connections of the virtual machine.
By default, the virtual machine is not connected to a network and the virtual machine memory size equals that of the original machine.
 - e. Click **Done**.
10. [Optional] In the validation plan template, click **Schedule**, and then configure it.
 11. [If the backup sets selected in **Items to validate** are encrypted] Enable the **Backup password** switch, and then provide the encryption password.
 12. [Optional] To modify the plan options, click the gear icon.
 13. Click **Create**.

As a result, your validation plan is ready and will run according to the schedule that you configured. To run the plan immediately, select it in **Management > Validation**, and then click **Run now**.

After the plan starts, you can check the running activities and drill down to their details in Cyber Protect console, under **Monitoring > Activities**.

A validation plan might include multiple backups and one backup can be validated by multiple validation plans.

Note

All backups are processed sequentially, one by one, by a single validation task.

Only one validation task can run at a time on a given agent. Multiple validation tasks can run in parallel if they are executed by different agents: two simultaneous tasks require two agents, three tasks - three agents, and so on.

The following table summarizes the possible statuses of the validation activity.

Activity result	Plan with one backup	Plan with multiple backups
Success	All validation methods succeeded	All validation methods succeeded in all backups
Success with warnings	N/A	At least one validation method failed in at least one backup
Fail	At least one validation method failed	At least one validation method failed in all backups

Validation methods

In a validation plan, the following validation methods are available:

- Checksum verification
- Run as virtual machine
 - VM heartbeat
 - Screenshot validation

Checksum verification

Validation via checksum verification calculates a checksum for every data block that can be recovered from the backup, and then compares it against the original checksum for that data block, which was written during the backup process. The only exception is validation of file-level backups that are located in the cloud storage. These backups are validated by checking the consistency of the metadata saved in the backup.

Validation via checksum verification is a time-consuming process, even for an incremental or a differential backup, which are small in size. The reason is that the validation operation checks not only the data that is physically contained in a particular backup, but all of the data that needs to be recovered – that is, previous backups might also need to be validated.

A successful validation via checksum verification means a high probability of data recovery. However, the validation via this method does not check all factors that influence the recovery process.

If you back up an operating system, we recommend that you use some of the following additional operations:

- [Test recovery](#) under the bootable media to a spare hard drive.
- [Running a virtual machine from the backup](#) in an ESXi or Hyper-V environment.
- [Running a validation plan](#) in which the **Run as virtual machine** validation method is enabled.

Run as virtual machine

This method works only for disk-level backups that contain an operating system. To use it, you need an ESXi or Hyper-V host and a protection agent (Agent for VMware or Agent for Hyper-V) that manages this host.

The **Run as virtual machine** validation method is available in the following variants:

- VM heartbeat
- Screenshot validation

You must select at least one of them.

VM heartbeat

With this validation method, the agent runs a virtual machine from the backup, connects to VMware Tools or Hyper-V Integration Services, and then checks the heartbeat response to ensure that the operating system has started successfully. If the connection fails, the agent attempts to connect every two minutes, a total of five times. If none of the attempts are successful, the validation fails.

Regardless of the number of validation plans and validated backups, the agent that performs validation runs one virtual machine at a time. As soon as the validation result becomes clear, the agent deletes the virtual machine and runs the next one.

Note

Use this method only when you validate backups of VMware virtual machines by running these backups as virtual machines on an ESXi host, and backups of Hyper-V virtual machines by running them as virtual machines on a Hyper-V host.

Screenshot validation

With this validation method, the agent runs a virtual machine from the backup, and while the virtual machine is booting, screenshots are made. A machine intelligence (MI) module checks the screenshots and if there is a login screen on them, it marks the backup as validated.

The screenshot is attached to the recovery point and you can download it in the Cyber Protect console within one year of the validation. For more information on how to check the screenshot, refer to "Checking the validation status of a backup" (p. 203).

If notifications are enabled for your user account, you will receive an email about the validation status of the backup, in which the screenshot is attached. For more information about the notifications, refer to [Changing the notification settings for a user](#).

Screenshot validation is supported by agent version 15.0.30971 (released in November, 2022) and later.

Note

Screenshot validation works best with backups of Windows and Linux systems with GUI-based login screen. This method is not optimized for Linux systems with console login screen.

Changing the timeout for VM heartbeat and screenshot validation

When you validate a backup by running it as a virtual machine, you can configure the timeout between booting the virtual machine, and sending the heartbeat request or taking a screenshot.

The default period is as follows:

- One minute – for backups stored on a local folder or a network share
- Five minutes – for backups stored in the cloud

You can change this by editing the configuration file for Agent for VMware or Agent for Hyper-V.

To change the timeout

1. Open the configuration file for editing. You can find the file in the following locations:
 - For Agent for VMware or Agent for Hyper-V running in Windows: C:\Program Files\BackupClient\BackupAndRecovery\settings.config
 - For Agent for VMware (Virtual appliance): /bin/mms_settings.configFor more information on how to access the configuration file on a virtual appliance, see "SSH connections to a virtual appliance" (p. 170).
2. Go to <validation>, and then change the values for local backups and cloud backups as needed:

```
<validation>
<run_vm>
<initial_timeout_minutes>
<local_backups>1</local_backups>
<cloud_backups>5</cloud_backups>
</initial_timeout_minutes>
</run_vm>
</validation>
```

3. Save the configuration file.
4. Restart the agent.
 - [For Agent for VMware or Agent for Hyper-V running in Windows] Run the following commands at the command prompt:

```
net stop mms
```

```
net start mms
```

- [For Agent for VMware (Virtual appliance)] Restart the virtual machine with the agent.

Configuring the number of retries in case of an error

To maximize the number of successful validations, you can configure automatic retries for validation operations that end with an error.

To configure automatic retries

1. When creating a validation plan, click the gear icon.
2. In the **Options** pane, select **Error handling**.
3. Under **Re-attempt, if an error occurs**, click **Yes**.
4. In **Number of attempts**, configure the maximum number of retries if an error occurs.
The validation operation will run again until it finishes successfully or until the maximum number of retries is reached.
5. In **Interval between attempts**, configure the timeout between two consecutive retries.
6. Click **Done**.

Checking the validation status of a backup

You can check the validation status of a backup in the **Devices** tab or in the **Backup storage** tab.

You can also see the status for each validation method and download the screenshot taken by the screenshot validation method.

For more information about how the statuses work, refer to "Validation status" (p. 197).

To check the validation status of a backup

Devices

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the workload whose backup validation status you want to check, and then click **Recovery**.
3. [If more than one backup location is available] Select the backup location.
4. Select the backup whose status you want to check.

Backup storage

1. In the Cyber Protect console, go to **Backup storage**.
2. Select the location where your backup set is stored.
3. Select the backup set, and then click **Show backups**.
4. Select the backup whose validation status you want to check.

Cleanup

Cleanup is an operation that deletes outdated backups according to the retention rules. This operation is only applicable to agents and workloads, and not cloud to cloud backups (which can only be manually deleted).

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

Supported locations

Cleanup plans support all backup locations, except for NFS folders and Secure Zone.

To create a cleanup plan

1. In the Cyber Protect console, click **Management > Cleanup**.
2. Click **Create plan**.
3. In **Agent**, select the agent that will perform the cleanup.
You can select any agent that has access to the backup location.
4. In **Items to clean up**, select the archives or backup locations to clean up.

To switch between archives and locations, use the **Locations / Backups** switch in the upper-right corner.

If you select multiple encrypted archives, their encryption password must be the same. For archives that use different encryption passwords, create separate plans.

5. In **Schedule**, configure the cleanup schedule.
6. In **Retention rules**, specify the retention rules.
The following options are available:
 - **By number of backups**
 - **By backup age** (separate settings for monthly, weekly, daily, and hourly backups)
 - **By total size of backups**
7. [If you selected encrypted archives in **Items to replicate**] Enable the **Backup password** switch, and then provide the encryption password.
8. [Optional] To modify the plan options, click the gear icon, and then configure the options as required.
9. Click **Create**.

Conversion to a virtual machine

Conversion to a virtual machine is available only for disk-level backups. If a backup includes the system volume and contains all of the information necessary for the operating system to start, the resulting virtual machine can start on its own. Otherwise, you can add its virtual disks to another virtual machine.

Note

VMs replicated via native Scale Computing VM replication functionality cannot be backed up.

You can create a separate plan for conversion to a virtual machine and run this plan manually or on a schedule.

For information about prerequisites and limitations, refer to "What you need to know about conversion" (p. 206).

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

To create a plan for conversion to a virtual machine

1. Click **Management > Conversion to VM**.
2. Click **Create plan**.
The software displays a new plan template.
3. [Optional] To modify the plan name, click the default name.

4. In **Convert to**, select the type of the target virtual machine. **You can select one of the following:**

- **VMware ESXi**
- **Microsoft Hyper-V**
- **Scale Computing HC3**
- **VMware Workstation**
- **VHDX files**

Note

To save storage space, each conversion to VHDX files or VMware Workstation overwrites the VHDX/VMDK files in the target location that were created during the previous conversion.

5. Do one of the following:

- [For VMware ESXi, Hyper-V, and Scale Computing HC3] Click **Host**, select the target host, and then specify the new machine name template.
- [For other virtual machine types] In **Path**, specify where to save the virtual machine files and the file name template.

The default name is **[Machine Name]_converted**.

6. Click **Agent**, and then select the agent that will perform the conversion.

7. Click **Items to convert**, and then select the backups that this plan will convert to virtual machines.

You can switch between selecting backups and selecting entire locations by using the **Locations / Backups** switch in the upper-right corner.

If the selected backups are encrypted, all of them must use the same encryption password. For backups that use different encryption passwords, create separate plans.

8. [Only for VMware ESXi and Hyper-V] Click **Datastore** for ESXi or **Path** for Hyper-V, and then select the datastore (storage) for the virtual machine.

9. [Only for VMware ESXi and Hyper-V] Select the disk provisioning mode. The default setting is **Thin** for VMware ESXi and **Dynamically expanding** for Hyper-V.

10. [Optional] [For VMware ESXi, Hyper-V, and Scale Computing HC3] Click **VM settings** to modify the memory size, the number of processors, or the network connections of the virtual machine.

11. [Optional] Click **Schedule**, and then change the schedule.

12. If the backups selected in **Items to convert** are encrypted, enable the **Backup password** switch, and then provide the encryption password. Otherwise, skip this step.

13. [Optional] To modify the plan options, click the gear icon.

14. Click **Create**.

What you need to know about conversion

Supported virtual machine types

Conversion of a backup to a virtual machine can be done by the same agent that created the backup or by another agent.

To perform a conversion to VMware ESXi, Hyper-V, or Scale Computing HC3, you need an ESXi, Hyper-V, or Scale Computing HC3 host respectively and a protection agent (Agent for VMware, Agent for Hyper-V, or Agent for Scale Computing HC3) that manages this host.

Conversion to VHDX files assumes that the files will be connected as virtual disks to a Hyper-V virtual machine.

The following table summarizes the types of virtual machines that you can create with the **Convert to VM** operation. The rows in the table show the type of converted virtual machines. The columns show the agents that perform the conversion.

VM type	Agent for VMware	Agent for Hyper-V	Agent for Windows	Agent for Linux	Agent for Mac	Agent for Scale Computing HC3	Agent for oVirt (KVM)	Agent for Virtuozzo Hybrid Infrastructure	Agent for Virtuozzo
VMware ESXi	+	-	-	-	-	-	-	-	-
Microsoft Hyper-V	-	+	-	-	-	-	-	-	-
VMware Workstation	+	+	+	+	-	-	-	-	-
VHDX files	+	+	+	+	-	-	-	-	-
Scale Computing HC3	-	-	-	-	-	+	-	-	-

Limitations

- Backups stored on NFS cannot be converted.
- Backups stored in Secure Zone can be converted only by the agent running on the same machine.

- Backups that contain Linux logical volumes (LVM) can be converted only if they were created by Agent for VMware, Agent for Hyper-V, and Agent for Scale Computing HC3 and are directed to the same hypervisor. Cross-hypervisor conversion is not supported.
- When backups of a Windows machine are converted to VMware Workstation or VHDX files, the resulting virtual machine inherits the CPU type from the machine that performs the conversion. As a result, the corresponding CPU drivers are installed in the guest operating system. If started on a host with a different CPU type, the guest system displays a driver error. Update this driver manually.

Regular conversion to virtual machine vs. running a virtual machine from a backup

Both operations provide you with a virtual machine that can be started in seconds if the original machine fails.

Regular conversion to virtual machine takes CPU and memory resources. Files of the virtual machine constantly occupy space on the datastore (storage). This may be not practical if a production host is used for conversion. However, the virtual machine performance is limited only by the host resources.

Running a virtual machine from a backup consumes resources only while the virtual machine is running. The datastore (storage) space is required only to keep changes to the virtual disks. However, the virtual machine may run slower, because the host does not access the virtual disks directly, but communicates with the agent that reads data from the backup. In addition, the virtual machine is temporary.

How the regular conversion to a virtual machine works

The way the regular conversion works depends on where you choose to create the virtual machine.

- **If you choose to save the virtual machine as a set of files:** each conversion re-creates the virtual machine from scratch.
- **If you choose to create the virtual machine on a virtualization server:** when converting an incremental or differential backup, the software incrementally updates the existing virtual machine instead of re-creating it. Such conversion is normally faster. It saves network traffic and CPU resource of the host that performs the conversion. If updating the virtual machine is not possible, the software re-creates it from scratch.

The following is a detailed description of both cases.

If you choose to save the virtual machine as a set of files

As a result of the first conversion, a new virtual machine will be created. Every subsequent conversion will re-create this machine from scratch. First, the old machine is temporarily renamed. Then, a new virtual machine is created that has the previous name of the old machine. If this operation succeeds, the old machine is deleted. If this operation fails, the new machine is deleted and the old machine is given its previous name. This way, the conversion always ends up with a

single machine. However, extra storage space is required during conversion to store the old machine.

If you choose to create the virtual machine on a virtualization server

The first conversion creates a new virtual machine. Any subsequent conversion works as follows:

- If there has been *a full backup* since the last conversion, the virtual machine is re-created from scratch, as described earlier in this section.
- Otherwise, the existing virtual machine is updated to reflect changes since the last conversion. If updating is not possible (for example, if you deleted the intermediate snapshots, see below), the virtual machine is re-created from scratch.

Intermediate snapshots

To be able to update the converted virtual machine securely, the software stores an intermediate hypervisor snapshot of this machine. The snapshot is named **Replica...** and must be kept.

The **Replica...** snapshot corresponds to the result of the latest conversion. You can go to this snapshot if you want to return the machine to that state; for example, if you worked with the machine and now you want to discard the changes made to it.

For converted Scale Computing HC3 virtual machines, an additional **Utility Snapshot** is created. Only Cyber Protection service uses it.

Protection plans and modules

To protect your data, you must create protection plans, and then apply them to your workloads.

A protection plan consists of different protection modules. Enable the modules that you need and configure their settings to create protection plans that meet your specific needs.

The following modules are available:

- **Backup**. Backs up your data sources to a local or cloud storage.
- "Implementing disaster recovery" (p. 679). Launches exact copies of your machines in the cloud site and switches the workload from corrupted original machines to the recovery servers in the cloud.
- **Antivirus and Antimalware protection**. Checks your workloads by using a built-in antimalware solution.
- **Endpoint Detection and Response (EDR)**. Detects suspicious activity on the workload, including attacks that have gone unnoticed, and generates incidents to help you understand how an attack happened and how to prevent it from happening again.
- **URL filtering**. Protects your machines from threats originating from the Internet, by blocking access to malicious URLs and downloadable content.
- **Windows Defender Antivirus**. Manages the settings of Windows Defender Antivirus to protect your environment.

- [Microsoft Security Essentials](#). Manages the settings of Microsoft Security Essentials to protect your environment.
- [Vulnerability assessment](#). Checks Windows, Linux, macOS, Microsoft third-party products, and macOS third-party products installed on your machines and notifies you about vulnerabilities.
- [Patch management](#). Installs patches and updates for Windows, Linux, macOS, Microsoft third-party products, and macOS third-party products on your machines, to resolve the detected vulnerabilities.
- [Data protection map](#). Discovers data in order to monitor the protection status of important files.
- [Device control](#). Specifies devices that users are allowed or prohibited to use on your machines.
- [Advanced Data Loss Prevention](#). Prevents leakage of sensitive data via peripheral devices (such as printers or removable storage), or through internal and external network transfers, based on a data flow policy.

Creating a protection plan

You can create a protection plan in the following ways:

- On the **Devices** tab. Select one or more workloads to protect, and then create a protection plan for them.
- On the **Management > Protection plans** tab. Create a protection plan, and then select one or more workloads to which to apply the plan.

When you create a protection plan, only the modules that are applicable to your type of workload are shown.

You can apply a protection plan to more than one workload. You can also apply multiple protection plans to the same workload. To learn more about possible conflicts, see "Resolving plan conflicts" (p. 214).

To create a protection plan

Devices

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the workloads that you want to protect, and then click **Protect**.
3. [If there are already applied plans] Click **Add plan**.
4. Click **Create plan > Protection**.
The protection plan panel opens.
5. [Optional] To modify the protection plan name, click the pencil icon.
6. [Optional] To enable or disable a module in the plan, toggle the switch next to the module name.
7. [Optional] To configure a module, click it to expand it, and then change the settings according to your needs.
8. When ready, click **Create**.

Note

To create a protection plan with encryption, specify an encryption password. For more information, see "Encryption" (p. 417).

Management > Protection plans

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Click **Create plan**.
The template for a protection plan opens.
3. [Optional] To modify the protection plan name, click the pencil icon next to the name.
4. [Optional] To enable or disable a module in the plan, toggle the switch next to the module name.
5. [Optional] To configure a module, click it to expand it, and then change the settings according to your needs.
6. [Optional] To select the workloads to which you want to apply the plan, click **Add devices**.

Note

You can create a plan without applying it to any workloads. You can add workloads later, by editing the plan. For more information about how to add a workload to a plan, see "Applying a protection plan to a workload" (p. 211).

7. When ready, click **Create**.

Note

To create a protection plan with encryption, specify an encryption password. For more information, see "Encryption" (p. 417).

To run a module on demand (such as **Backup, Antivirus and Antimalware protection, Vulnerability assessment, Patch management, or Data protection map**), click **Run now**.

Watch the how-to video [Creating the first protection plan](#).

For more information on the Disaster recovery module, see "Create a disaster recovery protection plan" (p. 684).

For more information on the Device control module, see "Working with the Device control module" (p. 342).

Actions with protection plans

After creating a protection plan, you can perform the following actions with it:

- Apply a plan to a workload or a device group.
- Rename a plan.
- Edit a plan.

You can enable and disable the modules in a plan, and change their settings.

- Enable or disable a plan.
A disabled plan will not run on the workloads to which it is applied.
This action is convenient for administrators who intend to protect the same workload with the same plan later. The plan is not revoked from the workload and you can quickly restore the protection by re-enabling the plan.
- Revoke a plan from a workload.
A revoked plan is not applied to the workload anymore.
This action is convenient for administrators who do not need rapid protection for the same workload with the same plan again. To restore the protection provided by a revoked plan, you must know the name of this plan, select it from the list of available plans, and then re-apply it to the respective workload.
- Stop a plan.
This action stops all running backup operations on all workloads to which the plan is applied. Backups will start again according to the plan schedule.
Antimalware scanning is not affected by this action and will proceed as configured in the schedule.
- Clone a plan.
You can create an exact copy of an existing plan. The new plan is not assigned to any workloads.
- Export and import a plan.
You can export a plan as a JSON file, which you can import back later. Thus, you do not need to create a new plan manually and configure its settings.

Note

You can import protection plans created in Cyber Protection 9.0 (released in March 2020) and later. Plans created in earlier versions are not compatible with Cyber Protection 9.0 and later.

- Check the details of a plan.
- Check the activities and alerts related to a plan.
- Delete a plan.

Applying a protection plan to a workload

To protect a workload, you must apply a protection plan to it.

You can apply a plan from the **Devices** tab and from the **Management > Protection plans** tab.

Devices

1. Select one or more workloads that you want to protect.
2. Click **Protect**.
3. [If another protection plan was already applied to the selected workloads] Click **Add plan**.
4. A list of available protection plans is shown.
5. Select the protection plan that you want to apply, and then click **Apply**.

Management > Protection plans

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select the protection plan that you want to apply.
3. Click **Edit**.
4. Click **Manage devices**.
5. In the **Devices** window, click **Add**.
6. Select the workloads to which you want to apply the plan, and then click **Add**.
7. In the **Devices** window, click **Done**.
8. In the protection plan panel, click **Save**.

To learn how to apply a protection plan to a device group, see "Applying a plan to a group" (p. 340).

Editing a protection plan

When you edit a plan, you can enable and disable the modules in it, and change their settings.

You can edit a protection plan for all workloads to which it is applied or only for selected workloads.

You can edit a plan from the **Devices** tab and from the **Management > Protection plans** tab.

Devices

1. Select one or more workloads to which the plan is applied.
2. Click **Protect**.
3. Select the protection plan that you want to edit.
4. Click the ellipsis icon (...) next to the plan name, and then click **Edit**.
5. Click a module that you want to edit, and then configure its settings as needed.
6. Click **Save**.
7. [If you have not selected all workloads to which the plan is applied] Select the scope of the edit:
 - To edit the plan for all workloads to which it is applied, click **Apply the changes to this protection plan (this will affect other devices)**.
 - To change the plan only for selected workloads, click **Create a new protection plan only for the selected devices**.

As a result, the existing plan will be revoked from the selected workloads. A new protection plan with the settings that you configured will be created and applied to these workloads.

Management > Protection plans

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select the protection plan that you want to edit.
3. Click **Edit**.
4. Click the modules that you want to edit, and then configure their settings as needed.
5. Click **Save**.

Note

Editing a plan from the **Management > Protection plans** tab affects all workloads to which that plan is applied.

Revoking a protection plan

When you revoke a plan, you remove it from one or more workloads. The plan still protects the other workloads to which it is applied.

You can revoke a plan from the **Devices** tab and the **Management > Protection plans** tab.

Devices

1. Select the workloads from which you want to revoke the plan.
2. Click **Protect**.
3. Select the protection plan that you want to revoke.
4. Click the ellipsis icon (...) next to the plan name, and then click **Revoke**.

Management > Protection plans

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select the protection plan that you want to revoke.
3. Click **Edit**.
4. Click **Manage devices**.
5. In the **Devices** window, select the workloads from which you want to revoke the plan.
6. Click **Remove**.
7. In the **Devices** window, click **Done**.
8. In the protection plan template, click **Save**.

Enabling or disabling a protection plan

An enabled plan is active and runs on the workloads to which it is applied. A disabled plan is inactive – it is still applied to workloads but it does not run on them.

When you enable or disable a protection plan from the **Devices** tab, your action affects only the selected workloads.

When you enable or disable a protection plan from the **Management > Protection plans** tab, your action affects all workloads to which this plan is applied. Also, you can enable or disable multiple protection plans.

Devices

1. Select the workload whose plan you want to disable.
2. Click **Protect**.
3. Select the protection plan that you want to disable.
4. Click the ellipsis icon (...) next to the plan name, and then click **Enable** or **Disable**, respectively.

Management > Protection plans

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select one or more protection plans that you want to enable or disable.

3. Click **Edit**.
4. Click **Enable** or **Disable**, respectively.

Note

This action does not affect protection plans that were already in the target state. For example, if your selection includes both enabled and disabled plans, and you click **Enable**, all selected plans will be enabled.

Deleting a protection plan

When you delete a plan, it is revoked from all workloads and removed from the Cyber Protect console.

You can delete a plan from the **Devices** tab and the **Management > Protection plans** tab.

Devices

1. Select any workload to which the protection plan that you want to delete is applied.
2. Click **Protect**.
3. Select the protection plan that you want to delete.
4. Click the ellipsis icon (...) next to the plan name, and then click **Delete**.

Management > Protection plans

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select the protection plan that you want to delete.
3. Click **Delete**.
4. Confirm your choice by selecting the **I confirm the deletion of plan** check box, and then click **Delete**.

Resolving plan conflicts

You can apply multiple protection plans to the same workload. For example, you may apply one protection plan in which you enabled and configured only the **Antivirus and Antimalware** module, and another protection plan in which you enabled and configured only the **Backup** module.

You can combine protection plans in which different modules are enabled. You can also combine multiple protection plans in which only the **Backup** module is enabled. However, if any other module is enabled in more than one plan, a conflict occurs. To apply the plan, first you must resolve the conflict.

Conflict between a new and existing plan

If a new plan conflicts with an existing plan, you can resolve the conflict in one of the following ways:

- Create a new plan, apply it, and then disable the existing plan that conflicts with the new one.
- Create a new plan, and then disable it.

Conflict between an individual and group plan

If an individual protection plan conflicts with a group plan that is applied to a device group, you can resolve the conflict in one of the following ways:

- Remove the workload from the device group, and then apply the individual protection plan to it.
- Edit the existing group plan or apply a new group plan to the device group.

License issue

A protection plan module might require that a specific service quota is assigned to the protected workload. If the assigned service quota is not appropriate, you will not be able to run, update, or apply the protection plan in which the respective module is enabled.

To resolve a license issue, do one of the following:

- Disable the module that is not supported by the currently assigned service quota, and then continue using the protection plan.
- Change the assigned service quota manually. To learn how to do this, see "Changing the service quota of machines" (p. 181).

Default protection plans

A default protection plan is a preconfigured template that you can apply to your workloads, thus ensuring quick protection. By using a default protection plan, you do not have to create new protection plans from scratch.

When you apply a default protection plan for the first time, the template is copied to your tenant and you can edit the modules in the plan and their settings.

The following default plans are available:

- Cyber Protect Essentials
This plan provides basic protection functionality and file-level backup.
- Remote workers
This plan is optimized for users who work remotely. It provides more frequent tasks (such as backup, antimalware protection, and vulnerability assessment), stricter protection actions, and optimized performance and power options.
- Office workers (third-party Antivirus)
This plan is optimized for users who work at the office and prefer third-party antivirus software. In this plan, the **Antivirus and Antimalware protection** module is disabled.
- Office workers (Acronis Antivirus)
This plan is optimized for users who work at the office and prefer the Acronis antivirus software.

Comparison of the default protection plans

Modules and options	Default protection plans			
	Cyber Protect Essentials	Remote workers	Office workers (third-party Antivirus)	Office workers (Acronis Antivirus)
Backup	Available	Available	Available	Available
What to back up	Files/folders	Entire machine	Entire machine	Entire machine
Items to back up	[All Profiles Folder]			
Continuous data protection (CDP)	Disabled	Enabled	Disabled	Disabled
Where to back up	Cloud storage	Cloud storage	Cloud storage	Cloud storage
Schedule	Monday to Friday at 11:00 PM	Monday to Friday at 12:00 AM Additionally enabled options and start conditions: <ul style="list-style-type: none"> • If the machine is turned off, run missed tasks at the machine startup • Wake up from the sleep or hibernate mode to start a scheduled backup • Save battery power: Do not start when on battery • Do not start when on metered connection 	Monday to Friday at 11:00 PM	Monday to Friday at 11:00 PM
Backup scheme	Always incremental	Always incremental	Always incremental	Always incremental

Modules and options	Default protection plans			
	Cyber Protect Essentials	Remote workers	Office workers (third-party Antivirus)	Office workers (Acronis Antivirus)
How long to keep	Keep backups infinitely	Monthly: 12 months Weekly: 4 weeks Daily: 7 days	Monthly: 12 months Weekly: 4 weeks Daily: 7 days	Monthly: 12 months Weekly: 4 weeks Daily: 7 days
Backup options	Default options	Default options, plus: <ul style="list-style-type: none"> Performance and backup window (the green set): CPU priority: Low Output speed: 50% 	Default options	Default options
Antivirus and Antimalware protection	Available	Available	Not available	Available
Active Protection	Off	Off	-	Off
Advanced Antimalware	On	On	-	On
Network folder protection	On	On	-	On
Server-side protection	Off	Off	-	Off
Self protection	On	On	-	On
Cryptomining process detection	On	On	-	On
Quarantine	Remove quarantined files after 30 days	Remove quarantined files after 30 days	-	Remove quarantined files after 30 days
Behavior engine	Quarantine	Quarantine	-	Quarantine
Exploit prevention	Notify and stop the process	Notify and stop the process	-	Notify and stop the process

Modules and options	Default protection plans			
	Cyber Protect Essentials	Remote workers	Office workers (third-party Antivirus)	Office workers (Acronis Antivirus)
Real-time protection	Quarantine	Quarantine	–	Quarantine
Schedule scan	Quick scan:Quarantine At 02:20 PM, Sunday to Saturday Full scan:Off	Quick scan: Off Full scan: Quarantine At 01:55 PM, Sunday to Saturday Additionally enabled options and start conditions: <ul style="list-style-type: none"> • If the machine is turned off, run missed tasks at the machine startup • Wake up from the sleep or hibernate mode to start a scheduled backup • Save battery power: Do not start when on battery 	–	Quick scan: Quarantine At 02:20 PM, Sunday to Saturday Full scan: Off
Exclusions	None	None	–	None
URL filtering	Available	Available	Available	Available
Malicious website access	Always ask user	Block	Always ask user	Always ask user
Categories to filter	Default options	Default options	Default options	Default options
Exclusions	None	None	None	None
Vulnerability	Available	Available	Available	Available

Modules and options	Default protection plans			
	Cyber Protect Essentials	Remote workers	Office workers (third-party Antivirus)	Office workers (Acronis Antivirus)
assessment				
Vulnerability assessment scope	Microsoft products, Windows third-party products	Microsoft products, Windows third-party products	Microsoft products, Windows third-party products	Microsoft products, Windows third-party products
Schedule	At 01:15 PM, only on Monday	At 02:20 PM, only on Monday	At 01:15 PM, only on Monday	At 01:15 PM, only on Monday
Patch management	Available	Available	Available	Available
Microsoft products	All updates	All updates	All updates	All updates
Windows third-party products	Only major updates	Only major updates	Only major updates	Only major updates
Schedule	At 03:10 PM, only on Monday	At 02:20 PM, Monday to Friday	At 03:10 PM, only on Monday	At 03:10 PM, only on Monday
Pre-update backup	Off	On	Off	Off
Data protection map	Not available	Available	Available	Available
Extensions and exception rules	-	Default options and the following additional extensions: Images <ul style="list-style-type: none"> • .jpeg • .jpg • .png • .gif • .bmp • .ico • .wbmp • .xcf • .psd • .tiff • .dwg 	Default options (66 extensions to detect)	Default options (66 extensions to detect)

Modules and options	Default protection plans			
	Cyber Protect Essentials	Remote workers	Office workers (third-party Antivirus)	Office workers (Acronis Antivirus)
		Audio and video <ul style="list-style-type: none"> • .avi, • .mov, • .mpeg, • .mpg, • .mkv • .wav • .aif • .aifc • .aiff • .au • .snd • .mid • .midi • .mpga • .mp3 • .oga • .flac • .opus • .spx • .ogg • .ogx • .mp4 		
Schedule	–	At 03:35 PM, Monday to Friday	At 03:40 PM, Monday to Friday	At 03:40 PM, Monday to Friday

Note

The number of modules in a default protection plan may vary according to your Cyber Protection license.

Applying a default protection plan

The initial default protection plans are templates the settings of which you cannot edit. When you apply a default plan for the first time, the template is copied to your tenant as a preconfigured protection plan and is enabled on the selected workloads.

The protection plan appears in the **Management > Protection plans** tab, and then you can manage it there.

To apply a default protection plan for the first time

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the workloads that you want to protect.
3. Click **Protect**.
4. Select one of the default plans, and then click **Apply**.

Editing a default protection plan

You can edit a default protection plan after you apply it for the first time.

To edit an applied default protection plan


1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select the plan that you want to edit, and then click **Edit**.
3. Modify the modules that are included in this plan, or their options, and then click **Save**.

Important

Some of the options cannot be modified.

Individual protection plans for hosting control panel integrations

When you enable hosting control panel integrations on your [web hosting servers](#) that use DirectAdmin, cPanel, or Plesk, the Cyber Protection service automatically creates an individual protection plan under your user account for each workload. This protection plan is associated with the particular workload that initiated the protection plan creation, and cannot be revoked or assigned to other workloads.

To stop using an individual protection plan, you can delete it from the Cyber Protect console. You can identify individual protection plans by the  sign next to their name.

If you want a protection plan to protect multiple web hosting servers that use hosting control panel integrations, you can create a regular protection plan in the Cyber Protect console and assign these workloads to it. However, any modifications to a protection plan that is shared by multiple web hosting control panels, can only be made in the Cyber Protect console, and not from within the integrations.

#CyberFit Score for machines

#CyberFit Score provides you with a security assessment and scoring mechanism that evaluates the security posture of your machine. It identifies security gaps in the IT environment and open attack vectors to endpoints and provides recommended actions for improvements in the form of a report. This feature is available in all Cyber Protect editions.

The #CyberFit Score functionality is supported on:

- Windows 7 (first version) and later versions
- Windows Server 2008 R2 and later versions

How it works

The protection agent that is installed on a machine performs a security assessment and calculates the #CyberFit Score for the machine. The #CyberFit Score of a machine is automatically periodically recalculated.

#CyberFit scoring mechanism

The #CyberFit Score for a machine is calculated, based on the following metrics:

- Antimalware protection 0-275
- Backup protection 0-175
- Firewall 0-175
- Virtual private network (VPN) 0-75
- Full disk encryption 0-125
- Network security 0-25

The maximum #CyberFit Score for a machine is 850.

Metric	What is assessed?	Recommendations to users	Scoring
Antimalware	The agent checks whether antimalware software is installed on a machine.	<p>Findings:</p> <ul style="list-style-type: none"> • You have antimalware protection enabled (+275 points) • You don't have antimalware protection, your system may be at risk (0 points) <p>Recommendations provided by #CyberFit Score:</p> <p>You should have an antimalware solution installed and enabled on your machine to stay protected from security risks.</p> <p>You should refer to websites such as AV-Test or AV-Comparatives for a list of recommended antimalware solutions.</p>	<p>275 - antimalware software is installed on a machine</p> <p>0 - no antimalware software is installed on a machine</p>
Backup	The agent checks if a backup solution is installed on a machine.	<p>Findings:</p> <ul style="list-style-type: none"> • You have a backup solution protecting your data (+175 points) • No backup solution was found, your data may be at risk (0 points) <p>Recommendations provided by #CyberFit Score:</p>	<p>175 - a backup solution is installed on a machine</p> <p>0 - no backup solution is installed on a</p>

		<p>We recommend that you back up your data regularly to prevent data loss or ransomware attacks. Below are some backup solutions that you should consider using:</p> <ul style="list-style-type: none"> • Acronis Cyber Protect / Cyber Backup / True Image • Windows Server Backup (Windows Server 2008 R2 and later) 	machine
Firewall	<p>The agent checks whether a firewall is available and enabled in your environment.</p> <p>The agent does the following:</p> <ol style="list-style-type: none"> 1. Checks Windows Firewall and Network Protection whether a public firewall is turned on. 2. Checks Windows Firewall and Network Protection whether a private firewall is turned on. 3. Checks for a 3-rd party firewall solution/agent if Windows public and private firewalls are disabled. 	<p>Findings:</p> <ul style="list-style-type: none"> • You have a firewall enabled for public and private networks, or a 3-rd party firewall solution is found (+175 points) • You have a firewall enabled only for public networks (+100 points) • You have a firewall enabled only for private networks (+75 points) • You have no firewall enabled, your network connection is not secure (0 points) <p>Recommendations provided by #CyberFit Score:</p> <p>We recommend that you enable firewall for your public and private networks to improve your security protection against malicious attacks on your system. Below, detailed guides are provided on setting-up your Windows firewall, depending on your security needs and network architecture:</p> <p>Guides for end-users/employees:</p> <p>How to set up Windows Defender Firewall on your PC</p> <p>How to set up Windows Firewall on your PC</p> <p>Guides for system administrators and engineers:</p> <p>How to deploy Windows Defender Firewall with Advanced Security</p> <p>How to create Advanced Rules in Windows Firewall</p>	<p>100 - Windows public firewall is enabled</p> <p>75 - Windows private firewall is enabled</p> <p>175 - Windows public and private firewall are enabled</p> <p>OR</p> <p>a third-party firewall solution is enabled</p> <p>0 - neither a Windows firewall, nor a third-party firewall solution are enabled</p>
Virtual Private Network (VPN)	The agent checks whether a VPN solution is installed on a	<p>Findings:</p> <ul style="list-style-type: none"> • You have a VPN solution and can safely receive and send data across public and shared networks (+75 points) 	<p>75 - VPN is enabled and running</p> <p>0 - VPN is not</p>

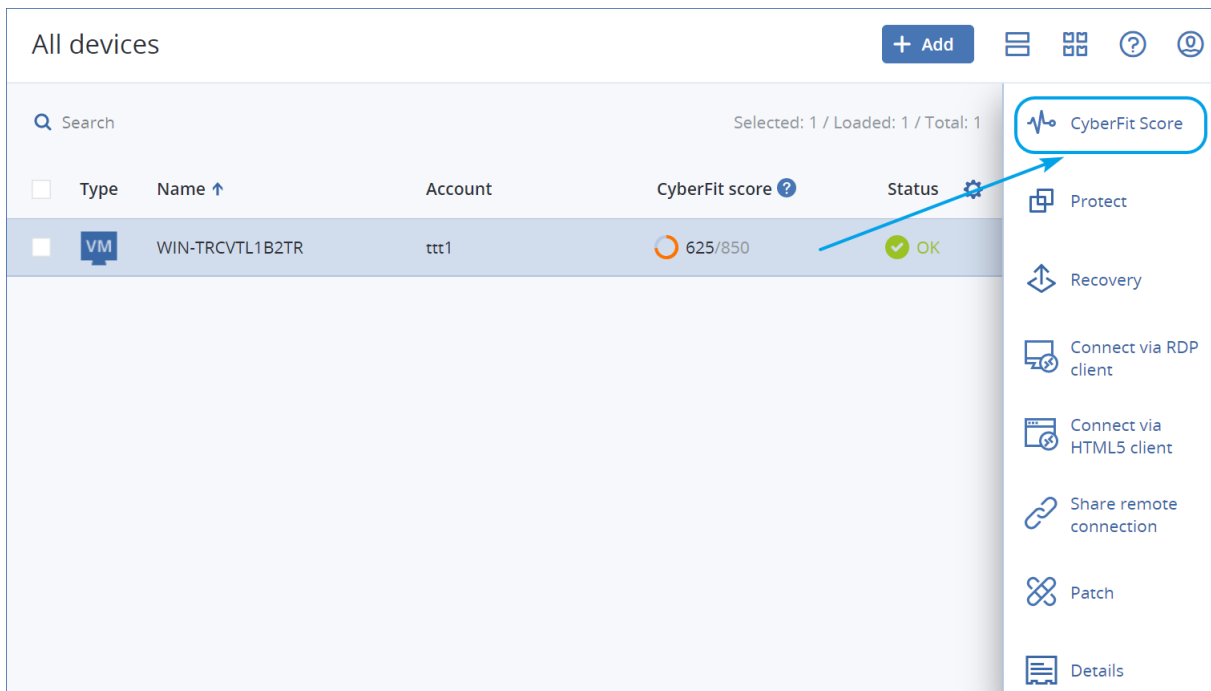
	<p>machine and whether the VPN is enabled and running.</p>	<ul style="list-style-type: none"> No VPN solution was found, your connection to public and shared networks is not secure (0 points) <p>Recommendations provided by #CyberFit Score:</p> <p>We recommend that you use VPN to access your corporate network and confidential data. It is critical to use a VPN to keep your communications safe and private, especially if you use complimentary Internet access from a cafe, library, airport, or elsewhere. Below are some VPN solutions that you should consider using:</p> <ul style="list-style-type: none"> Acronis Business VPN OpenVPN Cisco AnyConnect NordVPN TunnelBear ExpressVPN PureVPN CyberGhost VPN Perimeter 81 VyprVPN IPVanish VPN Hotspot Shield VPN Fortigate VPN ZYXEL VPN SonicWall GVPN LANCOM VPN 	<p>enabled</p>
<p>Disk encryption</p>	<p>The agent checks whether a machine has disk encryption enabled.</p> <p>The agent checks whether Windows BitLocker is turned on.</p>	<p>Findings:</p> <ul style="list-style-type: none"> You have full disk encryption enabled, your machine is protected against physical tampering (+125 points) Only some hard drives are encrypted, your machine may be at risk from physical tampering (+75 points) No disk encryption was found, your machine is at risk from physical tampering (0 points) <p>Recommendations provided by #CyberFit Score:</p> <p>We recommend that you turn on Windows BitLocker to improve protection of your data and files.</p> <p>Guide: How to turn on device encryption on</p>	<p>125 - all disks are encrypted</p> <p>75 - at least one of your disks is encrypted but there are also unencrypted disks</p> <p>0 - no disks are encrypted</p>

		Windows	
Network security (outgoing NTLM traffic to remote servers)	The agent checks whether a machine has restricted outgoing NTLM traffic to remote servers.	<p>Findings:</p> <ul style="list-style-type: none"> Outgoing NTLM traffic to remote servers is denied, your credentials are protected (+25 points) Outgoing NTLM traffic to remote servers is not denied, your credentials may be vulnerable to exposure (0 points) <p>Recommendations provided by #CyberFit Score:</p> <p>For better security protection, we recommend that you deny all outgoing NTLM traffic to remote servers. You can find information on how to change the NTLM settings and add exceptions by following the link below.</p> <p>Guide: Restrict outgoing NTLM traffic to remote servers</p>	<p>25 - outgoing NTLM traffic is set to DenyAll</p> <p>0 - outgoing NTLM traffic is set to another value</p>

Based on the summed points awarded to each metric, the total #CyberFit Score of a machine can fit one of the following ratings that reflect the endpoint's level of protection:

- 0 - 579 - Poor
- 580 - 669 - Fair
- 670 - 739 - Good
- 740 - 799 - Very good
- 800 - 850 - Excellent

You can see the #CyberFit Score for your machines in the Cyber Protect console: go to **Devices > All devices**. In the list of devices, you can see the **#CyberFit Score** column. You can also [run the #CyberFit Score scan](#) for a machine to check its security posture.




You can also get information about the #CyberFit Score in the corresponding [widget](#) and [report](#) pages.

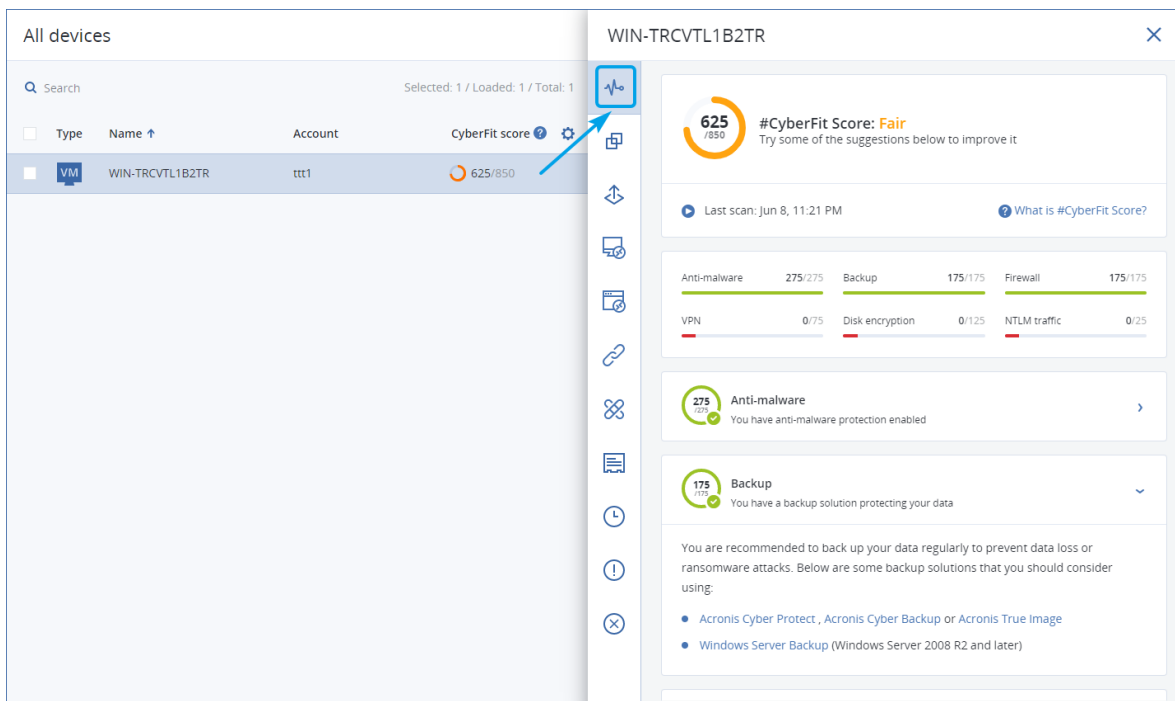
Running a #CyberFit Score scan

To run a #CyberFit Score scan

1. In the Cyber Protect console, go to **Devices**.
2. Select the machine and click **#CyberFit Score**.
3. If the machine has never been scanned before, then click **Run a first scan**.
4. After the scan is completed, you will see the total #CyberFit Score for the machine along with the scores of each of the six assessed metrics - Antimalware, Backup, Firewall, Virtual Private Network (VPN), Disk encryption, and NT LAN Manager (NTLM) traffic.



- To check how to increase the score of each metric for which the security configurations could be improved, expand the corresponding section and read the recommendations.



- After addressing the recommendations, you can always recalculate the #CyberFit Score of the machine by clicking on the arrow button right under the total #CyberFit Score.

Cyber Scripting

With Cyber Scripting, you can automate routine operations on Windows and macOS machines in your environment, such as installing software, modifying configurations, starting or stopping services, and creating accounts. Thus, you can decrease the time that you spend on such operations and reduce the risk of error when you perform them manually.

Cyber Scripting is available for administrators and users on the customer level, as well as to partner administrators (service providers). For more information about the different levels of administration, refer to "Multitenancy support" (p. 307).

The scripts that you can use must be approved in advance. Only the administrators with the Cyber administrator role can approve and test new scripts.

Performing operations with scripts and scripting plans depend on your user role. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

Prerequisites

- Cyber Scripting functionality requires the Advanced Management pack.
- To use all the features of Cyber Scripting such as script editing, script run, creation of scripting plans, and so on, you must enable two-factor authentication for your account.

Limitations

- The following scripting languages are supported:
 - PowerShell
 - Bash
- Cyber Scripting operations can only run on target machines that have an installed protection agent.

Supported platforms

Cyber Scripting is available for Windows and macOS workloads.

The following table summarizes the supported versions.

Operating system	Version
Windows	Windows 7 SP1 and later – all editions
	Windows 8/8.1 – all editions (x86, x64), except for the Windows RT editions
	Windows 10 – Home, Pro, Education, Enterprise, IoT Enterprise editions
	Windows 11
	Windows Server 2008 R2 SP1 and later – Standard, Enterprise, Datacenter, Foundation, and Web editions
	Windows Server 2012/2012 R2 – all editions
	Windows Server 2016
	Windows Server 2019
	Windows Server 2022
	Windows Storage Server (2008 R2, 2012, 2012 R2, 2016)
macOS	macOS Mojave 10.14
	macOS Catalina 10.15
	macOS Big Sur 11
	macOS Monterey 12

Scripts

A script is a set of instructions that are interpreted at runtime and executed on a target machine. It provides a convenient solution for automating repetitive or complex tasks.

With Cyber Scripting, you can run a predefined script or create a custom script. You can find all scripts that are available to you in **Management > Script repository**. The predefined scripts are located in the **Library** section. The scripts that you created or cloned to your tenant are located in the **My Scripts** section.

You can use a script by including it in a scripting plan or by starting a **Script quick run** operation.

Note

You can only use scripts that are created in your tenant or were cloned to it. If a script was removed from the script repository or its status was changed to **Draft**, it will not run. You can check the details of a scripting operation or cancel it in **Monitoring > Activities**.

The following table summarizes the possible actions with a script, depending on its status.

Status	Possible actions
Draft	All new scripts and the scripts that you clone in your repository are in the Draft status. These scripts cannot be run or included in scripting plans.
Testing	The scripts in the Testing status can be run and included in a scripting plan only by an administrator with the Cyber administrator role.
Approved	These scripts are available for running and including in scripting plans.

Only an administrator with the Cyber administrator role can change the state of a script or delete an approved script. For more information about the administrator rights, refer to "User roles and Cyber Scripting rights" (p. 245).

Creating a script

Note

Performing operations with scripts and scripting plans depend on your user role. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

To create a script

1. In the Cyber Protect console, go to **Management > Script repository**.
2. In **My Scripts**, click **Create**.
3. In the main pane, write the body of the script.

Important

When you create a script, include exit code checks for each operation. Otherwise, a failed operation might be ignored and the scripting activity status in **Monitoring > Activities** might be incorrectly shown as **Succeeded**.

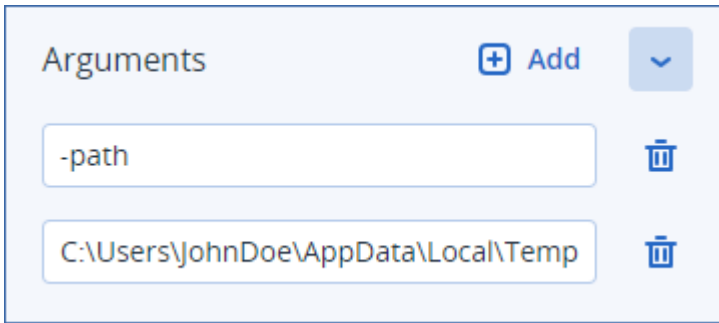
4. Specify the script properties that will help you find the script when you need it later:
 - a. **Script name**
 - b. [Optional] **Description**
 - c. **Language**
 - d. **Operating system**
 - e. Status. In the **Status** drop-down list, select one of the following statuses:
 - **Draft**
 - **Testing**
 - **Approved**
 - f. [Optional] Tags

The tags are not case-sensitive and can be up to 32 characters long. You cannot use round and angle brackets, commas, and spaces.
5. [Only for scripts that require credentials] Specify the credentials.

You can use a single credential (for example, a token) or a pair of credentials (for example, a user name and a password).

6. [Only for scripts that require arguments] Specify the arguments and their values, as follows:
 - a. Click **Add** twice.
 - b. In the first field, specify the argument.
 - c. In the second field, specify the argument value.

For example:



- d. Repeat the steps above if you need to add more than one argument.

You can only specify arguments that you have already defined in the script body.

```
Delete temporary files ● Approved
1 <#
2 .DESCRIPTION
3 Deletes all files in the specified temporary folder. If no arguments are specified, deletes the files in the folder specified in the "TEMP" environment variable.
4
5 .PARAMETER path
6 Optional. A path to folder with temporary files.
7 By default, uses the path specified in the "TEMP" environment variable.
8
9 .PARAMETER help
10 Displays a detailed usage description of this script.
11
12 .EXAMPLE
13 PS> .\Delete-Temporary-Files.ps1
14
15 .EXAMPLE
16 PS> .\Delete-Temporary-Files.ps1 -path "path-to-temp"
17
18 .EXAMPLE
19 PS> .\Delete-Temporary-Files.ps1 -help
20 #>
21
22 # Getting command line parameters
23 param (
24     [parameter(Mandatory = $false)][string]$path,
25     [parameter(Mandatory = $false)][switch]$help
26 )
```

7. Click **Save**.

As a result, you created a new script and you saved it to your repository. To use this script, an administrator with the Cyber administrator role must change its status to **Approved**. For more information about how to do this, refer to "Changing the script status" (p. 233).

To use a script in another tenant that you manage, you must clone the script to that tenant. For more information about how to do this, refer to "Cloning a script" (p. 231).

Cloning a script

Cloning a script is necessary in the following cases:

- Before using a script from **Library**. In this case, first you must clone the script to your **My Scripts** section.
- When you want to clone scripts that you created in a parent tenant to its child tenants or units.

To clone a script

1. In **Script repository**, find the script that you want to clone.
2. Do one of the following:
 - [If you clone a script from **My Scripts**] Click the ellipsis (...) next to the script name, and then click **Clone**.
 - [If you clone a script from **Library**] Click **Clone** next to the name of the script that you have selected.
3. In the **Clone script** pop-up, select one of the following script statuses from the **Status** drop-down list:
 - **Draft** (by default) — this status does not allow you to execute the script right away.
 - **Testing** — this status allows you to execute the script.
 - **Approved** — this status allows you to execute the script.
4. [If you manage more than one tenant or unit] Select where you want to clone the script.
In the **Clone script** dialog box, you see only the tenants that you can manage and which have the Advanced Management pack applied.

As a result, the script is cloned to the **My Scripts** section of the tenant or unit that you selected. If you manage only one tenant with no units in it, the script is automatically copied to your **My Scripts** section.

Important

Credentials that a script uses are not copied when you clone a script to a non-original tenant.

Editing or deleting a script

Note

Performing operations with scripts and scripting plans depend on your user role. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

To edit a script

1. In **Script repository**, go to **My Scripts**, and then find the script that you want to edit.
2. Click the ellipsis (...) next to the script name, and then click **Edit**.
3. Edit the script, and then click **Save**.
4. [If you edit a script that is used by a scripting plan] Confirm your choice by clicking **Save script**.

Note

The latest version of the script will be used next time the scripting plan runs.

Script versions

A new version of the script is created if you edit any of the following script attributes:

- script body
- script name
- description

- script language
- credentials
- arguments

If you change other attributes, your edits will be added to the current script version. To learn more about versions and how to compare them, refer to "Comparing script versions" (p. 234).

Note

The script status is updated only when you modify the value in the **Status** field. Only administrators with the Cyber administrator role can change a script status.

To delete a script

1. In **Script repository**, go to **My Scripts**, and then find the script that you want to delete.
2. Click the ellipsis (...) next to the script name, and then click **Delete**.
3. Click **Delete**.
4. [If you want to delete a script that is used by a scripting plan] Confirm your choice by clicking **Save script**.

Note

Scripting plans that use the deleted script will fail to run.

Changing the script status

Note

Performing operations with scripts and scripting plans depend on your user role. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

To change the script status

1. In **Script repository**, go to **My Scripts**, and then find the script whose status you want to change.
2. Click the ellipsis (...) next to the script name, and then click **Edit**.
3. In the **Status** drop-down list, select one of the following statuses:
 - **Draft**
 - **Testing**
 - **Approved**
4. Click **Save**.
5. [If you change the status of an approved script] Confirm your choice by clicking **Save script**.

Note

If the script status was downgraded to **Draft**, the scripting plans that use it will fail to run.

Only administrators with the Cyber administrator role can run scripts in the **Testing** state and scripting plans with such scripts.

Comparing script versions

You can compare two versions of a script and revert to an earlier version. You can also check who created a specific version, and when.

To compare script versions

1. In **Script repository**, go to **My Scripts**, and then find the script whose versions you want to compare.
2. Click the ellipsis (...) next to the script name, and then click **Version history**.
3. Select two versions that you want to compare, and then click **Compare versions**.

Any changes in the body text of the script, its arguments or credentials are highlighted.

To revert to an earlier version

1. In the **Compare script versions** window, click **Revert to this version**.
2. In the Revert to a previous version pop-up, select one of the following script statuses from the **Status** drop-down list:

- **Draft** (by default) — this status does not allow you to execute the script right away.
- **Testing** — this status allows you to execute the script.
- **Approved** — this status allows you to execute the script.

The selected version is restored and saved as the latest one in the version history.

To restore a script, you can also select a version from the **Version history** window, and then click on the **Restore** button.

Important

You can execute scripts only with the **Testing** or **Approved** statuses. For more information, refer to "Changing the script status" (p. 233).

Downloading the output of a scripting operation

You can download the output of a scripting operation as a .zip file. It contains two text files – stdout and stderr. In stdout, you can see the results of a successfully completed scripting operation. The stderr file contains information about the errors that occurred during the scripting operation.

To download the output file

1. In the Cyber Protect console, go to **Monitoring > Activities**.
2. Click the Cyber Scripting activity whose output you want to download.
3. On the **Activity details** screen, click **Download output**.

Script repository

You can locate the script repository under the **Management** tab. In the repository, you can search the scripts by their name and description. You can also use filters, or sort the scripts by their name

or status.

To manage a script, click the ellipsis (...) next to its name, and then select the desired action. Alternatively, click the script and use the buttons on the screen that opens.

The script repository contains the following sections:

- **My scripts**

Here, you can find the scripts that you can directly use in your environment. These are the scripts that you created from scratch and the scripts that you cloned here.

You can filter the scripts in this section by the following criteria:

- Tags
- Status
- Language
- Operating system
- Script owner

- **Library**

The library contains predefined scripts that you can use in your environment after cloning them to the **My scripts** section. You can only inspect and clone these scripts.

You can filter the scripts in this section by the following criteria:

- Tags
- Language
- Operating system

For more information, refer to [Vendor-Approved Scripts \(70595\)](#).

Scripting plans

A scripting plan allows you to run a script on multiple workloads, to schedule the running of a script, and to configure additional settings.

You can find the scripting plans that you created and the ones that are applied to your workloads in **Management > Scripting plans**. Here, you can check the plan execution location, owner, or status.

A clickable bar shows the following color-coded statuses for scripting plans:

- Running (Blue)
- Checking for compatibility (Dark gray)
- Disabled (Light gray)
- OK (Green)
- Critical alert (Red)
- Error (Orange)
- Warning (Yellow)

By clicking the bar, you can see which status a plan has and on how many workloads. Each status is also clickable.

On the **Scripting plans** tab, you can manage the plans by performing the following actions:

- Run
- Stop
- Edit
- Rename
- Disable
- Enable
- Clone
- Export. The plan configuration will be exported in a JSON format to the local machine.
- Delete

The visibility of a scripting plan and the available actions with it depend on the plan owner and your user role. For example, company administrators can only see the partner-owned scripting plans that are applied to their workloads, and cannot perform any actions with these plans.

For more information about who can create and manage scripting plans, refer to "User roles and Cyber Scripting rights" (p. 245).

To manage a scripting plan

1. In the Cyber Protect console, go to **Management > Scripting plans**.
2. Find the plan that you want to manage, and then click the ellipsis (...) next to it.
3. Select the desired action, and then follow the instructions on the screen.

Creating a scripting plan

You can create a scripting plan in the following ways:

- On the **Devices** tab
Select workloads, and then create a scripting plan for them.
- On the **Management > Scripting plans** tab
Create a scripting plan, and then select the workloads to which to apply the plan.

To create a scripting plan on the Devices tab

1. In the Cyber Protect console, go to **Devices > Machine with agents**.
2. Select the workloads or the device groups to which you want to apply a scripting plan, and then click **Protect** or **Protect group**, respectively.
3. [If there are already applied plans] Click **Add plan**.
4. Click **Create plan > Scripting plan**.
A template for the scripting plan opens.
5. [Optional] To modify the scripting plan name, click the pencil icon.
6. Click **Choose script**, select the script that you want to use, and then click **Done**.

Note

You can only use your own scripts from **Script repository > My scripts**. Only an administrator with the Cyber administrator role can use scripts in the **Testing** status. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

7. Configure the schedule and the start conditions for the scripting plan.
8. Choose under which account the script will run on the target workload. The following options are available:
 - System account (in macOS, this is the root account)
 - Currently logged-in account
9. Specify how long the script can run on the target workload.

If the script cannot finish running within the set time frame, the Cyber Scripting operation will fail.

The minimum value that you can specify is one minute and the maximum is 1440 minutes.
10. [Only for PowerShell scripts] Configure the PowerShell execution policy.

For more information about this policy, refer to the [Microsoft documentation](#).
11. Click **Create**.

To create a scripting plan on the Scripting plans tab

1. In the Cyber Protect console, go to **Management > Scripting plans**.
2. Click **Create plan**.

A template for the scripting plan opens.
3. [Optional] To select the workloads or the device groups to which you want to apply the new plan, click **Add workloads**.
 - a. Click **Machines with agents** to expand the list, and then select the desired workloads or device groups.
 - b. Click **Add**.

For more information about how to create device groups on the partner level, refer to "Devices tab" (p. 306).

Note

You can also select workloads or device groups after you create the plan.

4. [Optional] To modify the scripting plan name, click the pencil icon.
 5. Click **Choose script**, select the script that you want to use, and then click **Done**.
-

Note

You can only use your own scripts from **Script repository > My scripts**. Only an administrator with the Cyber administrator role can use scripts in the **Testing** status. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

6. Configure the schedule and the start conditions for the scripting plan.

7. Choose under which account the script will run on the target workload. The following options are available:
 - System account (in macOS, this is the root account)
 - Currently logged-in account
8. Specify how long the script can run on the target workload.

If the script cannot finish running within the set time frame, the Cyber Scripting operation will fail.

The minimum value that you can specify is one minute and the maximum is 1440 minutes.
9. [Only for PowerShell scripts] Configure the PowerShell execution policy.

For more information about this policy, refer to the [Microsoft documentation](#).
10. Click **Create**.

Schedule and start conditions

Schedule

You can configure a scripting plan to run once or repeatedly, and to start on a schedule or to be triggered by a certain event.

The following options are available:

- Run once
For this option, you must configure the date and time when the plan will run.
- Schedule by time
With this option, you can configure scripting plans that run hourly, daily, or monthly.
To make the schedule effective only temporarily, select the **Run within a date range** check box, and then configure the period during which the scheduled plan will run.
- When user logs in to the system
You can choose whether a specific user or any user who logs in triggers the scripting plan.
- When user logs off the system
You can choose whether a specific user or any user who logs off triggers the scripting plan.
- On the system startup
- When system is shut down

Note

This scheduling option only works with scripts that run under the system account.

- When system goes online

Start conditions

Start conditions add more flexibility to your scheduled plans. If you configure multiple conditions, all of them must be met simultaneously in order for the plan to start.

Start conditions are not effective if you run the plan manually, by using the **Run now** option.

Condition	Description
Run only if workload is online	The script will run when the target workload is connected to the Internet.
User is idle	This condition is met when a screen saver is running on the machine or the machine is locked.
User logged off	With this condition, you can postpone a scheduled scripting plan until the user of the target workload logs off.
Fits time interval	With this condition, a scripting plan can only start within the specified time interval. For example, you can use this condition to limit the User is logged off condition.
Save battery power	<p>With this condition, you can ensure that the scripting plan would not be interrupted because of a low battery. The following options are available:</p> <ul style="list-style-type: none"> • Do not start when on battery The plan will start only if the machine is connected to a power source. • Start when on battery if the battery level is higher than The plan will start if the machine is connected to a power source or if the battery level is higher than the specified value.
Do not start on metered connection	This condition prevents the plan from starting if the target workload accesses the Internet via a metered connection.
Do not start when connected to the following Wi-Fi networks	<p>This condition prevents the plan from starting if the target workload is connected to any of the specified wireless networks. To use this condition, you must specify the SSID of the forbidden network.</p> <p>The restriction applies to all networks that contain the specified name as a substring in their name, case-insensitive. For example, if you specify phone as the network name, the plan will not start when the device is connected to any of the following networks: John's iPhone, phone_wifi, or my_PHONE_wifi.</p>
Check device IP address	<p>This condition prevents the plan from starting if any of the IP addresses of the target workload are within or outside of the specified IP address range.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> • Start if outside IP range • Start if within IP range <p>Only IPv4 addresses are supported.</p>
If start conditions are not met, run the task anyway	<p>This option allows you to set the time interval after which the plan will run, irrespective of any other conditions. The plan will start as soon as the other conditions are met or the specified period ends, depending on which comes first.</p> <p>This option is not available if you configured the scripting plan to run only once.</p>

Managing the target workloads for a plan

You can select the workloads or the device groups to which to apply a scripting plan while you create the plan, or later.

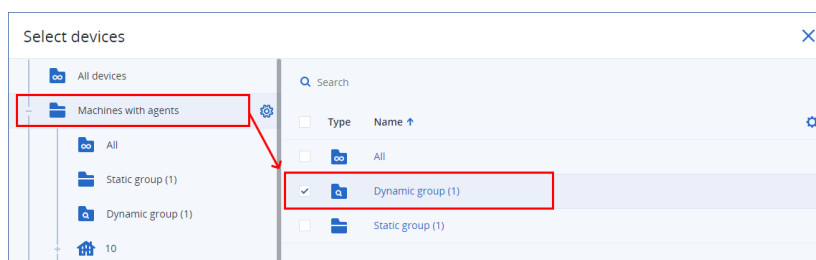
Partner administrators can apply the same plan to workloads from different customers, and can create device groups that contain workloads from different customers. To learn how to create a static or a dynamic device group on the partner level, refer to the "Devices tab" (p. 306).

To add initial workloads to a plan

1. In the Cyber Protect console, go to **Management > Scripting plans**.
2. Click the name of the plan for which you want to specify target workloads.
3. Click **Add workloads**.
4. Select the desired workloads or device groups, and then click **Add**.

Note

To select a device group, click its parent level, and then, in the main pane, select the check box next to its name.



5. To save the edited plan, click **Save**.

To manage existing workloads for a plan


1. In the Cyber Protect console, go to **Management > Scripting plans**.
2. Click the name of the plan whose target workloads you want to change.
3. Click **Manage workloads**.

The **Devices** screen lists the workloads to which the scripting plan is currently applied. If you manage more than one tenant, the workloads are sorted by tenant.

- To add new workloads or device groups, click **Add**.
 - a. Select the desired workloads or device groups. You can add workloads from all tenants that you manage.

Note

To select a device group, click its parent level, and then, in the main pane, select the check box next to its name.



b. Click **Add**.

- To remove workloads or device groups, select them, and then click **Remove**.

4. Click **Done**.

5. To save the edited plan, click **Save**.

Plans on different administration levels

The following table summarizes which plans administrators from different levels can see and manage.

Administrator	Administration level	Plans	Rights
Partner administrator	Partner level	Own plans	Full access
		Customer plans (including plans in units)	Full access
		Unit plans	Full access
	Customer level (for customers that are managed by the service provider)	Partner plans that are applied to workloads of this customer	Read-only
		Customer plans (including plans in units)	Full access
		Unit plans	Full access
	Unit level (for customers that are managed by the service provider)	Partner plans that are applied to workloads of this unit	Read-only
		Customer plans that are applied to workloads of this unit	Read-only
		Unit plans	Full access

Administrator	Administration level	Plans	Rights
Company administrator	Customer level	Partner plans that are applied to workloads of this customer or unit	Read-only
		Customer plans (including plans in units)	Full access
		Unit plans	Full access
	Unit level	Partner plans that are applied to workloads of this unit	Read-only
		Customer plans that are applied to workloads of this unit	Read-only
		Unit plans	Full access
Unit administrator	Unit level	Partner plans that are applied to workloads of this unit	Read-only
		Customer plans that are applied to workloads of this unit	Read-only
		Unit plans	Full access

Important

The owner of a plan is the tenant in which the plan was created. Thus, if a partner administrator created a plan on the customer tenant level, the customer tenant is the owner of that plan.

Compatibility issues with scripting plans

In some cases, applying a scripting plan on a workload might cause compatibility issues. You might observe the following compatibility issues:

- Incompatible operating system – this issue appears when the workload's operating system is not supported.
- Unsupported agent – this issue appears when the version of the protection agent on the workload is outdated and does not support the Cyber Scripting functionality.
- Insufficient quota – this issue appears when there is not enough service quota in the tenant to assign to the selected workloads.

If the scripting plan is applied to up to 150 individually selected workloads, you will be prompted to resolve the existing conflicts before saving the plan. To resolve a conflict, remove the root cause for it or remove the affected workloads from the plan. For more information, see "Resolving compatibility issues with scripting plans" (p. 243). If you save the plan without resolving the conflicts, it will be automatically disabled for the incompatible workloads, and alerts will be shown.

If the scripting plan is applied to more than 150 workloads or to device groups, it will be saved, and then checked for compatibility. The plan will be automatically disabled for the incompatible workloads, and alerts will be shown.

Resolving compatibility issues with scripting plans

Depending on the cause of the compatibility issues, you can perform different actions to resolve the compatibility issues as a part of the process of creating a new scripting plan.

Note

When resolving a compatibility issue by removing workloads from a plan, you cannot remove workloads that are part of a device group.

To resolve the compatibility issues

1. Click **Review issues**.
2. [To resolve compatibility issues with incompatible operating systems]
 - a. On the **Incompatible operating system** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
3. [To resolve compatibility issues with unsupported agents by removing workloads from the plan]
 - a. On the **Unsupported agents** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
4. [To resolve compatibility issues with unsupported agents by updating the agent version] Click **Go to the Agents list**.

Note

This option is available only for customer administrators.

5. [To resolve compatibility issues with insufficient quota by removing workloads from the plan]
 - a. On the **Insufficient quota** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
6. [To resolve compatibility issues with insufficient quota by increasing the quota of the tenant]

Note

This option is available only for partner administrators.

- a. On the **Insufficient quota** tab, click **Go to the Management portal**.
- b. Increase the service quota for the customer.

Script quick run

The **Script quick run** operation allows you to run a script immediately, without including it in a scripting plan. You cannot use this operation on more than 150 workloads, on offline workloads, or on device groups.

The target workload must be assigned a service quota that supports the Script quick run functionality, and the Advanced Management pack must be enabled for its tenant. An appropriate service quota will be automatically assigned if it is available in the tenant.

Note

You can only use your own scripts from **Script repository > My scripts**. Only an administrator with the Cyber administrator role can use scripts in the **Testing** status. For more information about the roles, refer to "User roles and Cyber Scripting rights" (p. 245).

You can start a quick run in the following ways:

- From the **Devices** tab
Select one or more workloads, and then select which script to run on it.
- From the **Management > Scripting repository** tab
Select a script, and then select one or more target workloads.

To run a script from the Devices tab

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the workload on which you want to run the script, and then click **Protect**.
3. Click **Script quick run**.
4. Click **Choose script**, select the script that you want to use, and then click **Done**.
5. Choose under which account the script will run on the target workload. The following options are available:
 - System account (in macOS, this is the root account)
 - Currently logged-in account
6. Specify how long the script can run on the target workload.
If the script cannot finish running within the set time frame, the Cyber Script operation will fail.
The minimum value that you can specify is one minute and the maximum is 1440 minutes.
7. [Only for PowerShell scripts] Configure the PowerShell execution policy.
For more information about this policy, refer to the [Microsoft documentation](#).
8. Click **Run now**.

To run a script from the Scripting repository tab

1. In the Cyber Protect console, go to **Management > Scripting repository**.
2. Select the script that you want to run, and then click **Script quick run**.
3. Click **Add workloads** to select the target workloads, and then click **Add**.
4. Click **Choose script**, select the script that you want to use, and then click **Done**.

5. Choose under which account the script will run on the target workload. The following options are available:
 - System account (in macOS, this is the root account)
 - Currently logged-in account
6. Specify how long the script can run on the target workload.

If the script cannot finish running within the set time frame, the Cyber Script operation will fail. The minimum value that you can specify is one minute and the maximum is 1440 minutes.
7. [Only for PowerShell scripts] Configure the PowerShell execution policy.

For more information about this policy, refer to the [Microsoft documentation](#).
8. Click **Run now**.

User roles and Cyber Scripting rights

The available actions with scripts and scripting plans depend on the script status and your user role.

Administrators can manage objects in their own tenant and in its child tenants. They cannot see or access objects on an upper administration level, if any.

Lower-level administrators have only read-only access to the scripting plans applied to their workloads by an upper-level administrator.

The following roles provide rights with regard to Cyber Scripting:

- **Company administrator**

This role grants full administrator rights in all services. With regard to Cyber Scripting, it grants the same rights as the Cyber administrator role.
- **Cyber administrator**

This role grants full permissions, including approval of scripts that can be used in the tenant, and the ability to run scripts with the **Testing** status.
- **Administrator**

This role grants partial permissions, with the ability to run approved scripts as well as create and run scripting plans that use approved scripts.
- **Read-only administrator**

This role grants limited permissions, with the ability to view scripts and protection plans that are used in the tenant.
- **User**

This role grants partial permissions, with the ability to run approved scripts as well as create and run scripting plans that use approved scripts, but only on the user's own machine.

The following table summarizes all available actions, depending on the script status and the user role.

Role	Object	Script status		
		Draft	Testing	Approved
Cyber administrator Company administrator	Scripting plan	Edit (Remove a draft script from a plan) Delete Revoke Disable Stop	Create Edit Apply Enable Run Delete Revoke Disable Stop	Create Edit Apply Enable Run Delete Revoke Disable Stop
	Script	Create Edit Change status Clone Delete Cancel running	Create Edit Change status Run Clone Delete Cancel running	Create Edit Change status Run Clone Delete Cancel running
Administrator User (for their own workloads)	Scripting plan	View Revoke Disable Stop	View Cancel run	Create Edit Apply Enable Run Delete Revoke Disable Stop
	Script	Create Edit Clone Delete	View Clone Cancel running	Run Clone Cancel running

		Cancel running		
Read-only administrator	Scripting plan	View	View	View
	Script	View	View	View

Protection of collaboration and communication applications

Zoom, Cisco Webex Meetings, Citrix Workspace, and Microsoft Teams are now widely used for video/web conferencing and communications. The Cyber Protection service allows you to protect your collaboration tools.

The protection configuration for Zoom, Cisco Webex Meetings, Citrix Workspace, and Microsoft Teams is similar. In the example below, we will consider configuration for Zoom.

To set up Zoom protection

1. [Install the protection agent](#) on the machine where the collaboration application is installed.
2. Log in to the Cyber Protect console and [apply a protection plan](#) that has one of the following modules enabled:
 - **Antivirus and Antimalware protection** (with the **Self-Protection** and **Active Protection** settings enabled) – if you have one of the Cyber Protect editions.
 - **Active Protection** (with the **Self-Protection** setting enabled) – if you have one of the Cyber Backup editions.
3. [Optional] For automatic update installation, configure the [Patch management module](#) in the protection plan.

As a result, your Zoom application will be under protection that includes the following activities:

- Installing Zoom client updates automatically
- Protecting Zoom processes from code injections
- Preventing suspicious operations by Zoom processes
- Protecting the "hosts" file from adding the domains related to Zoom

Understanding your current level of protection

Monitoring

The **Monitoring** tab provides important information about your current level of protection, and includes the following dashboards:

- **Overview**
- **Activities**
- **Alerts**
- **Threat feed** (for more information, see "Threat feed" (p. 288))

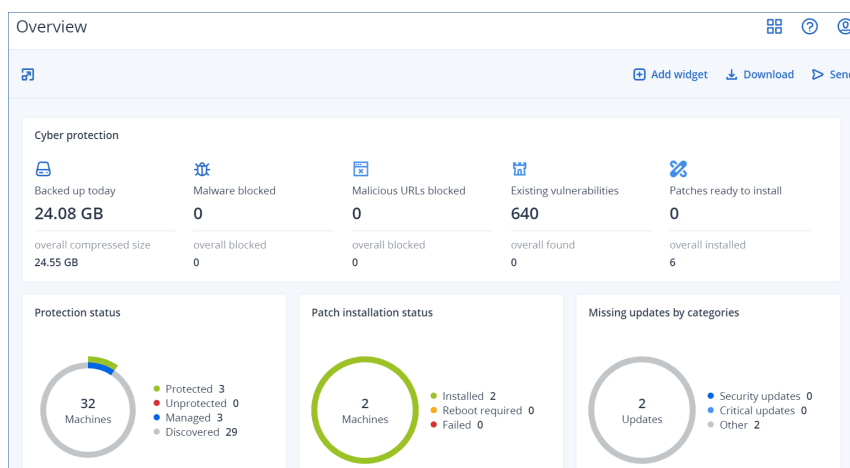
The Overview dashboard

The **Overview** dashboard provides a number of customizable widgets that give an overview of operations related to the Cyber Protection service. Widgets for other services will be available in future releases.

The widgets are updated every five minutes. The widgets have clickable elements that enable you to investigate and troubleshoot issues. You can download the current state of the dashboard or send it via email in the .pdf or/and .xlsx format.

You can choose from a variety of widgets, presented as tables, pie charts, bar charts, lists, and tree maps. You can add multiple widgets of the same type with different filters.

The buttons **Download** and **Send** in **Monitoring > Overview** are not available in the Standard editions of the Cyber Protection service.



To rearrange the widgets on the dashboard

Drag and drop the widgets by clicking on their names.

To edit a widget

Click the pencil icon next to the widget name. Editing a widget enables you to rename it, change the time range, set filters, and group rows.

To add a widget

Click **Add widget**, and then do one of the following:

- Click the widget that you want to add. The widget will be added with the default settings.
- To edit the widget before adding it, click **Customize** when the widget is selected. After editing the widget, click **Done**.

To remove a widget

Click the X sign next to the widget name.

The Activities dashboard

The **Activities** dashboard provides an overview of the current and past activities. By default, the retention period is 90 days.

To customize the view of the **Activities** dashboard, click the gear icon, and then select the columns that you want to see.

To see the activity progress in real time, select the **Refresh automatically** check box. However, frequent updating of multiple activities degrades the performance of the management server.

You can search the listed activities by the following criteria:

- **Device name**
This is the machine on which the activity is carried out.
- **Started by**
This is the account who started the activity.

You can also filter the activities by the following properties:

- **Status**
For example, succeeded, failed, in progress, canceled.
- **Type**
For example, applying plan, deleting backups, installing software updates.
- **Time**
For example, the most recent activities, the activities from the past 24 hours, or the activities during a specific period within the default retention period.

To see more details about an activity, select this activity from the list, and then, in the **Activity details** panel, click **All properties**. For more information about the available properties, refer to the [Activity](#) and [Task](#) API references in the Developer Network Portal.

The Alerts dashboard

The **Alerts** dashboard displays all your current alerts. Alerts listed are critical or error alerts, and are typically related to tasks such as backup that have failed for some reason.

To filter alerts on the dashboard

1. From the **View** drop-down list, select one of the following criteria:
 - **Alert severity**
 - **Alert category**
 - **Alert type**
 - **Monitoring type**
 - **Date range: from ... to ...**
 - **Workload**
 - **Plan**
 - **Customer**
2. If you have selected the **Alert category**, from the **Category** drop-down list, select the category of alerts that you want to view.
3. If you want to view all the alerts without filtering them, click **All alert types**.

Within each alert you can do the following:

- Access the relevant device the alert relates to by clicking the **Devices** link.
- Read and try to follow some advice in the **Troubleshooting** section of the alert.
- Access the relevant documentation and knowledge base article by clicking **Search for solution**.
The **Search for solution** functionality will pre-fill your request with the current alert details to assist you the most effectively.

To sort alerts on the dashboard

On the alerts table, click on the arrow button next to one of the following column names:

- **Alert severity**
- **Alert type**
- **Created**
- **Alert category**
- **Workload**
- **Plan**

If the Advanced Automation service is enabled for your account, you can also create a new service desk ticket directly from the alert.

To create a service desk ticket

1. In the relevant alert, click **Create a new ticket**.
Alternatively, when working in the table view mode, select an alert and then select **Create a new ticket** in the right pane.
2. Define the following:
 - In the header section, select the **Billable** check box if you want the time recorded on the ticket to be billed to the customer. In addition, select the **Email the customer** check box if you want to send ticket updates to the customer.
 - In the **General information** section, define a ticket title. This field is pre-filled with an alert summary but can be edited.
 - In the **Customer information** section, the fields are pre-filled with the relevant information from the alert.
 - In the **Configuration item or service** section, the fields are pre-filled with the device linked to the alert. You can reassign a device, as required.
 - In the **Support agent** section, the fields are pre-filled with the default support agent, category, and support group. You can reassign a different agent, as required.
 - In the **Ticket update** section, the fields are pre-filled with the alert description and details. The **Status** field is set as **New** by default, and can be changed.
 - In the **Attachments**, **Billable items**, and **Internal notes** sections, add the relevant items as required.
3. Click **Done**. When the ticket is created, a link to the ticket is added to the alert.
If an alert is closed, the related ticket is also automatically closed.

Note

You can only create one ticket per alert.

Alert types

Alerts will be generated for the following alert types:

- [Backup alerts](#)
- [Disaster recovery alerts](#)
- [Antimalware protection alerts](#)
- [Licensing alerts](#)
- [URL Filtering alerts](#)
- [EDR alerts](#)
- [Device Control alerts](#)
- [System alerts](#)

Backup alerts

Alert	Description	How to resolve the alert
Backup failed	An alert is generated when	Check the log of the faulty backup

Alert	Description	How to resolve the alert
	the backup failed with a resolvable error while running or it was interrupted due to system shut down.	operation: click the workload to select it, click Activities , and then find the warning in the log. The message should point you to the root cause of the issue the software notifies you about.
Backup succeeded with warnings	An alert is generated when the backup succeeded with warnings.	Check logs of conversion to VM, replication, or validation plans. Issues during these operations generate an "Activity failed" or "Activity finished with warning" alerts.
Backup is canceled	An alert is generated every time a backup activity is manually canceled by the user.	You can either start the backup manually by clicking Run now or wait until it runs at the next scheduled time.
Backup canceled due to closed backup window	An alert is generated when the backup activity was missed because it did not fit in the window specified in the backup options.	Re-configure schedule or edit options of the backup plan in Performance and backup window. Expand the section with your product for instructions.
Backup is waiting	This alert is generated anytime you have a scheduling conflict and two backups tasks are initiated at the same time. In this case, the second backup task is queued until the first one is finished or stopped.	Make sure that your backups are running in the expected time windows and according to their schedule, and avoid scheduling conflicts where possible.
Backup is not responding	An alert is generated when the running backup has not shown any progress for some time, and may be frozen.	The issue might be caused by a lockup. Follow this article to collect the necessary troubleshooting information.
Backup did not start	An alert is generated when the scheduled backup failed to start for unknown reason.	Make sure you are using the latest build of your Acronis Backup product. <ul style="list-style-type: none"> If the agent machine was available during the backup start time: <ol style="list-style-type: none"> Edit the backup task start time. If the alert appears again, recreate the backup task. If the newly created backup task also triggers the alert, contact

Alert	Description	How to resolve the alert
		<p>Acronis Support for assistance.</p> <ul style="list-style-type: none"> • If the agent was offline: <ol style="list-style-type: none"> 1. Do not turn off machine during backup time. 2. If machine was not turned off, make sure Acronis Managed Machine Service is running: Start -> Search -> services.msc -> locate Acronis Managed Machine Service. Contact Acronis Support, if you need assistance.
Backup status is unknown	An alert is generated when the backup agent was offline at a scheduled backup time. The status of the resource backups will be unknown until the backup agent becomes online.	<ol style="list-style-type: none"> 1. Check if the agent is expected to be offline (for example, it is a notebook that is outside the Management Server network). 2. If the agent should not be offline, make sure Acronis Managed Machine Service is running: Start -> Search -> services.msc -> locate Acronis Managed Machine Service and check its status. Start the service, if it is stopped.
Backup is missing	An alert is generated when there is not a successful backup for more than [Days from last backup] days.	
Backup is corrupted	An alert is generated when the validation activity is successful and shows that the backup is corrupted.	<p>Follow steps from the article Troubleshooting Issues with Corrupt Backups.</p> <p>If you need assistance with identifying the root cause for archive corruption, contact Acronis Support.</p>
Continuous Data Protection failed	An alert is generated if the continuous protection of backup failed.	<p>Verify the following limitations:</p> <ol style="list-style-type: none"> 1. Continuous data protection is supported only for the NTFS file system and the following operating systems: <ul style="list-style-type: none"> • Desktop: Windows 7 and later • Server: Windows Server 2008 R2 and later

Alert	Description	How to resolve the alert
		<ol style="list-style-type: none"> 2. CDP doesn't support Acronis Secure Zone as a destination. 3. NFS folders that are mounted on Windows are not supported. 4. Continuous replication is not supported: if there are two locations in the protection plan, CDP slices are created only in the first destination, and then the changes are replicated to the second one with the next backup. 5. If changes in a local protected folder are applied from a network source (e.g. when users access the folder from network), CDP doesn't detect them. 6. If a file is being used, e.g. some changes are being made in an Excel file, CDP doesn't detect the changes. For the changes to be detected by CDP save them and close the file.
Hyper-V hosts configuration is not valid	An alert is generated when there are 2 or more Agents for Hyper-V installed on Hyper-V hosts with the same host name, which is not supported on the same account level.	You should register these Agents for Hyper-V under different child units of this account to avoid conflicts.
Validation failed	An alert is generated when the validation process of your backup cannot be completed.	Check the log of the faulty operation: click the machine to select it, click Activities , and then find the warning in the log. The message should point you to the root cause of the issue the software notifies you about.
Failed to migrate the backups in the cloud storage to the new format	An alert is generated when it failed to migrate the backups in the cloud storage to the new format.	<p>Migration of Acronis Cyber Backup Advanced archives is described here.</p> <p>Migration of Acronis Cyber Backup archives is described here.</p> <p>Before contacting Acronis Support, please collect the following reports using the migrate_archives tool:</p>

Alert	Description	How to resolve the alert
		<pre>migrate_archives.exe -- account=<Acronis Account> -- password=<password> -- subaccounts=All > report1.txt migrate_archives.exe -- cmd=finishUpgrade -- account=<Acronis Account> -- password=<password> > report2.txt</pre>
Encryption password is missing	An alert is generated when the database encryption key is incorrect, corrupt, or missing.	There is no way to recover encrypted backups if you lose or forget the password. You must set the encryption password locally, on the protected device. You cannot set the encryption password in the protection plan. For more information, see Setting the encryption password .
Upload is pending	An alert is generated if scheduled check finds that Physical Data Shipping to cloud archive for this backup plan is not uploaded to storage.	
Backup recovery failed	An alert is generated when the recovery operation fails when you try to recover files or system backups.	Determine the exact date of the backup failure and attempt recovery with the last successful backup.

Disaster recovery alerts

Alert	Description	How to resolve the alert
Storage quota exceeds	An alert is generated when the soft quota is exceeded for disaster recovery storage	Increase the quota or remove some archives from the cloud storage.
Quota is reached	An alert is generated when: <ul style="list-style-type: none"> • Soft quota is exceeded for cloud servers. • Soft quota is exceeded for compute point. • Soft quota is exceeded for public IP addresses. 	

Alert	Description	How to resolve the alert
Storage quota is exceeded	<p>An alert is generated when hard quota is exceeded for disaster recovery storage.</p> <p>This storage is used by primary and recovery servers. If the overage for this quota is reached, it is not possible to create primary and recovery servers, or add/extend disks of the existing primary servers. If the overage for this quota is exceeded, it is not possible to initiate a failover or just start a stopped server. Running servers continue to run.</p>	
Quota is exceeded	<p>An alert is generated when:</p> <ul style="list-style-type: none"> • Hard quota is exceeded for cloud servers. • Hard quota is exceeded for compute point. • Hard quota is exceeded for public IP addresses. 	Consider purchasing additional device quotas or disable backup tasks for the devices you no longer need to protect.
Failover error	<p>An alert is generated when a system problem occurred after the fail-over action was submitted.</p>	<ol style="list-style-type: none"> 1. Click Edit on the recovery server. For more information, see Creating a recovery server. 2. Decrease CPU/RAM for the recovery server. 3. Try the failover again.
Test failover error	<p>An alert is generated when a system problem occurred after the test action was submitted.</p>	<ol style="list-style-type: none"> 1. Click Edit on the recovery server. For more information, see Creating a recovery server. 2. Decrease CPU/RAM for the recovery server. 3. Try the failover again. <hr/> <p>Note Make sure that there is the same IP address in IP address in production network as the one configured in the DHCP server.</p>
Failback error	<p>An alert is generated when a system problem occurred after</p>	<p>You can see the erroneous location in the list of backup</p>

Alert	Description	How to resolve the alert
	the fail-back was initiated.	<p>storages: it has a number instead of a name (normally, a location name matches one of the existing end users names) and you have not created this location. Remove the erroneous location:</p> <ol style="list-style-type: none"> 1. In the Cyber Protect console, navigate to Backup Storage. 2. Find the location and click the cross (x) icon to delete it. 3. Confirm your choice by clicking Delete. 4. Retry the failover.
Failback is canceled	An alert is generated when the failback was canceled by the user.	Manually dismiss the alert from the console.
VPN connection error	An alert is generated when the VPN connection failure occurs due to reasons not depend on the user's actions. Status report from VPN appliance is outdated.	<p>In case you have faced an issue with deploying or connecting Acronis VPN appliance, please contact Acronis Support.</p> <p>Please send the following information with your email:</p> <ul style="list-style-type: none"> • Screenshots of the error messages (if there are any) • Screenshot of the Acronis VPN Appliance CLI interface • Your Acronis Backup Cloud data center and group name.
(Vpn Unreachable) Connectivity gateway is not reachable	An alert is generated when the DR service can't reach connectivity gateway. Status report from connectivity gateway is outdated.	<p>In case you have faced an issue with deploying or connecting Acronis VPN appliance, please contact Acronis Support.</p> <p>Please send the following information with your email:</p> <ul style="list-style-type: none"> • Screenshots of the error messages (if there are any) • Screenshot of the Acronis VPN Appliance CLI interface • Your Acronis Backup Cloud data center and group name

Alert	Description	How to resolve the alert
DR IP reassignment required	An alert is generated if VPN appliance detects network changes.	Reassign the IP address. For more information, see Reassigning IP addresses .
Connectivity gateway failure	An alert is generated when it failed to deploy VPN server in the cloud.	Use Connection Verification Tool and check its output for errors. Allow Acronis software through application control of your firewalls and antimalware software.
Primary server creation failure	An alert is generated when the primary server was not created due to error.	
Recovery server creation failure	An alert is generated when the recovery server was not created due to error.	Make sure the recovery server matches the Software requirements .
Delete Primary Server	An alert is generated when a primary server is deleted.	
Server recovery failure	An alert is generated when the primary or recovery server failed to recover.	Find the details. If the error message is generic or unclear, for example "Internal error", navigate to Disaster Recovery → Servers , click to select the affected machine and click Activities . Click an activity, hold ctrl and left-click the activity. Now you will be able to see the ellipsis (...) sign near every activity. Click and select Task activity info .
Backup failed	An alerts is generated when the backup of cloud server (primary or server in production failover state) failed.	<ol style="list-style-type: none"> 1. Verify the connection of the backup location. 2. Check the backup storage device (local backups).
Network limit exceeds	An alert is generated when the maximum number of cloud networks is reached (5 networks).	
Runbook failure	An alert is generated when the runbook execution failed.	It does not affect the product functionality, and it can be safely ignored. For more information, see Creating a runbook .

Alert	Description	How to resolve the alert
Runbook warning	An alert is generated when the runbook execution is completed with warnings.	It does not affect the product functionality, and it can be safely ignored. For more information, see Creating a runbook .
Runbook User Interaction Required	An alert is generated when the runbook is waiting for user interaction.	It does not affect the product functionality, and it can be safely ignored. For more information, see Creating a runbook .
Internet traffic blocked	An alert is generated when the internet traffic was blocked by the administrator.	
Internet traffic unblocked	An alert is generated when the internet traffic was unblocked by the administrator.	
Local networks overlap	An alert is generated when identical or overlapping local networks is detected.	
Licensing switch insufficient server quota	An alert is generated when the cloud servers quota is not enough.	<ul style="list-style-type: none"> • Make sure the tenant and user have Web hosting servers quota or Servers quota available for a physical server. • Make sure the tenant and user have Web hosting servers quota or Virtual machines quota available for a virtual server. A virtual server cannot use Servers quota.
Licensing switch insufficient offering item	An alert is generated when the disaster recovery storage offering item is disabled.	For more information, see Disaster recovery quotas .
Licensing switch error	An alert is generated when the disaster recovery upgrade encountered an error.	
Licensing switch insufficient compute points	An alert is generated when there are no compute points available.	In the management portal, check and increase hard quota for Compute points.
Licensing switch insufficient servers offering items	An alert is generated when the cloud servers offering item is disabled.	

Alert	Description	How to resolve the alert
Policy failed to create recovery server	An alert is generated when an error occurred while setting up the disaster recovery infrastructure.	Manually create Recovery Server without the Internet Access property. For more information, see Creating recovery server
Backup processor auto test failover rescheduled	An alert is generated when the automated test failover run was rescheduled.	
Backup processor auto test failover timeout reached	<p>An alert is generated when the automated test failover operation expired.</p> <hr/> <p>Note Each Automated Test Failover run will consume chargeable compute points.</p> <hr/>	
Backup processor auto test failover overall failure	An alert is generated when the last scheduled automated test failover of the recovery server failed.	<ol style="list-style-type: none"> 1. Start a test failover of the recover server manually. For more information, see Performing a test failover. 2. Wait for the next scheduled date when automatic test failover will be performed
Failback data transfer error	An alert is generated when failback data transfer fails.	
Failback failed	An alert is generated when there is an error in the failback.	<p>You can see the erroneous location in the list of backup storages: it has a number instead of a name (normally, a location name matches one of the existing end users names) and you have not created this location. Remove the erroneous location:</p> <ol style="list-style-type: none"> 1. In Cyber Protection, navigate to backup storage. 2. Find the location and click the cross (x) icon to delete it. 3. Confirm your choice by clicking Delete. <p>Retry the failover.</p>
Failback confirming failed	An alert is generated when the	

Alert	Description	How to resolve the alert
	failback confirmation failed.	
Failback machine is ready for switchover	An alert is generated when the machine is ready for switchover.	
Failback switchover finished	An alert is generated when the switchover is successful.	Manually dismiss the alert from the console.
Failback target agent offline	An alert is generated when the agent is offline.	

Antimalware protection alerts

Alert	Description	How to resolve the alert
Suspicious remote connection activity is detected	An alert is generated when ransomware coming from a remote connection is detected.	Manually dismiss the alert from the console.
Suspicious activity is detected	An alert is generated when ransomware is detected in the workload.	<p>Manually dismiss the alert from the console. to deactivate the alert.</p> <p>Depending on the option you have specified in Active Protection plan, the malicious process is stopped, the changes made by the process are reverted or none actions have been taken yet and you need to resolve this issue manually.</p> <p>Read details of the alert to find out which process is encrypting files and which files are affected.</p> <p>If you decide that the process encrypting the files is sanctioned (false-positive alert), add this process to Trusted processes:</p> <ol style="list-style-type: none"> 1. Open Active Protection plan. 2. Click Edit to modify the settings. 3. In Trusted processes, specify trusted processes that will never be considered ransomware. Specify the full path to the process executable, starting with the drive letter. For example: C:\Windows\Temp\er76s7sdkh.exe.
Cryptomining activity is detected	An alert is generated when Illicit cryptominers are detected in the workload	Manually dismiss the alert from the console.

Alert	Description	How to resolve the alert
MBR defence: Suspicious activity is detected and suspended	An alert is generated when ransomware is detected in the workload (specifically MBR / GPT partition is modified by ransomware).	Manually dismiss the alert from the console.
Unsupported network path is specified	An alert is generated when the recovery path provided by the administrator is not a local folder path.	Specify the local path for network folder protection (recovery path). Manually dismiss the alert from the console
Critical process is added as harmful to the Active Protection plan	An alert is generated when a critical process is added as a blocked process in the Protection exclusions list.	Manually dismiss the alert from the console.
Failed to apply Active Protection policy	An alert is generated when Active Protection policy failed to be applied.	Check the error message to see why Active Protection policy cannot be applied.
Secure Zone: Unauthorized operation is detected and blocked	An alert is generated when ransomware is detected in the workload (ASZ partition is modified by ransomware).	Manually dismiss the alert from the console.
Active Protection service is not running	An alert is generated when the Active Protection service crashed / is not running.	Check the error message to see why Active Protection service is not running.
Active Protection service is not available	An alert is generated when the Active Protection service is not available because a driver is incompatible or missing.	Check Windows event logs for crashes of Acronis Active Protection service (acronis_protection_service.exe).
Conflict with another security solution	An alert is generated if Active Protection is not available for machine '{{resourceName}}' because a conflict with another security solution was detected. To enable Active Protection, disable or uninstall the conflicting security solution.	Solution 1: If you want to use Acronis real-time protection then uninstall third-party antivirus on the machine. Solution 2: If you want to use the third-party antivirus, disable Acronis real-time protection, URL filtering and Windows defender antivirus in the protection plan.
Quarantine action failed	An alert is generated when antimalware failed to quarantine a detected	Check the error message to see why quarantine failed.

Alert	Description	How to resolve the alert
	malware.	
Malicious process is detected	An alert is generated when a malware (process type) is detected by Behavior engine. The detected malware is quarantined.	Manually dismiss the alert from the console.
Malicious process is detected, but not quarantined	An alert is generated when a malware (process type) is detected by Behavior engine. The detected malware is not quarantined.	Manually dismiss the alert from the console.
Malware is detected and blocked (ODS)	An alert is generated when a malware is detected by scheduled scan. The detected malware is quarantined.	Manually dismiss the alert from the console.
Malware is detected and blocked (RTP)	An alert is generated when a malware is detected by Real-Time protection. The detected malware is quarantined.	Manually dismiss the alert from the console.
Malware is detected in a backup	An alert is generated when a malware is detected during backup scanning.	Manually dismiss the alert from the console.
Conflict detected between Real-time antimalware protection and a security product	An alert is generated when antimalware failed to register with Windows Security Center.	Disable or uninstall 3rd party security product, or disable Real-time antimalware protection in the protection plan.
Failed to run the Microsoft Security Essentials module	An alert is generated when it failed to run the Microsoft Security Essentials module.	Check the error message to see why Microsoft Security Essentials module failed to run.
Real-time protection is not available because third-party antivirus software is installed	An alert is generated when Real-time protection failed to turn on, because 3rd party antivirus still have Real-time protection enabled.	Disable or uninstall 3rd party security product, or disable Real-time antimalware protection in the protection plan.
Real-time protection is not available due to incompatible or missing driver	An alert is generated when Real-time protection is not available due to	Check the error message to see why Acronis failed to install driver on workload.

Alert	Description	How to resolve the alert
	incompatible or missing driver.	
Cyber Protection (or Active Protection) service is not responding	An alert is generated when Cyber Protection Service responds to health check ping from console.	Manually dismiss the alert from the console.
Security definition update failed	An alert is generated when security definition update failed.	Check the error message to see why security definition update failed.
Tamper Protection is enabled	An alert is generated when Microsoft Defender settings cannot be changed because tamper protection is enabled.	Disable Tamper Protection settings on the Windows workload.
Windows Defender module execution failed	An alert is generated when Windows Defender module execution failed.	Check the error message to see why Windows defender module failed to run.
Windows Defender is blocked by a third-party antivirus software	An alert is generated when Windows Defender is blocked because a third party Antivirus is installed on the machine.	Disable or uninstall 3rd party security product.
Group policy conflict	An alert is generated when Microsoft Defender settings cannot be changed because it is controlled by a group policy.	Disable group policy settings on the Windows workload.
Microsoft Security Essentials took action to protect this machine from malware	An alert is generated when Microsoft Security Essential deleted / quarantined a malware.	Manually dismiss the alert from the console.
Microsoft Security Essentials detected malware	An alert is generated when Microsoft Security Essentials detected malware or other potentially unwanted software.	Manually dismiss the alert from the console.

Licensing alerts

Alert	Description	How to resolve the alert
Storage quota almost reached	An alert is generated when the usage drops below 80% (after cleanup or quota upgrade).	Consider purchasing additional storage or freeing up space in your cloud storage.
Storage quota exceeded	An alert is generated when all 100% of the storage quota is used.	Buy more storage space. For more information on how to do that, verify the how to purchase more cloud storage .
Workload quota reached	An alert is generated when usage for offering item > 0 and usage > quota, but usage <= quota + overage.	
Workload quota exceeded	An alert is generated when the usage for offering item > quota + overage.	
The workload has no quota to apply a backup plan (resource has no service quota)	An alert is generated when: <ul style="list-style-type: none"> The quota was removed manually: Device > Details > Service quota, and then click Change and select the No quota option. The Management Console offering item is disabled. The Management Console quota+overage value of the offering item is decreased below current usage. 	
Cannot protect a workload with assigned quota	An alert is generated when the offering item is not sufficient, and you need to have: <ul style="list-style-type: none"> a dynamic group. a backup plan assigned to that group. you added a resource that falls to that dynamic group, but has some qualities that forbid applying the same backup plan to it. 	
Subscription license expired	An alert is generated when the daily check for license/maintenance expiration alerts, asked the license server, and got the response that the	After a subscription expires, all product functionality except recovery is blocked until further subscription renewal. Backed

Alert	Description	How to resolve the alert
	license is expired.	<p>up data is still accessible for recovery. Purchase a new license.</p> <hr/> <p>Note If you have recently purchased a new subscription but still receive the message that subscription is expired, you need to import new subscription from Acronis Account: in Management Console, go to Settings -> Licenses and click Sync in the top right corner. Subscriptions will be synchronized.</p>
Subscription license will expire soon	An alert is generated when the daily check for license/maintenance expiration alerts, asked the license server, and got the response that the license will expire in less than 30 days.	Consider purchasing a new subscription.

URL Filtering alerts

Alert	Description	How to resolve the alert
Malicious URL was blocked	An alert is generated when a malicious URL is blocked by URL filtering.	Check the URL filtering settings. URL filtering is blocking pages which are supposed to be blocked according to the URL filtering settings.
A malicious URL warning was ignored	An alert is generated when you selected to proceed with the malicious URL blocked by URL filtering.	Check the URL filtering settings.
Conflict detected between URL filtering and a security product	An alert is generated when the URL filtering cannot be enabled due to a conflict with another security product.	Check the URL filtering settings.
Website URL is blocked	An alert is generated when a URL meets all the criteria specified in the blocked category for URL	Check the URL filtering settings.

Alert	Description	How to resolve the alert
	filtering.	

EDR alerts

Alert	Description	How to resolve the alert
Incident Detected	An alert is generated when an incident is created or when the status for an existing incident is updated.	This alert informs you about a new incident or if an old incident has been updated. You can view the alert and close it. You can choose to open the incident for further investigation if required.
Indicator of compromise (IOCs) detected	An alert is generated when a new indicator of compromise was detected by EDR IOC threat search service.	This alert is to inform you that an IOC has been detected on one or many workloads. You will view the alert and then you can click on the link in the alert to view details about the IOC.
Failed to isolate the workload from the network	An alert is generated when the user triggers the action to isolate the machine from network, and isolation action fails.	Take the necessary actions.
Failed to reconnect the workload to the network	An alert is generated when the user triggers the action to reconnect the machine back to network, and the action failed.	Take the necessary actions.
Windows Defender Firewall settings was modified	An alert is generated when the settings to the firewall were modified on isolated machine.	This alert is to inform you that firewall details were modified on the isolated machine. It is informative only and you can close the alert after viewing it.

Device Control alerts

Alert	Description	How to resolve the alert
Device control and Data loss prevention will run with limited functionality (Incompatible CPU detected)	An alert is generated when the DeviceLock agent started on physical machine with CPU which has supporting for CET technology.	Disable the option on the affected machines to avoid alerts.
Device control functionality is not yet supported on macOS Ventura	An alert is generated when DeviceLock agent started on	

Alert	Description	How to resolve the alert
	physical macOS Ventura machine, and the protection plan with Device Control is applied to the agent. Applicable only for versions when there is a problem with the kernel panic due DeviceLock driver.	
Allowed transfer of sensitive data	An alert is generated when transferring for sensitivity content is allowed.	
Justified transfer of sensitive data	An alert is generated when transferring sensitivity content is justified.	
Denied transfer of sensitive data	An alert is generated when transferring sensitivity content is blocked.	
Review the results of Data Loss Prevention observation mode	<p>An alert is generated when it is time to review the Observation results:</p> <ul style="list-style-type: none"> • Advanced DLP Pack license is not applied. • A month passed since the Observation mode was enabled in any Protection plan applied to at least one workload. • A month passed since the last similar alert was raised and some usage of DLP in Observation mode is detected. 	
Security identifier was changed for user	An alert is generated when we have the situation when a SID is updated for known username. This might happen when OS is re-installed on a non-domain PC.	
Peripheral device access is blocked	An alert is generated when some actions (read/write operations) for supported devices are blocked.	
Unable to connect to a remote SSL resource.	An alert is generated when the access to a remote SSL resource is blocked due to additional handshake prevention used at the resource.	Add the resource to the allowlist for remote hosts.

System alerts

Alert	Description	How to resolve the alert
Agent is outdated	An alert is generated when the agent version is outdated.	Go to Agents list and initiate updating the agent.
Automatic update failed	An alert is generated when the agent auto update failed.	Try to perform a manual update.
You need to restart device after installing a new agent	An alert is generated when a reboot is required after remote install was successful.	Restart the workload.
Activity failed	An alert is generated when an activity failed.	Restart all Acronis services on the machine.
Activity succeeded with warnings	An alert is generated when an activity was successful but some warnings were generated.	
Activity is not responding	An alert is generated when an activity in progress is not responding.	
Plan deployment failed	An alert is generated when the protection plan deployment failed.	
Failed to convert user name to SID	An alert is generated when the schedule SID conversion failed.	






Alert widgets

In the alert widgets, you can see the following details of alerts related to your workload:

Field	Description
5 latest alerts widget	A list of five latest alerts.
Historical alerts summary	A graphical widget showing alerts by alert severity, alert type and the time range.
Active alerts summary	A graphical widget showing active alerts by alert severity and alert type, as well as the sum of active alerts.
Alerts history	A table view of historical alerts.
Active alerts details	A table view of active alerts.

Cyber Protection

This widget shows the overall information about the size of backups, blocked malware, blocked URLs, found vulnerabilities, and installed patches.

Cyber Protection				
 Backed up today 1.60 GB	 Malware blocked 0	 Malicious URLs blocked 0	 Existing vulnerabilities 347	 Patches ready to install 114
overall compressed size 2.43 GB	overall blocked 14	overall blocked 4	overall found 819	overall installed 5

The upper row shows the current statistics:

- **Backed up today** – the sum of recovery point sizes for the last 24 hours
- **Malware blocked** – the number of currently active alerts about malware blocked
- **URLs blocked** – the number of currently active alerts about URLs blocked
- **Existing vulnerabilities** – the number of currently existing vulnerabilities
- **Patches ready to install** – the number of currently available patches to be installed

The lower row shows the overall statistics:

- The compressed size of all backups
- The accumulated number of blocked malware across all machines
- The accumulated number of blocked URLs across all machines
- The accumulated number of discovered vulnerabilities across all machines
- The accumulated number of installed updates/patches across all machines


Protection status

This widget shows the current protection status for all machines.

A machine can be in one of the following statuses:

- **Protected** – machines with applied protection plan.
- **Unprotected** – machines without applied protection plan. These include both discovered machines and managed machines with no protection plan applied.
- **Managed** – machines with installed protection agent.
- **Discovered** – machines without installed protection agent.

If you click on the machine status, you will be redirected to the list of machines with this status for more details.



Discovered machines

This widget shows the list of discovered machines during the specified time range.

Discovered machines					
Device name ↑	IP address	OS	Organizational unit	Discovery type	⚙️
▼ Windows Server 2012 R2					
win-6g34mv70qa3	10.248.90.221	Windows Server 2012 R2	-	Local Network	
▼ Windows 10 Enterprise 2016 LTSB					
device1	10.248.90.238	Windows 10 Enterprise 2016 LTSB	OU1	Active Directory, Local Network	
device2	-	Windows 10 Enterprise 2016 LTSB	OU1	Active Directory	
device3	10.248.91.243	Windows 10 Enterprise 2016 LTSB	OU1	Active Directory, Manual, Loc...	
device4	10.248.91.125	Windows 10 Enterprise 2016 LTSB	-	Active Directory, Local Network	
device5	-	Windows 10 Enterprise 2016 LTSB	-	Active Directory, Manual	
▼ -					
-	10.250.41.189	-	-	Manual	
-	10.248.44.199	-	-	Manual	

Endpoint Detection and Response (EDR) widgets

Endpoint Detection and Response (EDR) includes seven widgets, all of which can be accessed from the **Overview** dashboard; three of these widgets are also displayed by default within the EDR functionality (see "Reviewing incidents" (p. 837)).

The seven widgets available are:

- Top incident distribution per workload
- Threat status (displayed in EDR)
- Incident severity history (displayed in EDR)
- Security incident MTTR

- Security incident burndown
- Detection by tactics (displayed in EDR)
- Workload network status

Top incident distribution per workload

This widget displays the top five workloads with the most incidents (click **Show all** to redirect to the incident list, which is filtered according to the widget settings).


Hover over a workload row to view a breakdown of the current investigation state for the incidents; the investigation states are **Not started**, **Investigating**, **Closed**, and **False positive**. Then click on the workload you want to analyze further; the incident list is refreshed according to the widget settings.

Top Incident distribution per workload		
 SCRANTON		123
 qa-gw3t68hh		41
 RG_345		32
 Georgy_Win_64		11
 w_35jf_4		12
Show all		

Threat status

This widget displays the current threat status for all workloads, highlighting the current number of incidents that are not mitigated and that need investigating. The widget also indicates the number of incidents that were mitigated (manually and/or automatically by the system).


Click on the **Not mitigated** number to display the incident list filtered to show incidents that are not mitigated.

Threat status	
Not Mitigated	
 2	
Automatically mitigated	2
Manually mitigated	2
Total	6

Incident severity history

This widget displays the evolution of attacks by severity, and can help indicate attack campaigns. When spikes are visible, this can indicate that the organization is under attack.

Hover over the graph to view a breakdown of the incident history at a specific point in the previous 24 hours (the default period). Click on the severity level (**Critical**, **High**, or **Medium**) if you want to view the list of related incidents; you are redirected to the incident list pre-filtered with incidents matching the selected severity level.




Security incident MTTR

This widget displays the average resolution time for security incidents. It indicates how quickly incidents are being investigated and resolved.

Click on a column to view a breakdown of the incidents according to severity (**Critical**, **High**, and **Medium**), and an indication of how long it took to resolve the different severity levels. The % value shown in parentheses indicates the increase or decrease in comparison to the previous time period.

Incident MTTR




Security incident burndown

This widget shows the efficiency rate in closing incidents; the number of open incidents are measured against the number of closed incidents over a period of time.

Hover over a column to view a breakdown of the closed and open incidents for the selected day. If you click the Open value, the incident list is displayed, and filtered to display incidents currently open (in the **Investigating** or **Not started** states). If you click the Closed value, the incident list is displayed, and filtered to display incidents that are no longer open (in the **Closed** or **False positive** states).

The % value shown in parentheses indicates the increase or decrease in comparison to the previous time period.



Detection by tactics

This widget displays the number of times specific attack techniques have been found in incidents during the selected period.

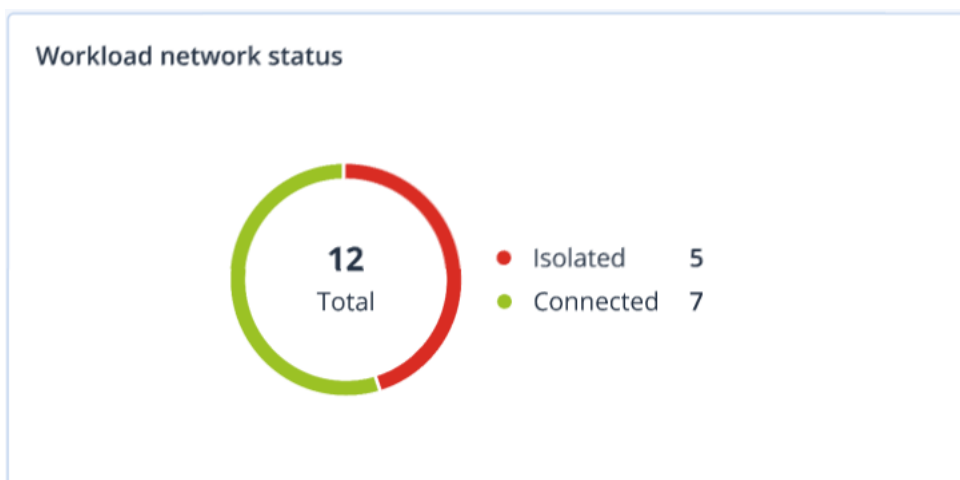
The values in green and red indicate if there has been an increase or decrease over the previous time period. In the example below, Privilege Escalation and Command and Control attacks have seen an increase over the previous time period; this could indicate that your credential management needs to be analyzed and security enhanced.

Detection by tactics			
Initial Access	3	Discovery	3
Execution	7	Lateral Movement	0
Persistence	15	Collection	15
Privilege Escalation	31	Command and Control	31
Defense Evasion	23	Exfiltration	23
Credential Access	7	Discovery	0
Impact	0	Resource Development	0

Workload network status

This widget displays the current network status of your workloads, and indicates how many workloads are isolated and how many are connected.

Click the Isolated value to view the Workload with agents list (under the **Workloads** menu in the Cyber Protect console), which is filtered to display isolated workloads. Click the Connected value to view the Workload with agents list filtered to display connected workloads.



#CyberFit Score by machine

This widget shows for each machine the total #CyberFit Score, its compound scores, and findings for each of the assessed metrics:

- Antimalware
- Backup
- Firewall
- VPN

- Encryption
- NTLM traffic

To improve the score of each of the metrics, you can view the recommendations that are available in the report.

For more details about the #CyberFit Score, refer to "[#CyberFit Score for machines](#)".

#CyberFit Score by machine ?			
Metric	#CyberFit Score	Findings	
▼ DESKTOP-2N2TRE8	625 / 850		
Anti-malware	275 / 275	You have anti-malware protection enabled	
Backup	175 / 175	You have a backup solution protecting your data	
Firewall	175 / 175	You have a firewall enabled for public and private networks	
VPN	0 / 75	No VPN solution was found, your connection to public and shared networks is n...	
Encryption	0 / 125	No disk encryption was found, your device is at risk from physical tampering	
NTLM traffic	0 / 25	Outgoing NTLM traffic to remote servers is not denied, your credentials may be ...	

Disk health monitoring

Disk health monitoring provides information about the current disk health status and a forecast about it, so that you can prevent data loss that might be related to a disk failure. Both HDD and SSD disks are supported.

Limitations

- Disk health forecast is supported only for machines running Windows.
- Only disks of physical machines are monitored. Disks of virtual machines cannot be monitored and are not shown in the disk health widgets.
- RAID configurations are not supported. The disk health widgets do not include any information about machines with RAID implementation.
- NVMe SSDs are not supported.


The disk health is represented by one of the following statuses:

- **OK**
Disk health is between 70% and 100%.
- **Warning**
Disk health is between 30% and 70%.
- **Critical**
Disk health is between 0% and 30%.
- **Calculating disk data**
The current disk status and forecast are being calculated.

How it works

The Disk Health Prediction Service uses an AI-based prediction model.

1. The protection agent collects the SMART parameters of the disks and passes this data to the Disk Health Prediction Service:
 - SMART 5 – Reallocated sectors count.
 - SMART 9 – Power-on hours.
 - SMART 187 – Reported uncorrectable errors.
 - SMART 188 – Command timeout.
 - SMART 197 – Current pending sector count.
 - SMART 198 – Offline uncorrectable sector count.
 - SMART 200 – Write error rate.
2. The Disk Health Prediction Service processes the received SMART parameters, makes forecasts, and then provides the following disk health characteristics:
 - Disk health current state: OK, warning, critical.
 - Disk health forecast: negative, stable, positive.
 - Disk health forecast probability in percentage.The prediction period is one month.
3. The Monitoring Service receives these characteristics, and then shows the relevant information in the disk health widgets in the Cyber Protect console.




Disk health widgets

The results of the disk health monitoring are presented in the following widgets that are available in the Cyber Protect console.

- **Disk health overview** is a treemap widget with two levels of detail that can be switched by drilling down.
 - Machine level
 - Shows summarized information about the disk health status of the selected customer machines. Only the most critical disk status is shown. The other statuses are shown in a tooltip when you hover over a particular block. The machine block size depends on the total size of all disks of the machine. The machine block color depends on the most critical disk status found.


Disk health overview

Resources




- Disk level
Shows the current disk health status of all disks for the selected machine. Each disk block shows one of the following disk health forecasts and its probability in percentage:
 - Will be degraded
 - Will stay stable

- Will be improved



- **Disk health status** is a pie chart widget that shows the number of disks for each status.



Disk health status alerts

The disk health check runs every 30 minutes, while the corresponding alert is generated once a day. When the disk health changes from **Warning** to **Critical**, an alert always is generated.

Alert name	Severity	Disk health status	Description
Disk failure is possible	Warning	(30 - 70)	The <disk name> disk on this machine is likely to fail in the future. Run a full image backup of this disk as soon as possible, replace it, and then recover the image to the new disk.
Disk failure is imminent	Critical	(0 - 30)	The <disk name> disk on this machine is in a critical state, and will most likely fail very soon. We do not recommend an image backup of this disk at this point, as the added stress can cause the disk to fail. Back up the most important files on this disk immediately and replace it.

Data protection map

Note

This feature is available with the Advanced Backup pack.

The data protection map feature allows you to discover all data that are important for you and get detailed information about number, size, location, protection status of all important files in a treemap scalable view.


Each block size depends on the total number/size of all important files that belong to a customer/machine.

Files can have one of the following protection statuses:

- **Critical** – there are 51-100% of unprotected files with the extensions specified by you that are not being backed up and will not be backed up with the existing backup settings for the selected machine/location.
- **Low** – there are 21-50% of unprotected files with the extensions specified by you that are not being backed up and will not be backed up with the existing backup settings for the selected machine/location.
- **Medium** – there are 1-20% of unprotected files with the extensions specified by you that are not being backed up and will not be backed up with the existing backup settings for the selected machine/location.
- **High** – all files with the extensions specified by you are protected (backed up) for the selected machine/location.

The results of the data protection examination can be found on the monitoring dashboard, in the Data Protection Map widget, a treemap widget that shows details on a machine level:

- Machine level – shows information about the protection status of important files per machines of the selected customer.



To protect files that are not protected, hover over the block and click **Protect all files**. In the dialog window, you can find information about the number of unprotected files and their location. To protect them, click **Protect all files**.

You can also download a detailed report in CSV format.

Vulnerability assessment widgets

Vulnerable machines

This widget shows the vulnerable machines by the vulnerability severity.

The found vulnerability can have one of the following severity levels according to the [Common Vulnerability Scoring System \(CVSS\) v3.0](#):

- Secured: no vulnerabilities are found
- Critical: 9.0 - 10.0 CVSS
- High: 7.0 - 8.9 CVSS
- Medium: 4.0 - 6.9 CVSS

- Low: 0.1 - 3.9 CVSS
- None: 0.0 CVSS



Existing vulnerabilities

This widget shows currently existing vulnerabilities on machines. In the **Existing vulnerabilities** widget, there are two columns showing timestamps:

- **First detected** – date and time when a vulnerability was detected initially on the machine.
- **Last detected** – date and time when a vulnerability was detected the last time on the machine.

Existing vulnerabilities							
Machine name	Vendor	Product	Vulnerability name/ID	Severity ↓	Last detected	First detected	⚙
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-7096	Critical	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0856	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0688	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0739	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0752	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0753	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0806	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0810	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0812	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0829	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	

[More](#)

Patch installation widgets

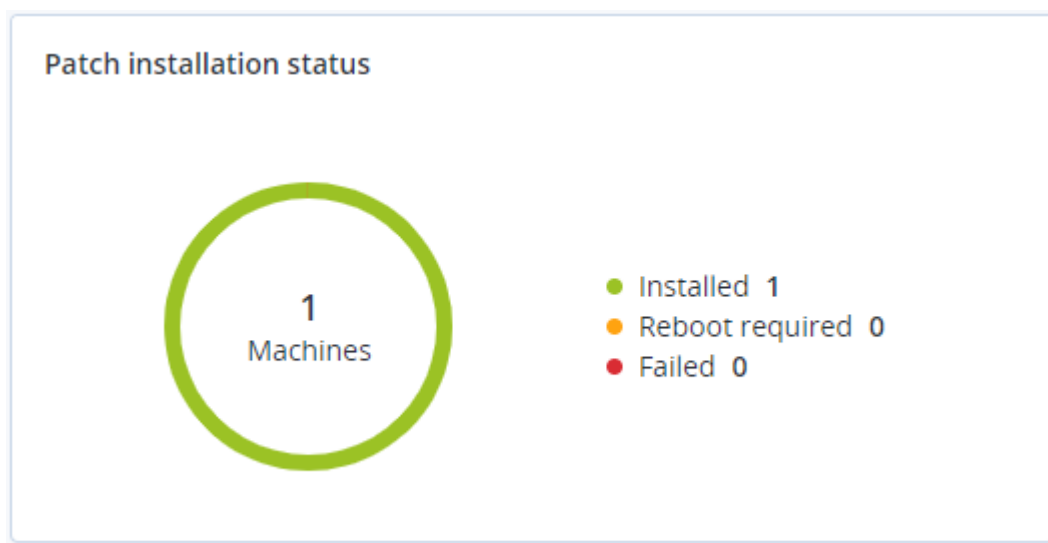
There are four widgets related to the patch management functionality.

Patch installation status

This widget shows the number of machines grouped by the patch installation status.

- **Installed** – all available patches are installed on a machine
- **Reboot required** – after patch installation reboot is required for a machine

- **Failed** – patch installation failed on a machine



Patch installation summary

This widget shows the summary of patches on machines by the patch installation status.

Patch installation summary							
Installation status	Total number of machines	Total number of updates	Microsoft updates	Application updates	Critical severity	High severity	Medium severity
Installed	1	2	1	1	2	0	0

Patch installation history

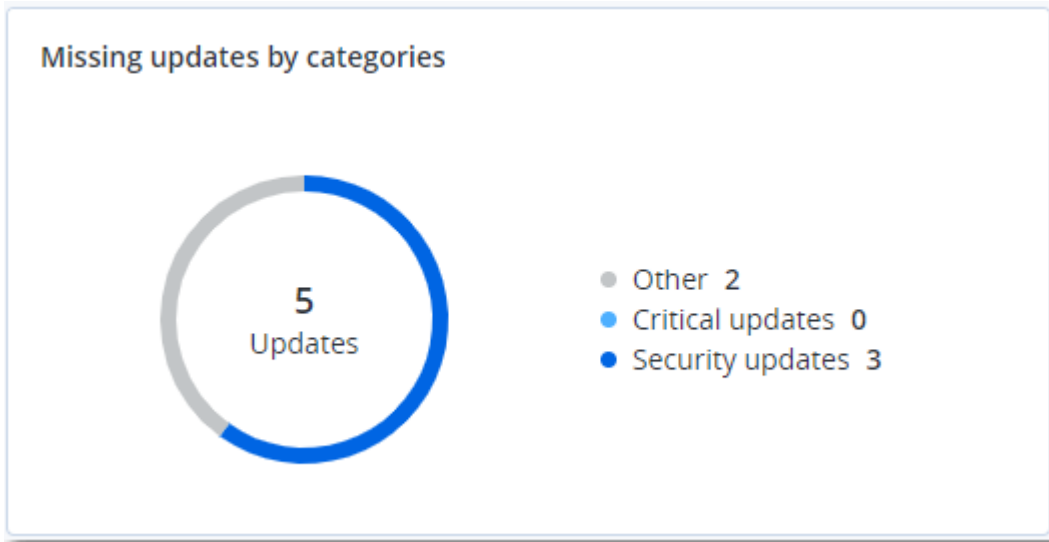
This widget shows the detailed information about patches on machines.

Machine name	Update name	Version	Severity	Approval status	Installation status	Installation date
NIKITATIKHOC4E5	FastStone Soft FastStone I...	5.9	Medium	New	Installed	02/05/2020
NIKITATIKHOB524	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020
NIKITATIKHOB524	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Oracle Java Runtime Envir...	8.0.2410.7	High	New	Failed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020

Missing updates by categories

This widget shows the number of missing updates per category. The following categories are shown:

- Security updates
- Critical updates
- Other



Backup scanning details

This widget shows the detailed information about the detected threats in backups.

Backup scanning details (threats)							
Device name	Plan name	Backup Date and time	Contents type	Location	Threat name	Affected files	Date and time
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Gen.Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 11:40 AM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 11:40 AM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Gen.Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 11:45 AM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 11:45 AM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Gen.Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 11:50 AM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 11:50 AM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Gen.Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 1:10 PM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 1:10 PM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Gen.Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 1:33 PM
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	████████	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 1:33 PM

Recently affected

This widget shows detailed information about workloads that were affected by threats, such as viruses, malware, and ransomware. You can find information about the detected threats, the time when the threats were detected, and how many files were affected.

Recently affected					
Machine name	Protection plan	Threat	Affected files	Detection time	
Ubuntu_14.04_x64-1	Total protection	Adware.DealPlylgen2	15	27.12.2	
Ubuntu_14.04_x64-1	Protection plan	Bloodhound.MalMacroIg1	274	27.12.2	
dc_w2k12_r2	Protection plan	Backdoor:Win32/Caphaw...	13	27.12.2	
Win2012_r2-Hyper-V	Protection plan	W97M.DownloaderIg32	5	27.12.2	
HyperV_for12A	Total protection	Miner.XMRigIgen1	68	27.12.2	
vm-sql_2012	Total protection	Backdoor:Win32/Caphaw...	61	27.12.2	
vm-sql_2012	Protection plan	Adware.DealPlylgen2	9	27.12.2	
MF_2012_R2	Total protection	MSH.DownloaderIgen8	73	27.12.2	
MF_2012_R2	Protection plan	Bloodhound.MalMacroIg1	182	27.12.2	
MF_2012_R2	Protection plan	Bloodhound.MalMacroIg1	18	27.12.2	
ESXirestore	Protection plan	MSH.DownloaderIgen8	682	27.12.2017 11:23 AM	
MF_2012_R2	Protection plan	Miner.XMRigIgen1	13	27.12.2017 11:23 AM	
Ubuntu_14.04_x64-1	Total protection	Adware.DealPlylgen2	3	27.12.2017 11:23 AM	
Win2012_r2-Hyper-V	Total protection	W97M.DownloaderIg32	27	27.12.2017 11:23 AM	

- Folder
- Customer
- ✓ Machine name
- ✓ Protection plan
- Detected by
- ✓ Threat
- File name
- File path
- ✓ Affected files
- ✓ Detection time

Downloading data for recently affected workloads

You can download the data for the recently affected workloads, generate a CSV file, and send it to the recipients that you specify.

To download the data for the recently affected workloads

1. In the **Recently affected** widget, click **Download data**.
2. In the **Time period** field, enter the number of days for which you want to download data. The maximum number of days that you can enter is 200.
3. In the **Recipients** field, enter the email addresses of all the people who will receive an email with a link for downloading the CSV file.
4. Click **Download**.

The system starts generating the CSV file with the data for the workloads that were affected in the time period that you specified. When the CSV file is complete, the system sends an email to the recipients. Each recipient can then download the CSV file.

Cloud applications

This widget shows detailed information about cloud-to-cloud resources:

- Microsoft 365 users (mailbox, OneDrive)
- Microsoft 365 groups (mailbox, group site)
- Microsoft 365 public folders
- Microsoft 365 site collections
- Microsoft 365 Teams

- Google Workspace users (Gmail, Google Drive)
- Google Workspace shared drives

Cloud applications				
Device name	Protection status ↑	Last successful backup	Next backup	Number of backups
HR - Onboarding	OK	06/17/2020 10:48 AM	06/18/2020 7:34 AM	1
Sales and Marketing	OK	06/17/2020 10:49 AM	06/18/2020 4:48 AM	1
HR Leadership Team	OK	06/17/2020 10:48 AM	06/18/2020 6:51 AM	1
Retail	OK	06/17/2020 10:47 AM	06/18/2020 2:53 AM	1
Contoso	OK	06/17/2020 10:47 AM	06/17/2020 3:23 PM	1
U.S. Sales	OK	06/17/2020 10:48 AM	06/18/2020 3:30 AM	1
IT	OK	06/17/2020 10:48 AM	06/17/2020 10:35 PM	1
Mark 8 Project Team	Warning	06/17/2020 10:49 AM	06/18/2020 3:06 AM	1
Finance	OK	06/17/2020 10:47 AM	06/17/2020 4:38 PM	1
Sales	Warning	06/17/2020 10:47 AM	06/17/2020 2:06 PM	1

Additional information about cloud-to-cloud resources is also available in the following widgets:


- Activities
- Activity list
- 5 latest alerts
- Alerts history
- Active alerts summary
- Historical alerts summary
- Active alert details
- Locations summary

Software inventory widgets

The **Software inventory** table widget shows detailed information about the all the software that is installed on Windows and macOS devices in your organization.

Machine name	Software name	Software version	Vendor name	Status	Date installed	Last run	Scan time	Location	User
~ Ivelins-Mac-mini-2.local	-	15.0.26046	-	No change	-	12/12/2020, 3:28 AM	12/14/2020, 10:24 AM	/Library/Application Supp...	root
Ivelins-Mac-mini-2.local	-	5989	Apple	No change	-	12/04/2020, 10:59 AM	12/14/2020, 10:24 AM	/Applications/Pages.app	root
Ivelins-Mac-mini-2.local	-	5989	Apple	No change	-	12/04/2020, 10:59 AM	12/14/2020, 10:24 AM	/Applications/Keynote.app	root
Ivelins-Mac-mini-2.local	-	5989	Apple	No change	-	12/04/2020, 10:59 AM	12/14/2020, 10:24 AM	/Applications/Numbers.a...	root
Ivelins-Mac-mini-2.local	Canon iJScanner2	4.0.0	Canon Inc. (XE2XNR9KZ5)	No change	-	12/04/2020, 10:35 AM	12/14/2020, 10:24 AM	/Library/Image Capture/D...	root
Ivelins-Mac-mini-2.local	Canon iJScanner4	4.0.0	Canon Inc. (XE2XNR9KZ5)	No change	-	12/04/2020, 10:35 AM	12/14/2020, 10:24 AM	/Library/Image Capture/D...	root
Ivelins-Mac-mini-2.local	Canon iJScanner6	4.0.0	Canon Inc. (XE2XNR9KZ5)	No change	-	12/04/2020, 10:35 AM	12/14/2020, 10:24 AM	/Library/Image Capture/D...	root
Ivelins-Mac-mini-2.local	commandFilter	1.71	EPSON (TXAEAV5RN4)	No change	-	12/04/2020, 10:35 AM	12/14/2020, 10:24 AM	/Library/Printers/EPSON/...	root
Ivelins-Mac-mini-2.local	Cyber Protect Agent Assis...	1	Acronis International Gm...	No change	-	12/12/2020, 10:01 AM	12/14/2020, 10:24 AM	/Applications/Utilities/Cy...	root
Ivelins-Mac-mini-2.local	Cyber Protect Agent Unin...	1	Acronis International Gm...	No change	-	12/12/2020, 3:28 AM	12/14/2020, 10:24 AM	/Library/Application Supp...	root

The **Software overview** widget shows the number of new, updated, and deleted applications on Windows and macOS devices in your organization for a specified time period (7 days, 30 days, or the current month).



When you hover over a certain bar on the chart, a tooltip with the following information shows:

New - the number of newly installed applications.

Updated - the number of updated applications.

Removed - the number of removed applications.

When you click the part of the bar for a certain status, you are redirected to the **Software Management** -> **Software Inventory** page. The information in the page is filtered for the corresponding date and status.

Hardware inventory widgets

The **Hardware inventory** and **Hardware details** table widgets show information about all the hardware that is installed on physical and virtual Windows and macOS devices in your organization.

Machine name	OS name	OS version	CPU cores	Disks total size	RAM total (Gb)	Motherboard name	Motherboard seria...	BIOS version	Domain	Registered owner	Registered organiz...	Scan date and time
Ivelins-Mac-mini...	macOS 11.0.1	10.16	6	233.47 GB	8.00 GB			0.1	-	-	-	12/14/2020 10:23 ...
00003079.corp...	Microsoft Window...	10.0.16299	2	476.94 GB	11.83 GB	Base Board	L1HF6AC08PY	N1CET81W (1.49)	corp.acronis.com	User	Acronis Inc.	12/13/2020 8:18 PM

Machine name	Hardware category	Hardware name	Hardware details	Manufacturer	Status	Scan date
Ivelins-Mac-mini-2.local	Motherboard	Macmini8,1	Mac-7BA5B2DFE22DD8C	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Ethernet	Ethernet, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Wi-Fi	IEEE80211, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Bluetooth PAN	Ethernet, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Thunderbolt 1	Ethernet, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Thunderbolt 2	Ethernet, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Thunderbolt 3	Ethernet, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Thunderbolt 4	Ethernet, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Network adapter	Thunderbolt Bridge	Bridge, 00:00:00:00:00:00	-	-	12/14/2020, 10:23 AM
Ivelins-Mac-mini-2.local	Disk	disk1	APPLE SSD AP0256M, SSD, 250685575...	-	-	12/14/2020, 10:23 AM

The **Hardware changes** table widget shows information about the added, removed, and changed hardware on physical and virtual Windows and macOS devices in your organization for a specified time period (7 days, 30 days, or the current month).

Hardware changes						
Machine name	Hardware category	Status	Old value	New value	Modification date and time	
▼ DESKTOP-0FF9TTF						
DESKTOP-0FF9TTF	Network adapter	Changed	Oracle Corporation, Ethernet 802.3, ...	Oracle Corporation, Ethernet 802.3, ...	01/11/2021 9:28 AM	
DESKTOP-0FF9TTF	Network adapter	New	-	Realtek Semiconductor Corp., Ether...	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	Motherboard	New	-	LENOVO, Toronto 5C1, PF0PJB10	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	GPU	New	-	Intel(R) HD Graphics Family	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	Disk	New	-	(Standard disk drives), WDC WD10JP...	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	Network adapter	New	-	Cisco Systems, Ethernet 802.3, 00:0...	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	Network adapter	New	-	Oracle Corporation, Ethernet 802.3, ...	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	RAM	New	-	Samsung, 985D7122, 4.00 GB	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	Network adapter	New	-	TAP-NordVPN Windows Provider V9...	01/04/2021 2:37 PM	
DESKTOP-0FF9TTF	RAM	New	-	Micron, 00000000, 4.00 GB	01/04/2021 2:37 PM	

[More](#)

Remote sessions widget

This widget shows the detailed information about the remote desktop and file transfer sessions.

Remote sessions							
Start time	End time	Duration	Connection type	Protocol	Connection sou...	Accessed by	Connection des...
12/15/2022 4:...	12/15/2022 4:4...	a few seco...	Direct	Screen Sharing	RU-PC0YHMZL	sk-part	. i.1.4
12/15/2022 4:...	12/15/2022 4:4...	a few seco...	Cloud	NEAR	RU-PC0YHMZL	sk-part	fiat-virtual-mac...
12/15/2022 4:...	12/15/2022 4:4...	2 minutes	Cloud	NEAR	RU-PC0YHMZL	sk-part	ACPM-Sveta
12/15/2022 4:...	12/15/2022 4:1...	16 minutes	Cloud	NEAR	BG-PF3EJ2GZ	Boryana-part	ACPM-Sveta
12/15/2022 3:...	12/15/2022 4:0...	a minute	Cloud	NEAR	BG-PF3EJ2GZ	Boryana-part	ACPM-Sveta
12/15/2022 3:...	12/15/2022 3:5...	a few seco...	Direct	RDP	RU-PC0YHMZL	sk-part	.35.112.
12/15/2022 3:...	12/15/2022 3:4...	a few seco...	Direct	Screen Sharing	RU-PC0YHMZL	sk-part	. i.1.
12/15/2022 3:...	12/15/2022 3:4...	a few seco...	Direct	Screen Sharing	RU-PC0YHMZL	sk-part	. i.1.4
12/15/2022 1:...	12/15/2022 12:...	a few seco...	Direct	RDP	RU-PC0YHMZL	sk-part	.35.112.
12/15/2022 1:...	12/15/2022 12:...	a few seco...	Cloud	NEAR	RU-PC0YHMZL	sk-part	fiat-virtual-mac...

[More](#)

Smart protection

Threat feed

Acronis Cyber Protection Operations Center (CPOC) generates security alerts that are sent only to the related geographic regions. These security alerts provide information about malware, vulnerabilities, natural disasters, public health, and other types of global events that may affect your data protection. The threat feed informs you about all the potential threats and allows you to prevent them.

Some security alerts can be resolved by following a set of specific actions that are provided by the security experts. Other security alerts just notify you about the upcoming threats but no recommended actions are available.

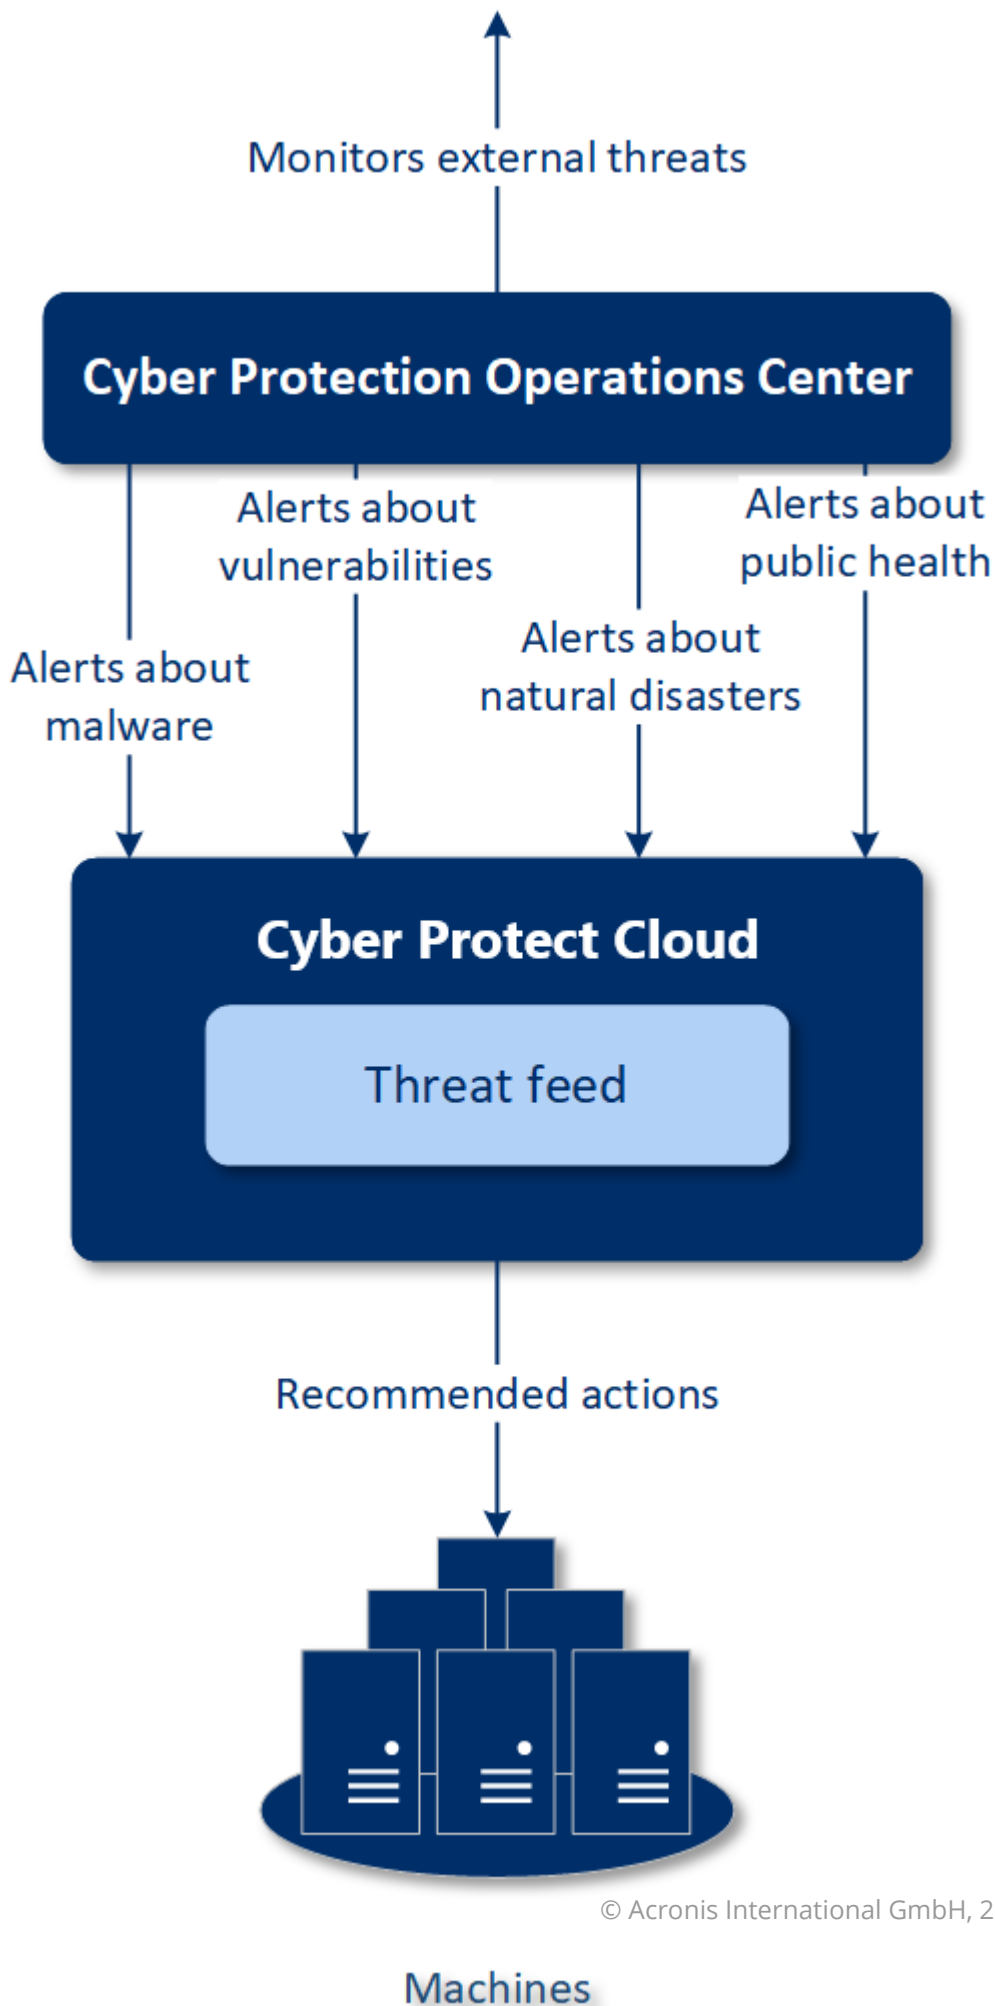
Note

Malware alerts are generated only for machines that have the agent for Antimalware protection installed.

How it works

Acronis Cyber Protection Operations Center monitors external threats and generates alerts about malware, vulnerability, natural disaster, and public health threats. You will be able to see all these alerts in the Cyber Protect console, in the **Threat feed** section. You can perform respective recommended actions depending on the type of alert.

The main workflow of the threat feed is illustrated in the diagram below.



To run the recommended actions on received alerts from Acronis Cyber Protection Operations Center, do the following:

1. In the Cyber Protect console, go to **Monitoring > Threat feed** to review if there are any existing security alerts.
2. Select an alert in the list and review the provided details.
3. Click **Start** to launch the wizard.
4. Enable the actions that you want to be performed and machines to which these actions must be applied. The following actions can be suggested:
 - **Vulnerability assessment** – to scan machines for vulnerabilities
 - **Patch management** – to install patches on the selected machines
 - **Antimalware Protection** – to run full scan of the selected machines

Note

This action is available only for machines that have the agent for Antimalware protection installed.

- **Backup of protected or unprotected machines** – to back up protected and unprotected workloads.

If there are no backups yet for the workload (in all accessible locations, cloud and local), or the existing backups are encrypted, the system creates a full backup with the following name format:

%workload_name%-Remediation

By default, the destination for the backup is the Cyber Protect Cloud storage, but you can configure another location before you start the operation.

If a non-encrypted backup already exists, the system will create an incremental backup in the existing archive.

5. Click **Start**.
6. On the **Activities** page, verify that the activity was successfully performed.

Name	Severity	Type	Date
Warning over powerful Smominru crypto mining botnet	MEDIUM	Malware	Dec 13, 2019
Acronis discovers new Autoit Cryptominer campaign injecting Windows process	HIGH	Malware	Dec 11, 2019
Manila vulnerable to major earthquake	LOW	Natural Disaster	Dec 11, 2019
Snatch ransomware reboots PCs into Safe Mode to bypass protection	HIGH	Malware	Dec 10, 2019
Caution! Ryuk ransomware decrypter damages larger files, even if you pay	MEDIUM	Malware	Dec 10, 2019
5.3 earthquake shakes New Zealand's North Island	LOW	Natural Disaster	Dec 10, 2019
Town hit by ransomware: System shut down to limit damage	MEDIUM	Malware	Dec 9, 2019
5.0M earthquake strikes Gunungkidul, Yogyakarta	LOW	Natural Disaster	Dec 9, 2019
Beware: Windows 10 update email is a ransomware trap	LOW	Malware	Dec 4, 2019
Dexphot malware uses fileless techniques to install cryptominer	LOW	Malware	Dec 4, 2019
New Chrome Password Stealer Sends Stolen Data to a MongoDB Database	LOW	Malware	Dec 2, 2019
New Malware Campaign Targets the Hospitality Industry	LOW	Malware	Dec 2, 2019
New DeathRansomware started encrypting files for real	HIGH	Malware	Nov 28, 2019
Docker platforms are targeted by hackers to deliver cryptominer malware	MEDIUM	Malware	Nov 28, 2019
Fake software update tries to download malware	MEDIUM	Malware	Nov 25, 2019
New malware DePrinMon registers as Default Print Monitor	MEDIUM	Malware	Nov 22, 2019

Deleting all alerts

Automatic clean-up from the threat feed is made after the following time periods:

- Natural disaster – 1 week
- Vulnerability – 1 month
- Malware – 1 month
- Public health – 1 week

Data protection map

The Data protection map functionality allows you

- To get detailed information about stored data (classification, locations, protection status, and additional information) on your machines.
- To detect whether data are protected or not. The data are considered protected if they are protected with backup (a protection plan with the backup module enabled).
- To perform actions for data protection.

How it works

1. First, you create a protection plan with the [Data protection map module](#) enabled.
2. Then, after the plan was performed and your data were discovered and analyzed, you will get the visual representation of data protection on the [Data protection map](#) widget.
3. You can also go to **Devices > Data protection map** and find there information about unprotected files per device.
4. You can take actions to protect the detected unprotected files on devices.

Managing the detected unprotected files

To protect the important files that were detected as unprotected, do the following:

1. In the Cyber Protect console, go to **Devices > Data protection map**.
In the list of devices, you can find general information about the number of unprotected files, size of such files per device, and the last data discovery.
To protect files on a particular machine, click the Ellipsis icon and then **Protect all files**. You will be redirected to the list of plans where you can create a protection plan with the backup module enabled.
To delete the particular device with unprotected files from the list, click **Hide until next data discovery**.
2. To view a more detailed information about the unprotected files on a particular device, click on the name of the device.
You will see the number of unprotected files per extension and per location. Define the extensions in the search field, for which you want to get the information about unprotected files.

- To protect all unprotected files, click **Protect all files**. You will be redirected to the list of plans where you can create a protection plan with the backup module enabled.

To get the information about the unprotected files in the form of report, click **Download detailed report in CSV**.

Data protection map settings

To learn how to create a protection plan with the Data protection map module, refer to "[Creating a protection plan](#)".

The following settings can be specified for the Data protection map module.

Schedule

You can define different settings to create the schedule according to which the task for data protection map will be performed.

Field	Description
Schedule the task run using the following events	<p>This setting defines when the task will run.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Schedule by time – This is the default setting. The task will run according to the specified time. • When user logs in to the system – By default, a login of any user will trigger the task. You can modify this setting so that only a specific user account can trigger the task. • When user logs off the system – By default, a logoff of any user will trigger the task. You can modify this setting so that only a specific user account can trigger the task. <hr/> <p>Note</p> <p>The task will not run at system shutdown. Shutting down and logging off are different events in the scheduling configuration.</p> <hr/> <ul style="list-style-type: none"> • On the system startup – The task will run when the operating system starts. • On the system shutdown – The task will run when the operating system shuts down.
Schedule type	<p>The field appears if in Schedule the task run using the following events you have selected Schedule by time.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Monthly – Select the months and the weeks or days of the month when the task will run. • Daily – This is the default setting. Select the days of the week when the task will run.

Field	Description
	<ul style="list-style-type: none"> • Hourly – Select the days of the week, repetition number, and the time interval in which the task will run.
Start at	<p>The field appears if in Schedule the task run using the following events you have selected Schedule by time</p> <p>Select the exact time when the task will run.</p>
Run within a date range	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time.</p> <p>Set a range in which the configured schedule will be effective.</p>
Specify a user account whose login to the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs in to the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the login of any user to trigger the task. • The following user - Use this option if you want only the login of a specific user account to trigger the task.
Specify a user account whose logout from the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs off the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the logout of any user to trigger the task. • The following user - Use this option if you want only the logout of a specific user account to trigger the task.
Start conditions	<p>Defines all conditions that must be met simultaneously for the task to run.</p> <p>Start conditions for antimalware scans are similar to the start conditions for the Backup module that are described in "Start conditions".</p> <p>You can define the following additional start conditions:</p> <ul style="list-style-type: none"> • Distribute task start time within a time window – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM. • If the machine is turned off, run missed tasks at the machine startup • Prevent the sleep or hibernate mode during task running – This option is effective only for machines running Windows. • If start conditions are not met, run the task anyway after –

Field	Description
	<p>Specify the period after which the task will run, regardless of the other start conditions.</p> <hr/> <p>Note Start conditions are not supported for Linux.</p> <hr/>

Extensions and exception rules

On the **Extensions** tab, you can define the list of file extensions that will be considered as important during data discovery and checked whether they are protected. Use the following format for defining extensions:

```
.html, .7z, .docx, .zip, .pptx, .xml
```

On the **Exception rules** tab, you can define which files and folders not to check on protection status during data discovery.

- **Hidden files and folders** – if selected, hidden files and folders will be skipped during data examination.
- **System files and folders** – if selected, system files and folders will be skipped during data examination.

The Activities tab

The **Activities** tab provides an overview of activities from the past 90 days.

To filter activities on the dashboard

1. In the **Device name** field, specify the machine on which the activity is carried out.
2. From the **Status** dropdown list, select the status. For example, succeeded, failed, in progress, canceled.
3. From the **Remote actions** dropdown list, select the action. For example, applying plan, deleting backups, installing software updates.
4. In the **Most recent** field, set the period of activities. For example, the most recent activities, the activities from the past 24 hours, or the activities during a specific period within the past 90 days.
5. If you are accessing the **Activities** tab as a partner administrator, you can filter the activities for a specific customer that you manage.

To customize the view of the **Activities** tab, click the gear icon, and then select the columns that you want to see. To see the activity progress in real time, select the **Refresh automatically** check box.

To cancel a running activity, click its name, and then, on the **Details** screen, click **Cancel**.

You can search the listed activities by the following criteria:

- **Device name**
This is the machine on which the activity is carried out.

- Started by
This is the account that started the activity.

Remote desktop activities can be filtered by the following properties:

- Creating plan
- Applying plan
- Revoking plan
- Deleting plan
- Remote connection
 - Cloud remote desktop connection via RDP
 - Cloud remote desktop connection via NEAR
 - Cloud remote desktop connection via Apple Screen Sharing
 - Remote desktop connection via web client
 - Remote desktop connection via Quick Assist
 - Direct remote desktop connection via RDP
 - Direct remote desktop connection via Apple Screen Sharing
 - File transfer
 - File transfer via Quick Assist
- Remote action
 - Shutting down a workload
 - Restarting a workload
 - Logging out remote user on the workload
 - Emptying recycle bin for user on the workload
 - Putting to sleep a workload

Cyber Protect Monitor

Cyber Protect Monitor provides a graphical user interface for Agent for Windows, Agent for Mac, and Agent for File Sync & Share. It shows information about the protection status of the machine on which Agent for Windows or Agent for Mac is installed, and allows its users to configure the backup encryption and proxy server settings. With Agent for File Sync & Share, it provides access to the File Sync & Share service.

The agents are registered in the account of the user who installs them. However, the File Sync & Share functionality is accessible after a mandatory onboarding during which the users sign in to their own File Sync & Share account and select a personal sync folder. For more information about Agent for File Sync & Share, refer to the [Cyber Files Cloud user guide](#).

Cyber Protect Monitor is accessible to users who might not have administrative rights for the Cyber Protection or the File Sync & Share service.

Cyber Protection users without administrative rights can perform the following tasks:

- Apply a default protection plan to their machines
- Check the protection status of their machines
- Temporarily pause the backups of their machines

They cannot apply custom protection plans or manage protection plans that are already applied.

File Sync & Share users without administrative rights can perform the following tasks:

- Sync content between their local sync folder and their File Sync & Share account
- Pause their sync operations
- Change their sync folder
- Check the file types whose syncing is restricted

All Cyber Protect Monitor users can change the backup encryption settings or configure the proxy server settings.

Warning!

Changing the encryption settings in Cyber Protect Monitor overwrites the settings in the protection plan and affects all backups of the machine. This operation can make some protection plans fail. For more information, refer to "Encryption" (p. 417).

There is no way to recover encrypted backups if you lose or forget the password.

Configuring proxy server settings in Cyber Protect Monitor

You can configure the proxy server settings in Cyber Protect Monitor. The configuration will affect all agents that are installed on the machine.

To configure the proxy server settings

1. Open Cyber Protect Monitor, and then click the gear icon in the top right corner.
2. Click **Settings**, and then click **Proxy**.
3. Enable the **Use a proxy server** switch, and then enter the proxy server address and port.
4. [If the proxy server access is password-protected] Enable the **Password required** switch, and then enter the user name and password to access the proxy server.
5. Click **Save**.

The proxy server settings are saved in the `http-proxy.yaml` file.

Reports

Note

The availability of this feature depends on the service quotas that are enabled for your account.

A report about operations can include any set of [dashboard widgets](#). All widgets show summary information for the entire company.

Depending on the widget type, the report includes data for a time range or for the moment of browsing or report generation. See "Reported data according to widget type" (p. 301).

All historical widgets show data for the same time range. You can change this range in the report settings.

You can use default reports or create a custom report.

You can download a report or send it via email in XLSX (Excel) or PDF format.

The set of default reports depends on the Cyber Protection service edition that you have. The default reports are listed below:

Report name	Description
#CyberFit Score by machine	Shows the #CyberFit Score, based on the evaluation of security metrics and configurations for each machine, and recommendations for improvements.
Alerts	Shows alerts that occurred during a specified time period.
Backup scanning details	Shows the detailed information about detected threats in the backups.
Daily activities	Shows the summary information about activities performed during a specified time period.
Data protection map	Shows the detailed information about the number, size, location, protection status of all important files on machines.
Detected threats	Shows the details of the affected machines by number of blocked threats and the healthy and vulnerable machines.
Discovered machines	Shows all found machines in the organization network.
Disk health prediction	Shows predictions when your HDD/SSD will break down and current disk status.
Existing vulnerabilities	Shows the existing vulnerabilities for OS and applications in your organization. The report also displays the details of the affected machines in your network for every product that is listed.
Software inventory	Shows information about the software that is installed on your company devices.
Hardware inventory	Shows information about the hardware that is available on your company devices.
Patch management summary	Shows the number of missing patches, installed patches, and applicable patches. You can drill down the reports to get the missing/installed patch information and details of all the systems.
Summary	Shows the summary information about the protected devices for a specified time

	period.
Weekly activities	Shows the summary information about activities performed during a specified time period.
Remote sessions	Shows information about the remote desktop and file transfer sessions.

Actions with reports

To view a report, click its name.

To add a new report

1. In the Cyber Protect console, go to **Reports**.
2. Under the list of available reports, click **Add report**.
3. [To add a predefined report] Click the name of the predefined report.
4. [To add a custom report] Click **Custom**, and then add widgets to the report.
5. [Optional] Drag and drop the widgets to rearrange them.

To edit a report

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report that you want to edit.

You can do the following:

- Rename the report.
- Change the time range for all widgets in the report.
- Specify the report recipients and when the report will be send to them. The available formats are PDF and XLSX.

To delete a report

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report that you want to delete.
3. Click the ellipsis icon (...), and then click **Delete**.
4. Confirm your choice by clicking **Delete**.

To schedule a report

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report that you want to schedule, and then click **Settings**.
3. Enable the **Scheduled** switch.
 - Specify the email addresses of the recipients.
 - Select the format of the report.

Note

You can export up to 1000 items in a PDF file and up to 10 000 items in a XLSX file. The timestamps in the PDF and XLSX files use the local time of your machine.

- Select the language of the report.
- Configure the schedule.

4. Click **Save**.

To download a report

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report, and then click **Download**.
3. Select the format of the report.

To send a report

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report, and then click **Send**.
3. Specify the email addresses of the recipients.
4. Select the format of the report.
5. Click **Send**.

To export the report structure

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report.
3. Click ellipsis icon (...), and then click **Export**.

As a result, the report structure is saved on your machine as a JSON file.

To dump the report data

By using this option, you can export all data for a custom period, without filtering it, to a CSV file and send the CSV file to an email recipient.

Note

You can export up to 150 000 items in a CSV file. The timestamps in the CSV file use Coordinated Universal Time (UTC).

1. In the Cyber Protect console, go to **Reports**.
2. In the list of reports, select the report whose data you want to dump.
3. Click the ellipsis icon (...), and then click **Dump data**.
4. Specify the email addresses of the recipients.
5. In **Time range**, specify the custom period for which you want to dump data.

Note

Preparing CSV files for longer periods takes more time.

6. Click **Send**.

Reported data according to widget type

According to the data range that they display, widgets on the dashboard are two types:

- Widgets that display actual data at the moment of browsing or report generation.
- Widgets that display historical data.

When you configure a date range in the report settings to dump data for a certain period, the selected time range will apply only for widgets that display historical data. For widgets that display actual data at the moment of browsing, the time range parameter is not applicable.

The following table lists the available widgets and their data ranges.

Widget name	Data displayed in widget and reports
#CyberFit Score by machine	Actual
5 latest alerts	Actual
Active alerts details	Actual
Active alerts summary	Actual
Activities	Historical
Activity list	Historical
Alerts history	Historical
Attack tactics statistics	Historical
Backup scanning details (threats)	Historical
Backup status	Historical - in columns Total runs and Number of successful runs Actual - in all other columns
Blocked URLs	Actual
Cloud applications	Actual
Cyber protection	Actual
Data protection map	Historical
Devices	Actual
Discovered machines	Actual
Disk health overview	Actual
Disk health status by physical	Actual

devices	
Existing vulnerabilities	Historical
Hardware changes	Historical
Hardware details	Actual
Hardware inventory	Actual
Historical alerts summary	Historical
Incident severity history	Historical
Locations summary	Actual
Missing updates by categories	Actual
Not protected	Actual
Patch installation history	Historical
Patch installation status	Historical
Patch installation summary	Historical
Protection status	Actual
Recently affected	Historical
Remote sessions	Historical
Security incident burndown	Historical
Security incident MTTR	Historical
Software inventory	Actual
Software overview	Historical
Threat status	Actual
Vulnerable machines	Actual
Workload network status	Actual

Managing workloads in the Cyber Protect console

This section describes how to manage your workloads in the Cyber Protect console.


The Cyber Protect console

In the Cyber Protect console, you can manage workloads and plans, change the protection settings, configure reports, or check the backup storage.

The Cyber Protect console provides access to additional services or features, such as File Sync & Share or Antivirus and Antimalware protection, Patch management, Device control, and Vulnerability assessment. The type and number of these services and features vary according to your Cyber Protection license.

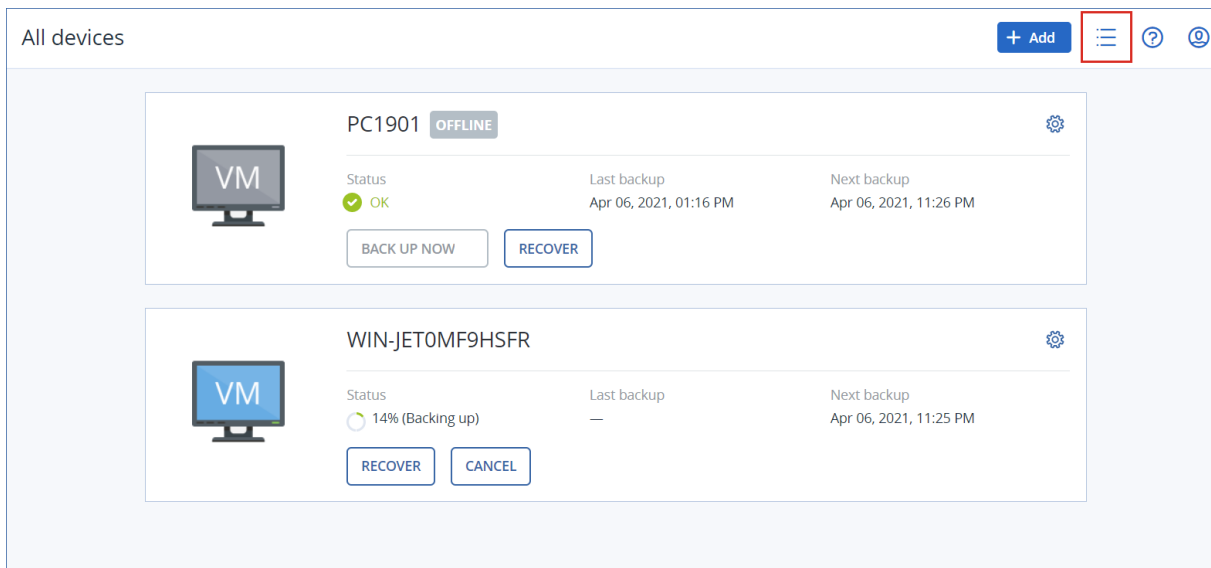
To check the dashboard with the most important information about your protection, go to **Monitoring > Overview**.

Depending on your access permissions, you can manage the protection for one or multiple customer tenants or units in a tenant. To switch the hierarchy level, use the drop-down list in the navigation menu. Only the levels to which you have access are shown. To go to the management portal, click **Manage**.

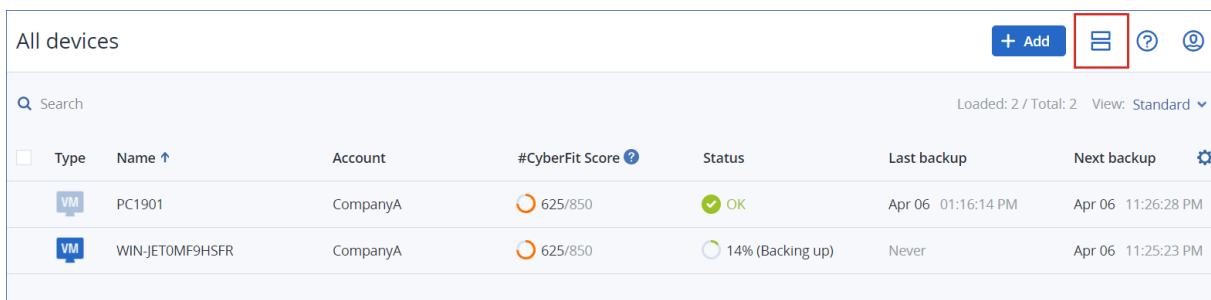


The **Devices** section is available in simple and table view. To switch between them, click the corresponding icon in the top right corner.

The simple view shows only a few workloads.



The table view is enabled automatically when the number of workloads becomes larger.



Both views provide access to the same features and operations. This document describes access to operations from the table view.

When a workload goes online or offline, it takes some time for its status to change in the Cyber Protect console. The workload status is checked every minute. If the agent installed on the corresponding machine is not transferring data, and there is no answer to five consecutive checks, the workload is shown as offline. The workload is shown as back online when it answers to a status check or starts transferring data.

What's new in the Cyber Protect console

When new features of Cyber Protect Cloud are available, you see a pop-up window with a brief description of these features upon logging in to the Cyber Protect console.

You can also view the description of the new features by clicking the **What's new** link in the bottom-left corner of the main Cyber Protect console window.

If there are no new features, the **What's new** link is not displayed.

Using the Cyber Protect console as a partner administrator

A partner administrator can use the Cyber Protect console at the following levels:

- Partner tenant (**All customers**) level

On this level, you can manage scripting plans for workloads from all your managed customer tenants.

You can apply the same scripting plan to workloads in from different customers, and can create device groups with workloads from different customers. To learn how to create a static or a dynamic device group on the partner level, refer to the "Devices tab" (p. 306). For more information about the scripts and scripting plans, refer to "Cyber Scripting" (p. 228).

- Customer tenant level

On this level, you have the same rights as the company administrator on whose behalf you act.

To change the level of administration, use the drop-down list in the navigation menu. The drop-down list is only available for administrators who can access both the Cyber Protect console and the management portal, and can manage more than one tenant or unit.

To work on the partner level, select **All customers**.

To work on the customer or unit level, select the name of that customer or unit.

The screenshot shows the Acronis Cyber Protect Cloud interface. The left-hand navigation menu is highlighted with a red border. At the top of the menu, it says "John Doe" and "Manage". Below that is a dropdown menu currently set to "All customers". Underneath is a search bar and a "RECENT" section with "All customers" and "Company B". The "CUSTOMERS" section lists "All customers", "Company A", "Unit 1", "Unit 2", and "Company B". The main content area is titled "Overview" and features a "Protection status" section with a donut chart. The chart shows 88 machines in total, with 7 protected (green), 81 unprotected (red), 88 managed (blue), and 0 discovered (grey). Below the chart is a "Patch installation status" section.

Cyber Protect console – partner level view

When you use the Cyber Protect console on the partner level, a customized view is available.

The **Alerts** and **Activities** tabs provide additional partner-related filters, while the **Devices** and the **Management** tabs provide access only to the features or objects that are accessible to partner administrators.

Alerts tab

Here, you can see the alerts from all your managed customers, search them, and filter them according to the following criteria:

- Device
- Customer
- Plan

You can select multiple items for each of these criteria.

Activities tab

Here, you can see the activities from all the tenants that you manage or the activities in a specific customer tenant.

You can filter the activities by customer, status, time, and type.

The following types of activities are automatically pre-selected on this level:

- Applying plan
- Creating the protection plan
- Protection plan
- Revoking plan
- Scripting

Devices tab

Only the **All devices**, **Machines with agents**, and virtualization host tabs are available under **Devices**.

In the **Machines with agents** tab, you can see all workloads from your managed customer tenants, and you can select workloads from one or more tenants. You can also create device groups that include workloads from different tenants.

Important


When you work on the partner (**All customers**) level, a limited number of operations with devices are available. For example, you cannot see and manage existing protection plans on customer devices, as well as create new protection plans, add new devices, recover backups, use Disaster Recovery, or access the Cyber Protection Desktop features. To perform any of these operations, switch to the customer level.

To see the workloads of a specific customer

1. In the Cyber Protect console, go **Devices > Machines with agents**.
2. In the tree, click **Machines with agents** to expand the list.
3. Click the name of the customer whose workloads you want to manage.

To create a static device group on the partner level

1. In the Cyber Protect console, go **Devices > Machines with agents**.
2. Click the gear icon next to **Machines with agents**, and then click **New group**.




3. Specify the group name.
4. [Optional] Add a description.
5. Click **OK**.

To create a dynamic device group on the partner level

1. In the Cyber Protect console, go **Devices > Machines with agents**.
2. In the tree, click **Machines with agents** to expand the list.
3. Click **All**.
4. In the search field, specify the criteria according to which you want to create a dynamic device group, and then click **Search**.

To learn more about the available search criteria, refer to "Search attributes for non-cloud-to-cloud workloads" (p. 326) and "Search attributes for cloud-to-cloud workloads" (p. 325).

5. Click **Save as**, and then specify the group name.



6. [Optional] Add a description.
7. Click **OK**.

Software management tab

If the software inventory scanning is enabled for customer workloads, partner administrators can see the software scanning results.

Multitenancy support

The Cyber Protection service supports multitenancy, which implies administration on the following levels:

- [For service providers] Partner tenant (**All customers**) level
This level is only available for partner administrators who manage customer tenants.
- Customer tenant level
This level is managed by company administrators.
Partner administrators can also work on this level in the customer tenants that they manage. On this level, partner administrators have the same rights as the customer administrators on whose behalf they act.
- Unit level
This level is managed by unit administrators and by company administrators from the parent customer tenant.
Partner administrators who manage the parent customer tenant can also access the unit level.
On this level, they have the same rights as the customer administrators on whose behalf they act.

Administrators can manage objects in their own tenant and in its child tenants. They cannot see or access objects on an upper administration level, if any.

For example, company administrators can manage protection plans both on the customer tenant level and on the unit level. Unit administrators can manage only their own protection plans on the unit level. They cannot manage any protection plans on the customer tenant level and cannot manage the protection plans that are created by the customer administrator on the unit level.

Also, partner administrators can create and apply scripting plans in the customer tenants that they manage. The company administrators in such tenants have only read-only access to the scripting plans that are applied to their workloads by a partner administrator. However, customer administrators can create and apply their own scripting or protection plans.

Workloads

A workload is any type of protected resource – for example, a physical machine, a virtual machine, a mailbox, or a database instance. In the Cyber Protect console, the workload is shown as an object to which you can apply a plan (protection plan, backup plan, or scripting plan).

Some workloads require installing a protection agent or deploying a virtual appliance. You can install agents by using the graphical user interface or by using the command-line interface (unattended installation). You can use the unattended installation to automate the installation procedure. For more information about how to install protection agents, refer to "Installing and deploying Cyber Protection agents" (p. 61).

A virtual appliance (VA) is a ready-made virtual machine that contains a protection agent. With a virtual appliance, you can back up other virtual machines in the same environment without installing a protection agent on them (agentless backup). The virtual appliances are available in hypervisor-specific formats, such as .ovf, .ova, or .qcow. For more information about which virtualization platforms support agentless backup, refer to "Supported virtualization platforms" (p. 35).

Important

Agents must be online at least once every 30 days. Otherwise, their plans will be revoked and the workloads will become unprotected.

The table below summarizes the workload types and their respective agents.

Workload type	Agent	Examples (non-exhaustive list)
Physical machines	A protection agent is installed on every protected machine.	Workstation Laptop Server
Virtual machines	Depending on the virtualization platform, the following backup methods might be available: <ul style="list-style-type: none">• Agent-based backup – A protection agent is installed on every protected machine.• Agentless backup – A protection agent is installed only on the hypervisor host, on a dedicated virtual machine, or is deployed as a virtual appliance. This agent backs up all virtual machines in the environment.	VMware virtual machine Hyper-V virtual machine Kernel-based virtual machine (KVM) managed by oVirt
Microsoft 365 Business workloads Google Workspace workloads	These workloads are backed up by a cloud agent for which no installation is required. To use the cloud agent, you need to add your Microsoft 365 or Google Workspace organization to the Cyber Protect console. Additionally, a local Agent for Office 365 is also available. It requires installation and can only be used to back up Exchange Online mailboxes. For more information about the differences between the local and the cloud agent, refer to "Protecting Microsoft 365 data" (p. 553).	Microsoft 365 mailbox Microsoft 365 OneDrive Microsoft Teams SharePoint site Google mailbox Google Drive
Applications	The data of specific applications is backed up by dedicated agents, such as Agent for SQL, Agent for Exchange, Agent for MySQL/MariaDB, or Agent for Active Directory.	SQL Server databases MySQL/MariaDB databases Oracle databases Active Directory
Mobile devices	A mobile app is installed on the protected devices.	Android or iOS devices

Workload type	Agent	Examples (non-exhaustive list)
Websites	The websites are backed up by a cloud agent for which no installation is required.	Websites accessed via the FTP protocol

For more information about which agent you need and where to install it, refer to "Which agent do I need?" (p. 63)

Adding workloads to the Cyber Protect console

To start protecting your workloads, add them to the Cyber Protect console first.

Note

The workload types that you can add depend on the service quotas for your account. If a specific workload type is missing, it is grayed out in the **Add devices** pane.

A partner administrator can enable the required service quotas in the Management portal. For details, refer to "Information for partner administrators" (p. 314).

To add a workload

1. Log in to the Cyber Protect console.
2. Go to **Devices > All devices**, and then click **Add**.
The **Add devices** pane opens on the right.
3. Select the release channel.
4. Click the workload type that you want to add, and then follow the instructions for the specific workload that you selected.

The following table summarizes the workload types and required actions.

Workloads to add	Required action	Procedure to follow
Multiple Windows machines	Perform autodiscovery in your environment. To perform autodiscovery, you need at least one machine with an installed protection agent in your local network or Active Directory domain. This agent is used as a discovery agent.	"Autodiscovery and manual discovery" (p. 127)
Windows workstations Windows servers	Install Agent for Windows.	"Installing protection agents in Windows" (p. 78) or

Workloads to add	Required action	Procedure to follow
		"Unattended installation or uninstallation in Windows" (p. 87)
macOS workstations	Install Agent for macOS.	"Installing protection agents in macOS" (p. 83) or "Unattended installation and uninstallation in macOS" (p. 110)
Linux servers	Install Agent for Linux.	"Installing protection agents in Linux" (p. 80) or "Unattended installation or uninstallation in Linux" (p. 104)
Mobile devices (iOS, Android)	Install the mobile app.	"Protecting mobile devices" (p. 547)
Cloud-to-cloud workloads		
Microsoft 365 Business	Add your Microsoft 365 organization to the Cyber Protect console and use the cloud agent to protect Exchange online mailboxes, OneDrive files, Microsoft Teams, and SharePoint sites. Alternatively, you can install the local Agent for Office 365. It only provides backup of Exchange Online mailboxes. For more information on the differences between the local and the cloud agent, refer to "Protecting Microsoft 365 data" (p. 553).	"Protecting Microsoft 365 data" (p. 553)
Google Workspace	Add your Google Workspace organization to the Cyber Protect console and use the cloud agent to protect Gmail mailboxes and Google Drive files.	"Protecting Google Workspace data" (p. 593)
Virtual machines		
VMware ESXi	Deploy Agent for VMware (Virtual Appliance) in your environment.	"Deploying Agent for VMware (Virtual Appliance)" (p. 134)
	Install Agent for VMware (Windows).	"Installing protection agents in

Workloads to add	Required action	Procedure to follow
		Windows" (p. 78) or "Unattended installation or uninstallation in Windows" (p. 87)
Virtuozzo Hybrid infrastructure	Deploy Agent for Virtuozzo Hybrid Infrastructure (Virtual appliance) in your environment.	"Deploying Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance)" (p. 143)
Hyper-V	Install Agent for Hyper-V.	"Installing protection agents in Windows" (p. 78) or "Unattended installation or uninstallation in Windows" (p. 87)
Virtuozzo	Install Agent for Virtuozzo.	"Installing protection agents in Linux" (p. 80) or "Unattended installation or uninstallation in Linux" (p. 104)
KVM	Install Agent for Windows.	"Installing protection agents in Windows" (p. 78) or "Unattended installation or uninstallation in Windows" (p. 87)
	Install Agent for Linux.	"Installing protection agents in Linux" (p. 80) or "Unattended installation or uninstallation in Linux" (p. 104)
Red Hat Virtualization (oVirt)	Deploy Agent for oVirt (Virtual Appliance) in your environment.	"Deploying Agent for oVirt (Virtual Appliance)" (p. 151)
Citrix XenServer	Install Agent for Windows.	"Installing protection agents in Windows" (p. 78) or "Unattended installation or uninstallation in Windows" (p. 87)
	Install Agent for Linux.	"Installing protection agents in Linux" (p. 80)

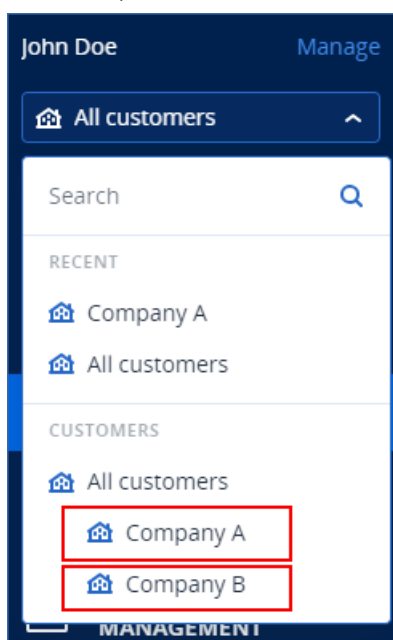
Workloads to add	Required action	Procedure to follow
		or "Unattended installation or uninstallation in Linux" (p. 104)
Nutanix AHV	Install Agent for Windows.	"Installing protection agents in Windows" (p. 78) or "Unattended installation or uninstallation in Windows" (p. 87)
	Install Agent for Linux.	"Installing protection agents in Linux" (p. 80) or "Unattended installation or uninstallation in Linux" (p. 104)
Oracle VM	Install Agent for Windows.	"Installing protection agents in Windows" (p. 78) or "Unattended installation or uninstallation in Windows" (p. 87)
	Install Agent for Linux.	"Installing protection agents in Linux" (p. 80) or "Unattended installation or uninstallation in Linux" (p. 104)
Scale Computing HC3	Deploy Agent for Scale Computing HC3 (Virtual Appliance) in your environment.	"Deploying Agent for Scale Computing HC3 (Virtual Appliance)" (p. 138)
Network-attached storage		
Synology	Deploy Agent for Synology (Virtual Appliance) in your environment.	"Deploying Agent for Synology" (p. 157)
Applications		

Workloads to add	Required action	Procedure to follow
Microsoft SQL Server	Install Agent for SQL.	"Installing protection agents in Windows" (p. 78) or
Microsoft Exchange Server	Install Agent for Exchange.	
Microsoft Active Directory	Install Agent for Active Directory.	"Unattended installation or uninstallation in Windows" (p. 87)
Oracle Database	Install Agent for Oracle.	"Protecting Oracle Database" (p. 618)
Website	Configure the connection to the website.	"Protecting websites and hosting servers" (p. 624)

For more information about the available protection agents and where to install them, refer to "Which agent do I need?" (p. 63)

Information for partner administrators

- A workload type might be missing in the **Add devices** pane if a required service quota is not enabled in the Management portal. For more information about which service quotas are required for which workloads, refer to [Enabling or disabling offering items](#) in the Partner administrator guide.
- As a partner administrator, you cannot add workloads on the **All customers** level. To add a workload, select an individual customer tenant.



Removing workloads from the Cyber Protect console

You can remove from the Cyber Protect console the workloads that you do not need to protect anymore. The procedure depends on the workload type.

Alternatively, you can uninstall the agent on the protected workload. When you uninstall an agent, the protected workload is automatically removed from the Cyber Protect console.

Important

When you remove a workload from the Cyber Protect console, all plans that are applied to that workload are revoked. Removing a workload does not delete any plans or backups, and does not uninstall the protection agent.

The following table summarizes the workload types and required actions.

Workloads to remove	Required actions	Procedure to follow
Physical and virtual machines		
Physical or virtual machines on which a protection agent is installed	<ol style="list-style-type: none">1. Remove the workload from the Cyber Protect console.2. [Optional] Uninstall the protection agent.	"To remove a workload from the Cyber Protect console" (p. 317) (Workload with protection agent)
Virtual machines that are backed up on the hypervisor level (agentless backup)	<ol style="list-style-type: none">1. In the Cyber Protect console, remove the machine on which the protection agent is installed. All virtual machines that are backed up by this agent will be automatically removed from the console.2. [Optional] Uninstall the protection agent.	"To remove a workload from the Cyber Protect console" (p. 317) (Workload without a protection agent)

Workloads to remove	Required actions	Procedure to follow
Cloud-to-cloud workloads		
Microsoft 365 Business workloads Google Workspace workloads	Delete the Microsoft 365 or the Google Workspace organization from the Cyber Protect console. All resources in that organization will be automatically removed from the console.	"To remove a workload from the Cyber Protect console" (p. 317) (Cloud-to-cloud workload)
Mobile devices		
Android devices iOS devices	<ol style="list-style-type: none"> 1. Remove the mobile device from the Cyber Protect console. 2. [Optional] On the mobile device, uninstall the app. 	"To remove a workload from the Cyber Protect console" (p. 317) (Mobile device)
Network-attached storage		
Synology	<ol style="list-style-type: none"> 1. Remove the workload from the Cyber Protect console. 2. [Optional] Uninstall the protection agent. 	"To remove a workload from the Cyber Protect console" (p. 317) (Workload with a protection agent)
Applications		
Microsoft SQL Server Microsoft Exchange Server Microsoft Active Directory	<ol style="list-style-type: none"> 1. In the Cyber Protect console, remove the machine on which the protection agent is installed. The 	"To remove a workload from the Cyber Protect console" (p. 317) (Workload without a protection agent)

Workloads to remove	Required actions	Procedure to follow
Oracle Database	<p>objects that are backed up by this agent will be automatically removed from the console.</p> <p>2. [Optional] Uninstall the protection agent.</p>	
Websites	Remove the website from the Cyber Protect console.	"To remove a workload from the Cyber Protect console" (p. 317) (Website)

To remove a workload from the Cyber Protect console

Workload with a protection agent

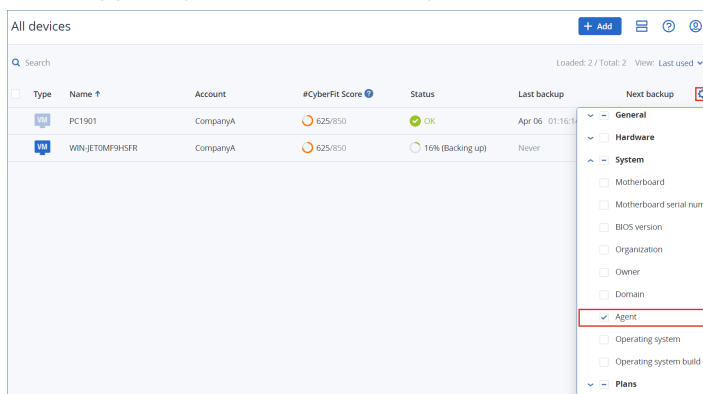
You can remove this type of workload directly.

1. In the Cyber Protect console, navigate to **Devices > All devices**.
2. Select the check box next to one or more workloads that you want to remove.
3. In the **Actions** pane, click **Delete**.
4. Confirm your choice by clicking **Delete**.
5. [Optional] Uninstall the agent as described in "Uninstalling agents" (p. 177).

Workload without a protection agent

To remove this type of workload, you need to remove the machine on which the protection agent is installed.

1. In the Cyber Protect console, go to **Devices > All devices**.
2. In the upper right corner, click the gear icon, and then select the **Agent** check box.



The **Agent** column appears.

3. In the **Agent** column, check the name of the machine where the protection agent is installed.
4. In the Cyber Protect console, select the check box next to the machine on which the protection agent is installed.
5. In the **Actions** pane, click **Delete**.
6. Confirm your choice by clicking **Delete**.
7. [Optional] Uninstall the agent as described in "Uninstalling agents" (p. 177).

Cloud-to-cloud workload

To remove workloads that are backed up by the cloud agent, delete your Microsoft 365 or Google Workspace organization from the Cyber Protect console.

1. In the Cyber Protect console, navigate to **Devices > Microsoft 365** or **Devices > Google Workspace**.
2. Click the name of your Microsoft 365 or Google Workspace organization.
3. In the **Actions** pane, click **Delete group**.
4. Click **Delete** to confirm your action.

Mobile device

1. In the Cyber Protect console, navigate to **Devices > All devices**.
2. Select the check box next to the workload that you want to delete.
3. In the **Actions** pane, click **Delete**.
4. Confirm your choice by clicking **Delete**.
5. [Optional] Uninstall the app from the mobile device.

Website

1. In the Cyber Protect console, navigate to **Devices > All devices**.
2. Select the check box next to the workload that you want to delete.
3. In the **Actions** pane, click **Delete**.
4. Confirm your choice by clicking **Delete**.

Device groups

With device groups, you can protect multiple similar workloads with a group plan. The plan is applied to the group as a whole and cannot be revoked from a member of the group.

A workload can be a member of more than one group. A workload that is included in a device group can still be protected by individual plans.

You can add only workloads of the same type to a device group. For example, under **Hyper-V**, you can only create groups of Hyper-V virtual machines. Under **Machines with agents**, you can only create groups of machines with installed agents.

You cannot create device groups within any **All**-type group, such as the root group **All devices**, or built-in groups like **Machines with agents > All, Microsoft 365 > your organization > Users > All users**.

Built-in groups and custom groups

Built-in groups

After you register a workload in the Cyber Protect console, the workload appears in one of the built-in root groups on the **Devices** tab, such as **Machines with agents**, **Microsoft 365**, or **Hyper-V**.

All registered non-cloud-to-cloud workloads are also listed in the **All devices** root group. A separate built-in root group named after your tenant contains all non-cloud-to-cloud workloads and all units in this tenant.

You cannot delete or edit the root groups, or apply plans to them.

Some of the root groups contain one or more levels of built-in subgroups, for example, **Machines with agents** > **All**, **Microsoft 365** > your organization > **Teams** > **All teams**, **Google Workspace** > your organization > **Shared Drives** > **All Shared Drives**.

You cannot edit or delete built-in subgroups.

Custom groups

Protecting all workloads in a built-in group might not be convenient, because there might be workloads that need different protection settings or a different protection schedule.

In some of the root groups, for example in **Machines with agents**, **Microsoft 365**, or **Google Workspace**, you can create custom subgroups. These subgroups can be static or dynamic.

You can edit, rename, or delete any custom group.

Static groups and dynamic groups

You can create the following type of custom groups:

- Static
- Dynamic

Static groups

Static groups contain manually added workloads.

The content of a static group changes only when you explicitly add or remove a workload.

Example: You create a static group for the accounting department in your company, and then manually add the accountants' machines to this group. When you apply a group plan, the machines in that group become protected. If a new accountant is hired, you will have to add the accountant's machine to the static group manually.

Dynamic groups

Dynamic groups contain workloads that match specific criteria. You define these criteria in advance by creating a search query that includes attributes (for example, `osType`), their values (for example, `Windows`), and search operators (for example, `IN`).

Thus, you can create a dynamic group for all machines whose operating system is Windows or a dynamic group that contains all users in your Microsoft 365 organization whose email addresses begin with `john`.

All workloads that have the required attributes and values are automatically added to the group and any workload that loses a required attribute or value is automatically removed from the group.

Example 1: The host names of the machines that belong to the accounting department contain the word `accounting`. You search for the machines whose names contain `accounting`, and then you save the search results as a dynamic group. Then, you apply a protection plan to the group. If a new accountant is hired, the accountant's machine will have `accounting` in its name and will be automatically added to the dynamic group as soon as you register that machine in the Cyber Protect console.

Example 2: The accounting department forms a separate Active Directory organizational unit (OU). You specify the `accounting` OU as a required attribute, and then you save the search results as a dynamic group. Then, you apply a protection plan to the group. If a new accountant is hired, the accountant's machine will be added to the dynamic group as soon as it is added to the Active Directory OU and is registered in the Cyber Protect console (regardless of which comes first).

Cloud-to-cloud groups and non-cloud-to-cloud groups

Cloud-to-cloud groups contain Microsoft 365 or Google Workspace workloads that are backed up by a cloud agent.

Non-cloud-to-cloud groups contain all other workload types.

Supported plans for device groups

The following table summarizes the plans that you can apply to a device group.

Group	Available plans	Plan location
Cloud-to-cloud workloads (Microsoft 365 and Google Workspace workloads)	Backup plan	Management > Cloud applications backup
Non-cloud-to-cloud workloads	Protection plan	Management > Protection plans
	Remote management plan	Management > Remote management plans
	Scripting plan	Management > Scripting plans

Cloud resources, such as Microsoft 365 or Google Workspace users, OneDrive and Google Drive shares, Microsoft Teams, or Azure AD groups are synchronized to the Cyber Protect console right after you add a Microsoft 365 or Google Workspace organization to the console. Any further changes in an organization are synchronized once a day.

If you need to synchronize a change immediately, in the Cyber Protect console, navigate to **Devices** > **Microsoft 365** or **Devices** > **Google Workspace** respectively, select the required organization, and then click **Refresh**.

Creating a static group

You can create an empty static group and add workloads to it.

Alternatively, you can select workloads and create a new static group from your selection.

You cannot create device groups within any **All**-type group, such as the root group **All devices**, or built-in groups like **Machines with agents** > **All**, **Microsoft 365** > your organization > **Users** > **All users**.

To create a static group

In the main window

1. Click **Devices**, and then select the root group that contains the workloads for which you want to create a static group.
2. [Optional] To create a nested group, navigate to an existing static group.

Note

Creating nested static groups is not available for cloud-to-cloud workloads.

3. Click **+ New static group** below the group tree or click **New static group** in the **Actions** pane.
4. Specify a name for the new group.
5. [Optional] Add a comment for the group.
6. Click **OK**.

In the group tree

1. Click **Devices**, and then select the root group that contains the workloads for which you want to create a static group.
2. Click the gear icon next to the name of the group in which you want to create a new static group.

Note

Creating nested static groups is not available for cloud-to-cloud workloads.

3. Click **New static group**.
4. Specify a name for the new group.
5. [Optional] Add a comment for the group.
6. Click **OK**.

From selection

1. Click **Devices**, and then select the root group that contains the workloads for which you want to create a static group.

Note

You cannot create device groups within any **All**-type group, such as the root group **All devices**, or built-in groups like **Machines with agents > All, Microsoft 365 > your organization > Users > All users**.

2. Select the check boxes next to workloads for which you want to create a new group, and then click **Add to group**.
3. In the folder tree, select the parent level for the new group, and then click **New static group**.

Note

Creating nested static groups is not available for cloud-to-cloud workloads.

4. Specify a name for the new group.
5. [Optional] Add a comment for the group.
6. Click **OK**.
The new group appears in the folder tree.
7. Click **Done**.

Adding workloads to a static group

You can select the target group first, and then add workloads to it.

Alternatively, you can select the workloads first, and then add them to a group.

To add workloads to a static group

Selecting the target group first

1. Click **Devices**, and then navigate to your target group.
2. Select the target group, and then click **Add devices**.
3. In the folder tree, select the group that contains the required workloads.
4. Select the check boxes next to the workloads that you want to add, and then click **Add**.

Selecting the workloads first

1. Click **Devices**, and then select the root group that contains the required workloads.
2. Select the check boxes next to the workloads that you want to add, and then click **Add to group**.
3. In the folder tree, select the target group, and then click **Done**.

Creating a dynamic group

You create a dynamic group by searching for workloads that have specific attributes whose values you define in a search query. Then you save the search results as a dynamic group.

The attributes that are supported for searching and creating dynamic groups differ for cloud-to-cloud workloads and non-cloud-to-cloud workloads. For more information on supported attributes, see "Search attributes for non-cloud-to-cloud workloads" (p. 326) and "Search attributes for cloud-to-cloud workloads" (p. 325).

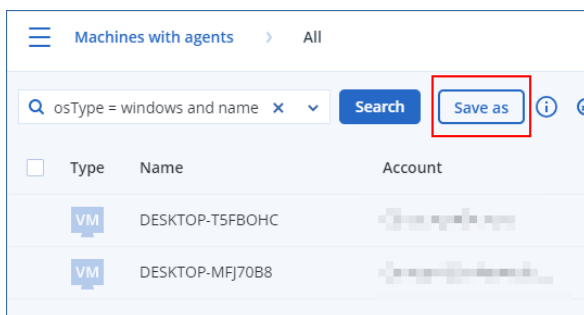
Dynamic groups are created in their respective root groups. Nested dynamic groups are not supported.

You cannot create device groups within any **All**-type group, such as the root group **All devices**, or built-in groups like **Machines with agents > All, Microsoft 365 > your organization > Users > All users**.

To create a dynamic group

Non-cloud-to-cloud workloads

1. Click **Devices**, and then select the group that contains the workloads for which you want to create a new dynamic group.
2. Search for workloads by using the supported search attributes and operators.
You can use multiple attributes and operators in a single query. For more information about the supported attributes, see "Search attributes for non-cloud-to-cloud workloads" (p. 326).
3. Click **Save as** next to the search field.



Note

The **Save as** button is not available when you are not allowed to create a dynamic group on a specific level. For example, in the root group **Devices > All devices**.

Select another level (for example, **Devices > Machines with agents > All**), and then repeat the steps above. With this search, you can create a dynamic group within **Machines with agents**, and not within **Machines with agents > All**.

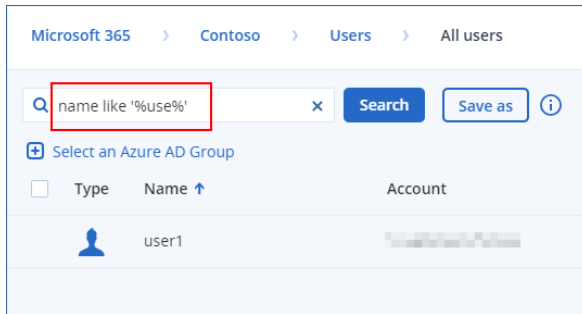
4. Specify a name for the new group.
5. [Optional] In the **Comment** field, add a description for the new group.
6. Click **OK**.

Cloud-to-cloud workloads

1. Click **Devices**, and then select **Microsoft 365** or **Google Workspace**.
2. Select the group that contains the workloads for which you want to create a new dynamic group. For example, **Users > All users**.

3. Search for workloads by using the supported search attributes and operators or by selecting Microsoft 365 users from a specific Active Directory group.

You can use multiple attributes and operators in a single query. For more information about the supported attributes, see "Search attributes for cloud-to-cloud workloads" (p. 325).



4. [Only for **Microsoft 365 > Users**] To select users from a specific Active Directory group, do the following:

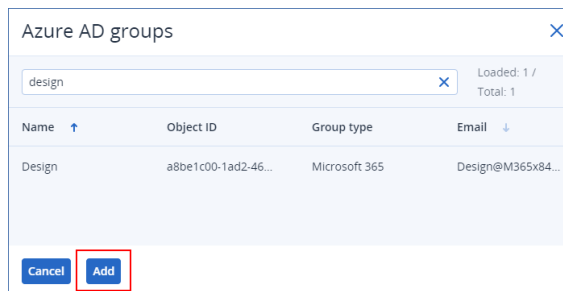
- a. Navigate to **Users > All users**.

- b. Click **Select an Azure AD Group**.

A list of the Active Directory groups in your organization opens.

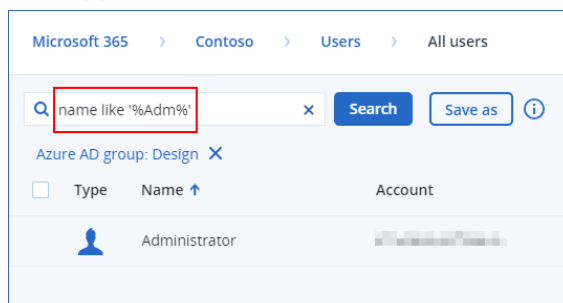
In this list, you can search for a specific group or sort the groups by name or email.

- c. Select the Active Directory group that you want, and then click **Add**.



- d. [Optional] To include or exclude specific users from the selected Active Directory group, create a search query by using the supported search attributes and operators.

You can use multiple attributes and operators in a single query. For more information about the supported attributes, see "Search attributes for cloud-to-cloud workloads" (p. 325).



5. Click **Save as** next to the search field.

Note

The **Save as** button is not available when you are not allowed to create a dynamic group on a specific level. For example, in **Microsoft 365** > your organization > **Users**.

Select another level (for example, **Microsoft 365** > your organization > **Users** > **All**), and then repeat the steps above. With this search, you can create a dynamic group within **Microsoft 365** > your organization > **Users** >, and not within **Users** > **All**.

6. Specify a name for the new group.
7. [Optional] In the **Comment** field, add a description for the new group.
8. Click **OK**.

Search attributes for cloud-to-cloud workloads

The following table summarizes the attributes that you can use in your search queries for Microsoft 365 and Google Workspace workloads.

To see which attributes you can use in search queries for other types of workloads, refer to "Search attributes for non-cloud-to-cloud workloads" (p. 326).

Attribute	Meaning	Can be used in	Search query examples	Supported for group creation
name	Display name of a Microsoft 365 or Google Workspace workload	All cloud-to-cloud resources	name = 'My Name' name LIKE '*nam*'	Yes
email	Email address for a Microsoft 365 user or group, or a Google Workspace user	Microsoft 365 > Groups Microsoft 365 > Users Google Workspace > Users	email = 'my_group_email@mycompany.com' email LIKE '*@company*' email NOT LIKE '*enterprise.com'	Yes
siteName	Name of a site that is associated with a Microsoft 365 group	Microsoft 365 > Groups	siteName = 'my_site' siteName LIKE '*company.com*support*'	Yes
url	Web address for a Microsoft 365 group or SharePoint site	Microsoft 365 > Groups Microsoft 365 > Site collections	url = 'https://www.mycompany.com/' url LIKE '*www.mycompany.com*'	Yes

Search attributes for non-cloud-to-cloud workloads

The following table summarizes the attributes that you can use in your search queries for non-cloud-to-cloud workloads.

To see which attributes you can use in search queries for cloud-to-cloud workloads, refer to "Search attributes for cloud-to-cloud workloads" (p. 325).

Attribute	Meaning	Search query examples	Supported for group creation
General			
name	Workload name, such as: <ul style="list-style-type: none"> • Host name for physical machines • Name for virtual machines • Database name • Email address for mailboxes 	name = 'en-00'	Yes
id	Device ID. To see the device ID, under Devices , select the device, click Details > All properties . The ID is shown in the id field.	id != '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	Yes
resourceType	Workload type. Possible values: <ul style="list-style-type: none"> • 'machine' • 'exchange' • 'mssql_server' • 'mssql_instance' • 'mssql_database' • 'mssql_database_folder' • 'msexchange_database' • 'msexchange_storage_group' • 'msexchange_mailbox.msexchange' 	resourceType = 'machine' resourceType in ('mssql_aag_database', 'mssql_database')	Yes

Attribute	Meaning	Search query examples	Supported for group creation
	<ul style="list-style-type: none"> • 'msexchange_mailbox.office365' • 'mssql_aag_group' • 'mssql_aag_database' • 'virtual_machine.vmww' • 'virtual_machine.vmwesx' • 'virtual_host.vmwesx' • 'virtual_cluster.vmwesx' • 'virtual_appliance.vmwesx' • 'virtual_application.vmwesx' • 'virtual_resource_pool.vmwesx' • 'virtual_center.vmwesx' • 'datastore.vmwesx' • 'datastore_cluster.vmwesx' • 'virtual_network.vmwesx' • 'virtual_data_center.vmwesx' • 'virtual_machine.vmww' • 'virtual_cluster.mshyperv' • 'virtual_machine.mshyperv' • 'virtual_host.mshyperv' • 'virtual_network.mshyperv' • 'virtual_folder.mshyperv' • 'virtual_data_center.mshyperv' • 'datastore.mshyperv' • 'virtual_machine.msvs' • 'virtual_machine.parallelsw' 		

Attribute	Meaning	Search query examples	Supported for group creation
	<ul style="list-style-type: none"> • 'virtual_host.parallelsw' • 'virtual_cluster.parallelsw' • 'virtual_machine.rhev' • 'virtual_machine.kvm' • 'virtual_machine.xen' • 'bootable_media' 		
chassis	<p>Chassis type.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • laptop • desktop • server • other • unknown 	<p>chassis = 'laptop'</p> <p>chassis IN ('laptop', 'desktop')</p>	Yes
ip	IP address (only for physical machines).	<p>ip RANGE</p> <p>('10.250.176.1', '10.250.176.50')</p>	Yes
comment	<p>Comment for a device. It can be specified automatically or manually.</p> <p>Default value:</p> <ul style="list-style-type: none"> • For physical machines running Windows, the computer description in Windows is automatically copied as a comment. This value is synchronized every 15 minutes. • Empty for other devices. 	<p>comment = 'important machine'</p> <p>comment = '' (all machines without a comment)</p>	Yes

Attribute	Meaning	Search query examples	Supported for group creation
	<p>Note The automatic synchronization is disabled if there is manually added text in the comment field. To enable the synchronization again, clear this text.</p> <hr/> <p>To refresh the automatically synchronized comments for your workloads, restart the Managed Machine Service in Windows Services or run the following commands at the command prompt:</p> <pre>net stop mms</pre> <pre>net start mms</pre> <p>To view a device comment, under Devices, select the device, click Details, and then locate the Comment section.</p> <p>To add or change a comment manually, click Add or Edit.</p> <p>For devices on which a protection agent is installed, there are two separate comment fields:</p> <ul style="list-style-type: none"> • Agent comment <ul style="list-style-type: none"> ◦ For physical machines running Windows, the computer description in Windows is 		

Attribute	Meaning	Search query examples	Supported for group creation
	<p>automatically copied as a comment. This value is synchronized every 15 minutes.</p> <ul style="list-style-type: none"> ◦ Empty for other devices. <hr/> <p>Note The automatic synchronization is disabled if there is manually added text in the comment field. To enable the synchronization again, clear this text.</p> <hr/> <ul style="list-style-type: none"> • Device comment <ul style="list-style-type: none"> ◦ If the agent comment is specified automatically, it is copied as a device comment. Manually added agent comments are not copied as device comments. ◦ Device comments are not copied as agent comments. <p>A device can have one or both of these comments specified, or have the both of them blank. If the both comments are specified, the device comment has priority.</p> <p>To view an agent comment, under Settings > Agents, select the device with the agent, click Details, and then locate the Comment</p>		

Attribute	Meaning	Search query examples	Supported for group creation
	<p>section.</p> <p>To view a device comment, under Devices, select the device, click Details, and then locate the Comment section.</p> <p>To add or change a comment manually, click Add or Edit.</p>		
isOnline	<p>Workload availability.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • true • false 	isOnline = true	No
hasAsz	<p>Secure Zone availability.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • true • false 	hasAsz = true	Yes
tzOffset	<p>Timezone offset from Coordinated Universal Time (UTC), in minutes.</p>	<p>tzOffset = 120</p> <p>tzOffset > 120</p> <p>tzOffset < 120</p>	Yes
CPU, memory, disks			
cpuArch	<p>CPU architecture.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • 'x64' • 'x86' 	cpuArch = 'x64'	Yes
cpuName	CPU name.	cpuName LIKE '%XEON%'	Yes
memorySize	RAM size in megabytes.	memorySize < 1024	Yes
diskSize	Hard drive size in gigabytes or megabytes (only for physical machines).	<p>diskSize < 300GB</p> <p>diskSize >= 3000000MB</p>	No
Operating system			

Attribute	Meaning	Search query examples	Supported for group creation
osName	Operating system name.	osName LIKE '%Windows XP%'	Yes
osType	Operating system type. Possible values: <ul style="list-style-type: none"> 'windows' 'linux' 'macosx' 	osType = 'windows' osType IN ('linux', 'macosx')	Yes
osArch	Operating system architecture. Possible values: <ul style="list-style-type: none"> 'x64' 'x86' 	cpuArch = 'x86'	Yes
osProductType	Operating system product type. Possible values: <ul style="list-style-type: none"> 'dc' Stands for Domain Controller. <hr/> <p>Note When the domain controller role is assigned on a Windows server, the osProductType changes from server to dc. Such machines will be not included in the search results for osProductType='server'.</p> <ul style="list-style-type: none"> 'server' 'workstation' 	osProductType = 'server'	Yes
osSp	Service pack of the operating system.	osSp = 1	Yes
osVersionMajor	Major version of the operating system.	osVersionMajor = 1	Yes

Attribute	Meaning	Search query examples	Supported for group creation
osVersionMinor	Minor version of the operating system.	osVersionMinor > 1	Yes
Agent			
agentVersion	Version of the installed protection agent.	agentVersion LIKE '12.0.*'	Yes
hostId	Internal ID of the protection agent. To see the protection agent ID, under Devices , select the device, click Details > All properties . Check the id value of the agent property.	hostId = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	Yes
virtualType	Virtual machine type. Possible values: <ul style="list-style-type: none"> 'vmwexx' VMware virtual machines. 'mshyperv' Hyper-V virtual machines. 'pcs' Virtuozzo virtual machines. 'hci' Virtuozzo Hybrid Infrastructure virtual machines. 'scale' Scale Computing HC3 virtual machines. 'ovirt' oVirt virtual machines 	virtualType = 'vmwexx'	Yes
insideVm	Virtual machine with an agent inside. Possible values:	insideVm = true	Yes

Attribute	Meaning	Search query examples	Supported for group creation
	<ul style="list-style-type: none"> • true • false 		
Location			
tenant	The name of the tenant to which the device belongs.	tenant = 'Unit 1'	Yes
tenantId	<p>The identifier of the tenant to which device belongs.</p> <p>To see the tenant ID, under Devices, select the device, click Details > All properties. The ID is shown in the ownerId field.</p>	tenantId = '3bfe6ca9-9c6a-4953-9cb2-a1323f454fc9'	Yes
ou	Devices that belong to the specified Active Directory organizational unit.	ou IN ('RnD', 'Computers')	Yes
Status			
state	<p>Device state.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • 'idle' • 'interactionRequired' • 'canceling' • 'backup' • 'recover' • 'install' • 'reboot' • 'failback' • 'testReplica' • 'run_from_image' • 'finalize' • 'failover' • 'replicate' • 'createAsz' • 'deleteAsz' • 'resizeAsz' 	state = 'backup'	No
status	Protection status.	status = 'ok'	No

Attribute	Meaning	Search query examples	Supported for group creation
	Possible values: <ul style="list-style-type: none"> • ok • warning • error • critical • protected • notProtected 	status IN ('error', 'warning')	
protectedByPlan	Devices that are protected by a protection plan with a given ID. To see the plan ID, in Management > Protection plans , select a plan, click the bar in the Status column, and then click the status name. A new search with the plan ID will be created.	protectedByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
okByPlan	Devices that are protected by a protection plan with a given ID and have an OK status.	okByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
errorByPlan	Devices that are protected by a protection plan with a given ID and have an Error status.	errorByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
warningByPlan	Devices that are protected by a protection plan with a given ID and have a Warning status.	warningByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
runningByPlan	Devices that are protected by a protection plan with a given ID and have a Running status.	runningByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
interactionByPlan	Devices that are protected by a protection plan with a given ID and have an Interaction Required	interactionByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No

Attribute	Meaning	Search query examples	Supported for group creation
	status.		
lastBackupTime*	The date and time of the last successful backup. The format is 'YYYY-MM-DD HH:MM'.	lastBackupTime > '2023-03-11' lastBackupTime <= '2023-03-11 00:15' lastBackupTime is null	No
lastBackupTryTime*	The time of the last backup attempt. The format is 'YYYY-MM-DD HH:MM'.	lastBackupTryTime >= '2023-03-11'	No
nextBackupTime*	The time of the next backup. The format is 'YYYY-MM-DD HH:MM'.	nextBackupTime >= '2023-08-11'	No
lastVAScanTime*	The date and time of the last successful vulnerability assessment. The format is 'YYYY-MM-DD HH:MM'.	lastVAScanTime > '2023-03-11' lastVAScanTime <= '2023-03-11 00:15' lastVAScanTime is null	Yes
lastVAScanTryTime*	The time of the last vulnerability assessment attempt. The format is 'YYYY-MM-DD HH:MM'.	lastVAScanTryTime >= '2022-03-11'	Yes
nextVAScanTime*	The time of the next vulnerability assessment. The format is 'YYYY-MM-DD HH:MM'.	nextVAScanTime <= '2023-08-11'	Yes
network_status	Network isolation status for Endpoint detection and response (EDR). Possible values: <ul style="list-style-type: none">connectedisolated	network_status= 'connected'	Yes

Note

If you skip the hour and minutes value, the start time is considered to be YYYY-MM-DD 00:00, and the end time is considered to be YYYY-MM-DD 23:59:59. For example, `lastBackupTime = 2023-01-20`, means that the search results will include all backups from the interval

`lastBackupTime >= 2023-01-20 00:00` and `lastBackup time <= 2023-01-20 23:59:59`.

Search operators

The following table summarizes the operators that you can use for your search queries.

You can use more than one operator in a single query.

Operator	Supported for	Meaning	Examples
AND	All workloads	Logical conjunction operator	<code>name like 'en-00' AND tenant = 'Unit 1'</code>
OR	All workloads	Logical disjunction operator	<code>state = 'backup' OR state = 'interactionRequired'</code>
NOT	All workloads	Logical negation operator	<code>NOT(osProductType = 'workstation')</code>
IN (<code><value1>, ... <valueN></code>)	All workloads	This operator checks if an expression matches any value in a list of values.	<code>osType IN ('windows', 'linux')</code>
NOT IN	All workloads	This operator is the opposite of the IN operator.	<code>NOT osType IN ('windows', 'linux')</code>
LIKE 'wildcard pattern'	All workloads	This operator checks if an expression matches the wildcard pattern. You can use the following wildcard operators: <ul style="list-style-type: none">• * or % The asterisk and the percent sign represent zero, one, or multiple	<code>name LIKE 'en-00'</code> <code>name LIKE '*en-00'</code> <code>name LIKE '*en-00*'</code> <code>name LIKE 'en-00_'</code>

Operator	Supported for	Meaning	Examples
		characters <ul style="list-style-type: none"> _ The underscore represents a single character 	
NOT LIKE 'wildcard pattern'	All workloads	This operator is the opposite of the LIKE operator. You can use the following wildcard operators: <ul style="list-style-type: none"> * or % The asterisk and the percent sign represent zero, one, or multiple characters _ The underscore represents a single character 	<pre>NOT name LIKE 'en-00'</pre> <pre>NOT name LIKE '*en-00'</pre> <pre>NOT name LIKE '*en-00*'</pre> <pre>NOT name LIKE 'en-00_'</pre>
RANGE (<starting_value>, <ending_value>)	All workloads	This operator checks if an expression is within a range of values (inclusive). Search queries with alphanumeric strings use the ASCII sort order but are case-insensitive.	<pre>ip RANGE('10.250.176.1', '10.250.176.50')</pre> <pre>name RANGE('a', 'd')</pre> <p>With this query, you can filter all names that begin with A, B, and C, such as Alice, Bob, Claire. However, only the single letter D meets the requirements, so names with more letters, such as Diana or Don will not be included.</p> <p>To achieve the same result, you can also use the following query:</p> <pre>name >= 'a' AND name <= 'd'</pre>
= or ==	All workloads	<i>Equal to operator</i>	<pre>osProductType = 'server'</pre>
!= or <>	All workloads	<i>Not equal to operator</i>	<pre>id != '4B2A7A93-A44F-4155-BDE3-A023C57C9431'</pre>
<	Non-cloud-to-cloud workloads	<i>Less than operator</i>	<pre>memorySize < 1024</pre>

Operator	Supported for	Meaning	Examples
>	Non-cloud-to-cloud workloads	<i>Greater than operator.</i>	diskSize > 300GB
<=	Non-cloud-to-cloud workloads	<i>Less than or equal to operator</i>	lastBackupTime <= '2022-03-11 00:15'
>=	Non-cloud-to-cloud workloads	<i>Greater than or equal to operator</i>	nextBackupTime >= '2022-08-11'

Editing a dynamic group

You edit a dynamic group by changing the search query that defines the group content.

In dynamic groups that are based on Active Directory, you can also change the Active Directory group.

To edit a dynamic group

By changing the search query

1. Click **Devices**, navigate to the dynamic group that you want to edit, and then select it.
2. Click the gear icon next to the name of the group, and then click **Edit**. Alternatively, click **Edit** in the **Actions** pane.
3. Change the search query by modifying the search attributes, their values, or the search operators, and then click **Search**.
4. Click **Save** next to the search field.

By changing the Active Directory group

Note

This procedure applies to dynamic groups based on Active Directory. Active Directory-based dynamic groups are available only in **Microsoft 365 > Users**.

1. Click **Devices**, navigate to **Devices > Microsoft 365 > your organization > Users**.
2. Select the dynamic group that you want to edit.
3. Click the gear icon next to the name of the group, and then click **Edit**. Alternatively, click **Edit** in the **Actions** pane.
4. Change the group content by doing any of the following:
 - Change the already selected Active Directory group by clicking its name, and then selecting a new Active Directory group from the list that opens.

- Edit the search query, and then click **Search**.

The search query is limited to the currently selected Active Directory group.

5. Click **Save** next to the search field.

You can also save your edits without overwriting the current group. To save the edited configuration as a new group, click the arrow button next to the search field, and then click **Save as**.

Deleting a group

When you delete a device group, all plans that are applied to that group will be revoked. The workloads in the group will become unprotected if no other plans are applied to them.

To delete a device group

1. Click **Devices**, and then navigate to the group that you want to delete.
2. Click the gear icon next to the name of the group, and then click **Delete**.
3. Confirm your choice by clicking **Delete**.

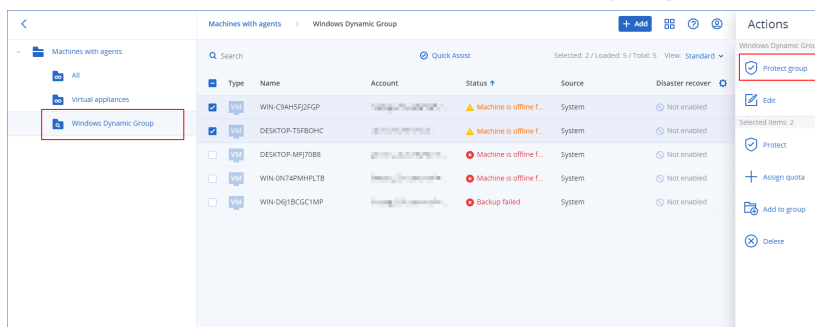
Applying a plan to a group

You can apply a plan to a group by selecting the group first, and then assigning a plan to it.

Alternatively, you can open a plan for editing, and then add a group to it.

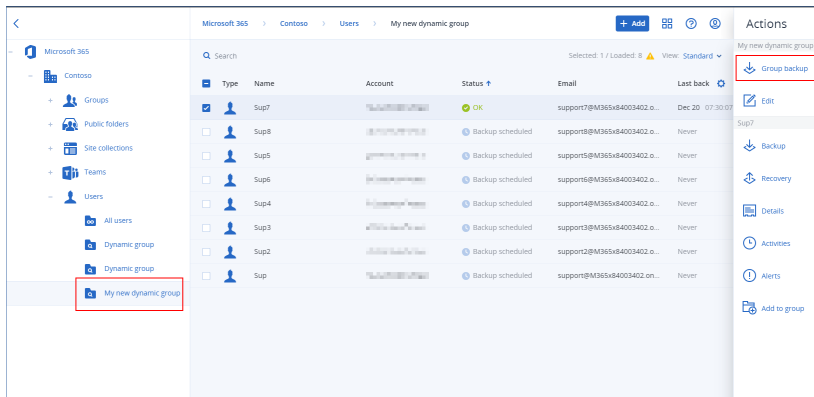
To apply a plan to a group

1. Click **Devices**, and then navigate to the group to which you want to apply a plan.
2. [For non-cloud-to-cloud workloads] Click **Protect group**.



A list of plans that can be applied is shown.

3. [For cloud-to-cloud workloads] Click **Group backup**.



A list of backup plans that can be applied is shown.

4. [To apply an existing plan] Select the plan, and then click **Apply**.
5. [To create a new plan] Click **Create plan**, select the plan type, and then create the new plan.
For more information about the available types of plans and how to create them, refer to "Supported plans for device groups" (p. 320).

Note

Backup plans that are applied to cloud-to-cloud device groups are automatically scheduled to run once a day. You cannot run these plans on demand by clicking **Run now**.

Revoking a plan from a group

You can revoke a plan from a group by selecting the group first, and then revoking the plan from it. Alternatively, you can open the plan for editing, and then remove the group from it.

To revoke a plan from a group

1. Click **Devices**, and then navigate to the group from which you want to revoke a plan.
2. [For non-cloud-to-cloud workloads] Click **Protect group**.
A list of plans that are applied to the group is shown.
3. [For cloud-to-cloud workloads] Click **Group backup**.
A list of backup plans that are applied to the group is shown.
4. Select the plan that you want to revoke.
5. [For non-cloud-to-cloud workloads] Click the ellipsis icon (...), and then click **Revoke**.
6. [For cloud-to-cloud workloads] Click the gear icon, and then click **Revoke**.

Working with the Device control module

A part of the Cyber Protection service protection plans, the device control module¹ leverages a functional subset of the agent for Data Loss Prevention² on each protected computer to detect and prevent unauthorized access and transmission of data over local computer channels. It provides fine-grained control over a wide range of data leakage pathways including data exchange using removable media, printers, virtual and redirected devices, and the Windows clipboard.

The module is available for Cyber Protect Essentials, Cyber Protect Standard, and Cyber Protect Advanced editions that are licensed per workload.

Note

On Windows machines, the device control features require the installation of Agent for Data Loss Prevention. It will be installed automatically for protected workloads if the **Device control** module is enabled in their protection plans.

The device control module relies on the data loss prevention³ functions of the agent to enforce contextual control over data access and transfer operations on the protected computer. These include user access to peripheral devices and ports, document printing, clipboard copy / paste operations, media format and eject operations, as well as synchronizations with locally connected mobile devices. The agent for Data Loss Prevention includes a framework for all central management and administration components of the device control module, and therefore it must be installed on every computer to be protected with the device control module. The agent allows, restricts, or denies user actions based on the device control settings it receives from the protection plan that is applied to the protected computer.

The device control module controls access to various peripheral devices, whether used directly on protected computers or redirected in virtualization environments hosted on protected computers. It recognizes devices redirected in Microsoft Remote Desktop Server, Citrix XenDesktop / XenApp / XenServer, and VMware Horizon. It can also control data copy operations between the clipboard of the guest operating system running on VMware Workstation / Player, Oracle VM VirtualBox, or

¹As part of a protection plan, the device control module leverages a functional subset of the data loss prevention agent on each protected computer to detect and prevent unauthorized access and transmission of data over local computer channels. These include user access to peripheral devices and ports, document printing, clipboard copy/paste operations, media format and eject operations, as well as synchronizations with locally connected mobile devices. The device control module provides granular, contextual control over the types of devices and ports that users are allowed to access on the protected computer and the actions that users can take on those devices.

²A data loss prevention system's client component that protects its host computer from unauthorized use, transmission, and storage of confidential, protected, or sensitive data by applying a combination of context and content analysis techniques and enforcing centrally managed data loss prevention policies. Cyber Protection provides a fully featured data loss prevention agent. However, the functionality of the agent on a protected computer is limited to the set of data loss prevention features available for licensing in Cyber Protection, and depends upon the protection plan applied to that computer.

³A system of integrated technologies and organizational measures aimed at detecting and preventing accidental or intentional disclosure / access to confidential, protected, or sensitive data by unauthorized entities outside or inside the organization, or the transfer of such data to untrusted environments.

Windows Virtual PC, and the clipboard of the host operating system running on the protected computer.

The device control module can protect computers running the following operating systems:

Device control

- Microsoft Windows 7 Service Pack 1 and later
- Microsoft Windows Server 2008 R2 and later
- macOS 10.15 (Catalina)
- macOS 11.2.3 (Big Sur)
- macOS 12 (Monterey)
- macOS 13 (Ventura)

Note

Agent for Data Loss Prevention for macOS supports only x64 processors. Apple silicon ARM-based processors are not supported.

Data loss prevention

- Microsoft Windows 7 Service Pack 1 and later
- Microsoft Windows Server 2008 R2 and later

Note

Agent for Data Loss Prevention might be installed on unsupported macOS systems because it is an integral part of Agent for Mac. In this case, the Cyber Protect console will indicate that Agent for Data Loss Prevention is installed on the computer, but the device control and data loss prevention functionality will not work. Device control functionality will only work on macOS systems that are supported by Agent for Data Loss Prevention.

Limitation on the use of the agent for Data Loss Prevention with Hyper-V

Do not install Agent for Data Loss Prevention on Hyper-V hosts in Hyper-V clusters because it might cause BSOD issues, mainly in Hyper-V clusters with Clustered Shared Volumes (CSV).

If you use any of the following versions of Agent for Hyper-V, you need to manually remove Agent for Data Loss Prevention:



- 15.0.26473 (C21.02)
- 15.0.26570 (C21.02 HF1)
- 15.0.26653 (C21.03)
- 15.0.26692 (C21.03 HF1)
- 15.0.26822 (C21.04)

To remove Agent for Data Loss Prevention, on the Hyper-V host, run the installer manually and clear the Agent for Data Loss Prevention check box, or run the following command:

```
<installer_name> --remove-components=agentForDlp -quiet
```

You can enable and configure the device control module in the **Device control** section of your protection plan in the Cyber Protect console. For instructions, see [steps to enable or disable device control](#).

The **Device control** section displays a summary of the module's configuration:

Device control Access to 7 device types is limited. Allowlists are configured	 
Access settings	Restricted: USB, Removable, Printers and 4 more
Device types allowlist	1 allowed
USB devices allowlist	1 allowed
Exclusions	2 excluded

- [Access settings](#) - Shows a summary of device types and ports with restricted (denied or read-only) access, if any. Otherwise, indicates that all device types are allowed. Click this summary to view or change the access settings (see [steps to view or change access settings](#)).
- [Device types allowlist](#) - Shows how many device subclasses are allowed by excluding from device access control, if any. Otherwise, indicates that the allowlist is empty. Click this summary to view or change the selection of allowed device subclasses (see [steps to exclude device subclasses from access control](#)).
- [USB devices allowlist](#) - Shows how many USB devices/models are allowed by excluding from device access control, if any. Otherwise, indicates that the allowlist is empty. Click this summary to view or change the list of allowed USB devices/models (see [steps to exclude individual USB devices from access control](#)).
- [Exclusions](#) - Shows how many access control exclusions have been set for Windows clipboard, screenshot capture, printers, and mobile devices.

Using device control

This section covers step-by-step instructions for basic tasks when using the device control module.

Enable or disable device control

You can enable device control when [creating a protection plan](#). You can change an existing protection plan to enable or disable device control.

To enable or disable device control

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Do one of the following to open the protection plan panel:
 - If you are going to create a new protection plan, select a machine to protect, click **Protect**, and then click **Create plan**.
 - If you are going to change an existing protection plan, select a protected machine, click **Protect**, click the ellipsis (...) next to the name of the protection plan, and then click **Edit**.
3. In the protection plan panel, navigate to the **Device control** area, and enable or disable **Device control**.
4. Do one of the following to apply your changes:
 - If creating a protection plan, click **Create**.
 - If editing a protection plan, click **Save**.

You might also access the protection plan panel from the [Management tab](#). However, this capability is not available in all editions of the Cyber Protection service.

Enabling the use of the device control module on macOS

The device control settings of a protection plan become effective only after loading the device control driver on the protected workload. This section describes how to load the device control driver to enable the use of the device control module on macOS. This is a one-time operation that requires administrator privileges on the endpoint machine.


Supported macOS versions:

- macOS 10.15 (Catalina) and later
- macOS 11.2.3 (Big Sur) and later
- macOS 12.2 (Monterey) and later
- macOS 13.2 (Ventura) and later

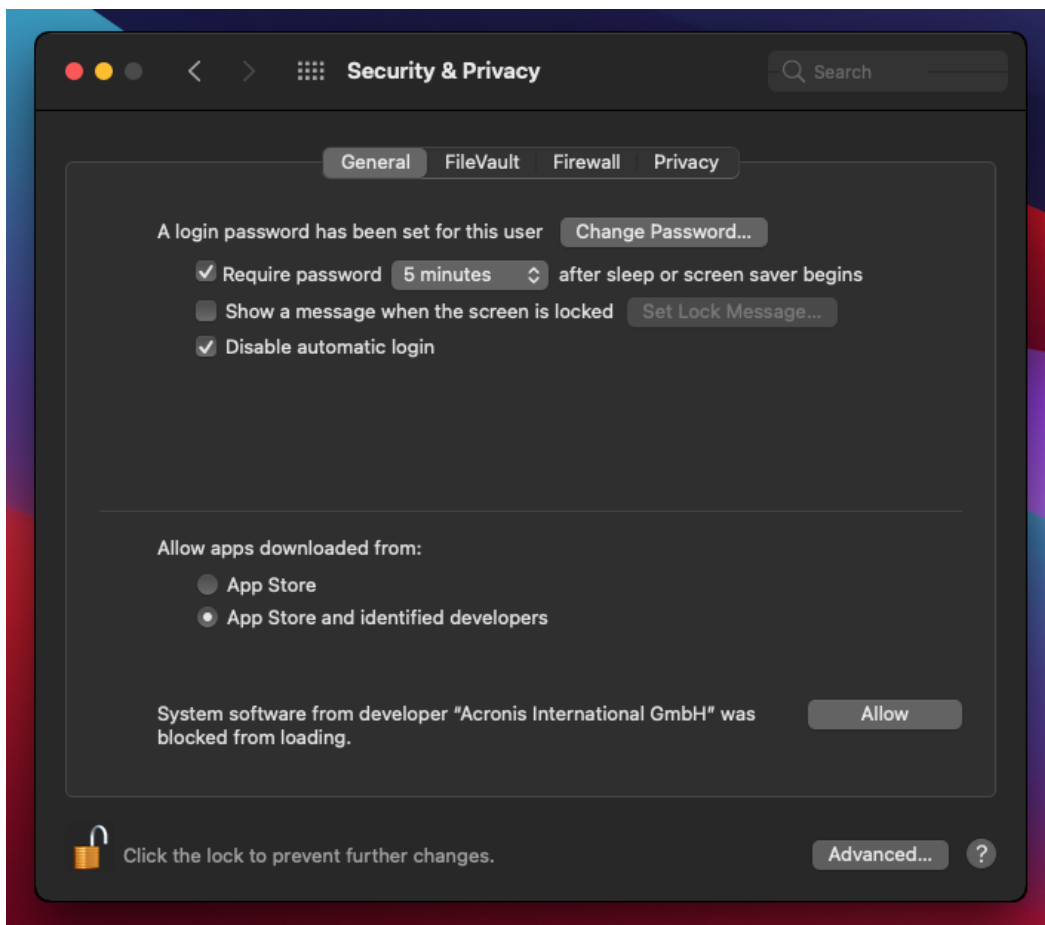
To enable the use of device control module on macOS

1. Install Agent for Mac on the machine that you want to protect.
2. Enable device control settings in the protection plan.
3. Apply the protection plan.

4. The "System Extension Blocked" warning will appear on the protected workload. Click **Open Security Preferences**.



5. In the **Security & Privacy** pane that appears, select **App Store and identified developers** and then click **Allow**.



6. In the dialog that appears, click **Restart** to restart the workload and activate the device control settings.

Note

You do not have to repeat these steps if the device control settings are disabled and then enabled again.

View or change access settings

From the protection plan panel, you can manage access settings for the device control module. In this way, you can allow or deny access to certain types of devices, as well as enable or disable notifications and alerts.

To view or change access settings

1. Open the protection plan panel for a protection plan and enable device control in that plan (see [steps to enable or disable device control](#)).
2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the link next to **Access settings**.

3. On the [page for managing access settings](#) that appears, view or change access settings as appropriate.

Note

The access settings configured in Device control might be overridden when using both Device control and Advanced DLP to protect a workload. See "Enabling Advanced Data Loss Prevention in protection plans" (p. 815).

Enable or disable OS notification and service alerts

When managing access settings, you can enable or disable [OS notification and service alerts](#), informing of user attempts to perform actions that are not allowed.

To enable or disable OS notification

1. Follow the [steps to view or change access settings](#).
2. On the [page for managing access settings](#), select or clear the **Show OS notification to end users if they try to use a blocked device type or port** check box.

To enable or disable service alerts

1. Follow the [steps to view or change access settings](#).
2. On the [page for managing access settings](#), select or clear the **Show alert** check box for the desired device type/s.

The **Show alert** check box is available only for device types with restricted access (Read-only or Denied access), except screenshot capture.

Exclude device subclasses from access control

From the protection plan panel, you can choose device subclasses to exclude from access control. As a result, access to those devices is allowed regardless of the device control access settings.

To exclude device subclasses from access control

1. Open the protection plan panel for a protection plan and enable device control in that plan (see [steps to enable or disable device control](#)).
2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the link next to **Device types allowlist**.
3. On the [page for managing the allowlist](#) that appears, view or change the selection of device subclasses to exclude from access control.

Exclude individual USB devices from access control

From the protection plan panel, you can specify individual USB devices or USB device models to exclude from access control. As a result, access to those devices is allowed regardless of the device control access settings.

To exclude a USB device from access control

1. Open the protection plan panel for a protection plan and enable device control in that plan (see [steps to enable or disable device control](#)).
2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the link next to **USB devices allowlist**.
3. On the [page for managing the allowlist](#) that appears, click **Add from database**.
4. On the [page for selecting USB devices](#) that appears, select the desired device/s from those registered with the [USB devices database](#).
5. Click the **Add to allowlist** button.

To stop excluding a USB device from access control

1. Open the protection plan panel for a protection plan and enable device control in that plan (see [steps to enable or disable device control](#)).
2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the link next to **USB devices allowlist**.
3. On the [page for managing the allowlist](#) that appears, click the delete icon at the end of list item representing the desired USB device.

Add or remove USB devices from the database

To exclude a particular USB device from access control, you need to add it to the [USB devices database](#). Then, you can add devices to the allowlist by selecting from that database.

The following procedures apply to protection plans that have the device control feature enabled.

To add USB devices to the database

1. Open the protection plan of a device for editing:
Click the ellipsis (...) next to the name of the protection plan and select **Edit**.

Note

Device control must be enabled in the plan, so you can access the Device control settings.

2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the link next to **USB devices allowlist**.
3. On the **USB devices allowlist** page that appears, click **Add from database**.
4. On the USB devices database management page that appears, click **Add to database**.
5. On the **Add USB device** dialog that appears, click the machine to which the USB device is connected.

Only machines that are online are displayed in the list of computers.

The list of USB devices is displayed only for machines that have the agent for Data Loss Prevention installed.

The USB devices are listed in tree view. The first level of the tree represents a device model. The second level represents a specific device of that model.

A blue icon next to the description of the device indicates that the device is currently attached to the computer. If the device is not attached to the computer, the icon is grayed out.

6. Select the check boxes for the USB devices that you want to add to the database, and then click **Add to database**.
The selected USB devices are added to the database.
7. Close or save the protection plan.

To add USB devices to the database from the computer Details panel

Note

This procedure applies only for devices that are online and have the agent for Data Loss Prevention installed on them. You cannot view the list of USB devices for a computer that is offline or does not have the Data Loss Prevention agent installed.

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select a computer to which the desired USB device has ever been connected, and, in the menu to the right, click **Inventory**.
The computer details panel opens.
3. On the computer details panel, click the **USB Devices** tab.
The list of USB devices that are known on the selected computer opens.
The USB devices are listed in tree view. The first level of the tree represents a device model. The second level represents a specific device of that model.
A blue icon next to the description of the device indicates that the device is currently attached to the computer. If the device is not attached to the computer, the icon is grayed out.
4. Select the check boxes for the USB devices that you want to add to the database and click **Add to database**.

To add USB devices to the database from service alerts

1. In the Cyber Protect console, go to **Monitoring > Alerts**.
2. [Locate a device control alert](#) that informs of denying access to the USB device.
3. In the alert simple view, click **Allow this USB device**.
This excludes the USB device from access control, and adds it to the database for further reference.

To add USB devices by importing a list of devices to the database

You can import a JSON file with a list of USB devices to the database. See "Import a list of USB devices to the database" (p. 360).

To remove USB devices from the database

1. Open the protection plan of a device for editing:
Click the ellipsis (...) next to the name of the protection plan and select **Edit**.

Note

Device control must be enabled in the plan, so you can access the Device control settings.

2. Click the arrow next to the **Device control** switch to expand the settings, and then click the **USB devices allowlist** row.
3. On the [page for managing the allowlist](#) that appears, click **Add from database**.
4. On the [page for selecting USB devices from the database](#), click ellipsis (...) at the end of the list item representing the device, click **Delete**, and confirm the deletion.
The USB devices are deleted from the database.
5. Close or save the protection plan.

View device control alerts

The device control module can be configured to raise alerts that inform of denied user attempts to use certain device types (see [Enable or disable OS notification and service alerts](#)). Use the following steps to view those alerts.

To view device control alerts

1. In the Cyber Protect console, go to **Monitoring > Alerts**.
2. Look for alerts with the following status: "Peripheral device access is blocked".

See [Device control alerts](#) for further details.

Access settings

On the **Access settings** page, you can allow or deny access to devices of certain types, as well as enable or disable OS notification and device control alerts.

Note

The access settings configured in Device control might be overridden when using both Device control and Advanced DLP to protect a workload. See "Enabling Advanced Data Loss Prevention in protection plans" (p. 815).

The access settings allow you to limit user access to the following device types and ports:

- **Removable** (access control by device type) - Devices with any interface for connecting to a computer (USB, FireWire, PCMCIA, IDE, SATA, SCSI, etc.) that are recognized by the operating system as removable storage devices (for example, USB sticks, card readers, magneto-optical drives, etc.). The device control classifies all hard drives connected via USB, FireWire, and PCMCIA as removable devices. It also classifies some hard drives (usually with SATA and SCSI) as removable devices if they support the hot-plug function and do not have the running operating system installed on them.

You can allow full access, read-only access, or deny access to removable devices to control data copy operations to and from any removable device on a protected computer. Access rights do not affect devices that are encrypted with BitLocker or FileVault (only HFS+ file system).

This device type is supported on both Windows and macOS.

- **Encrypted removable** (access control by device type) - Removable devices that are encrypted with BitLocker (on Windows) or FileVault (on macOS) drive encryption.

On macOS, only encrypted removable drives using the HFS+ (also known as HFS Plus or Mac OS Extended, or HFS Extended) file system are supported. Encrypted removable drives using the APFS file system are treated as removable drives.

You can allow full access, read-only access, or deny access to encrypted removable devices to control data copy operations to and from any encrypted removable device on a protected computer. Access rights affect only devices that are encrypted with BitLocker or FileVault (only HFS+ file system).

This device type is supported on both Windows and macOS.

- **Printers** (access control by device type) - Physical printers with any interface for connecting to a computer (USB, LPT, Bluetooth, etc.), as well as printers accessed from a computer on the network.

You can allow or deny access to printers to control the printing of documents on any printer on a protected computer.

Note

When you change the access setting for printers to **Deny**, the applications and processes accessing the printers must be restarted to enforce the newly configured access settings. To ensure that access settings are enforced correctly, restart the protected workloads.

This device type is supported only on Windows.

- **Clipboard** (access control by device type) - Windows clipboard.

You can allow or deny access to the clipboard to control the copy and paste operations through the Windows clipboard on a protected computer.

Note

When you change the access setting for clipboard to **Deny**, the applications and processes accessing the clipboard must be restarted to enforce the newly configured access settings. To ensure that access settings are enforced correctly, restart the protected workloads.

This device type is supported only on Windows.

- **Screenshot capture** (access control by device type) - Enables capturing of screenshots of the entire screen, the active window, or of selected portion of the screen.

You can allow or deny access to the screenshot capture to control the screenshot capturing on a protected computer.

Note

When you change the access setting for screenshot capture to **Deny**, the applications and processes accessing the screenshot capture must be restarted to enforce the newly configured access settings. To ensure that access settings are enforced correctly, restart the protected workloads.

This device type is supported only on Windows.

- **Mobile devices** (access control by device type) - Devices (such as Android-based smartphones, etc.) that communicate with a computer via Media Transfer Protocol (MTP), with any interface used for connecting to a computer (USB, IP, Bluetooth).

You can allow full access, allow read-only access, or deny access to mobile devices to control data copy operations to and from any MTP-based mobile device on a protected computer.

Note

When you change the access setting for mobile devices to **Read-only** or **Deny**, the applications and processes accessing the mobile devices must be restarted to enforce the newly configured access settings. To ensure that access settings are enforced correctly, restart the protected workloads.

This device type is supported only on Windows.

- **Bluetooth** (access control by device type) - External and internal Bluetooth devices with any interface for connecting to a computer (USB, PCMCIA, etc.). This setting controls the use of the devices of this type rather than data exchange using such devices.

You can allow or deny access to Bluetooth to control the use of any Bluetooth devices on a protected computer.

Note

On macOS, the access rights for Bluetooth do not affect Bluetooth HID devices. The access to these devices is always allowed to prevent wireless HID devices (mice and keyboards) from being disabled on iMac and Mac Pro hardware.

This device type is supported on both Windows and macOS.

- **Optical drives** (access control by device type) - External and internal CD/DVD/BD drives (including writers) with any interface for connecting to a computer (IDE, SATA, USB, FireWire, PCMCIA, etc.).

You can allow full access, allow read-only access, or deny access to optical drives to control data copy operations to and from any optical drive on a protected computer.

This device type is supported on both Windows and macOS.

- **Floppy drives** (access control by device type) - External and internal floppy drives with any interface for connecting to a computer (IDE, USB, PCMCIA, etc.). There are some models of floppy drives that the operating system recognizes as removable drives, in which case the device control also identifies these drives as removable devices.

You can allow full access, allow read-only access, or deny access to floppy drives to control data copy operations to and from any floppy drive on a protected computer.

This device type is supported only on Windows.

- **USB** (access control by device interface) - Any devices connected to a USB port, except hubs.

You can allow full access, allow read-only access, or deny access to USB port to control data copy operations to and from devices connected to any USB port on a protected computer.

This device type is supported on both Windows and macOS.

- **FireWire** (access control by device interface) - Any devices connected to a FireWire (IEEE 1394) port, except hubs.

You can allow full access, allow read-only access, or deny access to FireWire port to control data copy operations to and from devices connected to any FireWire port on a protected computer.

This device type is supported on both Windows and macOS.

- **Redirected devices** (access control by device interface) - Mapped drives (hard, removable and optical drives), USB devices, and the clipboard redirected to virtual application/desktop sessions. The device control recognizes devices redirected via the Microsoft RDP, Citrix ICA, VMware PCoIP, and HTML5/WebSockets remoting protocols in the Microsoft RDS, Citrix XenDesktop, Citrix XenApp, Citrix XenServer, and VMware Horizon virtualization environments hosted on protected Windows computers. It can also control data copy operations between the Windows clipboard of the guest operating system running on VMware Workstation, VMware Player, Oracle VM VirtualBox, or Windows Virtual PC, and the clipboard of the host operating system running on a protected Windows computer.

This device type is supported only on Windows.

You can configure access to redirected devices as follows:

- **Mapped drives** - Allow full access, allow read-only access, or deny access to control data copy operations to and from any hard drive, removable drive, or optical drive redirected to the session hosted on a protected computer.
- **Clipboard incoming** - Allow or deny access to control data copy operations through the clipboard to the session hosted on a protected computer.

Note

When you change the access setting for clipboard incoming to **Deny**, the applications and processes accessing the clipboard must be restarted to enforce the newly configured access settings. To ensure that access settings are enforced correctly, restart the protected workloads.

- **Clipboard outgoing** - Allow or deny access to control data copy operations through the clipboard from the session hosted on a protected computer.

Note

When you change the access setting for clipboard outgoing to **Deny**, the applications and processes accessing the clipboard must be restarted to enforce the newly configured access settings. To ensure that access settings are enforced correctly, restart the protected workloads.

- **USB ports** - Allow or deny access to control data copy operations to and from devices connected to any USB port redirected to the session hosted on a protected computer.

Device control settings affect all users equally. For example, if you deny access to removable devices, you prevent any user from copying data to and from such devices on a protected computer. It is possible to selectively allow access to individual USB devices by excluding them from access control (see [Device types allowlist](#) and [USB devices allowlist](#)).

When access to a device is controlled by both its type and its interface, denying access at the interface level takes precedence. For example, if access to USB ports is denied (device interface), then access to mobile devices connected to a USB port is denied regardless of whether access to mobile devices is allowed or denied (device type). To allow access to such a device, you must allow both its interface and type.

Note

If the protection plan used on macOS has settings for device types that are supported only on Windows, then the settings for these device types will be ignored on macOS.

Important

When a removable device, an encrypted removable device, a printer, or a Bluetooth device is connected to a USB port, allowing access to that device overrides the access denial set at the USB interface level. If you allow such a device type, access to the device is allowed regardless of whether access to the USB port is denied.

OS notification and service alerts

You can configure the device control to display OS notification to end users if they try to use a blocked device type on protected computers. When the **Show OS notification to end users if they try to use a blocked device type or port** check box is selected in the access settings, the agent displays a pop-up message in the notification area of the protected computer if any of the following events occurs:

- A denied attempt to use a device on a USB or FireWire port. This notification appears whenever the user plugs in a USB or FireWire device that is denied at the interface level (for example, when denying access to the USB port) or at the type level (for example, when denying the use of removable devices). The notification informs that the user is not allowed to access the specified device/drive.
- A denied attempt to copy a data object (such as a file) from a certain device. This notification appears when denying read access to the following devices: floppy drives, optical drives, removable devices, encrypted removable devices, mobile devices, redirected mapped drives, and redirected clipboard incoming data. The notification informs that the user is not allowed to get the specified data object from the specified device.

The denied read notification is also displayed when denying read/write access to Bluetooth, FireWire port, USB port, and redirected USB port.

- A denied attempt to copy a data object (such as a file) to a certain device. This notification appears when denying write access to the following devices: floppy drives, optical drives, removable devices, encrypted removable devices, mobile devices, local clipboard, screenshot capture, printers, redirected mapped drives, and redirected clipboard outgoing data. The notification informs that the user is not allowed to send the specified data object to the specified device.

User attempts to access blocked device types on protected computers can raise alerts that are logged in the Cyber Protect console. It is possible to enable alerts for each device type (excluding

screenshot capture) or port separately by selecting the **Show alert** check box in the access settings. For example, if access to removable devices is restricted to read-only, and the **Show alert** check box is selected for that device type, an alert is logged every time a user on a protected computer attempts to copy data to a removable device. See [Device control alerts](#) for further details.

See also [steps to enable or disable OS notification and service alerts](#).

Device types allowlist

On the **Device types allowlist** page, you can choose device subclasses to exclude from device access control. As a result, access to those devices is allowed regardless of the access settings in the device control module.

The device control module provides the option to allow access to devices of certain subclasses within a denied device type. This option allows you to deny all devices of a certain type, except for some subclasses of devices of this type. It can be useful, for example, when you need to deny access to all USB ports while allowing the use of a USB keyboard and mouse at the same time.

When configuring the device control module, you can specify which device subclasses to exclude from device access control. When a device belongs to an excluded subclass, access to that device is allowed regardless of whether or not the device type or port is denied. You can selectively exclude the following device subclasses from device access control:

- **USB HID (mouse, keyboard, etc.)** - When selected, allows access to Human Interface Devices (mouse, keyboard, and so on) connected to a USB port even if USB ports are denied. By default, this item is selected so that denying access to the USB port does not disable the keyboard or mouse.
Supported on both Windows and macOS.
- **USB and FireWire network cards** - When selected, allows access to network cards connected to a USB or FireWire (IEEE 1394) port even if USB ports and/or FireWire ports are denied.
Supported on both Windows and macOS.
- **USB scanners and still image devices** - When selected, allows access to scanners and still image devices connected to a USB port even if USB ports are denied.
Supported only on Windows.
- **USB audio devices** - When selected, allows access to audio devices, such as headsets and microphones, connected to a USB port even if USB ports are denied.
Supported only on Windows.
- **USB cameras** - When selected, allows access to Web cameras connected to a USB port even if USB ports are denied.
Supported only on Windows.
- **Bluetooth HID (mouse, keyboard, etc.)** - When selected, allows access to Human Interface Devices (mouse, keyboard, and so on) connected via Bluetooth even if Bluetooth is denied.
Supported only on Windows.
- **Clipboard copy/paste within application** - When selected, allows copying/pasting of data through the clipboard within the same application even if the clipboard is denied.
Supported only on Windows.

Note

Settings for unsupported device subclasses are ignored if these settings are configured in the applied protection plan.

When allowlisting device types, consider the following:

- With the device types allowlist, you can only allow a whole subclass of device. You cannot allow a specific device model, while denying all other devices of the same subclass. For example, by excluding USB cameras from device access control, you allow the use of any USB camera, no matter their model and vendor. On how to allow individual devices/models, see [USB devices allowlist](#).
- Device types can only be selected from a closed list of device subclasses. If the device to allow is of a different subclass, then it cannot be allowed by using device types allowlist. For example, such a subclass as USB smartcard readers cannot be added to the allowlist. To allow a USB smartcard reader when USB ports are denied, follow the instructions in [USB devices allowlist](#).
- The device types allowlist only works for devices that use standard Windows drivers. The device control may not recognize the subclass of some USB devices with proprietary drivers. As a result, you cannot allow access to such USB devices by using the device types allowlist. In this case, you could allow access on a per-device/model basis (see [USB devices allowlist](#)).

USB devices allowlist

The allowlist is intended to allow using certain USB devices regardless of any other device control settings. You can add individual devices or device models to the allowlist to disable the access control for those devices. For example, if you add a mobile device with a unique ID to the allowlist, you allow the use of that particular device even though any other USB devices are denied.

On the **USB devices allowlist** page, you can specify individual USB devices or USB device models to exclude from device access control. As a result, access to those devices is allowed regardless of the access settings in the device control module.

There are two ways to identify devices in the allowlist:

- Model of device - Collectively identifies all devices of a certain model. Each device model is identified by vendor ID (VID) and product ID (PID), such as `USB\VID_0FCE&PID_E19E`.
This combination of VID and PID does not identify a specific device, but an entire device model. By adding a device model to the allowlist, you allow access to any device of that model. For example, this way you can allow the use of USB printers of a particular model.
- Unique device - Identifies a certain device. Each unique device is identified by vendor ID (VID), product ID (PID), and serial number, such as `USB\VID_0FCE&PID_E19E\D55E7FCA`.
Not all USB devices are assigned a serial number. You can add a device to the allowlist as a unique device only if the device has been assigned a serial number during production. For example, a USB stick that has a unique serial number.

To add a device to the allowlist, you first need to add it to the [USB devices database](#). Then, you can add devices to the allowlist by selecting from that database.

The allowlist is managed on a separate configuration page called **USB devices allowlist**. Each item in the list represents a device or device model and has the following fields:

- **Description** - The operating system assigns a certain description when connecting the USB device. You can modify the description of the device in the USB devices database (see [USB database management page](#)).
- **Device type** - Displays Unique if the list item represents a unique device, or Model if it represents a device model.
- **Read-only** - When selected, allows only receiving data from the device. If the device does not support read-only access, then access to the device is blocked. Clear this check box to allow full access to the device.
- **Reinitialize** - When selected, causes the device to simulate disconnecting/reconnecting when a new user logs in. Some USB devices require reinitializing in order to function, so we recommend that you select this check box for such devices (mouse, keyboard, etc.). We also recommend that you clear this check box for data storage devices (USB sticks, optical drives, external hard drives, etc.).

The device control may not be able to reinitialize some USB devices with proprietary drivers. If there is no access to such a device, you must remove the USB device from the USB port, and then insert it back.

Note

The **Reinitialize** field is hidden by default. To display it in the table, click the gear icon in the upper right corner of the table, and then select the **Reinitialize** check box.

Note

The **Read-only** and **Reinitialize** fields are not supported on macOS. If these fields are configured in the applied protection plan, they will be ignored.

You can add or remove devices/models from the allowlist as follows:

- Click **Add from database** above the list and then select the desired device/s from those registered with the [USB devices database](#). The selected device is added to the list, where you can configure its settings and confirm the changes.
- Click **Allow this USB device** in an alert informing that access to the USB device is denied (see [Device control alerts](#)). This adds the device to the allowlist and to the USB devices database.
- Click the delete icon at the end of a list item. This removes the respective device/model from the allowlist.

USB devices database

The device control module maintains a database of USB devices from which you can add devices to the list of exclusions (see [USB devices allowlist](#)). A USB device can be registered with the database in

any of these ways:

- Add a device on the page that appears when adding a device to the exclusion list (see [USB devices database management page](#)).
- Add a device from the USB Devices tab of a computer's Inventory pane in the Cyber Protect console (see [List of USB devices on a computer](#)).
- Allow the device from an alert on denying access to the USB device (see [Device control alerts](#)).

See also [steps to add or remove USB devices from the database](#).

USB devices database management page

When configuring the allowlist for USB devices, you have the option to add a device from the database. If you choose this option, a management page appears with a list of devices. On this page you can view the list of all devices that are registered with the database, you can select devices to add to the allowlist, and perform the following operations:

Register a device with the database

1. Click **Add to database** at the top of the page.
2. On the **Add USB device** dialog that appears, choose the machine to which the USB device is connected.

Only machines that are online are displayed in the list of computers.

The list of USB devices is displayed only for machines that have the agent for Data Loss Prevention installed.

The USB devices are listed in tree view. The first level of the tree represents a device model. The second level represents a specific device of that model.

A blue icon next to the description of the device indicates that the device is currently attached to the computer. If the device is not attached to the computer, the icon is grayed out.

3. Select the check box for the USB device that you want to register, and click **Add to database**.

Change the description of a device

1. On the **USB devices database** page click ellipsis (...) at the end of the list item representing the device and then click **Edit**.
2. Make changes to the description in the dialog box that appears.

Remove a device from the database

1. Click the ellipsis (...) at the end of the list item representing the device.
2. Click **Delete**, and confirm the deletion.

For each device, the list on the page provides the following information:

- **Description** - A readable identifier of the device. You can change the description as needed.
- **Device type** - Displays Unique if the list item represents a unique device, or Model if it represents a device model. A unique device must have a serial number along with a vendor ID (VID) and product ID (PID), whereas a device model is identified by a combination of VID and PID.

- **Vendor ID, Product ID, Serial number** - These values together make up the device ID in the form USB\VID_<vendor ID>&PID_<product ID>\<serial number>.
- **Account** - Indicates the tenant to which this device belongs. This is the tenant that contains the user account that was used to register the device with the database.

Note

This column is hidden by default. To display it in the table, click the gear icon in the upper right corner of the table, and then select **Account**.

The leftmost column is intended to select the devices to add to the allowlist: Select the check box for each device to add, and then click the **Add to allowlist** button. To select or clear all check boxes, click the check box in the column header.

You can search or filter the list of devices:

- Click **Search** at the top of the page and enter a search string. The list displays devices whose description matches the string you typed.
- Click **Filter**, and then configure and apply a filter in the dialog box that appears. The list is limited to devices with the type, vendor ID, product ID, and account that you selected when configuring the filter. To cancel the filter and list all devices, click **Reset to default**.

Export the list of USB devices in the database

You can export the list of USB devices that are added to the database.

1. Open the protection plan of a device for editing.
2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the **USB devices allowlist** row.
3. On the USB devices allowlist page, click **Add from database**.
4. On the USB devices database management page that appears, click **Export**.
The standard Browse dialog opens.
5. Select the location to which you want to save the file, enter a new file name if needed, and click **Save**.

The list of USB devices is exported to a JSON file.

You can edit the resulting JSON file to add or remove devices from it, and make mass changes of device descriptions.

Import a list of USB devices to the database

Instead of adding USB devices from the Cyber Protect console, you can import a list of USB devices. The list is a file in JSON format.

Note

You can import JSON files to a database that does not contain the devices described in the file. To import a modified file to the database from which it was exported, you must clear the database first because you cannot import duplicate entries. If you export the list of USB devices, modify it, and try to import to the same database without clearing it, the import will fail.

1. Open the protection plan of a device for editing.
2. Click the arrow icon next to the **Device control** switch to expand the settings, and then click the **USB devices allowlist** row.
3. On the USB devices allowlist page, click **Add from database**.
4. On the USB devices database management page that appears, click **Import**.
The dialog Import USB devices from file opens.
5. Use drag and drop (or browse) for the file that you want to import.

The Cyber Protect console checks if the list contains duplicate entries that already exist in the database and skips them. The USB devices that are not found in the database are appended to it.

List of USB devices on a computer

The Inventory panel of a computer in the Cyber Protect console includes the **USB Devices** tab. If the computer is online and the agent for Data Loss Prevention is installed on it, the **USB Devices** tab displays a list all USB devices that have ever been connected to that computer.

The USB devices are listed in tree view. The first level of the tree represents a device model. The second level represents a specific device of that model.

For each device, the list provides the following information:

- **Description** - The operating system assigns a description when connecting the USB device. This description can serve as a readable identifier of the device.
A blue icon next to the description of the device indicates that the device is currently attached to the computer. If the device is not attached to the computer, the icon is grayed out.
- **Device ID** - The identifier that the operating system assigned to the device. This identifier has the following format: USB\VID_<vendor ID>&PID_<product ID>\<serial number> where <serial number> is optional. Examples: USB\VID_0FCE&PID_ADDE\D55E7FCA (device with a serial number); USB\VID_0FCE&PID_ADDE (device without serial number).

To add devices to the USB devices database, select the check boxes of the desired devices, and then click the **Add to database** button.

Excluding processes from access control

The access to Windows clipboard, screenshot capture, printers, and mobile devices is controlled through hooks injected into processes. If processes are not hooked, the access to these devices will not be controlled.

Note

Excluding processes from access control is not supported on macOS. If a list of excluded processes is configured in the applied protection plan, it will be ignored.

On the **Exclusions** page, you can specify a list of processes that will not be hooked. This means that clipboard (local and redirected), screenshot capture, printer, and mobile device access controls will not be applied to such processes.

For example, you applied a protection plan that denies access to printers, then started the Microsoft Word application. An attempt to print from this application will be blocked. But if you add the Microsoft Word process to the list of exclusions, then the application will not be hooked. As a result, printing from Microsoft Word will not be blocked, while printing from other applications will still be blocked.

To add processes to exclusions

1. Open the protection plan of a device for editing:
Click the ellipsis (...) next to the name of the protection plan and select **Edit**.

Note

Device control must be enabled in the plan, so you can access the Device control settings.

2. Click the arrow next to the **Device control** switch to expand the settings, and then click the **Exclusions** row.
3. On the **Exclusions** page, in the **Processes and folders** row, click **+Add**.
4. Add the processes that you want to exclude from the access control.
For example, `C:\Folder\subfolder\process.exe`.
You can use wildcards:
 - * replaces any number of characters.
 - ? replaces one character.For example:
`C:\Folder*`
`*\Folder\SubFolder?*`
`*\process.exe`
5. Click the check mark, and then click **Done**.
6. In the protection plan, click **Save**.
7. Restart the processes that you excluded to ensure that the hooks are properly removed.

The excluded processes will have access to clipboard, screenshot capture, printers, and mobile devices regardless of the access settings for those devices.

To remove a process from exclusions

Open the protection plan of a device for editing:

Click the ellipsis (...) next to the name of the protection plan and select **Edit**.

Note

Device control must be enabled in the plan, so you can access the Device control settings.

1. Click the arrow next to the **Device control** switch to expand the settings, and then click the **Exclusions** row.
2. On the **Exclusions** page, click the trash can icon next to the process that you want to remove from the exclusions.
3. Click **Done**.
4. In the protection plan, click **Save**.
5. Restart the process to ensure that hooks are properly injected.

The access settings from the protection plan will be applied to the processes that you removed from the exclusions.

To edit a process in exclusions

1. Open the protection plan of a device for editing:
Click the ellipsis (...) next to the name of the protection plan and select **Edit**.

Note

Device control must be enabled in the plan, so you can access the Device control settings.

2. Click the arrow next to the **Device control** switch to expand the settings, and then click the **Exclusions** row.
3. On the **Exclusions** page, click the **Edit** icon next to the process that you want to edit.
4. Apply the changes and click the check mark to confirm.
5. Click **Done**.
6. In the protection plan, click **Save**.
7. Restart the affected processes to ensure that your changes are applied correctly.

Device control alerts

The device control maintains an event log by tracking user attempts to access controlled device types, ports, or interfaces. Certain events can raise alerts that are logged in the Cyber Protect console. For example, the device control module can be configured to prevent the use of removable devices, with an alert logged whenever a user tries to copy data to or from such a device.

When configuring the device control module, you can enable alerts for most items listed under device Type (except screenshot capture) or Ports. If alerts are enabled, each attempt by a user to perform an operation that is not allowed generates an alert. For example, if access to removable devices is restricted to read-only, and the **Show alert** option is selected for that device type, an alert is generated every time a user on a protected computer attempts to copy data to a removable device.

To view alerts in the Cyber Protect console, go to **Monitoring > Alerts**. Within each device control alert, the console provides the following information about the respective event:

- **Type**—Warning.
- **Status**—Displays “Peripheral device access is blocked”.
- **Message**—Displays “Access to '<device type or port>' on '<computer name>' is blocked”. For example, “Access to 'Removable' on 'accountant-pc' is blocked”.
- **Date and time**—The date and time that the event occurred.
- **Device**—The name of the computer on which the event occurred.
- **Plan name**—The name of the protection plan that caused the event.
- **Source**—The device type or port involved in the event. For example, in the event of a denied user attempt to access a removable device, this field reads Removable device.
- **Action**—The operation that caused the event. For example, in the event of a denied user attempt to copy data to a device, this field reads Write. For more information, see [Action field values](#).
- **Name**—The name of the event target object, such as the file the user attempted to copy or the device the user attempted to use. Not displayed if the target object cannot be identified.
- **Information**—Additional information about the event target device, such as the device ID for USB devices. Not displayed if no additional information about the target device is available.
- **User**—The name of the user who caused the event.
- **Process**—The fully qualified path to the executable file of the application that caused the event. In some cases, the process name might be displayed instead of the path. Not displayed if process information is not available.

If an alert applies to a USB device (including removable devices and encrypted removable devices), then, directly from the alert, the administrator can add the device to the allowlist, which prevents the device control module from restricting access to that particular device. Clicking **Allow this USB device** adds it to the USB devices allowlist in the device control module’s configuration, and also adds it to the [USB devices database](#) for further reference.

See also [steps to view device control alerts](#).

Action field values

Alert **Action** field can contain the following values:

- **Read** - Get data from the device or port.
- **Write** - Send data to the device or port.
- **Format** - Direct access (formatting, check disk, etc.) to the device. In the case of a port, applies to the device connected to that port.
- **Eject** - Remove the device from the system or eject the media from the device. In the case of a port, applies to the device connected to that port.
- **Print** - Send a document to the printer.
- **Copy audio** - Copy/paste audio data via the local clipboard.
- **Copy file** - Copy/paste a file via the local clipboard.

- **Copy image** - Copy/paste an image via the local clipboard.
- **Copy text** - Copy/paste text via the local clipboard.
- **Copy unidentified content** - Copy/paste other data via the local clipboard.
- **Copy RTF data (image)** - Copy/paste an image via the local clipboard using Rich Text Format.
- **Copy RTF data (file)** - Copy/paste a file via the local clipboard using Rich Text Format.
- **Copy RTF data (text, image)** - Copy/paste text along with an image via the local clipboard using Rich Text Format.
- **Copy RTF data (text, file)** - Copy/paste text along with a file via the local clipboard using Rich Text Format.
- **Copy RTF data (image, file)** - Copy/paste an image along with a file via the local clipboard using Rich Text Format.
- **Copy RTF data (text, image, file)** - Copy/paste text along with an image and a file via the local clipboard using Rich Text Format.
- **Delete** - Delete data from the device (for example, a removable device, a mobile device, and so on).
- **Device access** - Access to some device or port (for example, a Bluetooth device, a USB port, and so on).
- **Incoming audio** - Copy/paste audio data from the client computer to the hosted session via the redirected clipboard.
- **Incoming file** - Copy/paste a file from the client computer to the hosted session via the redirected clipboard.
- **Incoming image** - Copy/paste an image from the client computer to the hosted session via the redirected clipboard.
- **Incoming text** - Copy/paste text from the client computer to the hosted session via the redirected clipboard.
- **Incoming unidentified content** - Copy/paste other data from the client computer to the hosted session via the redirected clipboard.
- **Incoming RTF data (image)** - Copy/paste an image from the client computer to the hosted session via the redirected clipboard using Rich Text Format.
- **Incoming RTF data (file)** - Copy/paste a file from the client computer to the hosted session via the redirected clipboard using Rich Text Format.
- **Incoming RTF data (text, image)** - Copy/paste text along with an image from the client computer to the hosted session via the redirected clipboard using Rich Text Format.
- **Incoming RTF data (text, file)** - Copy/paste text along with a file from the client computer to the hosted session via the redirected clipboard using Rich Text Format.
- **Incoming RTF data (image, file)** - Copy/paste an image along with a file from the client computer to the hosted session via the redirected clipboard using Rich Text Format.
- **Incoming RTF data (text, image, file)** - Copy/paste text along with an image and a file from the client computer to the hosted session via the redirected clipboard using Rich Text Format.
- **Insert** - Connect a USB device or a FireWire device.

- **Outgoing audio** - Copy/paste audio data from the hosted session to the client computer via the redirected clipboard.
- **Outgoing file** - Copy/paste a file from the hosted session to the client computer via the redirected clipboard.
- **Outgoing image** - Copy/paste an image from the hosted session to the client computer via the redirected clipboard.
- **Outgoing text** - Copy/paste text from the hosted session to the client computer via the redirected clipboard.
- **Outgoing unidentified content** - Copy/paste other data from the hosted session to the client computer via the redirected clipboard.
- **Outgoing RTF data (image)** - Copy/paste an image from the hosted session to the client computer via the redirected clipboard using Rich Text Format.
- **Outgoing RTF data (file)** - Copy/paste a file from the hosted session to the client computer via the redirected clipboard using Rich Text Format.
- **Outgoing RTF data (text, image)** - Copy/paste text along with an image from the hosted session to the client computer via the redirected clipboard using Rich Text Format.
- **Outgoing RTF data (text, file)** - Copy/paste text along with a file from the hosted session to the client computer via the redirected clipboard using Rich Text Format.
- **Outgoing RTF data (image, file)** - Copy/paste an image along with a file from the hosted session to the client computer via the redirected clipboard using Rich Text Format.
- **Outgoing RTF data (text, image, file)** - Copy/paste text along with an image and a file from the hosted session to the client computer via the redirected clipboard using Rich Text Format.
- **Rename** - Rename files on a device (for example, on removable devices, mobile devices, and others).

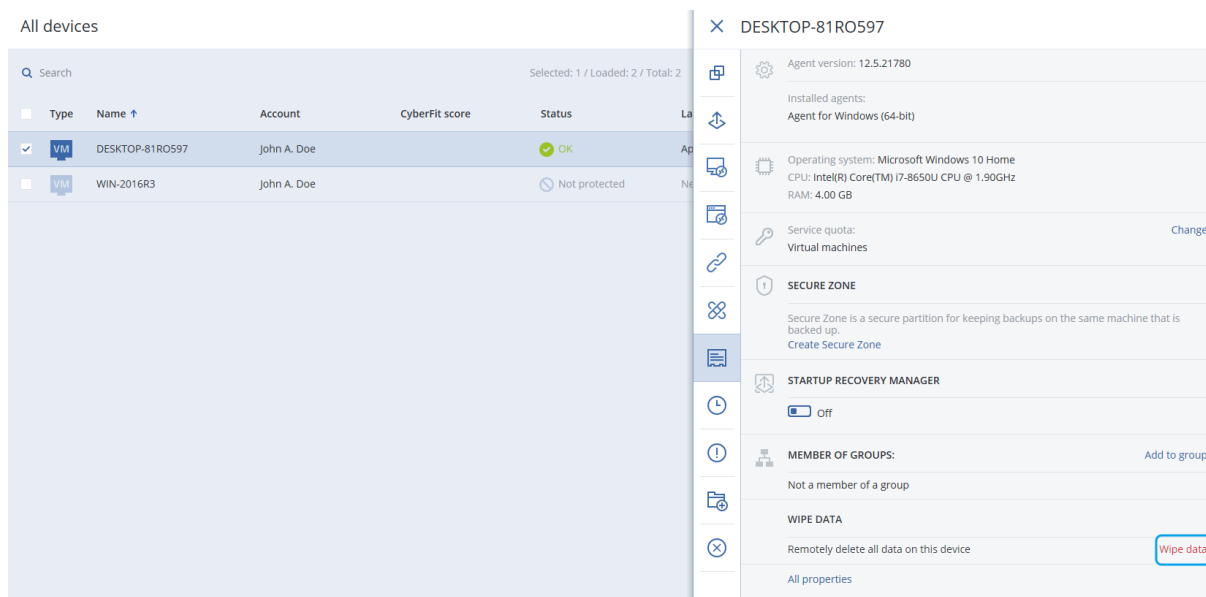
Wiping data from a managed workload

Note

Remote wipe is available with the Advanced Security pack.

Remote wipe allows a Cyber Protection service administrator and a machine owner to delete the data on a managed machine – for example, if it gets lost or stolen. Thus, any unauthorized access to sensitive information will be prevented.

Remote wipe is only available for machines running Windows versions 10 and later. To receive the wipe command, the machine must be turned on and connected to the Internet.



To wipe data from a machine

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the machine whose data you want to wipe.

Note

You can wipe data from one machine at a time.

3. Click **Details**, and then click **Wipe data**.
If the machine that you selected is offline, the **Wipe data** option is inaccessible.
4. Confirm your choice.
5. Enter the credentials of this machine's local administrator, and then click **Wipe data**.

Note

You can check the details about the wiping process and who started it in **Monitoring > Activities**.

Managing the isolation of workloads

Note

This functionality is part of the Advanced Security + EDR pack. In addition, it only works with workloads that have the Acronis agent installed.

Isolating a workload from the network enables you to mitigate the risk of malware present on a specific workload from spreading to other workloads.

Note that you can also isolate a workload when defining response actions to an incident. For more information, see "Manage the network isolation of a workload" (p. 865).

If you need to recover a backup for a workload, see "Recovery" (p. 469).

Isolating a workload from the network

As part of your response to an attack, you can isolate affected workloads in the Cyber Protect console. You can also reconnect a workload to the network, as and when required.

To isolate a workload from the network

1. In the Cyber Protect console, go to **Workloads > Workloads with agents**.
2. Locate and select the relevant workload(s). Use the search option to find the relevant workloads; for example, you can search for workloads according to their current network status, or according to a defined IP range.
Note that you can also select offline workloads to be isolated; the workload will become **Isolated** when the workload is back online. Similarly, a workload that is offline can also be reconnected to the network; the workload will exit the **Isolated** state when the workload is back online.
3. In the displayed right sidebar, click **Manage network isolation**. The following dialog is displayed.

The screenshot shows a dialog box with a light blue background. At the top, it says "Network status: **Connected**". Below that, a paragraph explains: "Workload network isolation restricts network access to all non-RDP and backup processes for infected/suspicious workload to reduce the risk of the infection spreading." This is followed by the question: "Do you want to isolate the network of workload qa-gw3t68hh?". A white-bordered box contains the text: "Message to display: The administrator is placing this workload in Isolated State. Network access in and out of this workload is restricted." At the bottom, there are two buttons: a red "Isolate" button and a blue "Manage network exclusions" button.

Note

The **Network status** value indicates if the workload is currently connected or not. If the value displays **Isolated**, you can reconnect the isolated workload to the network, as described in the procedure below.

4. In the **Message to display** field, add a message to display to end users to let them know that the workload is under investigation, and access to and from the workload is prohibited until further notice.
5. Click **Manage network exclusions** to add ports, URLs, hostnames, and IP addresses that will have access to the workload during the isolation. For more information, see [how to manage network exclusions](#).

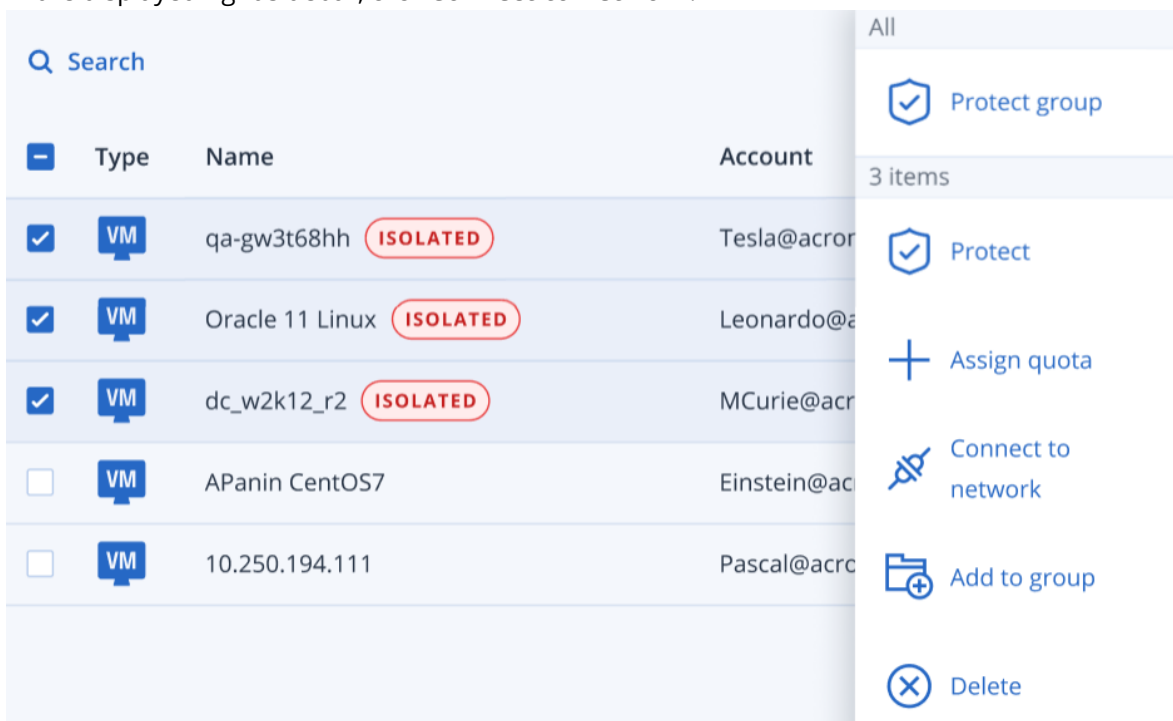
- Click **Isolate**. The workload is isolated and shows **ISOLATED** in the displayed list of workloads. Note that when you isolate a workload, or reconnect it to the network, this action is recorded in the Activities screen (go to **Monitoring > Activities**). You can also see these recorded actions for each individual workload (select a workload, and click **Activities** in the right side panel).

Note

Backups that occurred during the isolation will be marked as suspicious. If you try to recover this suspicious backup, a warning message is displayed.

To connect an isolated workload back to the network

- In the **Workloads with agents** screen, locate and select the required workload(s). Note that you can use the search option to find workloads with the **Isolated** network status.
- In the displayed right sidebar, click **Connect to network**.



- [Optional] In the displayed dialog, add a message in the **Message to display** field; this message is displayed to end users when they access the connected workload.
- Click **Connect** (or **Connect all** if you selected multiple workloads). The workload is reconnected to the network and all access to and from the workload is no longer restricted.

Managing network exclusions

You can manage the ports, DNS names, and IP addresses that will have access to a workload during isolation.

To manage network exclusions

Note

Even if all Acronis Cyber Protect technologies are working when the workload is in isolation, there may be scenarios in which you need additional network connections to be established (for example, you may need to upload a file from the workload to a shared directory). In these scenarios, you can add a network exclusion, but make sure any threats are removed before you add the exclusion.

1. Locate and select the relevant workload(s), as described in "Isolating a workload from the network" (p. 368).
 2. Click **Manage network exclusions**.
 3. For each of the options available (Ports, DNS names / IP addresses), do the following:
 - a. Click **Add** and then enter the relevant port(s), or DNS names / IP addresses. Note that if you selected multiple workloads, you cannot define ports or DNS names / IP addresses.
 - b. In the **Traffic direction** drop-down list, select one of **Incoming and outgoing connections**, **Incoming connections only**, or **Outgoing connections only**.
 - c. Click **Add**.
-

Note

For best practice, make the exclusion rule as restrictive as possible.

4. Click **Save**.

Viewing workloads managed by RMM integrations

Note

This feature is only available if the Advanced Automation service is enabled.

When you integrate an RMM platform as part of the Advanced Automation service, you can view and monitor information from devices that are managed by the RMM platform. This information is available in the Cyber Protect console by navigating to **Devices**.

To view workloads managed by RMM integrations

1. Go to **Devices > All devices**.
2. (Optional) Sort the **RMM integration** column to locate the relevant integrations.
3. Select the relevant workload.
4. In the **Actions** pane, select **Details**.
5. In the displayed pane, one of three options is displayed, according to your configured workload:
 - If Acronis services are defined for the workload without RMM integration: If the workload is configured to work only with Acronis services, no RMM integration information is displayed.
 - If Acronis services and an RMM integration is configured for the workload: The Acronis services and RMM integration details are located in two tabs, **Overview** and **RMM integration**. Click **RMM integration** to view the integration details, including the workload name and type (provided by the RMM platform), description and location. In addition, any

installed and enabled RMM agent add-ons are also shown.

- If the workload is configured with an RMM integration only: The RMM integration details are displayed, including the workload name and type (provided by the RMM platform), description and location. In addition, any installed and enabled RMM agent add-ons are also shown.

Note that when the workload is configured with RMM integration (either in tandem with Acronis services or with an RMM integration only), you can do the following:

- Initiate a remote connection (available for Datto RMM, N-able N-central, N-able RMM integrations)
- Review installed add-ons on the third party RMM device (available for N-able RMM only)
- Directly access the third party RMM device's details (available for Datto RMM, N-able N-central, NinjaOne)

CyberApp workloads

CyberApp workloads are created by ISVs (Independent software vendors) and appear in the Cyber Protect console after you enable a CyberApp integration. The following conditions must be met:

- The **Workloads and actions** extension point must be enabled in the CyberApp.
- At least one **Workload type** must be defined in the CyberApp.
- The connector service hosted by the ISV must ensure that the CyberApp workloads are added and updated to the Acronis platform.

For more information about Vendor Portal and creating CyberApps, see the Vendor Portal User Guide.

Aggregated workloads

A physical workload may have Cyber Protect agent and one or several CyberApp agents installed at the same time. In this case, the same workload will have more than one representation on the **All Devices** screen - a separate record will be shown for the Acronis workload, and for each CyberApp workload. If the automatic merging of workloads is enabled and configured from Vendor Portal or from the Cyber Protect console, the system will compare the host addresses and the MAC addresses of the Acronis workloads and the CyberApp workloads, and will merge all representations into a single aggregated workload. You can also manually merge and unmerge workloads in the Cyber Protect console.

Working with CyberApp workloads

Apart from the standard actions that are built-in to the Cyber Protect console, you can perform actions that become available after the CyberApp workloads appear in the console: manually merge workloads into an aggregated workload and perform custom actions that are configured in the CyberApp.

Merge

Prerequisites

- Workloads from different sources are available for the tenant.

You can manually merge an Acronis workload with one or several CyberApp workloads into a single aggregated workload.

To manually merge workloads into an aggregated workload

1. In the **All devices** screen, select the workloads that you want to merge.

Note

The merge action is displayed if you select workloads from different sources, such as an Acronis workload and a CyberApp workload.

2. Click **Merge workloads**.

Perform custom actions

Prerequisites

- A CyberApp integration that has **Workload actions** defined is enabled for the tenant.

Custom actions are actions that are configured in the CyberApp, and become available for the corresponding CyberApp workload when you enable the CyberApp integration for the tenant.

To perform custom actions

1. In the **All devices** screen, click the workload.
2. Click **Integrated App actions**.
3. Click the action.

Working with aggregated workloads

Apart from the standard actions that are built-in to the Cyber Protect console, you can perform the following operations with aggregated workloads: view details, unmerge source workloads, and perform custom actions that are configured in the CyberApps.

View details

Prerequisites

- At least one aggregated workload is available for the tenant.

To view the details of an aggregated workload

1. In the **All devices** screen, click the aggregated workload.
2. Click **Details**.

The details of the aggregated workload are separated into tabs. Each tab shows the details for each workload representation.

Unmerge

Prerequisites

- At least one aggregated workload is available for the tenant.

When you unmerge an aggregated workload, it will no longer be displayed in the devices list. Instead, you will view a separate entry for each source workload that has been merged into the aggregate workload.

To unmerge an aggregated workload

1. In the **All devices** screen, click the aggregated workload that you want to unmerge.
2. Click **Unmerge source workloads**.
3. In the confirmation window, click **Unmerge**.

Perform custom actions

Prerequisites

- At least one CyberApp integration that has **Workload actions** defined is enabled for the tenant.

Custom actions are actions that are configured in the CyberApps and become available for the corresponding CyberApp workload when you enable the CyberApp integration for the tenant.

To perform custom actions

1. In the **All devices** screen, click the workload.
2. Click **Integrated App actions**.
3. Depending on the available custom actions, do one of the following.
 - If the aggregated workload has one CyberApp workload, click the action.
 - If the aggregated workload has more than one CyberApp workload, click the name of the CyberApp, and then click the action.

Linking workloads to specific users

Note

This feature is only available if the Advanced Automation service is enabled.

By linking a workload to a specific user, you can automatically link the workload to new service desk tickets created by or assigned to the user.

To link a workload to a user

1. Go to **Devices > All devices**, and then select the relevant workload.
2. In the **Actions** pane, select **Link to a user**.
3. Select the relevant user.

You can also change the selected user for existing linked workloads, as required.
4. Click **Done**. The selected user is now displayed in the **Linked user** column.

To unlink a workload from a user

1. Go to **Devices > All devices**, and then select the relevant workload.
2. In the **Actions** pane, select **Link to a user**.
3. Click **Unlink user**.
4. Click **Done**.

Find the last logged in user

In order for the administrators to manage devices, they have to identify which user is and was logged in to a device. This information is displayed in the Dashboard or in the workloads details.

You can enable or disable displaying the Last login information in [Remote management plans](#).

In the Dashboard:

1. Click **Devices**. The **All devices** window is displayed.
2. In the **Last login** column, the name of the user who logged in the last time for each device is displayed.
3. In the **Last login time** column, the time when the user logged in the last time for each device is displayed.

In Device Details:

1. Click **Devices**. The **All devices** window is displayed.
2. Click on the device for which you want to verify the details.
3. Click the **Details** icon. The name of the user, the date and time of the last logins for the selected device is displayed in the **Last users logged in** section.

Note

In the **Last users logged in** section there will be displayed up to 5 different users who logged in to the device.

To show or hide Last login and Last login time columns In the Dashboard

1. Click **Devices**. The **All devices** window is displayed.
2. Click the gear icon in the upper right corner, and do one of the following in the **General** section:
 - Enable the **Last login** and **Last login time** columns, if you want to show them on the Dashboard.
 - Disable the **Last login** and **Last login time** columns, if you want to hide them from the Dashboard.

Managing the backup and recovery of workloads and files

The backup module enables backup and recovery of physical and virtual machines, files, and databases to local or cloud storage.

Backup

A protection plan with the Backup module enabled is a set of rules that specify how the given data will be protected on a given machine.

A protection plan can be applied to multiple machines at the time of its creation, or later.

To create the first protection plan with the Backup module enabled

1. Select the machines that you want to back up.
2. Click **Protect**.
Protection plans that are applied to the machine are shown. If the machine does not have any plans already assigned to it, then you will see the default protection plan that can be applied. You can adjust the settings as needed and apply this plan or create a new one.
3. To create a new plan, click **Create plan**. Enable the **Backup** module and unroll the settings.

New protection plan (2)

Cancel
Create

Backup

Entire machine to Cloud storage, Monday to Friday at 05:45 PM

▼

What to back up

Entire machine ▼

Continuous data protection (CDP)

Where to back up

Cloud storage

Schedule

Monday to Friday at 05:45 PM i

How long to keep

Monthly: 6 months

Weekly: 4 weeks

Daily: 7 days

Encryption

i

Application backup

Disabled i

Backup options


Change

4. [Optional] To modify the protection plan name, click the default name.
5. [Optional] To modify the Backup module parameters, click the corresponding setting of the protection plan panel.
6. [Optional] To modify the backup options, click **Change** next to **Backup options**.
7. Click **Create**.

To apply an existing protection plan

1. Select the machines that you want to back up.
2. Click **Protect**. If a common protection plan is already applied to the selected machines, click **Add plan**.

The software displays previously created protection plans.



3. Select a protection plan to apply.
4. Click **Apply**.

Protection plan cheat sheet

The following table summarizes the available protection plan parameters. Use the table to create a protection plan that best fits your needs.

WHAT TO BACK UP	ITEMS TO BACK UP Selection methods	WHERE TO BACK UP	SCHEDULE Backup schemes	HOW LONG TO KEEP
Disks/volumes (physical machines ¹)	Direct selection Policy rules File filters	Cloud Local folder Network folder	Always incremental (Single-file) Always full Weekly full, Daily incremental	By backup age (single rule/per backup set) By number of backups

¹A machine that is backed up by an agent installed in the operating system.

			NFS*		
			Secure Zone**		
Disks/volumes (virtual machines ¹)	Policy rules File filters		Cloud Local folder Network folder NFS*	Monthly full, Weekly differential, Daily incremental (GFS) Custom (F-D-I)	By total size of backups*** Keep indefinitely
Files (physical machines only ²)	Direct selection Policy rules File filters		Cloud Local folder Network folder NFS* Secure Zone**	Always incremental (Single-file) Always full Weekly full, Daily incremental Monthly full, Weekly differential, Daily incremental (GFS)	
ESXi configuration	Direct selection		Local folder Network folder NFS*	Custom (F-D-I)	
Websites (files and MySQL databases)	Direct selection		Cloud	—	
System state	Direct selection		Cloud	Always full Weekly full, daily incremental	
SQL databases			Local folder	Custom (F-I)	
Exchange databases			Network folder	Always incremental (Single-file) - only for SQL databases	
Microsoft 365	Mailboxes	Direct selection	Cloud	Always incremental (Single-file)	

¹A virtual machine that is backed up at a hypervisor level by an external agent such as Agent for VMware or Agent for Hyper-V. A virtual machine with an agent inside is treated as physical from the backup standpoint.

²A machine that is backed up by an agent installed in the operating system.

	(local Agent for Microsoft 365)		Local folder Network folder		
	Mailboxes (cloud Agent for Microsoft 365)	Direct selection	Cloud	Up to 6 backups per day	
	Public folders				
	Teams				
	OneDrive files	Direct selection	Cloud	Up to 6 backups per day	
	SharePoint Online data	Policy rules			
Google Workspace	Gmail mailboxes	Direct selection	Cloud	Up to 6 backups per day	
	Google Drive files	Direct selection			
	Shared drive files	Policy rules			

* Backup to NFS shares is not available in Windows.

** Secure Zone cannot be created on a Mac.

*** The **By total size of backups** retention rule is not available with the **Always incremental (single-file)** backup scheme or when backing up to the cloud storage.

Selecting data to back up

Selecting entire machine

A backup of an entire machine is a backup of all its non-removable disks. For more information about disk backup, refer to "Selecting disks or volumes" (p. 380).

Limitations

- Disk-level backups are not supported for encrypted APFS volumes that are locked. During a backup of an entire machine, such volumes are skipped.
- The OneDrive root folder is excluded from backup operations by default. If you select to back up specific OneDrive files and folders, they will be backed up. Files that are not available on the device will have invalid contents in the backup set.

Selecting disks or volumes

A disk-level backup contains a copy of a disk or a volume in a packaged form. From a disk-level backup, you can recover disks, volumes, folders, and files.

Note

The disk-level backup supports only plain file systems, such as .nfts, .ext, .hptfs, excluding from browsing .lvm and .md volumes.

You can select the disks or volumes to back up for each individual workload in the protection plan (direct selection) or you can configure policy rules for multiple workloads. Additionally, you can exclude specific files from a backup, or include only specific files to it, by configuring file filters. For more information, see "File filters (Inclusions/Exclusions)" (p. 435).

To select disks or volumes

Direct selection

Direct selection is available only for physical machines.

1. In **What to back up**, select **Disks/volumes**.
2. Click **Items to back up**.
3. In **Select items for backup**, select **Directly**.
4. For each of the workloads included in the protection plan, select the check boxes next to the disks or volumes to back up.
5. Click **Done**.

By policy rules

1. In **What to back up**, select **Disks/volumes**.
2. Click **Items to back up**.
3. In **Select items for backup**, select **Using policy rules**.
4. Select any of the predefined rules, type your own rules, or combine both.
For more information about the available policy rules, see "Policy rules for disks and volumes" (p. 382).
The policy rules will be applied to all workloads that are included in the protection plan.
If none of the specified rules can be applied to a workload, the backup of that workload fails.
5. Click **Done**.

Limitations

- Disk-level backups are not supported for encrypted APFS volumes that are locked. During a backup of an entire machine, such volumes are skipped.
- The OneDrive root folder is excluded from backup operations by default. If you select to back up specific OneDrive files and folders, they will be backed up. Files that are not available on the device will have invalid contents in the backup set.

- You cannot select individual Linux LVM volumes as backup source—neither by direct selection nor by using policy rules. You can back up workloads with such volumes only by selecting **Entire machine** in **What to back up**.
- You can back up disks that are connected via the iSCSI protocol to a physical machine. However, limitations apply if you use Agent for VMware or Agent for Hyper-V for backing up the iSCSI-connected disks. For more information, see "Limitations" (p. 36).

What does a disk or volume backup store?

A disk or volume backup stores a disk or a volume **file system** as a whole and includes all of the information necessary for the operating system to boot. It is possible to recover disks or volumes as a whole from such backups as well as individual folders or files.

With the **sector-by-sector (raw mode) backup option** enabled, a disk backup stores all the disk sectors. The sector-by-sector backup can be used for backing up disks with unrecognized or unsupported file systems and other proprietary data formats.

Windows

A volume backup stores all files and folders of the selected volume independent of their attributes (including hidden and system files), the boot record, the file allocation table (FAT) if it exists, the root and the zero track of the hard disk with the master boot record (MBR).

A disk backup stores all volumes of the selected disk (including hidden volumes such as the vendor's maintenance partitions) and the zero track with the master boot record.

The following items are *not* included in a disk or volume backup (as well as in a file-level backup):

- The swap file (pagefile.sys) and the file that keeps the RAM content when the machine goes into hibernation (hiberfil.sys). After recovery, the files will be re-created in the appropriate place with the zero size.
- If the backup is performed under the operating system (as opposed to bootable media or backing up virtual machines at a hypervisor level):
 - Windows shadow storage. The path to it is determined in the registry value **VSS Default Provider** which can be found in the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup**. This means that in operating systems starting with Windows Vista, Windows Restore Points are not backed up.
 - If the **Volume Shadow Copy Service (VSS) backup option** is enabled, files and folders that are specified in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot** registry key.

Linux

A volume backup stores all files and directories of the selected volume independent of their attributes, a boot record, and the file system super block.

A disk backup stores all disk volumes as well as the zero track with the master boot record.

Mac

A disk or volume backup stores all files and directories of the selected disk or volume, plus a description of the volume layout.

The following items are excluded:

- System metadata, such as the file system journal and Spotlight index
- The Trash
- Time machine backups

Physically, disks and volumes on a Mac are backed up at a file level. Bare metal recovery from disk and volume backups is possible, but the sector-by-sector backup mode is not available.

Policy rules for disks and volumes

When you select disks or volumes to back up, you can use the following policy rules, according to the operating system of the protected workload.

Windows

- [All Volumes] selects all volumes on the machine.
- Drive letter (for example, C:\) selects the volume with the specified drive letter.
- [Fixed Volumes (physical machines)] selects all volumes of a physical machine, other than removable media. Fixed volumes include volumes on SCSI, ATAPI, ATA, SSA, SAS, and SATA devices, and on RAID arrays.
- [BOOT+SYSTEM] selects the system and boot volumes. This is the minimal combination from which you can recover an operating system.
- [Disk 1] selects the first disk of the machine, including all volumes on that disk. To select another disk, type the corresponding number.

Linux

- [All Volumes] selects all mounted volumes on the machine.
- /dev/hda1 selects the first volume on the first IDE hard disk.
- /dev/sda1 selects the first volume on the first SCSI hard disk.
- /dev/md1 selects the first software RAID hard disk.
- To select other basic volumes, specify /dev/xdyN, where:
 - "x" corresponds to the disk type
 - "y" corresponds to the disk number (a for the first disk, b for the second disk, and so on)
 - "N" is the volume number.
- To select a logical volume, specify its path as it appears after running the `ls /dev/mapper` command under the root account.

For example:

```
[root@localhost ~]# ls /dev/mapper/  
control vg_1-lv1 vg_1-lv2
```

This output shows two logical volumes, lv1 and lv2, that belong to the volume group vg_1. To back up these volumes, specify:

```
/dev/mapper/vg_1-lv1  
/dev/mapper/vg_1-lv2
```

macOS

- [All Volumes] selects all mounted volumes on the machine.
- [Disk 1] Selects the first disk of the machine, including all volumes on that disk. To select another disk, specify the corresponding number.

Selecting files or folders

Use file-level backup to protect only specific data, for example, the files in your current project. File-level backups are smaller than disk-level backups and save storage space.

Important

You cannot recover an operating system from a file-level backup.

You can select the files and folders to back up for each individual workload in the protection plan (direct selection) or you can configure policy rules for multiple workloads. Additionally, you can exclude specific files from a backup, or include only specific files in it, by configuring the filters. For more information, see "File filters (Inclusions/Exclusions)" (p. 435).

To select files or folders

Direct selection

1. In **What to back up**, select **Files/folders**.
2. In **Items to back up**, click **Specify**.
3. In **Select items for backup**, select **Directly**.
4. Specify the files or folders to back up for each workload in the protection plan.
 - a. Click **Select files and folders**.
 - b. Click **Local folder** or **Network folder**.

Network folders must be accessible from the selected machine.

When you select **Network folder** as a source, you can back up data from network-attached storages (NAS), such as NetApp devices. NAS devices from all vendors are supported.
 - c. In the folder tree, navigate to the required files or folders.

Alternatively, specify the path to them, and then click the arrow button.
 - d. [For shared folders] When prompted, specify the access credentials to the shared folder.

Backing up folders with anonymous access is not supported.

- e. Select the required files and folders.
- f. Click **Done**.

By policy rules

1. In **What to back up**, select **Files/folders**.
2. In **Items to back up**, click **Specify**.
3. In **Select items for backup**, select **Using policy rules**.
4. Select any of the predefined rules, type your own rules, or combine both.
For more information about the available policy rules, see "Policy rules for files and folders" (p. 384).
The policy rules will be applied to all workloads that are included in the protection plan.
If none of the specified rules can be applied to a workload, the backup of that workload fails.
5. Click **Done**.

Limitations

- You can select files and folders when you back up physical machines or virtual machines on which an agent is installed (agent-based backup). File-level backup is not available for virtual machines that you back up in the agentless mode. For more information about the differences between these types of backup, see "Agent-based and agentless backup" (p. 67).
- The OneDrive root folder is excluded from backup operations by default. If you select to back up specific OneDrive files and folders, they will be backed up. Files that are not available on the device will have invalid contents in the backup set.
- You can back up files and folders that are located on disks connected via the iSCSI protocol to a physical machine. Some [limitations](#) apply if you use Agent for VMware or Agent for Hyper-V for backing up the data on the iSCSI-connected disks.

Policy rules for files and folders

When you select files or folders to back up, you can use the following policy rules, according to the operating system of the protected workload.

Windows

- Full path to a file or folder. For example, D:\Work\Text.doc or C:\Windows.
- Predefined rules:
 - [All Files] selects all files on all volumes of the machine.
 - [All Profiles Folder] selects the folder in which all user profiles are located. For example, C:\Users or C:\Documents and Settings.
- Environment variables:
 - %ALLUSERSPROFILE% selects the folder in which the common data of all user profiles is located. For example, C:\ProgramData or C:\Documents and Settings\All Users.
 - %PROGRAMFILES% selects the Program Files folder. For example, C:\Program Files.
 - %WINDIR% selects the Windows folder. For example, C:\Windows.

You can use other environment variables or a combination of environment variables and text. For example, to select the Java folder in the Program Files folder, specify: %PROGRAMFILES%\Java.

Linux

- Full path to a file or directory.
For example, to back up the file.txt file on volume /dev/hda3 that is mounted on /home/usr/docs, specify /dev/hda3/file.txt or /home/usr/docs/file.txt.
- Predefined rules:
 - [All Profiles Folder] selects /home. By default, all user profiles are stored in this folder.
 - /home selects the home directory of the common users.
 - /root selects the root user's home directory.
 - /usr selects the directory for all user-related programs.
 - /etc selects the directory for system configuration files.

macOS

- Full path to a file or directory.
For example:
 - To back up file.txt on a user's desktop, specify /Users/<user name>/Desktop/file.txt.
 - To back up the Desktop, the Documents, and the Downloads folders of a user, specify /Users/<user name>/Desktop, /Users/<user name>/Documents, and /Users/<user name>/Downloads.
 - To back up the home folders of all users who have an account on this machine, specify /Users.
 - To back up the folder in which the applications are installed, specify /Applications.
- Predefined rules
 - [All Profiles Folder] selects /Users. By default, all user profiles are stored in this folder.

Selecting system state

Note

System state backup is available for machines running Windows 7 or later on which Agent for Windows is installed. System state backup is not available for virtual machines that are backed up at the hypervisor level (agentless backup).

To back up system state, in **What to back up**, select **System state**.

A system state backup is comprised of the following files:

- Task scheduler configuration
- VSS Metadata Store
- Performance counter configuration information
- MSSearch Service
- Background Intelligent Transfer Service (BITS)
- The registry

- Windows Management Instrumentation (WMI)
- Component Services Class registration database

Selecting ESXi configuration

A backup of an ESXi host configuration enables you to recover an ESXi host to bare metal. The recovery is performed under bootable media.

The virtual machines running on the host are not included in the backup. They can be backed up and recovered separately.

A backup of an ESXi host configuration includes:

- The bootloader and boot bank partitions of the host.
- The host state (configuration of virtual networking and storage, SSL keys, server network settings, and local user information).
- Extensions and patches installed or staged on the host.
- Log files.

Prerequisites

- SSH must be enabled in the **Security Profile** of the ESXi host configuration.
- You must know the password for the 'root' account on the ESXi host.

Limitations

- ESXi configuration backup is not supported for hosts running VMware ESXi 7.0 and later.
- An ESXi configuration cannot be backed up to the cloud storage.

To select an ESXi configuration

1. Click **Devices > All devices**, and then select the ESXi hosts that you want to back up.
2. Click **Protect**.
3. In **What to back up**, select **ESXi configuration**.
4. In **ESXi 'root' password**, specify a password for the 'root' account on each of the selected hosts or apply the same password to all of the hosts.

Continuous data protection (CDP)

Continuous data protection (CDP) is part of the Advanced Backup pack. It backs up critical data immediately after this data is changed, ensuring that no changes will be lost if your system fails between two scheduled backups. You can configure Continuous data protection for the following data:

- Files or folders in specific locations
- Files modified by specific applications

Continuous data protection is supported only for the NTFS file system and the following operating systems:

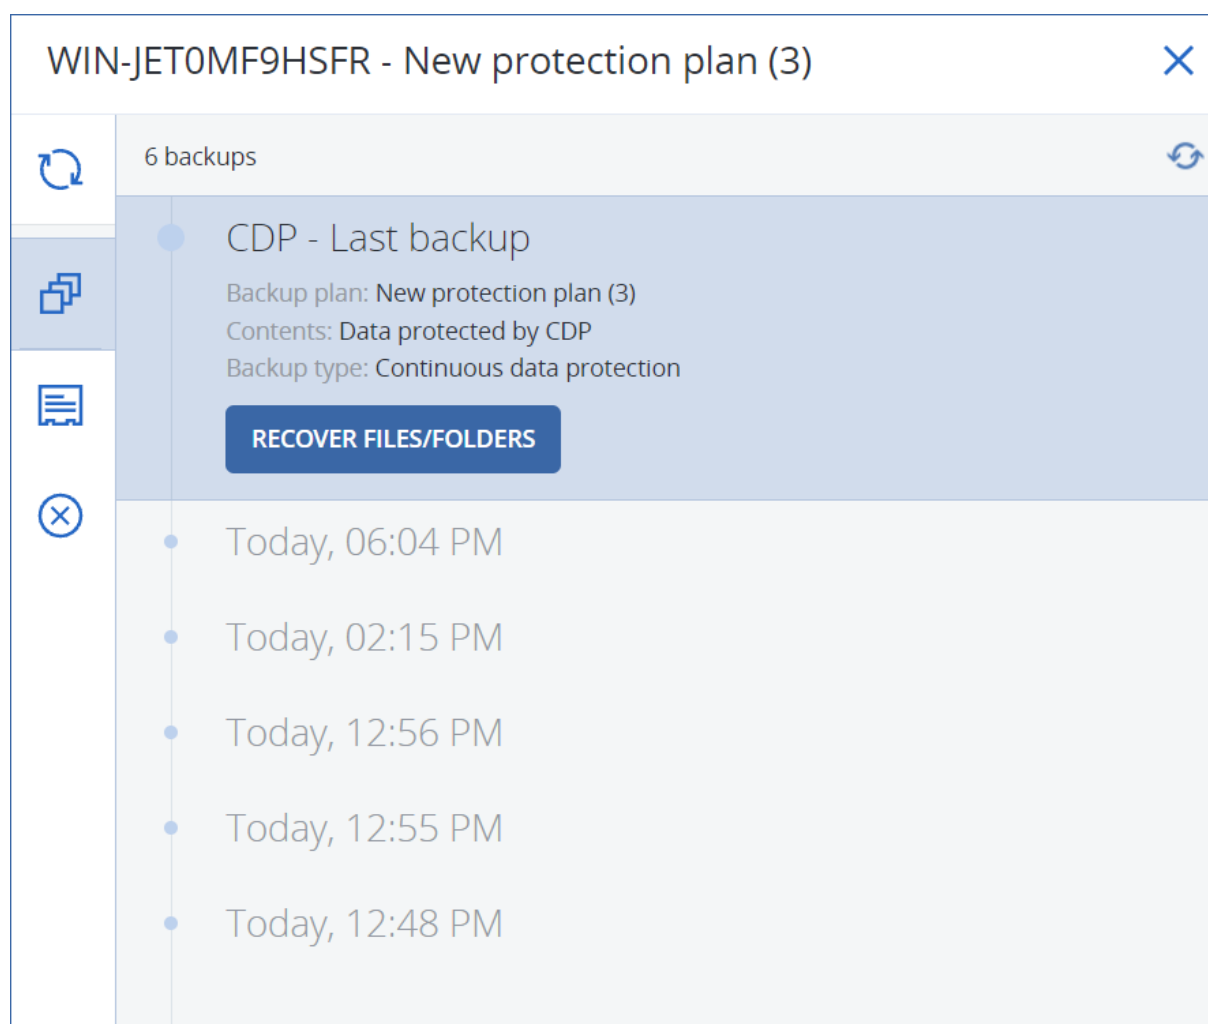
- Desktop: Windows 7 and later
- Server: Windows Server 2008 R2 and later

Only local folders are supported. Network folders cannot be selected for Continuous data protection.

Continuous data protection is not compatible with the **Application backup** option.

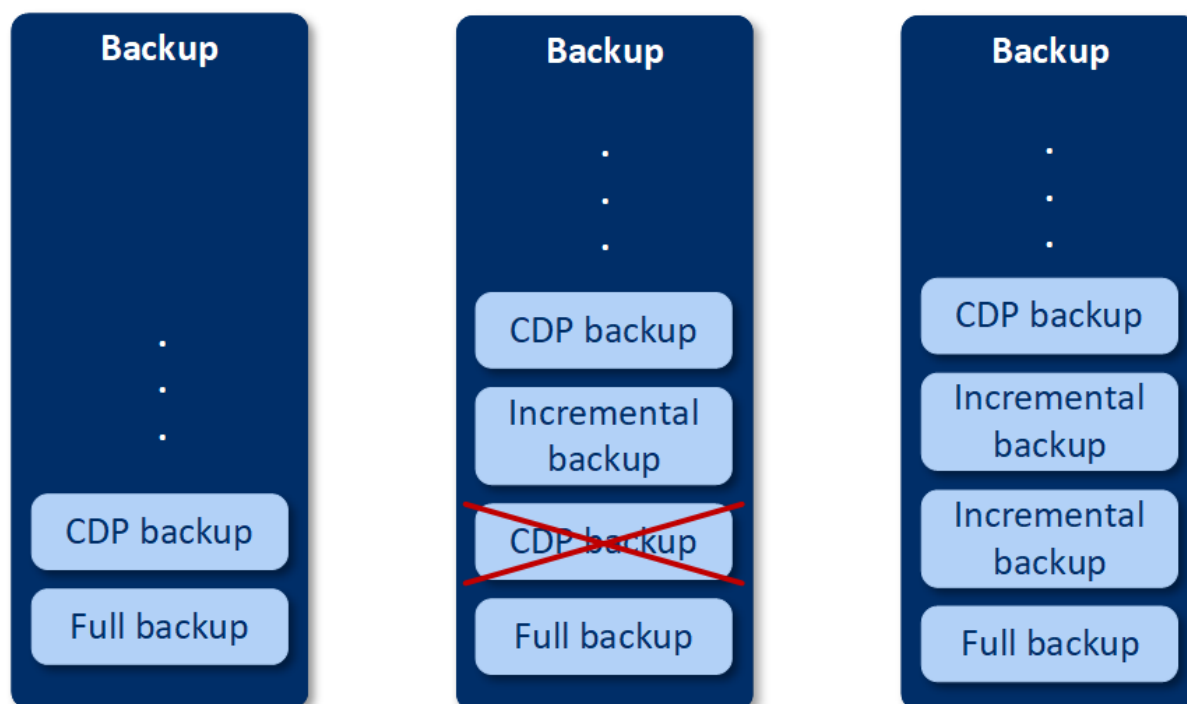
How it works

Changes in the files and folders that are tracked by Continuous data protection are immediately saved to a special CDP backup. There is only one CDP backup in a backup set, and it is always the most recent one.



When a scheduled regular backup starts, Continuous data protection is put on hold because the latest data is to be included in the scheduled backup. When the scheduled backup finishes, Continuous data protection resumes, the old CDP backup is deleted, and a new CDP backup is

created. Thus, the CDP backup always stays the most recent backup in the backup set and stores only the latest state of the tracked files or folders.



If your machine crashes during a regular backup, Continuous data protection resumes automatically after the machine restarts and creates a CDP backup on top of the last successful scheduled backup.

Continuous data protection requires that at least one regular backup is created before the CDP backup. That is why, when you run a protection plan with Continuous data protection for the first time, a full backup is created, and a CDP backup is immediately added on top of it. If you enable the **Continuous data protection** option for an existing protection plan, the CDP backup is added to the existing backup set.

Note

Continuous Data protection is enabled by default for protection plans that you create from the **Devices** tab, if the Advanced Backup functionality is enabled for you and you are not using other Advanced Backup features for the selected machines. If you already have a plan with Continuous data protection for a selected machine, Continuous data protection will not be enabled by default for that machine in newly created plans.

Continuous data protection is not enabled by default for plans created for device groups.

Supported data sources

You can configure Continuous data protection with the following data sources:

- Entire machine
- Disks/volumes

- Files/folders

After selecting the data source in **What to backup** section in the protection plan, in the **Items to protect continuously** section, select the files, folders, or applications for Continuous data protection. For more information on how to configure Continuous data protection, refer to "Configuring a CDP backup" (p. 389).

Supported destinations

You can configure Continuous data protection with the following destinations:

- Local folder
- Network folder
- Cloud storage
- Acronis Cyber Infrastructure
- Location defined by a script

Note

You can define by a script only the locations listed above.

Configuring a CDP backup

You can configure Continuous data protection in the **Backup** module of a protection plan. For more information on how to create a protection plan, refer to "Creating a protection plan" (p. 209).

To configure the Continuous data protection settings

1. In the **Backup** module of a protection plan, enable the **Continuous data protection (CDP)** switch.
This switch is available only for the following data sources:
 - Entire machine
 - Disk/volumes
 - Files/folders
2. In **Items to protect continuously**, configure Continuous data protection for **Applications** or **Files/folders**, or both.
 - Click **Applications** to configure CDP backup for files that are modified by specific applications. You can select applications from predefined categories or add other applications by specifying the path to their executable file, for example:
 - C:\Program Files\Microsoft Office\Office16\WINWORD.EXE
 - *:\Program Files (x86)\Microsoft Office\Office16\WINWORD.EXE
 - Click **Files/folders** to configure CDP backup for files in specific locations. You can define these locations by using selection rules or by selecting the files and folders directly.

- [For all machines] To create a selection rule, use the text box.
You can use the full paths to files or paths with wildcard characters (* and ?). The asterisk matches zero or more characters. The question mark matches a single character.

Important

To create a CDP backup for a folder, you must specify its content by using the asterisk wildcard character:

Correct path: D:\Data*

Incorrect path: D:\Data\

- [For online machines] To select files and folders directly:
 - In **Machine to browse from**, select the machine on which the files or folders reside.
 - Click **Select files and folders** to browse the selected machine.
Your direct selection creates a selection rule. If you apply the protection plan to multiple machines and a selection rule is not valid for a machine, it will be skipped on this machine.

3. In the protection plan pane, click **Create**.

As a result, the data that you specified will be backed up continuously between the scheduled backups.

Selecting a destination

Click **Where to back up**, and then select one of the following:

- **Cloud storage**

Backups will be stored in the cloud data center.

- **Local folders**

If a single machine is selected, browse to a folder on the selected machine or type the folder path.

If multiple machines are selected, type the folder path. Backups will be stored in this folder on each of the selected physical machines or on the machine where the agent for virtual machines is installed. If the folder does not exist, it will be created.

- **Network folder**

This is a folder shared via SMB/CIFS/DFS.

Browse to the required shared folder or enter the path in the following format:

- For SMB/CIFS shares: \\<host name>\<path>\ or smb://<host name>/<path>/
- For DFS shares: \\<full DNS domain name>\<DFS root>\<path>

For example, \\example.company.com\shared\files

Then, click the arrow button. If prompted, specify the user name and password for the shared folder. You can change these credentials at any time by clicking the key icon next to the folder name.

Backing up to a folder with anonymous access is not supported.

- **Public cloud**

This option is available as part of the Advanced Backup pack.

It enables you to configure a direct backup to a public cloud compatible storage, without the need to deploy additional components (such as Microsoft Azure or other virtual machines as gateways). Select and connect to the relevant public cloud, as required.

For more information, see "Backing up workloads to public clouds" (p. 510).

- **NFS folder** (available for machines running Linux or macOS)

Verify that the nfs-utils package is installed on the Linux server where the Agent for Linux is installed.

Browse to the required NFS folder or enter the path in the following format:

```
nfs://<host name>/<exported folder>:/<subfolder>
```

Then, click the arrow button.

Note

It is not possible to back up to an NFS folder protected with a password.

- **Secure Zone** (available if it is present on each of the selected machines)

Secure Zone is a secure partition on a disk of the backed-up machine. This partition has to be created manually prior to configuring a backup. For information about how to create Secure Zone, its advantages and limitations, refer to "About Secure Zone" (p. 392).

Advanced storage option

Note

This functionality is available only in the Advanced edition of the Cyber Protection service.

Defined by a script (available for machines running Windows)

You can store each machine's backups in a folder defined by a script. The software supports scripts written in JScript, VBScript, or Python 3.5. When deploying the protection plan, the software runs the script on each machine. The script output for each machine should be a local or network folder path. If a folder does not exist, it will be created (limitation: scripts written in Python cannot create folders on network shares). On the **Backup storage** tab, each folder is shown as a separate backup location.

In **Script type**, select the script type (**JScript**, **VBScript**, or **Python**), and then import, or copy and paste the script. For network folders, specify the access credentials with the read/write permissions.

Examples:

- The following JScript script outputs the backup location for a machine in the format

```
\\bkpsrv\<machine name>:
```

```
WScript.Echo("\\\\bkpsrv\\" + WScript.CreateObject("WScript.Network").ComputerName);
```

As a result, the backups of each machine will be saved in a folder of the same name on the server **bkpsrv**.

- The following JScript script outputs the backup location in a folder on the machine where the script runs:

```
WScript.Echo("C:\\Backup");
```

As a result, the backups of this machine will be saved in the folder C:\Backup on the same machine.

Note

The location path in these scripts is case-sensitive. Therefore, C:\Backup and C:\backup are displayed as different locations in the Cyber Protect console. Also, use upper case for the drive letter.

About Secure Zone

Secure Zone is a secure partition on a disk of the backed-up machine. It can store backups of disks or files of this machine.

Should the disk experience a physical failure, the backups located in the Secure Zone may be lost. That's why Secure Zone should not be the only location where a backup is stored. In enterprise environments, Secure Zone can be thought of as an intermediate location used for backup when an ordinary location is temporarily unavailable or connected through a slow or busy channel.

Why use Secure Zone?

Secure Zone:

- Enables recovery of a disk to the same disk where the disk's backup resides.
- Offers a cost-effective and handy method for protecting data from software malfunction, virus attack, human error.
- Eliminates the need for a separate media or network connection to back up or recover the data. This is especially useful for roaming users.
- Can serve as a primary destination when using replication of backups.

Limitations

- Secure Zone cannot be organized on a Mac.
- Secure Zone is a partition on a basic disk. It cannot be organized on a dynamic disk or created as a logical volume (managed by LVM).
- Secure Zone is formatted with the FAT32 file system. Because FAT32 has a 4-GB file size limit, larger backups are split when saved to Secure Zone. This does not affect the recovery procedure and speed.

How creating Secure Zone transforms the disk

- Secure Zone is always created at the end of the hard disk.
- If there is no or not enough unallocated space at the end of the disk, but there is unallocated space between volumes, the volumes will be moved to add more unallocated space to the end of the disk.
- When all unallocated space is collected but it is still not enough, the software will take free space from the volumes you select, proportionally reducing the volumes' size.
- However, there should be free space on a volume, so that the operating system and applications can operate; for example, create temporary files. The software will not decrease a volume where free space is or becomes less than 25 percent of the total volume size. Only when all volumes on the disk have 25 percent or less free space, will the software continue decreasing the volumes proportionally.

As is apparent from the above, specifying the maximum possible Secure Zone size is not advisable. You will end up with no free space on any volume, which might cause the operating system or applications to work unstably and even fail to start.

Important

Moving or resizing the volume from which the system is booted requires a reboot.

How to create Secure Zone

1. Select the machine that you want to create Secure Zone on.
2. Click **Details > Create Secure Zone** .
3. Under **Secure Zone disk**, click **Select**, and then select a hard disk (if several) on which to create the zone.

The software calculates the maximum possible size of Secure Zone.

4. Enter the Secure Zone size or drag the slider to select any size between the minimum and the maximum ones.

The minimum size is approximately 50 MB, depending on the geometry of the hard disk. The maximum size is equal to the disk's unallocated space plus the total free space on all of the disk's volumes.

5. If all unallocated space is not enough for the size you specified, the software will take free space from the existing volumes. By default, all volumes are selected. If you want to exclude some volumes, click **Select volumes**. Otherwise, skip this step.

✕ Create Secure Zone

Secure Zone disk

Disk 1, 60.0 GB

Maximum possible size of Secure Zone: 35.9 GB

Secure Zone size:

20 GB

There is not enough unallocated space. Free space will be taken from all volumes where it is present.

Select volumes

Password protection

Off

- [Optional] Enable the **Password protection** switch and specify a password.
The password will be required to access the backups located in Secure Zone. Backing up to Secure Zone does not require a password, unless the backup is performed under bootable media.
- Click **Create**.
The software displays the expected partition layout. Click **OK**.
- Wait while the software creates Secure Zone.

You can now choose Secure Zone in **Where to back up** when creating a protection plan.

How to delete Secure Zone

- Select a machine with Secure Zone.
- Click **Details**.
- Click the gear icon next to **Secure Zone**, and then click **Delete**.
- [Optional] Specify the volumes to which the space freed from the zone will be added. By default, all volumes are selected.
The space will be distributed equally among the selected volumes. If you do not select any volumes, the freed space will become unallocated.
Resizing the volume from which the system is booted requires a reboot.
- Click **Delete**.

As a result, Secure Zone will be deleted along with all backups stored in it.

Backup schedule

You can configure a backup to run automatically at a specific time, at specific intervals, or on a specific event.

Scheduled backups for non-cloud-to-cloud resources run according to the time zone settings of the workload on which the protection agent is installed. For example, if you apply the same protection plan to workloads with different time zones settings, the backups will start according to the local time zone of each workload.

Scheduling a backup includes the following actions:

- Selecting a backup scheme
- Configuring the time or selecting the event that triggers the backup
- Configuring optional setting and start conditions

Backup schemes

A backup scheme is a part of the protection plan schedule that defines which type of backup (full, differential, or incremental) is created and when. You can select one of the predefined backup schemes or create a custom scheme.

The available backup schemes and types depend on the backup location and source. For example, a differential backup is not available when you back up SQL data, Exchange data, or system state. The **Always incremental (single-file)** scheme is not supported for tape devices.

Backup scheme	Description	Configurable elements
Always incremental (single-file)	<p>The first backup is full and might be time-consuming. Subsequent backups are incremental and significantly faster.</p> <p>The backups use the single-file backup format^{1*}.</p> <p>By default, backups are performed on a daily basis, Monday to Friday.</p> <p>We recommend that you use this scheme when you store your backups in the cloud storage, because incremental backups are fast and involve less network traffic.</p>	<ul style="list-style-type: none">• Schedule type: monthly, weekly, daily, hourly• Backup trigger: time or event• Start time• Start conditions• Additional options

¹A backup format, in which the initial full and subsequent incremental backups are saved to a single .tibx file. This format leverages the speed of the incremental backup method, while avoiding its main disadvantage—difficult deletion of outdated backups. The software marks the blocks used by outdated backups as "free" and writes new backups to these blocks. This results in extremely fast cleanup, with minimal resource consumption. The single-file backup format is not available when backing up to locations that do not support random-access reads and writes.

Backup scheme	Description	Configurable elements
Always full	<p>All backups in the backup set are full.</p> <p>By default, backups are performed on a daily basis, Monday to Friday.</p>	<ul style="list-style-type: none"> • Schedule type: monthly, weekly, daily, hourly • Backup trigger: time or event • Start time • Start conditions • Additional options
Weekly full, Daily incremental	<p>A full backup is created once a week and other backups are incremental.</p> <p>The first backup is full and the other backups during the week are incremental, then the cycle repeats.</p> <p>To select the day on which the weekly full backup is created, in the protection plan, click the gear icon, and then go to Backup options > Weekly backup.</p> <p>By default, backups are performed on a daily basis, Monday to Friday.</p>	<ul style="list-style-type: none"> • Backup trigger: time or event • Start time • Start conditions • Additional options
Monthly full, Weekly differential, Daily incremental (GFS)	<p>By default, incremental backups are performed on a daily basis, Monday to Friday. Differential backups are performed every Saturday. Full backups are performed on the first day of each month.</p> <hr/> <p>Note This is a predefined custom scheme. In the protection plan, it is shown as Custom.</p> <hr/>	<ul style="list-style-type: none"> • Change the existing schedule per backup type: <ul style="list-style-type: none"> ◦ Schedule type: monthly, weekly, daily, hourly ◦ Backup trigger: time or event ◦ Start time ◦ Start conditions ◦ Additional options • Add new schedules per backup type
Custom	<p>You must select the backup types (full, differential, and incremental), and configure a separate schedule for each of them*.</p>	<ul style="list-style-type: none"> • Change the existing schedule per backup type: <ul style="list-style-type: none"> ◦ Schedule type: monthly, weekly, daily, hourly ◦ Backup trigger: time or event ◦ Start time ◦ Start conditions

Backup scheme	Description	Configurable elements
		<ul style="list-style-type: none"> ◦ Additional options • Add new schedules per backup type

* After you create a protection plan, you cannot switch between **Always incremental (single-file)** and the other backup schemes, and vice versa. **Always incremental (single-file)** is a single-file format scheme, and the other schemes are multi-file format. If you want to switch between formats, create a new protection plan.

Backup types

The following backup types are available:

- Full—a full backup contains all source data. This backup is self-sufficient. To recover data, you do not need access to any other backups.

Note

The first backup created by any protection plan is a full backup.

- Incremental—an incremental backup stores changes to the data since the latest backup, regardless of whether the latest backup is full, differential, or incremental. To recover data, you need the whole chain of backups on which the incremental backup depends, back to the initial full backup.
- Differential—a differential backup stores changes to the data since the latest full backup. To recover data, you need both the differential backup and the corresponding full backup on which the differential backup depends.

Running a backup on a schedule

To run a backup automatically at a specific time or on a specific event, enable a schedule in the protection plan.

To enable a schedule

1. In the protection plan, expand the **Backup** module.
2. Click **Schedule**.
3. Enable the schedule switch.
4. Select the backup scheme.
5. Configure the schedule as required, and then click **Done**.
For more information about the available scheduling options, see "Schedule by time" (p. 398) and "Schedule by events" (p. 400).
6. [Optional] Configure start conditions or additional scheduling options.
7. Save the protection plan.

As a result, a backup operation starts every time when the schedule conditions are met.

To disable a schedule

1. In the protection plan, expand the **Backup** module.
2. Click **Schedule**.
3. Disable the schedule switch.
4. Save the protection plan.

As a result, the backup runs only if you start it manually.

Note

If the schedule is disabled, the retention rules are not applied automatically. To apply them, run the backup manually.

Schedule by time

The following table summarizes the scheduling options that are based on time. The availability of these options depends on the backup scheme. For more information, see "Backup schemes" (p. 395).

Option	Description	Examples
Monthly	Select the months, days of the month or days of the week, and then select the backup start time.	Run a backup on January 1, and February 3, at 12:00 AM. Run a backup on the first day of each month, at 10:00 AM. Run a backup on March 1, March 5, April 1, and April 5, at 09:00 AM. Run a backup on the second and third Friday of each month, at 11:00 AM. Run a backup on the last Wednesday of the month, at 10:30 PM.
Weekly	Select the days of the week, and then select the backup start time.	Run a backup Monday to Friday, at 10:00 AM. Run a backup on Monday, at 11:00 PM. Run a backup on Tuesday and Saturday, at 08:00 AM.
Daily	Select the days (everyday or weekdays only), and then select the backup start time.	Run a backup every day, at 11:45 AM. Run a backup Monday to Friday, at 09:30 PM.
Hourly	Select the days of the week, and then select a time interval between two consecutive backups and the time range	Run a backup every hour between 08:00 AM and 06:00 PM, Monday to Friday. Run a backup every 3 hours between

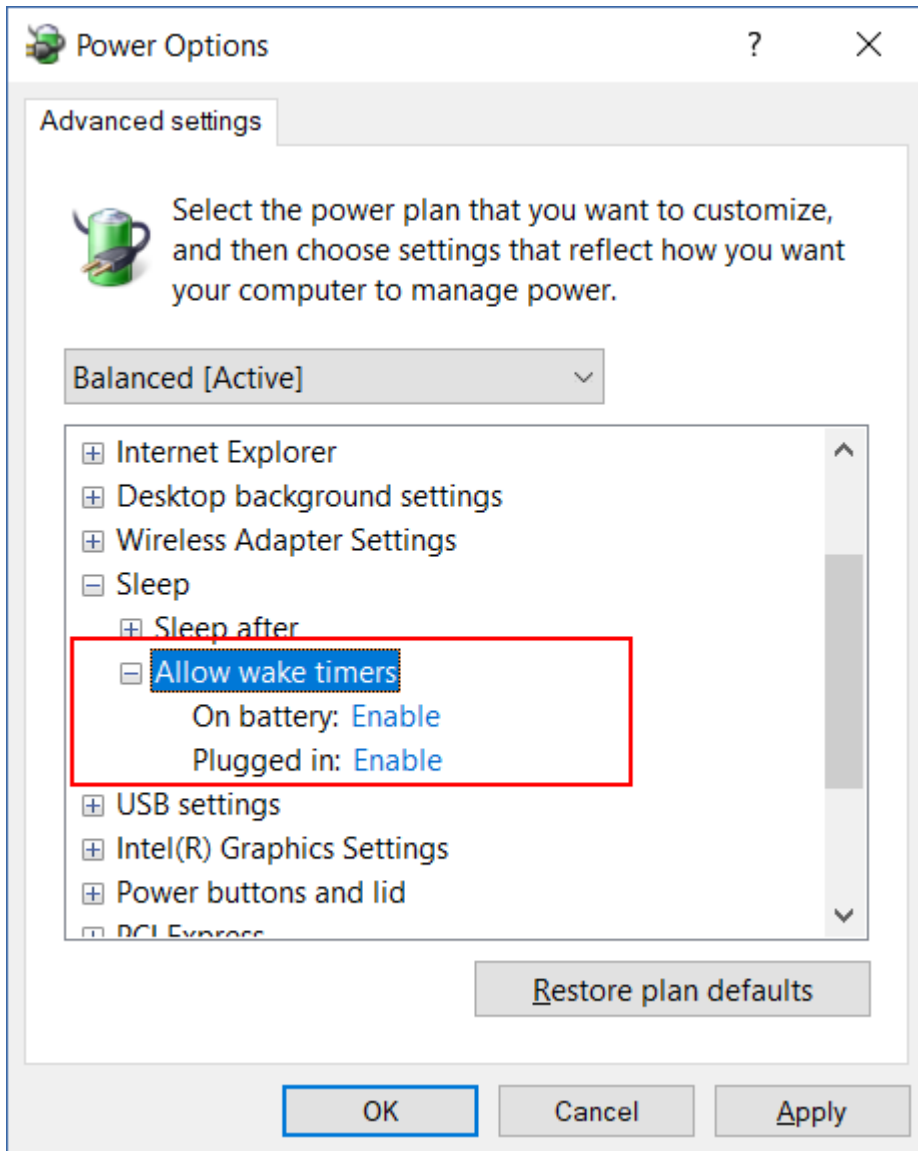
Option	Description	Examples
	<p>within which the backups run.</p> <p>When you configure the interval in minutes, you can select a suggested interval between 10 and 60 minutes, or specify a custom one, for example, 45 or 75 minutes.</p>	<p>01:00 AM and 06:00 PM, on Saturday and Sunday.</p>

Additional options

When you schedule a backup by time, the following additional scheduling options are available.

To access them, in the **Schedule** pane, click **Show more**.

- **If the machine is turned off, run missed tasks at the machine startup**
 Default setting: Disabled.
- **Prevent the sleep or hibernate mode during backup**
 This option is applicable only to machines running Windows.
 Default setting: Enabled.
- **Wake up from the sleep or hibernate mode to start a scheduled backup**
 This option is applicable only to machines running Windows, in the power plans for which the **Allow wake timers** option is enabled.



This option does not use the Wake-on-LAN functionality and is not applicable to powered-off machines.

Default setting: Disabled.

Schedule by events

To configure a backup that runs upon a specific event, select one of the following options.

Option	Description	Examples
Upon time since last backup	A backup starts after a specified period following the last successful backup.	Run a backup one day after the last successful backup. Run a backup four hours after the last successful backup.

Option	Description	Examples
	<p>Note This option depends on how the previous backup completed. If a backup fails, the next backup will not start automatically. In this case, you must run the backup manually and ensure that it completes successfully, in order to reset the schedule.</p>	
<p>When a user logs on to the system</p>	<p>A backup starts when a user logs in to the machine.</p> <p>You can configure this option for any login or for a login of a specific user.</p> <hr/> <p>Note Logging in with a temporary user profile will not start a backup.</p>	<p>Run a backup when user John Doe logs in.</p>
<p>When a user logs off the system</p>	<p>A backup starts when a user logs off the machine.</p> <p>You can configure this option for any logoff or for the logoff of a specific user.</p> <hr/> <p>Note Logging off from a temporary user profile will not start a backup.</p> <p>Shutting down a machine will not start a backup.</p>	<p>Run a backup when every user logs off.</p>
<p>On the system startup</p>	<p>A backup runs when the protected machine starts up.</p>	<p>Run a backup when a user starts the machine.</p>
<p>On the system shutdown</p>	<p>A backup runs when the protected machine shuts down.</p>	<p>Run a backup when a user shuts down the machine.</p>
<p>On Windows Event Log event</p>	<p>A backup runs upon a Windows event that you specify.</p>	<p>Run a backup when event 7 of type error and source disk is recorded in the Windows System log.</p>

The availability of these options depends on the backup source and the operating system of the protected workloads. The table below summarizes the available options for Windows, Linux, and macOS.

Event	Backup source (What to back up)					
	Entire machine, Disks/volumes, or Files/folders (physical machines)	Entire machines or Disk/volumes (virtual machines)	ESXi configuration	Microsoft 365 mailboxes	Exchange databases and mailboxes	SQL databases
Upon time since last backup	Windows, Linux, macOS	Windows, Linux	Windows, Linux	Windows	Windows	Windows
When a user logs on to the system	Windows	N/A	N/A	N/A	N/A	N/A
When a user logs off the system	Windows	N/A	N/A	N/A	N/A	N/A
On the system startup	Windows, Linux, macOS	N/A	N/A	N/A	N/A	N/A
On the system shutdown	Windows	N/A	N/A	N/A	N/A	N/A
On Windows Event Log event	Windows	N/A	N/A	Windows	Windows	Windows

On Windows Event Log event

You can automatically run a backup when a specific event is recorded in a Windows Event log, such as the Application log, Security log, or System log.

Note

You can browse the events and view their properties in **Computer Management > Event Viewer** in Windows. To open the Security log, you need administrator rights.

Event parameters

The following table summarizes the parameters that you must specify when configuring the **On Windows Event Log event** option.

Parameter	Description
Log name	The name of the log. Select the name of a standard log (Application, Security, System) or specify another log name. For example, Microsoft Office Sessions.
Event source	The event source indicates the program or the system component that caused the event. For example, disk. Any event source that contains the specified text string will trigger the scheduled backup. This option is not case-sensitive. For example, if you specify <i>service</i> , both Service Control Manager and Time-Service event sources will trigger a backup.
Event type	Type of the event: Error, Warning, Information, Audit success, or Audit failure.
Event ID	The event ID identifies a particular kind of event within an event source. For example, an Error event with event source disk and event ID 7 occurs when Windows discovers a bad block on a disk, while an Error event with event source disk and event ID 15 occurs when a disk is not ready for access.

Example: Emergency backup in case of bad blocks on the hard disk

One or more bad blocks on a hard disk drive might indicate an imminent fail. That is why you might want to create a backup when a bad block is detected.

When Windows detects a bad block on the disk, an error event with the event source disk and event number 7 is recorded to the system log. In the protection plan, configure the following schedule:

- Schedule: On Windows Event log event
- Log name: System
- Event source: disk
- Event type: Error
- Event ID: 7

Important

To ensure that the backup completes despite the bad blocks, in **Backup options**, go to **Error handling**, and then select the **Ignore bad sectors** check box.

Start conditions

To make a backup run only if specific conditions are met, configure one or more start conditions. If you configure multiple conditions, all of them must be met simultaneously for the backup to start. You can specify a period after which the backups will run, regardless of whether the conditions are met. For more information about this backup option, see "Task start conditions" (p. 465).

Start conditions are not applicable when you start a backup manually.

The table below lists the start conditions available for various data under Windows, Linux, and macOS.

Start condition	Backup source (What to back up)					
	Entire machine, Disks/volumes, or Files/folders (physical machines)	Entire machines or Disk/volumes (virtual machines)	ESXi configuration	Microsoft 365 mailboxes	Exchange databases and mailboxes	SQL databases
User is idle	Windows	N/A	N/A	N/A	N/A	N/A
The backup location's host is available	Windows, Linux, macOS	Windows, Linux	Windows, Linux	Windows	Windows	Windows
Users logged off	Windows	N/A	N/A	N/A	N/A	N/A
Fits the time interval	Windows, Linux, macOS	Windows, Linux	N/A	N/A	N/A	N/A
Save battery power	Windows	N/A	N/A	N/A	N/A	N/A
Do not start when on metered connection	Windows	N/A	N/A	N/A	N/A	N/A
Do not start when connected to the following Wi-Fi	Windows	N/A	N/A	N/A	N/A	N/A

Start condition	Backup source (What to back up)					
	Entire machine, Disks/volumes, or Files/folders (physical machines)	Entire machines or Disk/volumes (virtual machines)	ESXi configuration	Microsoft 365 mailboxes	Exchange databases and mailboxes	SQL databases
networks						
Check device IP address	Windows	N/A	N/A	N/A	N/A	N/A

User is idle

"User is idle" means that a screen saver is running on the machine or the machine is locked.

Example

Run a backup every day at 09:00 PM, preferably when the user is idle. If the user is still active by 11:00 PM, run the backup anyway.

- Schedule: **Daily, Run every day**. Start at: **09:00 PM**.
- Condition: **User is idle**.
- Backup start conditions: **Wait until the conditions are met, Start the task anyway after 2 hours**.

As a result:

- If the user is idle before 09:00 PM, the backup starts at 09:00 PM.
- If the user becomes idle between 09:00 PM and 11:00 PM, the backup starts immediately.
- If the user is still active at 11:00 PM, the backup starts at 11:00 PM.

The backup location's host is available

"The backup location's host is available" means that the machine that hosts the backup location is available over the network.

This condition is applicable to network folders, the cloud storage, and locations managed by a storage node.

This condition does not cover the availability of the location itself—only the host availability. For example, if the host is available, but the network folder on this host is not shared or the credentials for the folder are no longer valid, the condition is still considered met.

Example

You run backups to a network folder every workday at 09:00 PM. If the machine that hosts the folder is not available at that moment (for example, due to maintenance), you want to skip the backup and wait for the scheduled start on the next workday.

- Schedule: **Daily, Run Monday to Friday**. Start at: **09:00 PM**.
- Condition: **The backup location's host is available**.
- Backup start conditions: **Skip the scheduled backup**.

As a result:

- If the host is available at 09:00 PM, the backup starts immediately.
- If the host is not available at 09:00 PM, the backup starts the next workday (if the host is available at 09:00 PM on this day).
- If the host is never available on workdays at 09:00 PM, the backup never starts.

Users logged off

Use this start condition to postpone a backup until all users log off from a Windows machine.

Example

You run a backup every Friday at 08:00 PM, preferably when all users are logged off. If one of the users is still logged in at 11:00 PM, run the backup anyway.

- Schedule: **Weekly**, on Fridays. Start at: **08:00 PM**.
- Condition: **Users logged off**.
- Backup start conditions: **Wait until the conditions are met, Start the backup anyway after 3 hours**.

As a result:

- If all users are logged off at 08:00 PM, the backup starts at 08:00 PM.
- If the last user logs off between 08:00 PM and 11:00 PM, the backup starts immediately.
- If there are still logged-in users at 11:00 PM, the backup starts at 11:00 PM.

Fits the time interval

Use this start condition to restrict a backup start to a specified interval.

Example

A company backs up user data and servers to different locations on the same network-attached storage.

The workday starts at 08:00 AM and ends at 05:00 PM. User data should be backed up as soon as the users log off, but not earlier than 04:30 PM.

The company's servers are backed up every day at 11:00 PM. User data should preferably be backed up before 11:00 PM, in order to free network bandwidth for the server backups.

Backing up user data takes no more than one hour, so the latest backup start time is 10:00 PM. If a user is still logged in within the specified time interval, or logs off at any other time, the backup of the user data should be skipped.

- Event: **When a user logs off the system**. Specify the user account: **Any user**.
- Condition: **Fits the time interval** from **04:30 PM** to **10:00 PM**.
- Backup start conditions: **Skip the scheduled backup**.

As a result:

- If the user logs off between 04:30 PM and 10:00 PM, the backup starts immediately.
- If the user logs off at any other time, the backup is skipped.

Save battery power

Use this start condition to prevent a backup if a machine (for example, a laptop or a tablet) is not connected to a power source. Depending on the value of the [Backup start conditions](#) option, the skipped backup will or will not start after the machine is connected to a power source.

The following options are available:

- **Do not start when on battery**
A backup will start only if the machine is connected to a power source.
- **Start when on battery if the battery level is higher than**
A backup will start if the machine is connected to a power source or if the battery level is higher than a specified value.

Example

You back up your data every workday at 09:00 PM. If your machine is not connected to a power source, you want to skip the backup to save the battery power and wait until you connect the machine to a power source.

- Schedule: **Daily, Run Monday to Friday**. Start at: **09:00 PM**.
- Condition: **Save battery power, Do not start when on battery**.
- Backup start conditions: **Wait until the conditions are met**.

As a result:

- If the machine is connected to a power source at 09:00 PM, the backup starts immediately.
- If the machine is running on battery power at 09:00 PM, the backup starts when you connect the machine to a power source.

Do not start when on metered connection

Use this start condition to prevent a backup (including a backup to a local disk) if the machine is connected to the Internet through a connection that is set as metered in Windows. For more

information about metered connections in Windows, refer to <https://support.microsoft.com/en-us/help/17452/windows-metered-internet-connections-faq>.

The additional start condition **Do not start when connected to the following Wi-Fi networks** is automatically enabled when you enable the **Do not start when on metered connection** condition. This is an additional measure to prevent backups over mobile hotspots. The following network names are specified by default: android, phone, mobile, and modem.

To remove these names from the list, click the X sign. To add a new name, type it in the empty field.

Example

You back up your data every workday at 09:00 PM. If the machine is connected to the Internet by using a metered connection, you want to skip the backup to save the network traffic and wait for the scheduled start on the next workday.

- Schedule: **Daily, Run Monday to Friday**. Start at: **09:00 PM**.
- Condition: **Do not start when on metered connection**.
- Backup start conditions: **Skip the scheduled backup**.

As a result:

- At 09:00 PM, if the machine is not connected to the Internet through a metered connection, the backup starts immediately.
- At 09:00 PM, if the machine is connected to the Internet through a metered connection, the backup starts on the next workday.
- If the machine is always connected to the Internet through a metered connection on workdays at 09:00 PM, the backup never starts.

Do not start when connected to the following Wi-Fi networks

Use this start condition to prevent a backup (including a backup to a local disk) if the machine is connected to any of the specified wireless networks (for example, if you want to restrict backups through a mobile phone hotspot).

You can specify the Wi-Fi network names, also known as service set identifiers (SSID). The restriction applies to all networks that contain the specified name as a substring in their name, not case-sensitive. For example, if you specify phone as the network name, the backup will not start when the machine is connected to any of the following networks: John's iPhone, phone_wifi, or my_PHONE_wifi.

The start condition **Do not start when connected to the following Wi-Fi** is automatically enabled when you enable the **Do not start when on metered connection** condition. The following network names are specified by default: android, phone, mobile, and modem.

To remove these names from the list, click the X sign. To add a new name, type it in the empty field.

Example

You back up your data every workday at 09:00 PM. If the machine is connected to the Internet through a mobile hotspot, you want to skip the backup and wait for the scheduled start on the next

workday.

- Schedule: **Daily, Run Monday to Friday**. Start at: **09:00 PM**.
- Condition: **Do not start when connected to the following networks, Network name:** <SSID of the hotspot network>.
- Backup start conditions: **Skip the scheduled backup**.

As a result:

- If the machine is not connected to the specified network at 09:00 PM, the backup starts immediately.
- If the machine is connected to the specified network at 09:00 PM, the backup starts the next workday.
- If the machine is always connected to the specified network on workdays at 09:00 PM, the backup never starts.

Check device IP address

Use this start condition to prevent a backup (including a backup to a local disk) if any of the machine IP addresses are within or outside of the specified IP address range. Thus, for example, you can avoid large data transit charges when backing up machines of users who are overseas, or you can prevent backups over a Virtual Private Network (VPN) connection.

The following options are available:

- **Start if outside IP range**
- **Start if within IP range**

With either option, you can specify several ranges. Only IPv4 addresses are supported.

Example

You back up your data every workday at 09:00 PM. If the machine is connected to the corporate network by using a VPN tunnel, you want to skip the backup.

- Schedule: **Daily, Run Monday to Friday**. Start at **09:00 PM**.
- Condition: **Check device IP address, Start if outside IP range, From:** <beginning of the VPN IP address range>, **To:** <end of the VPN IP address range>.
- Backup start conditions: **Wait until the conditions are met**.

As a result:

- If the machine IP address is not in the specified range at 09:00 PM, the backup starts immediately.
- If the machine IP address is in the specified range at 09:00 PM, the backup starts when the machine obtains a non-VPN IP address.
- If the machine IP address is always in the specified range on workdays at 09:00 PM, the backup never starts.

Additional scheduling options

You can configure the backups to run only if specific conditions are met, to run only during a specified period, or to run with a delay compared to the schedule.

To configure start conditions

1. In the protection plan, expand the **Backup** module.
2. Click **Schedule**.
3. On the **Schedule** pane, click **Show more**.
4. Select the check boxes next to the start conditions that you want to include, and then click **Done**.
For more information about the available start conditions and how to configure them, see "Start conditions" (p. 404).
5. Save the protection plan.

To configure a time range

1. In the protection plan, expand the **Backup** module.
2. Click **Schedule**.
3. Select the **Run the plan within a date range** check box.
4. Specify the period according to your needs, and then click **Done**.
5. Save the protection plan.

As a result, the backups will run only during the specified period.

To configure a delay

To avoid excessive network load when you back up multiple workloads to a network location, a small random delay is configured as a backup option. You can disable it or change its setting.

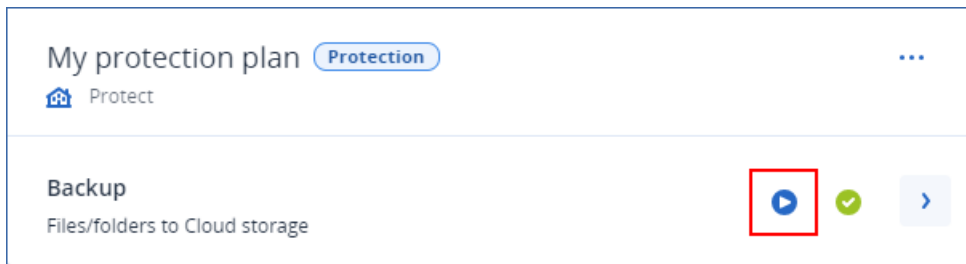
1. In the protection plan, expand the **Backup** module.
2. Click **Backup options**, and then select **Scheduling**.
The delay value for each workload is selected randomly between zero and the maximum value you specify. By default, the maximum value is 30 minutes.
For more information about this backup option, see "Scheduling" (p. 463).
The delay value for each workload is calculated when you apply the protection plan to that workload, and remains the same until you edit the maximum delay value.
3. Specify the period according to your needs, and then click **Done**.
4. Save the protection plan.

Running a backup manually

You can manually run scheduled and unscheduled backups.

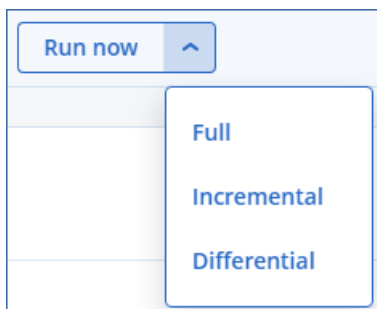
To run a backup manually

1. In the Cyber Protect console, go to **Devices**.
2. Select the workload for which you want to run a backup, and then click **Protect**.
3. Select the protection plan that you want to create the backup.
If no protection plan is applied to the workload, apply an existing plan or create a new one.
For more information about how to create a protection plan, see "Creating a protection plan" (p. 209).
4. [To create the default type of backup] In the protection plan, click the **Run now** icon.



Alternatively, in the protection plan, expand the **Backup** module, and then click the **Run now** button.

5. [To create a specific type of backup] In the protection plan, expand the **Backup** module, click the arrow next to the **Run now** button, and then select the backup type.



Note

Selecting the type is not available for backup schemes that use only one backup method, for example, **Always incremental (single-file)** or **Always full**.

As a result, the backup operation starts. You can check its progress and its result on the **Devices** tab, in the **Status** column.

Retention rules

To delete older backups automatically, configure the backup retention rules in the protection plan.

You can base the retention rules on any of the following backup properties:

- Number
- Age
- Size

The available retention rules and their options depend on the backup scheme. The rules are also relevant to agents, workloads, and cloud to cloud backups. For more information, see "Retention rules according to the backup scheme" (p. 412).

You can disable the automatic cleanup of older backups, by selecting the **Keeping backups infinitely** option while configuring the retention rules. This might result in increased storage usage, and you have to delete the unnecessary old backups manually.

Important tips

- Retention rules are part of the protection plan. If you revoke or delete a plan, the retention rules in that plan will no longer be applied. For more information about how to delete the backups that you no longer need, see "Deleting backups" (p. 504).
- If, according to the backup scheme and backup format, each backup is stored as a separate file, you cannot delete a backup on which other incremental or differential backups depend. This backup will be deleted according to the retention rules applied to the dependent backups. This configuration may result in increased storage usage because the deletion of some backups is postponed. Also, the backup age, number, or size of backups may exceed the values that you specified. For more information about how to change this behavior, see "Backup consolidation" (p. 425).
- By default, the newest backup that a protection plan creates is never deleted. However, if you configure a retention rule to clean up the backups before starting a new backup operation, and set the number of backups to keep to zero, the newest backup will also be deleted.

Warning!

If you apply this retention rule to a backup set with a single backup, and the backup operation fails, you will not be able to recover your data, because the existing backup will be deleted before a new one is created.

Retention rules according to the backup scheme

The available retention rules and their settings depend on the backup scheme that you use in the protection plan. For more information about the backup schemes, see "Backup schemes" (p. 395).

The following table summarizes the available retention rules and their settings.

Backup scheme	Schedule	Available retention rules and settings
Always incremental (single-file)	Monthly	By number of backups
	Weekly	By backup age (separate settings for monthly, weekly, daily, and hourly backups)
	Daily	
	Hourly	Keep backups indefinitely
	Event-triggered backups	
Always full	Monthly	By number of backups

Backup scheme	Schedule	Available retention rules and settings
	Weekly Daily Hourly Event-triggered backups	By backup age (separate settings for monthly, weekly, daily, and hourly backups) By total size of backups Keep backups indefinitely
Weekly full, Daily incremental	Daily Event-triggered backups	By number of backups By backup age (separate settings for weekly and daily backups) By total size of backups Keep backups indefinitely
Monthly full, Weekly differential, daily incremental	Monthly Weekly Daily Hourly Event-triggered backups	By number of backups By backup age (separate settings for full, differential, and incremental backups) By total size of backups Keep backups indefinitely
Custom	Monthly Weekly Daily Hourly Event-triggered backups	By number of backups By backup age (separate settings for full, differential, and incremental backups) By total size of backups Keep backups indefinitely

Why are there monthly backups with an hourly scheme?

Depending on the backup scheme, you can configure the **By backup age** option for one the following backups:

- Monthly, weekly, daily, and hourly backups.

These settings are available with all non-custom backup schemes, and are based on time. All these backups (monthly, weekly, daily, and hourly) are available, even if you configure your backups to run hourly. See the example below.

Backup	Description
Monthly	A monthly backup is the first backup each month.
Weekly	A weekly backup is the first backup on the day of the week that you specify in the Weekly backup option. This day is considered as the beginning of the week in terms of retention rules.

Backup	Description
	If a weekly backup is also the first backup of the month, it is considered a monthly backup. In this case, a weekly backup is created on the selected day the following week.
Daily	A daily backup is the first backup of the day, unless this backup falls within the definition of a monthly or weekly backup. In this case, a daily backup is created the following day.
Hourly	An hourly backup is the first backup of the hour, unless this backup falls within the definition of a monthly, weekly, or daily backup. In this case, an hourly backup is created the next hour.

- Full, differential, and incremental backups.

These settings are available for the **Custom** backup scheme, and are based on the backup method. The **Monthly full, Weekly differential, Daily incremental** is a pre-configured custom scheme.

Example

You use the **Always incremental (single-file)** backup scheme with the default setting for hourly backups:

- Scheduled by time.
- Backups run hourly: Monday to Friday, every 1 hour, from 08:00 AM to 06:00 PM.
- The **Weekly backup** option is set to Monday.

In the **How long to keep** section of the protection plan, you can apply retention rules to monthly, weekly, daily, and hourly backups.

The following table summarizes the backup types that are created during an 8-day period.

Date	Day of week	Description
July 1	Monday	The first backup each month is monthly, so the first backup today is a monthly backup. The other backups during the day are hourly. This week, the first backup is considered a monthly backup. That is why there is no weekly backup. The first backup next week will be a weekly backup.
July 2	Tuesday	The first backup is daily, the other backups during the day are hourly.
July 3	Wednesday	The first backup is daily, the other backups during the day are hourly.
July 4	Thursday	The first backup is daily, the other backups during the day are hourly.

Date	Day of week	Description
July 5	Friday	The first backup is daily, the other backups during the day are hourly.
July 6	Saturday	The first backup is daily, the other backups during the day are hourly.
July 7	Sunday	The first backup is daily, the other backups during the day are hourly.
July 8	Monday	The first backup is weekly, the other backups during the day are hourly.

Configuring retention rules

The retention rules are part of the protection plan, and their availability and options depend on the backup scheme. For more information, see "Retention rules according to the backup scheme" (p. 412).

To configure the retention rules

1. In the protection plan, expand the **Backup** module.
2. Click **How many to keep**.
3. Select one of the following options:
 - **By number of backups**
 - **By backup age**
Separate settings for monthly, weekly, daily, and hourly backups are available. The maximum value for all types is 99.
You can also use a single setting for all backup types.
 - **By total size of backups**
This setting is not available with the **Always incremental (single-file)** backup scheme.
 - **Keep backups indefinitely**
4. [If you did not select **Keep backups indefinitely**] Configure the values for the selected option.
5. [If you did not select **Keep backups indefinitely**] Select when the retention rules are applied:
 - After backup
 - Before backup
This option is not available when backing up Microsoft SQL Server clusters or Microsoft Exchange Server clusters.
6. Click **Done**.
7. Save the protection plan.

Replication

With replication, each new backup is automatically copied to a replication location. The backups in the replication location do not depend on the backups in the source location, and vice versa.

Only the last backup in the source location is replicated. However, if earlier backups are not replicated (for example, due to a network connection problem), the replication operation will include all backups that are created after the last successful replication.

If a replication operation is interrupted, the processed data will be used by the next replication operation.

Usage examples

- Ensuring reliable recovery
Store your backups both on-site (for immediate recovery) and off-site (to guarantee that the backups stay safe even in case of storage failure or a natural disaster that affects the primary location).
- Using the cloud storage to protect data from a natural disaster
Replicate the backups to the cloud storage by transferring only the data changes.
- Keeping only the latest recovery points
Configure retention rules to delete the older backups from a fast storage, in order to save on storage costs.

Supported locations

Location	As source location	As replication location
Local folder	+	+
Network folder	+	+
Cloud storage	-	+
Secure Zone	+	-
Public cloud	+	+

To enable replication

1. In a protection plan, expand the **Backup** module, and then click **Add location**.

Note

The **Add location** option is not available when you select **Cloud storage** in **Where to back up**.

2. From the list of available locations, select the replication location.

The location appears in the protection plan as **2nd location**, **3rd location**, **4th location**, or **5th location**, depending on the number of locations you added for replication.

3. [Optional] Click the gear icon to configure the options for the replication location.
 - **Performance and backup window** – set the backup window for the selected location, as described in "Performance and backup window" (p. 453). These settings define the replication performance.
 - **Remove location** – delete the currently selected replication location.
 - [Only for the cloud storage] **Physical Data Shipping** – save the initial backup on a removable storage device and ship it for upload to the cloud storage, instead of replicating it over the Internet.

This option is suitable for locations with slow network connection or when you want to save bandwidth on big file transfers over the network. Enabling the option does not require advanced Cyber Protect service quotas, but you will need a Physical Data Shipping service quota to create a shipping order and track it. See "Physical Data Shipping" (p. 457).

Note

This option is supported with protection agent version from release C21.06 or later.

4. [Optional] In the **How many to keep** row under the replication location, configure the retention rules for that location, as described in "Retention rules" (p. 411).
5. [Optional] Repeat steps 1 – 4 to add more replication locations.

You can configure up to four replication locations (**2nd location**, **3rd location**, **4th location**, and **5th location**). If you select **Cloud storage**, you cannot add more replication locations.

Important

If you enable backup and replication in the same protection plan, ensure that the replication completes before the next scheduled backup. If the replication is still in progress, the scheduled backup will not start—for example, a scheduled backup that runs once every 24 hours will not start if the replication takes 26 hours to complete.

To avoid this dependency, use a separate plan for backup replication. For more information about this specific plan, refer to "Backup replication" (p. 194).

Encryption

The Advanced Encryption Standard (AES) cryptographic algorithm operates in Galois/Counter mode (GCM) and uses a randomly generated key with a user-defined size of 128, 192 or 256 bits. The larger the key size, the more secure your data will be. The encryption key is then encrypted with AES-256 by using an SHA-2 (256-bit) hash of the password as a key. The password itself is not stored anywhere on the disk or in the backups, and the password hash is used for verification.

With this two-level security, the backup data is protected from unauthorized access, but recovering a lost password is not possible.

Note

Using the AES-256 algorithm with a strong password provides quantum-resistant encryption. It is safe against cryptanalytic attacks that rely on quantum computing.

We recommend that you encrypt all backups that are stored in the cloud storage, especially if your company is subject to regulatory compliance.

You can configure encryption in the following ways:

- In the protection plan
- As a machine property, by using the Cyber Protect Monitor or the command-line interface

Configuring encryption in the protection plan

This option is not available for accounts in the Enhanced security mode. In the Enhanced security mode, you can configure encryption only on the protected device. To learn how to do this, see "Configuring encryption as a machine property" (p. 418).

To configure encryption

1. In a protection plan, expand the **Backup** module.
2. In **Encryption**, click **Specify password**.
3. Specify and confirm the encryption password.
4. Select the key length for the Advanced Encryption Standard (AES) algorithm:
 - AES-256
 - AES-192
 - AES-128

Note

Using the AES-256 algorithm with a strong password provides quantum-resistant encryption. It is safe against cryptanalytic attacks that rely on quantum computing.

5. Click **OK**.

Warning!

There is no way to recover encrypted backups if you lose or forget the password.

You cannot change the encryption settings after you apply the protection plan. To use different encryption settings, create a new plan.

Configuring encryption as a machine property

You can configure backup encryption as a machine property. In this case, backup encryption is not configured in the protection plan, but on the protected workload. Encryption as a machine property uses the AES algorithm with a 256-bit key (AES-256).

Note

Using the AES-256 algorithm with a strong password provides quantum-resistant encryption. It is safe against cryptanalytic attacks that rely on quantum computing.

Configuring encryption as a machine property affects the protection plans in the following way:

- **Protection plans that are already applied to the machine.** If the encryption settings in a protection plan are different, the backups will fail.
- **Protection plans that will be applied to the machine later.** The encryption settings saved on the machine will override the encryption settings in the protection plan. Any backup will be encrypted, even if encryption is disabled in the Backup module settings.

For accounts in the Enhanced security mode, only encryption as a machine property is available.

If you have more than one Agent for VMware connected to the same vCenter Server, and you configure encryption as a machine property, you must use the same encryption password on all machines with Agent for VMware, because of the load balancing between the agents.

You can configure encryption as a machine property in the following ways:

- On the command line
- In Cyber Protect Monitor (Available for Windows and macOS)

To configure encryption**On the command line**

1. Log in as an administrator (in Windows) or the root user (in Linux).
2. On the command line, run the following command:
 - For Windows:

```
<installation_path>\PyShell\bin\acropsh.exe -m manage_creds --set-password  
<encryption_password>
```

By default, the installation path is %ProgramFiles%\BackupClient.

- For Linux:

```
/usr/sbin/acropsh -m manage_creds --set-password <encryption_password>
```

- For a virtual appliance:

```
./sbin/acropsh -m manage_creds --set-password <encryption_password>
```

Warning!

There is no way to recover encrypted backups if you lose or forget the password.

In Cyber Protect Monitor

1. Log in as an administrator.
2. Click the Cyber Protect Monitor icon in the notification area (in Windows) or the menu bar (in macOS).
3. Click the gear icon, and then click **Settings > Encryption**.
4. Select **Set a password for this machine**. Specify and confirm the encryption password.
5. Click **Save**.

Warning!

There is no way to recover encrypted backups if you lose or forget the password.

To reset the encryption settings

1. Log in as an administrator (in Windows) or root user (in Linux).
2. On the command line, run the following command:

- For Windows:

```
<installation_path>\PyShell\bin\acropsh.exe -m manage_creds --reset
```

By default, the installation path is %ProgramFiles%\BackupClient.

- For Linux:

```
/usr/sbin/acropsh -m manage_creds --reset
```

- For a virtual appliance:

```
./sbin/acropsh -m manage_creds --reset
```

Important

If you reset the encryption as a machine property or change the encryption password after a protection plan creates a backup, the next backup operation will fail. To continue backing up the workload, create a new protection plan.

Notarization

Note

This feature is available with the Advanced Backup pack.

Notarization enables you to prove that a file is authentic and unchanged since it was backed up. We recommend that you enable notarization when backing up your legal document files or other files that require proved authenticity.

Notarization is available only for file-level backups. Files that have a digital signature are skipped, because they do not need to be notarized.

Notarization is *not* available:

- If the backup format is set to **Version 11**
- If the backup destination is Secure Zone

How to use notarization

To enable notarization of all files selected for backup (except for the files that have a digital signature), enable the **Notarization** switch when creating a protection plan.

When configuring recovery, the notarized files will be marked with a special icon, and you can [verify the file authenticity](#).

How it works

During a backup, the agent calculates the hash codes of the backed-up files, builds a hash tree (based on the folder structure), saves the tree in the backup, and then sends the hash tree root to the notary service. The notary service saves the hash tree root in the Ethereum blockchain database to ensure that this value does not change.

When verifying the file authenticity, the agent calculates the hash of the file, and then compares it with the hash that is stored in the hash tree inside the backup. If these hashes do not match, the file is considered not authentic. Otherwise, the file authenticity is guaranteed by the hash tree.

To verify that the hash tree itself was not compromised, the agent sends the hash tree root to the notary service. The notary service compares it with the one stored in the blockchain database. If the hashes match, the selected file is guaranteed to be authentic. Otherwise, the software displays a message that the file is not authentic.

Default backup options

The default values of [backup options](#) exist at the company, unit, and user level. When a unit or a user account is created within a company or within a unit, it inherits the default values set for the company or for the unit.

Company administrators, unit administrators, and every user without the administrator rights can change a default option value against the pre-defined one. The new value will be used by default in all protection plans created at the respective level after the change takes place.

When creating a protection plan, a user can override a default value with a custom value that will be specific for this plan only.

To change a default option value

1. Do one of the following:
 - To change the default value for the company, sign in to the Cyber Protect console as a company administrator.
 - To change the default value for a unit, sign in to the Cyber Protect console as an administrator of the unit.

- To change the default value for yourself, sign in to the Cyber Protect console by using an account without the administrator rights.
2. Click **Settings > System settings**.
 3. Expand the **Default backup options** section.
 4. Select the option, and then make the necessary changes.
 5. Click **Save**.

Backup options

To modify the backup options of a protection plan, in the **Backup** module, in the **Backup options** field, click **Change**.

Availability of the backup options

The set of available backup options depends on:

- The environment the agent operates in (Windows, Linux, macOS).
- The type of the data being backed up (disks, files, virtual machines, application data).
- The backup destination (the cloud storage, local or network folder).

The following table summarizes the availability of the backup options.

	Disk-level backup			File-level backup			Virtual machines			SQL and Exchange
	Windows	Linux	macOS	Windows	Linux	macOS	ESXi	Hypervisor-V	Virtuozzo	Windows
Alerts	+	+	+	+	+	+	+	+	+	+
Backup consolidation	+	+	+	+	+	+	+	+	+	-
Backup file name	+	+	+	+	+	+	+	+	+	+
Backup format	+	+	+	+	+	+	+	+	+	+
Backup validation	+	+	+	+	+	+	+	+	+	+
Changed block tracking (CBT)	+	-	-	-	-	-	+	+	-	-
Cluster backup mode	-	-	-	-	-	-	-	-	-	+
Compression level	+	+	+	+	+	+	+	+	+	+

Error handling										
Re-attempt, if an error occurs	+	+	+	+	+	+	+	+	+	+
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+	+	+	+	+
Ignore bad sectors	+	-	+	+	-	+	+	+	+	-
Re-attempt, if an error occurs during VM snapshot creation	-	-	-	-	-	-	+	+	+	-
Fast incremental/differential backup	+	+	+	-	-	-	-	-	-	-
File-level backup snapshot	-	-	-	+	+	+	-	-	-	-
File filters	+	+	+	+	+	+	+	+	+	-
Forensic data	+	-	-	-	-	-	-	-	-	-
Log truncation	-	-	-	-	-	-	+	+	-	SQL only
LVM snapshotting	-	+	-	-	-	-	-	-	-	-
Mount points	-	-	-	+	-	-	-	-	-	-
Multi-volume snapshot	+	+	-	+	+	-	-	-	-	-
One-click recovery	+	+	-	-	-	-	-	-	-	-
Performance and backup window	+	+	+	+	+	+	+	+	+	+
Physical Data Shipping	+	+	+	+	+	+	+	+	+	-
Pre/Post commands	+	+	+	+	+	+	+	+	+	+

Pre/Post data capture commands	+	+	+	+	+	+	-	-	-	+
Scheduling										
Distribute start times within a time window	+	+	+	+	+	+	+	+	+	+
Limit the number of simultaneously running backups	-	-	-	-	-	-	+	+	+	-
Sector-by-sector backup	+	+	-	-	-	-	+	+	+	-
Splitting	+	+	+	+	+	+	+	+	+	+
Task failure handling	+	+	+	+	+	+	+	+	+	+
Task start conditions	+	+	-	+	+	-	+	+	+	+
Volume Shadow Copy Service (VSS)	+	-	-	+	-	-	-	+	-	+
Volume Shadow Copy Service (VSS) for virtual machines	-	-	-	-	-	-	+	+	-	-
Weekly backup	+	+	+	+	+	+	+	+	+	+
Windows event log	+	-	-	+	-	-	+	+	-	+

Alerts

No successful backups for a specified number of consecutive days

The preset is: **Disabled**.

This option determines whether to generate an alert if no successful backups were performed by the protection plan for a specified period of time. In addition to failed backups, the software counts backups that did not run on schedule (missed backups).

The alerts are generated on a per-machine basis and are displayed on the **Alerts** tab.

You can specify the number of consecutive days without backups after which the alert is generated.

Backup consolidation

This option defines whether to consolidate backups during cleanup or to delete entire backup chains.

The preset is: **Disabled**.

Consolidation is the process of combining two or more subsequent backups into a single backup.

If this option is enabled, a backup that should be deleted during cleanup is consolidated with the next dependent backup (incremental or differential).

Otherwise, the backup is retained until all dependent backups become subject to deletion. This helps avoid the potentially time-consuming consolidation, but requires extra space for storing backups whose deletion is postponed. The backups' age or number can exceed the values specified in the retention rules.

Important


Please be aware that consolidation is just a method of deletion, but not an alternative to deletion. The resulting backup will not contain data that was present in the deleted backup and was absent from the retained incremental or differential backup.

This option is *not* effective if any of the following is true:

- The backup destination is the cloud storage.
- The backup scheme is set to **Always incremental (single-file)**.
- The [backup format](#) is set to **Version 12**.

Backups stored in the cloud storage, as well as single-file backups (both version 11 and 12 formats), are always consolidated because their inner structure makes for fast and easy consolidation.

However, if version 12 format is used, and multiple backup chains are present (every chain being stored in a separate .tibx file), consolidation works only within the last chain. Any other chain is deleted as a whole, except for the first one, which is shrunk to the minimum size to keep the meta information (~12 KB). This meta information is required to ensure the data consistency during simultaneous read and write operations. The backups included in these chains disappear from the GUI as soon as the retention rule is applied, although they physically exist until the entire chain is deleted.

In all other cases, backups whose deletion is postponed are marked with the trash can icon () in the GUI. If you delete such a backup by clicking the X sign, consolidation will be performed.

Backup file name

This option defines the names of the backup files that are created by the protection plan or by the cloud applications backup plan.

For backup files that are created by protection plans, you can see these names in a file manager when you browse the backup location.

What is a backup file?

Each protection plan creates one or more files in the backup location, depending on which backup scheme and which [backup format](#) is used. The following table lists the files that can be created per machine or mailbox.

	Always incremental (single-file)	Other backup schemes
Version 11 backup format	One TIB file and one XML metadata file	Multiple TIB files and one XML metadata file
Version 12 backup format	One TIBX file per backup chain (a full or differential backup, and all incremental backups that depend on it). If the size of a file stored in a local or network (SMB) folder exceeds 200 GB, the file is split to 200-GB files by default.	

All files have the same name, with or without the addition of a timestamp or a sequence number. You can define this name (referred to as the backup file name) when you create or edit a protection plan or a cloud applications backup plan.

Note

Timestamp is added to the backup file name only in the version 11 backup format.

If you change a backup file name in a protection plan or a cloud applications backup plan, the next backup will be a full backup.

If you specify a file name of an existing backup of the same machine, a full, incremental, or differential backup will be created according to the plan schedule.

Note

If you move backup files (.tibx) from their original storage, do not rename them. Renamed files will appear corrupted and you will not be able to recover data from them.

It is possible to set backup file names for locations that cannot be browsed by a file manager (such as the cloud storage). In this case, you will see the custom names on the **Backup storage** tab.

Where can I see backup file names?

For protection plans, on the **Backup storage** tab, select the location, and then select the backup archive.

- The default backup file name is shown on the **Details** panel.
- If you set a non-default backup file name, it will be shown directly on the **Backup storage** tab, in the **Name** column.

For cloud applications backup plans, on the **Backup storage** tab, select the location, select the backup archive, and then click the gear icon.

Limitations for backup file names

- A backup file name cannot end with a digit.
In the default backup file name, to prevent the name from ending with a digit, the letter "A" is appended. When creating a custom name, always make sure that it does not end with a digit. When using variables, the name must not end with a variable, because a variable might end with a digit.
- A backup file name cannot contain the following symbols: **()&?*\${}<>":\|/##**, line endings (**\n**), and tabs (**\t**).

Note

Choose user-friendly backup file names. This will help you to easily distinguish backups when browsing the backup location with a file manager.

Default backup file name

The default backup file name for backups of entire physical and virtual machines, disks/volumes, files/folders, Microsoft SQL Server databases, Microsoft Exchange Server databases, and ESXi configuration is [Machine Name]-[Plan ID]-[Unique ID]A.

The default name for Exchange mailbox backups and Microsoft 365 mailbox backups created by a local Agent for Microsoft 365 is [Mailbox ID]_mailbox_[Plan ID]A.

The default name for Microsoft Azure backups is prefixed with [Mailbox ID]_. This prefix cannot be removed.


The default name for cloud application backups created by cloud agents is [Resource Name]_[Resource Type]_[Resource Id]_[Plan Id]A.

The default name consists of the following variables:


- [Machine Name] This variable is replaced with the name of the machine (the same name that is shown in the Cyber Protect console).
- [Plan ID], [Plan Id] These variables are replaced with the unique identifier of the protection plan. This value does not change if the plan is renamed.
- [Unique ID] This variable is replaced with the unique identifier of the selected machine. This value does not change if the machine is renamed.
- [Mailbox ID] This variable is replaced with the mailbox user's principal name (UPN).
- [Resource Name] This variable is replaced with the cloud data source name, such as the user's principal name (UPN), SharePoint site URL, or Shared drive name.
- [Resource Type] This variable is replaced with the cloud data source type, such as mailbox, 0365Mailbox, 0365PublicFolder, OneDrive, SharePoint, GDrive.

- [Resource ID] This variable is replaced with the unique identifier of the cloud data source. This value does not change if the cloud data source is renamed.
- "A" is a safeguard letter that is appended to prevent the name from ending with a digit.

The diagram below shows the default backup file name.



The diagram below shows the default backup file name for Microsoft 365 mailbox backups performed by a local agent.



Names without variables

If you change the backup file name to MyBackup, the backup files will look like the following examples. Both examples assume daily incremental backups scheduled at 14:40, starting from September 13, 2016.

For the version 12 format with the **Always incremental (single-file)** backup scheme:

```
MyBackup.tibx
```

For the version 12 format with other backup schemes:

```
MyBackup.tibx
MyBackup-0001.tibx
MyBackup-0002.tibx
...
```

Using variables

Besides the variables that are used by default, you can use the following variables:

- The [Plan name] variable, which is replaced with the name of the protection plan.
- The [Virtualization Server Type] variable, which is replaced with "vmwesx" if virtual machines are backed up by Agent for VMware or with "mshyperv" if virtual machines are backed up by Agent for Hyper-V.

If multiple machines or mailboxes are selected for backup, the backup file name must contain the [Machine Name], the [Unique ID], the [Mailbox ID], the [Resource Name], or the [Resource Id] variable.

Creating backups in an existing backup archive

You can configure the backups of a workload to be added to an existing backup archive.

This option might be useful, for example, when a protection plan is applied to a single machine, and you have to remove this machine from the Cyber Protect console, or uninstall the agent along with its configuration settings. After you add the machine again or reinstall the agent, you can force the protection plan to continue backing up to the original archive.

Backup file name

You can change the default backup file name or select an existing backup file to add backups to. If you change the backup file name, the next backup will be a full backup.

To configure the backups of a workload to be added to an existing backup archive

Non-cloud-to-cloud workloads

1. On the **All devices** screen, click the workload, and then click **Protect**.
2. In the protection plan settings, extend the **Backup** module.
3. Click **Backup options**, and then click **Change**.
4. On the **Backup file name** tab, click **Select**.

The **Select** button shows the backups in the location selected in the **Where to back up** section of the protection plan.

Note

The **Select** button is only available for protection plans that are created for and applied to a single workload.

5. Select an archive, and then click **Done**.
6. Click **Done**, and then click **Apply**.

Cloud-to-cloud workloads

1. On the **Management > Cloud applications backup** tab, select the plan.
2. Click **Edit**, and then click the gear icon next to the plan's name.
3. On the **File backup name** tab, click **Select**.

Note

The **Select** button is only available for backup plans that are created for and applied to a single workload.

4. Select a backup archive, and then click **Done**.
5. Click **Done**, and then click **Save changes**.

Backup format

The **Backup format** option defines the format of the backups created by the protection plan. This option is available only for protection plans that already use the version 11 backup format. If this is the case, you can change the backup format to version 12. After you switch the backup format to version 12, the option becomes unavailable.

- **Version 11**

The legacy format preserved for backward compatibility.

Note

You cannot back up Database Availability Groups (DAG) by using the backup format version 11. Backing up of DAG is supported only in the version 12 format.

- **Version 12**

The backup format that was introduced in Acronis Backup 12 for faster backup and recovery.

Each backup chain (a full or differential backup, and all incremental backups that depend on it) is saved to a single TIBX file.

Backup format and backup files

For backup locations that can be browsed with a file manager (such as local or network folders), the backup format determines the number of files and their extension. The following table lists the files that can be created per machine or mailbox.

	Always incremental (single-file)	Other backup schemes
Version 11 backup format	One TIB file and one XML metadata file	Multiple TIB files and one XML metadata file
Version 12 backup format	One TIBX file per backup chain (a full or differential backup, and all incremental backups that depend on it). If the size of a file stored in a local or network (SMB) folder exceeds 200 GB, the file is split to 200-GB files by default.	

Changing the backup format to version 12 (TIBX)

If you change the backup format from version 11 (TIB format) to version 12 (TIBX format):

- The next backup will be full.
- In backup locations that can be browsed with a file manager (such as local or network folders), a new TIBX file will be created. The new file will have the name of the original file, appended with the **_v12A** suffix.
- Retention rules and replication will be applied only to the new backups.
- The old backups will not be deleted and will remain available on the **Backup storage** tab. You can delete them manually.
- The old cloud backups will not consume the **Cloud storage** quota.
- The old local backups will consume the **Local backup** quota until you delete them manually.

In-archive deduplication

The TIBX backup format of version 12 supports in-archive deduplication that brings the following advantages:

- Significantly reduced backup size, with built-in block-level deduplication for any type of data
- Efficient handling of hard links ensures that there are no storage duplicates
- Hash-based chunking

Note

In-archive deduplication is enabled by default for all backups in the TIBX format. You do not have to enable it in the backup options, and you cannot disable it.

Backup format compatibility across different product versions

For information about backup format compatibility, see [Backup archive compatibility across different product versions \(1689\)](#).

Backup validation

Validation is an operation that checks the possibility of data recovery from a backup. When this option is enabled, each backup created by the protection plan is validated immediately after creation, by using the checksum verification method. This operation is performed by the protection agent.

The preset is: **Disabled**.

For more information about the validation via checksum verification, refer to "Checksum verification" (p. 200).

Note

Depending on the settings chosen by your service provider, validation might not be available when backing up to the cloud storage. Validation is also not available for backup locations on public clouds.

Changed block tracking (CBT)

This option is effective for the following backups:

- Disk-level backups of virtual machines
- Disk-level backups of physical machines running Windows
- Backups of Microsoft SQL Server databases
- Backups of Microsoft Exchange Server databases

The preset is: **Enabled**.

This option determines whether to use Changed Block Tracking (CBT) when performing an incremental or differential backup.

The CBT technology accelerates the backup process. Changes to the disk or database content are continuously tracked at the block level. When a backup starts, the changes can be immediately saved to the backup.

Cluster backup mode

Note

This feature is available with the Advanced Backup pack.

These options are effective for database-level backup of Microsoft SQL Server and Microsoft Exchange Server.

These options are effective only if the cluster itself (Microsoft SQL Server Always On Availability Groups (AAG) or Microsoft Exchange Server Database Availability Group (DAG)) is selected for backup, rather than the individual nodes or databases inside of it. If you select individual items inside the cluster, the backup will not be cluster-aware and only the selected copies of the items will be backed up.

Microsoft SQL Server

This option determines the backup mode for SQL Server Always On Availability Groups (AAG). For this option to be effective, Agent for SQL must be installed on all of the AAG nodes. For more information about backing up Always On Availability Groups, refer to "[Protecting Always On Availability Groups \(AAG\)](#)".

The preset is: **Secondary replica if possible**.

You can choose one of the following:

- **Secondary replica if possible**

If all secondary replicas are offline, the primary replica is backed up. Backing up the primary replica may slow down the SQL Server operation, but the data will be backed up in the most recent state.

- **Secondary replica**

If all secondary replicas are offline, the backup will fail. Backing up secondary replicas does not affect the SQL server performance and allows you to extend the backup window. However, passive replicas may contain information that is not up-to-date, because such replicas are often set to be updated asynchronously (lagged).

- **Primary replica**

If the primary replica is offline, the backup will fail. Backing up the primary replica may slow down the SQL Server operation, but the data will be backed up in the most recent state.

Regardless of the value of this option, to ensure the database consistency, the software skips databases that are *not* in the **SYNCHRONIZED** or **SYNCHRONIZING** states when the backup starts. If all databases are skipped, the backup fails.

Microsoft Exchange Server

This option determines the backup mode for Exchange Server Database Availability Groups (DAG). For this option to be effective, Agent for Exchange must be installed on all of the DAG nodes. For more information about backing up Database Availability Groups, refer to "Protecting Database Availability Groups (DAG)".

The preset is: **Passive copy if possible.**

You can choose one of the following:

- **Passive copy if possible**

If all passive copies are offline, the active copy is backed up. Backing up the active copy may slow down the Exchange Server operation, but the data will be backed up in the most recent state.

- **Passive copy**

If all passive copies are offline, the backup will fail. Backing up passive copies does not affect the Exchange Server performance and allows you to extend the backup window. However, passive copies may contain information that is not up-to-date, because such copies are often set to be updated asynchronously (lagged).

- **Active copy**

If the active copy is offline, the backup will fail. Backing up the active copy may slow down the Exchange Server operation, but the data will be backed up in the most recent state.

Regardless of the value of this option, to ensure the database consistency, the software skips databases that are *not* in the **HEALTHY** or **ACTIVE** states when the backup starts. If all databases are skipped, the backup fails.

Compression level

Note

This option is not available for cloud-to-cloud backups. Compression for these backups is enabled by default with a fixed level that corresponds to the **Normal** level below.

The option defines the level of compression applied to the data being backed up. The available levels are: **None, Normal, High, Maximum**.

The preset is: **Normal**.

A higher compression level means that the backup process takes more time, but the resulting backup occupies less space. Currently, the **High** and **Maximum** levels work similarly.

The optimal data compression level depends on the type of data being backed up. For example, even maximum compression will not significantly reduce the backup size if the backup contains essentially compressed files, such as .jpg, .pdf or .mp3. However, formats such as .doc or .xls will be compressed well.

Error handling

These options enable you to specify how to handle errors that might occur during backup.

Re-attempt, if an error occurs

The preset is: **Enabled. Number of attempts: 10. Interval between attempts: 30 seconds**.

When a recoverable error occurs, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds or the specified number of attempts are performed, depending on which comes first.

For example, if the backup destination on the network becomes unavailable or not reachable during a running backup, the software will attempt to reach the destination every 30 seconds, but no more than 30 times. The attempts will be stopped as soon as the connection is resumed or the specified number of attempts is performed, depending on which comes first.

However, if the backup destination is not available when the backup starts, only 10 attempts will be made.

Do not show messages and dialogs while processing (silent mode)

The preset is: **Enabled**.

With the silent mode enabled, the program will automatically handle situations requiring user interaction (except for handling bad sectors, which is defined as a separate option). If an operation cannot continue without user interaction, it will fail. Details of the operation, including errors, if any, can be found in the operation log.

Ignore bad sectors

The preset is: **Disabled**.

When this option is disabled, each time the program comes across a bad sector, the backup activity will be assigned the **Interaction required** status. In order to back up the valid information on a

rapidly dying disk, enable ignoring bad sectors. The rest of the data will be backed up and you will be able to mount the resulting disk backup and extract valid files to another disk.

Note

Skipping bad sectors is not supported on Linux. You can back up Linux systems with bad sectors in offline mode, by using the bootable media builder in the on-premises version of Cyber Protect. Using the on-premises bootable media builder requires a separate license. Contact support for assistance.

Re-attempt, if an error occurs during VM snapshot creation

The preset is: **Enabled. Number of attempts: 3. Interval between attempts: 5 minutes.**

When taking a virtual machine snapshot fails, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts are performed, depending on which comes first.

Fast incremental/differential backup

This option is effective for incremental and differential disk-level backup.

This option is not effective (always disabled) for volumes formatted with the JFS, ReiserFS3, ReiserFS4, ReFS, or XFS file systems.

The preset is: **Enabled.**

Incremental or differential backup captures only data changes. To speed up the backup process, the program determines whether a file has changed or not by the file size and the date/time when the file was last modified. Disabling this feature will make the program compare the entire file contents to those stored in the backup.

File filters (Inclusions/Exclusions)

Use file filters to include only specific files and folders in a backup, or to exclude specific files and folders from a backup.

File filters are available for entire machine backups, disk-level backups, and file-level backups, unless stated otherwise.

File filters are not available with the XFS, JFS, exFAT, and ReiserFS4 file systems. For more information, see "Supported file systems" (p. 58).

File filters are not applicable to dynamic disks (LVM or LDM volumes) of virtual machines that are backed up in the agentless mode, for example, by Agent for VMware, Agent for Hyper-V, or Agent for Scale Computing.

To enable file filters

1. In a protection plan, expand the **Backup** module.
2. In **Backup options**, click **Change**.
3. Select **File filters (Inclusions/Exclusions)**.
4. Use any of the options described below.

Inclusion and exclusion filters

There are two filters – inclusion filter and exclusion filter.

- **Include only the files that match the following criteria**

If you specify `C:\File.exe` in the inclusion filter, only this file will be backed up, even if you selected Entire machine backup.

Note

This filter is not supported for file-level backups when the backup format is **Version 11**, and the backup destination is not the cloud storage.

- **Exclude the files that match the following criteria**

If you specify `C:\File.exe` in the exclusion filter, this file will be skipped during a backup, even if you selected Entire machine backup.

You can use both filters in the same time. The exclusion filter takes precedence over the inclusion filter – that is, if you specify `C:\File.exe` in both fields, this file will be skipped during a backup.

Filter criteria

As filter criteria, you can use file and folder names, full paths to files and folders, and masks with wildcard symbols.

The filter criteria are case insensitive. For example, by specifying `C:\Temp`, you will select `C:\TEMP` and `C:\temp`.

- **Name**
Specify the name of the file or folder, such as `Document.txt`. All files and folders with that name will be selected.
- **Full path**
Specify the full path to the file or folder, starting with the drive letter (when backing up Windows) or the root directory (when backing up Linux or macOS). In Windows, Linux, and macOS, you can use forward slashes (as in `C:/Temp/File.tmp`). In Windows, you can also use the traditional backslashes (as in `C:\Temp\File.tmp`).

Important

If the operating system of the backed-up machine is not detected correctly during a disk-level backup, full path file filters will not work. For an exclusion filter, a warning will be shown. If there is an inclusion filter, the backup will fail.

For example, a full path to a file could be `C:\Temp\File.tmp`. A full path filter, which includes the drive letter or the root directory, such as `C:\Temp\File.tmp` or `C:\Temp*`, will result in a warning or failure.

A filter that does not use the drive letter or the root directory (for example, `Temp*` or `Temp\File.tmp`) or a filter that starts with an asterisk (for example, `*C:\`) will not result in warning or failure. However, if the operating system of the backed-up machine is not detected correctly, these filters will not work, either.

- Mask

You can use the following wildcard characters for the names and full paths: asterisk (*), double asterisk (**), and question mark (?).

The asterisk (*) represents zero or more characters. For example, the filter criterion **Doc*.txt** matches the files `Doc.txt` and `Document.txt`.

The double asterisk (**) represents zero or more characters, including the slash character. For example, ****/Docs/**/*.txt** matches all .txt files in all subfolders of all folders named `Docs`. You can use the double asterisk (**) wildcard only for backups in the Version 12 format.

The question mark (?) represents only one character. For example, **Doc?.txt** matches the files `Doc1.txt` and `Docs.txt`, but not the files `Doc.txt` or `Doc11.txt`.

File-level backup snapshot

This option is effective only for file-level backup.

This option defines whether to back up files one by one or by taking an instant data snapshot.

Note

Files that are stored on network shares are always backed up one by one.

The preset is:

- If only machines running Linux are selected for backup: **Do not create a snapshot.**
- Otherwise: **Create snapshot if it is possible.**

You can select one of the following:

- **Create a snapshot if it is possible**

Back up files directly if taking a snapshot is not possible.

- **Always create a snapshot**

The snapshot enables backing up of all files including files opened for exclusive access. The files will be backed up at the same point in time. Choose this setting only if these factors are critical,

that is, backing up files without a snapshot does not make sense. If a snapshot cannot be taken, the backup will fail.

- **Do not create a snapshot**

Always back up files directly. Trying to back up files that are opened for exclusive access will result in a read error. Files in the backup may be not time-consistent.

Forensic data

Viruses, malware, and ransomware can carry out malicious activities, such as stealing or changing data. These activities may need to be investigated, which is possible only if digital evidence is provided. However, pieces of digital evidence, such as files or activity traces, may be deleted or the machine on which the malicious activity happened may become unavailable.

Backups with forensic data allow investigators to analyze disk areas that are not usually included in a regular disk backup. The **Forensic data** backup option allows you to collect the following pieces of digital evidence that can be used in forensic investigations: snapshots of unused disk space, memory dumps, and snapshots of running processes.

Backups with forensic data are automatically notarized.

The **Forensic data** option is available only for entire machine backups of Windows machines that run the following operating systems:

- Windows 8.1, Windows 10
- Windows Server 2012 R2 – Windows Server 2019

Backups with forensic data are not available for the following machines:

- Machines that are connected to your network through VPN and do not have direct access to the Internet
- Machines with disks that are encrypted by BitLocker

Note

You cannot modify the forensic data settings after you apply a protection plan with enabled **Backup** module to a machine. To use different forensic data settings, create a new protection plan.

You can store backups with forensic data in the following locations:

- Cloud storage
- Local folder

Note

The local folder location is supported only for external hard disks connected via USB. Local dynamic disks are not supported as a location for backups with forensic data.

- Network folder

Forensic backup process

The system performs the following during a forensic backup process:

1. Collects raw memory dump and the list of running processes.
2. Automatically reboots a machine into the bootable media.
3. Creates the backup that includes both the occupied and unallocated space.
4. Notarizes the backed-up disks.
5. Reboots into the live operating system and continues plan execution (for example, replication, retention, validation and other).

To configure forensic data collection

1. In the Cyber Protect console, go to **Devices > All devices**. Alternatively, the protection plan can be created from the **Management** tab.
2. Select the device and click **Protect**.
3. In the protection plan, enable the **Backup** module.
4. In **What to back up**, select **Entire machine**.
5. In **Backup options**, click **Change**.
6. Find the **Forensic data** option.
7. Enable **Collect forensic data**. The system will automatically collect a memory dump and create a snapshot of running processes.

Note

Full memory dump may contain sensitive data such as passwords.

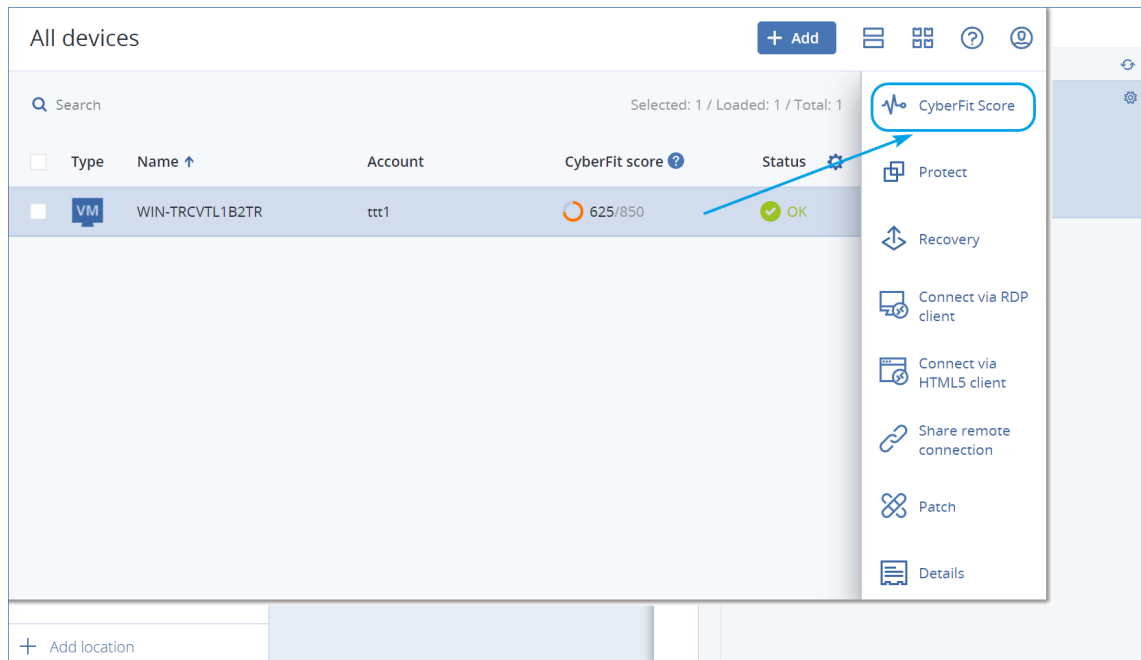
8. Specify the location.
9. Click **Run Now** to perform a backup with forensic data right away or wait until the backup is created according to the schedule.
10. Go to **Monitoring > Activities**, verify that the backup with forensic data was successfully created.

As a result, backups will include forensic data and you will be able to get them and analyze. Backups with forensic data are marked and can be filtered among other backups in **Backup storage > Locations** by using the **Only with forensic data** option.

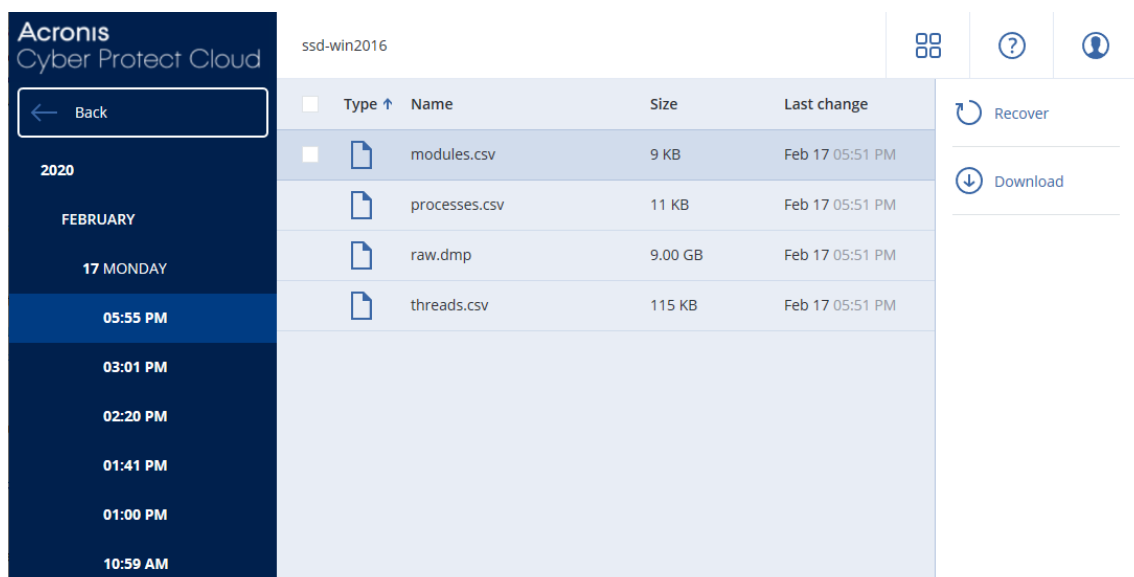
How to get forensic data from a backup?

1. In the Cyber Protect console, go to **Backup storage**, select the location with backups that include forensic data.
2. Select the backup with forensic data and click **Show backups**.
3. Click **Recover** for the backup with forensic data.

- To get only the forensic data, click **Forensic data**.



The system will show a folder with forensic data. Select a memory dump file or any other forensic file, and then click **Download**.



- To recover a full forensic backup, click **Entire machine**. The system will recover the backup without the boot mode. Thus, it will be possible to check that the disk was not changed.

You can use the provided memory dump with several of third-party forensic software, for example, use Volatility Framework at <https://www.volatilityfoundation.org/> for further memory analysis.

Notarization of backups with forensic data

To ensure that a backup with forensic data is exactly the image that was taken and it was not compromised, the backup module provides the notarization of backups with forensic data.

How it works

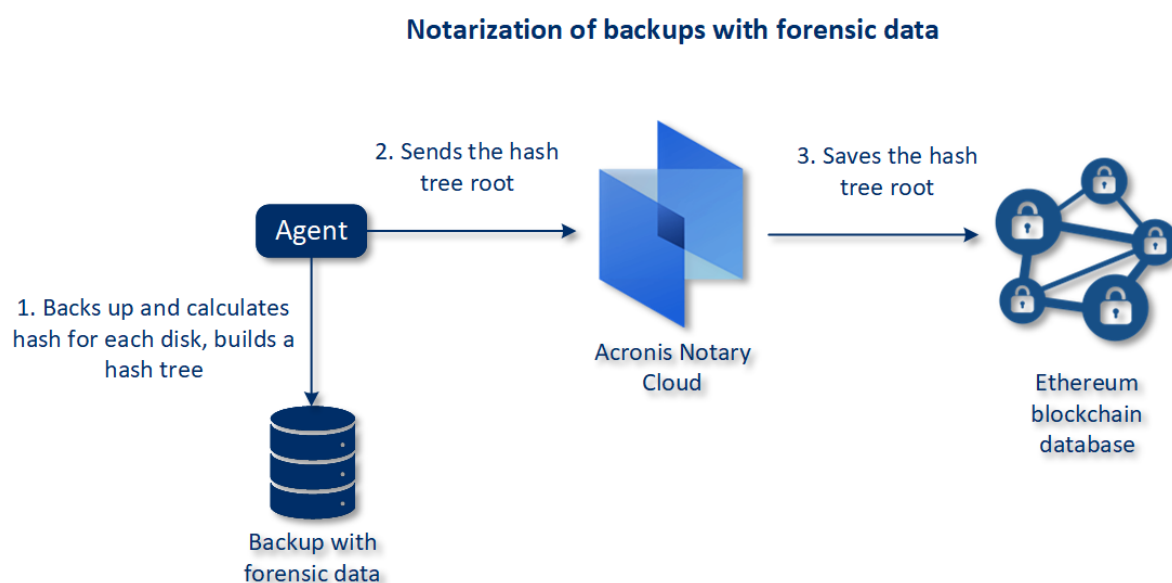
Notarization enables you to prove that a disk with forensic data is authentic and unchanged since it was backed up.

During a backup, the agent calculates the hash codes of the backed-up disks, builds a hash tree, saves the tree in the backup, and then sends the hash tree root to the notary service. The notary service saves the hash tree root in the Ethereum blockchain database to ensure that this value does not change.

When verifying the authenticity of the disk with forensic data, the agent calculates the hash of the disk, and then compares it with the hash that is stored in the hash tree inside the backup. If these hashes do not match, the disk is considered not authentic. Otherwise, the disk authenticity is guaranteed by the hash tree.

To verify that the hash tree itself was not compromised, the agent sends the hash tree root to the notary service. The notary service compares it with the one stored in the blockchain database. If the hashes match, the selected disk is guaranteed to be authentic. Otherwise, the software displays a message that the disk is not authentic.

The scheme below shows shortly the notarization process for backups with forensic data.



To verify the notarized disk backup manually, you can get the certificate for it and follow the verification procedure shown with the certificate by using the [tibxread](#) tool.

Getting the certificate for backups with forensic data

To get the certificate for a backup with forensic data from the console, do the following:

1. Go to **Backup storage** and select the backup with forensic data.
2. Recover the entire machine.

3. The system opens the **Disk Mapping** view.
4. Click the **Get certificate** icon for the disk.
5. The system will generate the certificate and open a new window in the browser with the certificate. Below the certificate you will see the instruction for manual verification of notarized disk backup.

The tool "tibxread" for getting the backed-up data

Cyber Protection provides the tool, called `tibxread`, for manual check of the backed-up disk integrity. The tool allows you to get data from a backup and calculate hash of the specified disk. The tool is installed automatically with the following components: Agent for Windows, Agent for Linux, and Agent for Mac.

The installation path: the same folder as the agent has (for example, `C:\Program Files\BackupClient\BackupAndRecovery`).

The supported locations are:

- The local disk
- The network folder (CIFS/SMB) that can be accessed without the credentials.
In case of a password-protected network folder, you can mount the network folder to the local folder by using the OS tools and then the local folder as the source for this tool.
- The cloud storage

You should provide the URL, port, and certificate. The URL and port can be obtained from the Windows registry key or configuration files on Linux/Mac machines.

For Windows:

```
HKEY_LOCAL_
MACHINE\SOFTWARE\Acronis\BackupAndRecovery\Settings\OnlineBackup\FesAddressCache\Default\<tenant_login>\FesUri
```

For Linux:

```
/etc/Acronis/BackupAndRecovery.config
```

For macOS:

```
/Library/Application Support/Acronis/Registry/BackupAndRecovery.config
```

The certificate can be found in the following locations:

For Windows:

```
%allusersprofile%\Acronis\BackupAndRecovery\OnlineBackup\Default
```

For Linux:

```
/var/lib/Acronis/BackupAndRecovery/OnlineBackup/Default
```

For macOS:

```
/Library/Application Support/Acronis/BackupAndRecovery/OnlineBackup/Default
```

The tool has the following commands:

- list backups
- list content
- get content
- calculate hash

list backups

Lists recovery points in a backup.

SYNOPSIS:

```
tibxread list backups --loc=URI --arc=BACKUP_NAME --raw
```

Options

```
--loc=URI  
--arc=BACKUP_NAME  
--raw  
--utc  
--log=PATH
```

Output template:

```
GUID    Date    Date timestamp  
----    -  
<guid> <date> <timestamp>
```

<guid> – a backup GUID.

<date> – a creation date of the backup. Format is “DD.MM.YYYY HH24:MM:SS”. In local timezone by default (can be changed by using the --utc option).

Output example:

```
GUID    Date    Date timestamp  
----    -  
516FCE73-5E5A-49EF-B673-A9EACB4093B8 18.12.2019 16:01:05 1576684865  
516FCE73-5E5A-49EF-B673-A9EACB4093B9 18.12.2019 16:02:05 1576684925
```

list content

Lists content in a recovery point.

SYNOPSIS:

```
tibxread list content --loc=URI --arc=BACKUP_NAME --password --backup=RECOVERY_POINT_ID
--raw --log=PATH
```

Options

```
--loc=URI
--arc=BACKUP_NAME
--password
--backup=RECOVERY_POINT_ID
--raw
--log=PATH
```

Output template:

```
Disk      Size      Notarization status
-----
<number> <size> <notarization_status>
```

<number> – identifier of the disk.

<size> – size in bytes.

<notarization_status> – the following statuses are possible: Without notarization, Notarized, Next backup.

Output example:

```
Disk      Size      Notary status
-----
1         123123465798 Notarized
2         123123465798 Notarized
```

get content

Writes content of the specified disk in the recovery point to the standard output (stdout).

SYNOPSIS:

```
tibxread get content --loc=URI --arc=BACKUP_NAME --password --backup=RECOVERY_POINT_ID -
-disk=DISK_NUMBER --raw --log=PATH --progress
```

Options

```
--loc=URI
--arc=BACKUP_NAME
--password
--backup=RECOVERY_POINT_ID
--disk=DISK_NUMBER
--raw
```



```
--log=PATH
--progress
```

calculate hash

Calculates the hash of the specified disk in the recovery point by using the SHA-2 (256-bit) algorithm and writes it to the stdout.

SYNOPSIS:

```
tibxread calculate hash --loc=URI --arc=BACKUP_NAME --password --backup=RECOVERY_POINT_ID --disk=DISK_NUMBER --raw --log=PATH --progress
```

Options

```
--loc=URI
--arc=BACKUP_NAME
--password
--backup=RECOVERY_POINT_ID
--disk=DISK_NUMBER
--raw
--log=PATH
```

Options description

Option	Description
--arc=BACKUP_NAME	The backup file name that you can get from the backup properties in the Cyber Protect console. The backup file must be specified with the extension .tibx.
--backup=RECOVERY_POINT_ID	The recovery point identifier
--disk=DISK_NUMBER	Disk number (the same as was written to the output of the "get content" command)
--loc=URI	A backup location URI. The possible formats of the "--loc" option are: <ul style="list-style-type: none">Local path name (Windows) c:/upload/backupsLocal path name (Linux) /var/tmpSMB/CIFS \\server\folderCloud storage --loc=<IP_address>:443 --cert=<path_to_certificate> [--storage_path=/1] <IP_address> - you can find it in the registry key in Windows: HKEY_LOCAL_

	<p>MACHINE\SOFTWARE\Acronis\BackupAndRecovery\Settings\OnlineBackup\FesAdd ressCache\Default\<tenant_login>\FesUri</tenant_login></p> <p><path_to_certificate> - a path to the certificate file to access Cyber Protect Cloud. For example, in Windows this certificate is located in C:\ProgramData\Acronis\BackupAndRecovery\OnlineBackup\Default\<username>.crt </username>.crt where <username> - is your account name to access Cyber Protect Cloud.</p>
--log=PATH	Enables writing the logs by the specified PATH (local path only, format is the same as for --loc=URI parameter). Logging level is DEBUG.
--password=PASS WORD	An encryption password for your backup. If the backup is not encrypted, leave this value empty.
--raw	<p>Hides the headers (2 first rows) in the command output. It is used when the command output should be parsed.</p> <p>Output example without "--raw":</p> <pre> GUID Date Date timestamp ---- - 516FCE73-5E5A-49EF-B673-A9EACB4093B8 18.12.2019 16:01:05 1576684865 516FCE73-5E5A-49EF-B673-A9EACB4093B9 18.12.2019 16:02:05 1576684925 </pre> <p>Output with "--raw":</p> <pre> 516FCE73-5E5A-49EF-B673-A9EACB4093B8 18.12.2019 16:01:05 1576684865 516FCE73-5E5A-49EF-B673-A9EACB4093B9 18.12.2019 16:02:05 1576684925 </pre>
--utc	Shows dates in UTC
--progress	<p>Shows progress of the operation.</p> <p>For example:</p> <pre> 1% 2% 3% 4% ... 100% </pre>

Log truncation

This option is effective for backup of Microsoft SQL Server databases and for disk-level backup with enabled Microsoft SQL Server application backup.

This option defines whether the SQL Server transaction logs are truncated after a successful backup.

The preset is: **Enabled**.

When this option is enabled, a database can be recovered only to a point in time of a backup created by this software. Disable this option if you back up transaction logs by using the native backup engine of Microsoft SQL Server. You will be able to apply the transaction logs after a recovery and thus recover a database to any point in time.

LVM snapshotting

This option is effective only for physical machines.

This option is effective for disk-level backup of volumes managed by Linux Logical Volume Manager (LVM). Such volumes are also called logical volumes.

This option defines how a snapshot of a logical volume is taken. The backup software can do this on its own or rely on Linux Logical Volume Manager (LVM).

The preset is: **By the backup software.**

- **By the backup software.** The snapshot data is kept mostly in RAM. The backup is faster, and unallocated space on the volume group is not required. Therefore, we recommend that you change the preset only if you are experiencing problems with backing up logical volumes.
- **By LVM.** The snapshot is stored on unallocated space of the volume group. If the unallocated space is missing, the snapshot will be taken by the backup software.

The snapshot is used only during the backup operation, and is automatically deleted when the backup operation completes. No temporary files are kept.

Mount points

This option is effective only in Windows for a file-level backup of a data source that includes [mounted volumes](#) or [cluster shared volumes](#).

This option is effective only when you select for backup a folder that is higher in the folder hierarchy than the mount point. (A mount point is a folder on which an additional volume is logically attached.)

- If such folder (a parent folder) is selected for backup, and the **Mount points** option is enabled, all files located on the mounted volume will be included in the backup. If the **Mount points** option is disabled, the mount point in the backup will be empty.
During recovery of a parent folder, the mount point content will or will not be recovered, depending on whether the [Mount points option for recovery](#) is enabled or disabled.
- If you select the mount point directly, or select any folder within the mounted volume, the selected folders will be considered as ordinary folders. They will be backed up regardless of the state of the **Mount points** option and recovered regardless of the state of the [Mount points option for recovery](#).

The preset is: **Disabled.**

Note

You can back up Hyper-V virtual machines residing on a cluster shared volume by backing up the required files or the entire volume with file-level backup. Just power off the virtual machines to be sure that they are backed up in a consistent state.

Example

Let's assume that the **C:\Data1** folder is a mount point for the mounted volume. The volume contains folders **Folder1** and **Folder2**. You create a protection plan for file-level backup of your data.

If you select the check box for volume C and enable the **Mount points** option, the **C:\Data1** folder in your backup will contain **Folder1** and **Folder2**. When recovering the backed-up data, be aware of proper using the [Mount points option for recovery](#).

If you select the check box for volume C, and disable the **Mount points** option, the **C:\Data1** folder in your backup will be empty.

If you select the check box for the **Data1**, **Folder1** or **Folder2** folder, the checked folders will be included in the backup as ordinary folders, regardless of the state of the **Mount points** option.

Multi-volume snapshot

This option is effective for backups of physical machines running Windows or Linux.

This option applies to disk-level backup. This option also applies to file-level backup when the file-level backup is performed by taking a snapshot. (The "[File-level backup snapshot](#)" option determines whether a snapshot is taken during file-level backup).

This option determines whether to take snapshots of multiple volumes at the same time or one by one.

The preset is:

- If at least one machine running Windows is selected for backup: **Enabled**.
- Otherwise: **Disabled**.

When this option is enabled, snapshots of all volumes being backed up are created simultaneously. Use this option to create a time-consistent backup of data spanning multiple volumes; for instance, for an Oracle database.

When this option is disabled, the volumes' snapshots are taken one after the other. As a result, if the data spans several volumes, the resulting backup may be not consistent.

One-click recovery

Note

This feature is available with the Advanced Backup pack.

With One-click recovery you can automatically recover a disk backup of your Windows or Linux machine. This backup can be a backup of the entire machine, or a backup of specific disks or volumes on this machine.

One-click recovery supports the following operations:

- Automatic recovery from the latest backup
- Recovery from a specific backup (also known as recovery point) within the backup archive

One-click recovery supports the following backup storages:

- Secure Zone
- Local folder
- Network folder
- Cloud storage

Important

Suspend the BitLocker encryption until the next restart of your machine when you perform any of the following operations:

- Creating, modifying, or deleting Secure Zone.
- Enabling or disabling Startup Recovery Manager.
- [Only if Startup Recovery Manager was not already enabled] Running the first backup after enabling One-click recovery in the protection plan. This operation enables Startup Recovery Manager automatically.
- Updating Startup Recovery Manager, for example by updating the protection.

If the BitLocker encryption was not suspended during these operations, you will need to specify your BitLocker PIN after restarting your machine.

Enabling One-click recovery

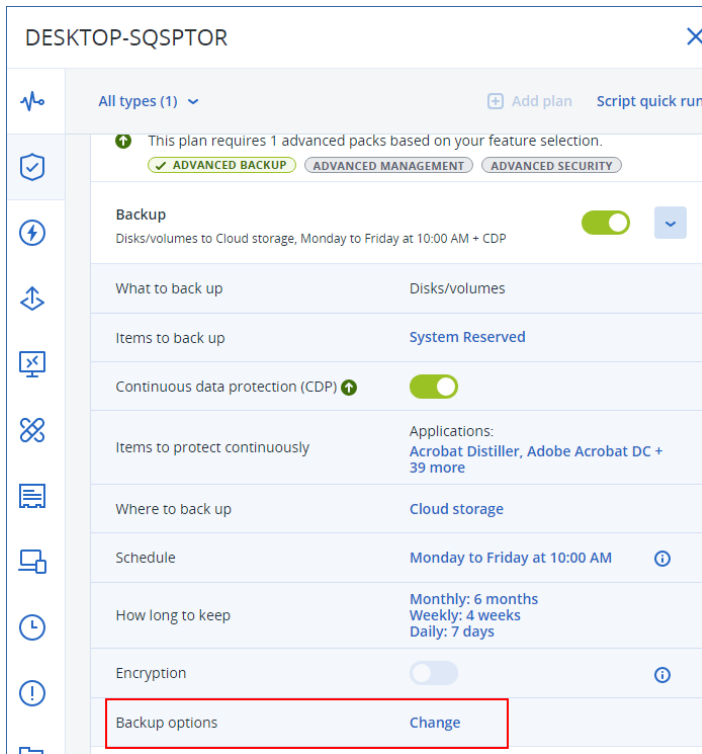
One-click recovery is a backup option in the protection plan. For more information on how to create a plan, see "Creating a protection plan" (p. 209).

Note

Enabling One-click recovery also enables Startup Recovery Manager on the target machine. If Startup Recovery Manager cannot be enabled, the backup operation that creates One-click recovery backups will fail. For more information about Startup Recovery Manager, see "Startup Recovery Manager" (p. 676).

To enable One-click recovery

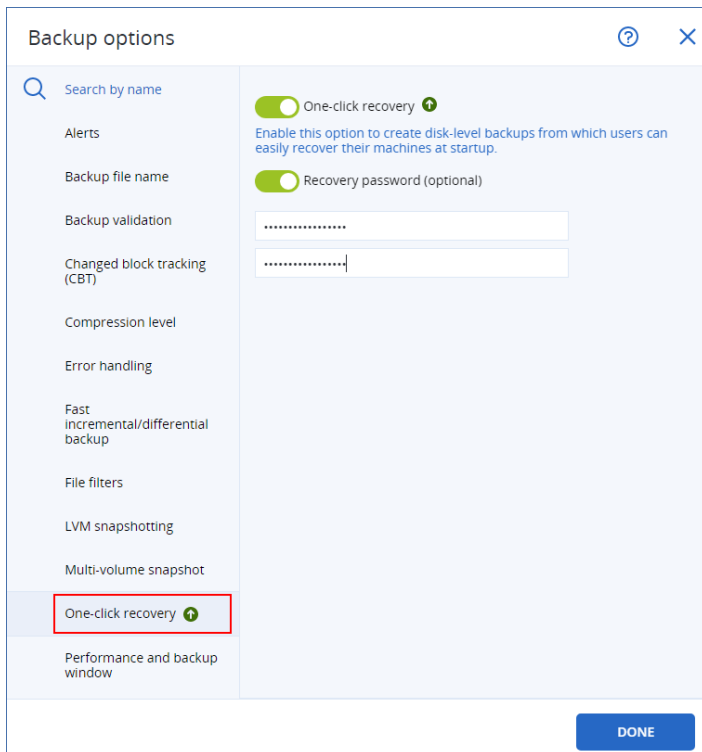
1. In the protection plan, expand the **Backup** module.
2. In **What to back up**, select **Entire machine** or **Disk/volumes**.
3. [If you selected **Disk/volumes**]. In **Items to back up**, specify the disk or volumes to back up.
4. In **Backup options**, click **Change**, and then select **One-click recovery**.



5. Enable the **One-click recovery** switch.
6. [Optional] Enable the **Recovery password** switch, and then specify a password.

Important

We strongly recommend that you specify a recovery password. Ensure that the user who performs One-click recovery on the target machine knows this password.



7. Click **Done**.
8. Configure the other elements of the protection plan according to your needs, and then save the plan.

As a result, after the protection plan runs and creates a backup, One-click recovery becomes accessible to the users of the protected machine.

Important

One-click recovery becomes temporarily unavailable after you update the protection agent. One backup after the agent update, one-click recovery becomes available again.

Disabling One-click recovery

You can disable One-click recovery for a specific workload in the following ways:

- Disable the **One-click recovery** option in the protection plan that is applied to the workload.
- Revoke the protection plan in which the **One-click recovery** option is enabled.
- Delete the protection plan in which the **One-click recovery** option is enabled.

Recovering a machine with One-click recovery

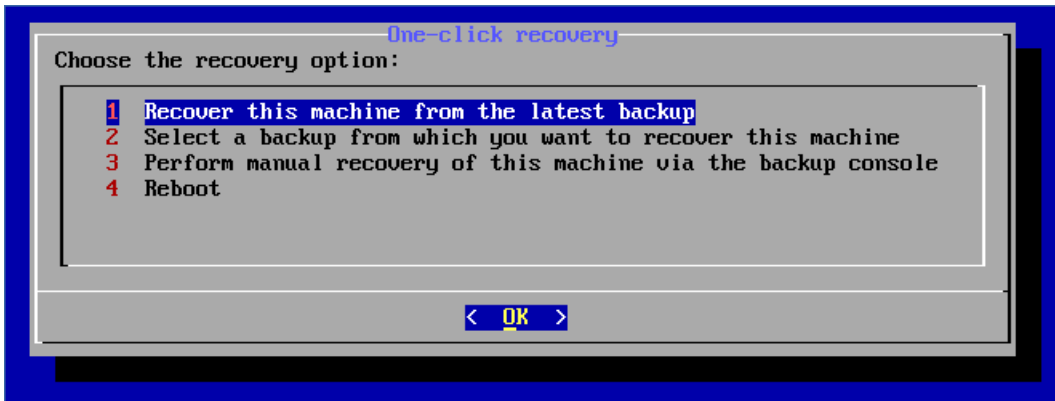
Prerequisites

- A protection plan with enabled **One-click recovery** backup option is applied to the machine.
- There is at least one disk backup of the machine.

To recover a machine

1. Reboot the machine that you want to recover.
2. During the reboot, press F11 to enter Startup Recovery Manager.
The rescue media window opens.
3. Select **Acronis Cyber Protect**.
4. [If a recovery password was specified in the protection plan] Enter the recovery password, and then click **OK**.
5. Select a One-click recovery option.
 - To recover the latest backup automatically, select the first option, and then click **OK**.
 - To recover another backup within the backup archive, select the second option, and then click

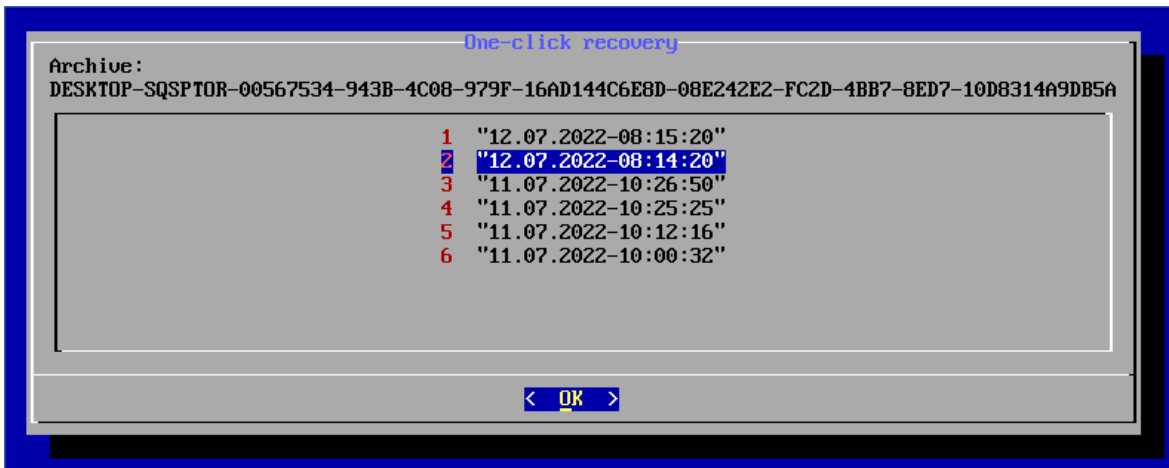
OK.



6. Confirm your choice by clicking **Yes**.

The rescue media window opens, and then disappears. The recovery procedure continues without it.

7. [If you chose to recover a specific backup] Select the backup that you want to recover, and then click **OK**.



After a while, the recovery starts and its progress is shown. When the recovery completes, your machine reboots.


```
One-click recovery
progress: 7%
elapsed time: 00:00:44
estimated time: 00:09:44
-----
progress: 8%
elapsed time: 00:00:48
estimated time: 00:09:11
-----
progress: 8%
elapsed time: 00:00:48
estimated time: 00:09:11
-----
progress: 9%
elapsed time: 00:00:53
estimated time: 00:08:55
-----
progress: 10%
elapsed time: 00:00:56
estimated time: 00:08:23
-----
progress: 10%
elapsed time: 00:01:00
estimated time: 00:08:59
-----
progress: 11%
elapsed time: 00:01:02
estimated time: 00:08:21
-----
```

Performance and backup window

This option enables you to set one of three levels of backup performance (high, low, prohibited) for every hour within a week. This way, you can define a time window when backups are allowed to start and run. The high and low performance levels are configurable in terms of the process priority and output speed.

This option is not available for backups executed by the cloud agents, such as website backups or backups of servers located on the cloud recovery site.

This option is effective only for the backup and backup replication processes. Post-backup commands and other operations included in a protection plan (for example, validation) will run regardless of this option.

The preset is: **Disabled**.

When this option is disabled, backups are allowed to run at any time, with the following parameters (no matter if the parameters were changed against the preset value):

- CPU priority: **Low** (in Windows, it corresponds to **Below normal**)
- Output speed: **Unlimited**

When this option is enabled, scheduled backups are allowed or blocked according to the performance parameters specified for the current hour. At the beginning of an hour when backups are blocked, a backup process is automatically stopped and an alert is generated. Even if scheduled

backups are blocked, a backup can be started manually. It will use the performance parameters of the most recent hour when backups were allowed.

Note

You can configure performance and backup window for every replication location individually. To access the settings of the replication location, in the protection plan, click the gear icon next to the location name, and then click **Performance and backup window**.

Backup window

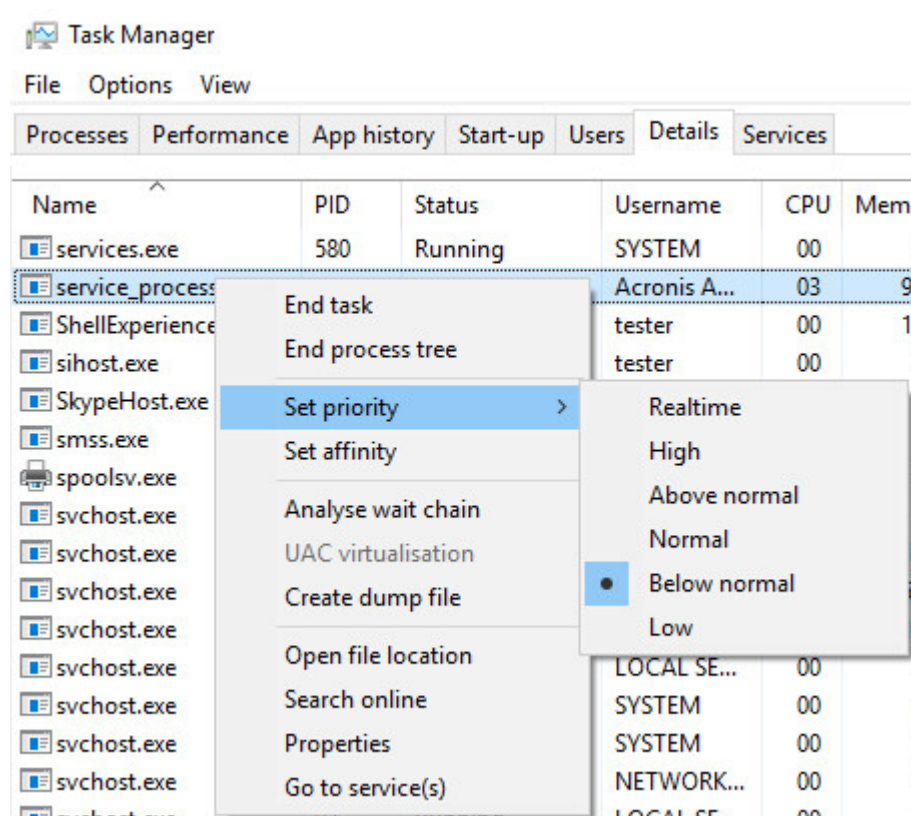
Each rectangle represents an hour within a week day. Click a rectangle to cycle through the following states:

- **Green:** backup is allowed with the parameters specified in the green section below.
- **Blue:** backup is allowed with the parameters specified in the blue section below.
This state is not available if the backup format is set to **Version 11**.
- **Gray:** backup is blocked.

You can click and drag to change the state of multiple rectangles simultaneously.

The priority of a process running in a system determines the amount of CPU and system resources allocated to that process. Decreasing the backup priority will free more resources for other applications. Increasing the backup priority might speed up the backup process by requesting the operating system to allocate more resources like the CPU to the backup application. However, the resulting effect will depend on the overall CPU usage and other factors like disk in/out speed or network traffic.

This option sets the priority of the backup process (**service_process.exe**) in Windows and the niceness of the backup process (**service_process**) in Linux and macOS.



The table below summarizes the mapping for this setting in Windows, Linux, and macOS.

Cyber Protection priority	Windows priority	Linux and macOS niceness
Low	Below normal	10
Normal	Normal	0
High	High	-10

Output speed during backup

This parameter enables you to limit the hard drive writing speed (when backing up to a local folder) or the speed of transferring the backup data through the network (when backing up to a network

share or to cloud storage).

When this option is enabled, you can specify the maximum allowed output speed:

- As a percentage of the estimated writing speed of the destination hard disk (when backing up to a local folder) or of the estimated maximum speed of the network connection (when backing up to a network share or cloud storage).

This setting works only if the agent is running in Windows.

- In KB/second (for all destinations).

Physical Data Shipping

This option is available if the backup or recovery destination is the cloud storage and the [backup format](#) is set to **Version 12**.

This option is effective for disk-level backups and file backups created by Agent for Windows, Agent for Linux, Agent for Mac, Agent for VMware, Agent for Hyper-V, and Agent for Virtuozzo.

Use this option to ship the first full backup created by a protection plan to the cloud storage on a hard disk drive by using the Physical Data Shipping service. The subsequent incremental backups are performed over the network.

For local backups that are replicated to cloud, incremental backups continue and are saved locally until the initial backup is uploaded in the cloud storage. Then all incremental changes are replicated to the cloud and the replication continues per the backup schedule.

The preset is: **Disabled**.

About the Physical Data Shipping service

The Physical Data Shipping service web interface is available only to administrators.

For detailed instructions about using the Physical Data Shipping service and the order creation tool, refer to the [Physical Data Shipping Administrator's Guide](#). To access this document in the Physical Data Shipping service web interface, click the question mark icon.

Overview of the physical data shipping process

1. [To ship backups that have cloud storage as the primary backup location]
 - a. Create a new protection plan with backup to cloud.
 - b. In the **Backup options** row, click **Change**.
 - c. In the list of available options, click **Physical Data Shipping**.

You can back up directly to a removable drive or back up to a local or a network folder, and then copy/move the backup(s) to the drive.

2. [To ship local backups that are replicated to cloud]

Note

This option is supported with protection agent version from release C21.06 or later.

- a. Create a new protection plan with backup to a local or network storage.
- b. Click **Add location** and select **Cloud storage**.
- c. In the **Cloud storage** location row, click the gear wheel and select **Physical Data Shipping**.
3. Under **Use Physical Data Shipping**, click **Yes** and **Done**.
The Encryption option is enabled automatically in the protection plan because all backups that are shipped must be encrypted.
4. In the **Encryption** row, click **Specify a password** and enter a password for encryption.
5. In the **Physical Data Shipping** row, select the removable drive where the initial backup will be saved.
6. Click **Create** to save the protection plan.
7. After the first backup is complete, use the Physical Data Shipping service web interface to download the order creation tool and create the order.
To access this web interface, log in to the management portal, click **Overview > Usage**, and then click **Manage service** under **Physical Data Shipping**.

Important

Once the initial full backup is done, the subsequent backups must be performed by the same protection plan. Another protection plan, even with the same parameters and for the same machine, will require another Physical Data Shipping cycle.

8. Package the drives and ship them to the data center.

Important

Ensure that you follow the packaging instructions provided in the [Physical Data Shipping Administrator's Guide](#).

9. Track the order status by using the Physical Data Shipping service web interface. Note that the subsequent backups will fail until the initial backup is uploaded to the cloud storage.

Pre/Post commands

The option enables you to define the commands to be automatically executed before and after the backup procedure.

The following scheme illustrates when pre/post commands are executed.

Pre-backup command	Backup	Post-backup command
--------------------	--------	---------------------

Examples of how you can use the pre/post commands:

- Delete some temporary files from the disk before starting backup.
- Configure a third-party antivirus product to be started each time before the backup starts.
- Selectively copy backups to another location. This option may be useful because the replication configured in a protection plan copies *every* backup to subsequent locations.

The agent performs the replication *after* executing the post-backup command.

The program does not support interactive commands, i.e. commands that require user input (for example, "pause").

Pre-backup command

To specify a command/batch file to be executed before the backup process starts

1. Enable the **Execute a command before the backup** switch.
2. In the **Command...** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the backup if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not back up until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the backup only after the command is successfully executed. Fail the backup if the command execution fails.	Perform the backup after the command is executed despite execution failure or success.	N/A	Perform the backup concurrently with the command execution and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

Note

If a script fails due to a conflict related to a required library version in Linux, exclude the LD_LIBRARY_PATH and LD_PRELOAD environmental variables, by adding the following lines in your script:

```
#!/bin/sh
unset LD_LIBRARY_PATH
unset LD_PRELOAD
```

Post-backup command

To specify a command/executable file to be executed after the backup is completed

1. Enable the **Execute a command after the backup** switch.
2. In the **Command...** field, type a command or browse to a batch file.
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field, specify the command execution arguments, if required.
5. Select the **Fail the backup if the command execution fails** check box if successful execution of the command is critical for you. The command is considered failed if its exit code is not equal to zero. If the command execution fails, the backup status will be set to **Error**.

When the check box is not selected, the command execution result does not affect the backup failure or success. You can track the command execution result by exploring the **Activities** tab.

6. Click **Done**.

Pre/Post data capture commands

The option enables you to define the commands to be automatically run before and after data capture (that is, taking the data snapshot). Data capture is performed at the beginning of the backup procedure.

The following scheme illustrates when the pre/post data capture commands are run.

	←----- Backup -----→				
Pre-backup command	Pre-data capture command	Data capture	Post-data capture command	Write data to the backup set	Post-backup command

Interaction with other backup options

Running of the pre/post data capture commands can be modified by other backup options.

If the **Multi-volume snapshot** option is enabled, the pre/post data capture commands will run only once, because the snapshots for all volumes are created simultaneously. If the **Multi-volume**

snapshot option is disabled, the pre/post data capture commands will run for every volume that is being backed up because the snapshots are created sequentially, one volume after another.

If the **Volume Shadow Copy Service (VSS)** option is enabled, the pre/post data capture commands and the Microsoft VSS actions will run as follows:

Pre-data capture commands > VSS Suspend > Data capture > VSS Resume > Post-data capture commands

By using the pre/post data capture commands, you can suspend and resume a database or application that is not compatible with VSS. Because the data capture takes seconds, the database or application idle time will be minimal.

Pre-data capture command

To specify a command/batch file to be executed before data capture

1. Enable the **Execute a command before the data capture** switch.
2. In the **Command...** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the backup if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not perform the data capture until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the data capture only after the command is successfully executed. Fail the backup if the command execution fails.	Perform the data capture after the command is executed despite execution failure or success.	N/A	Perform the data capture concurrently with the command and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

Note

If a script fails due to a conflict related to a required library version in Linux, exclude the LD_LIBRARY_PATH and LD_PRELOAD environmental variables, by adding the following lines in your script:

```
#!/bin/sh
unset LD_LIBRARY_PATH
unset LD_PRELOAD
```

Post-data capture command

To specify a command/batch file to be executed after data capture

1. Enable the **Execute a command after the data capture** switch.
2. In the **Command...** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the backup if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not back up until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Continue the backup only after the command is successfully executed.	Continue the backup after the command is executed despite command execution failure or success.	N/A	Continue the backup concurrently with the command execution and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

Scheduling

This option defines whether backups start exactly as scheduled or with a delay, and how many virtual machines are backed up simultaneously.

For more information about how to configure the backup schedule, see "Running a backup on a schedule" (p. 397).

The preset is: **Distribute backup start times within a time window. Maximum delay: 30 minutes.**

You can select one of the following:

- **Start all backups exactly as scheduled**

Backups of physical machines will start exactly as scheduled. Virtual machines will be backed up one by one.

- **Distribute start times within a time window**

Backups of physical machines will start with a delay from the scheduled time. The delay value for each machine is selected randomly and ranges from zero to the maximum value you specify. You may want to use this setting when backing up multiple machines to a network location, to avoid excessive network load. The delay value for each machine is determined when the protection plan is applied to the machine and remains the same until you edit the protection plan and change the maximum delay value.

Virtual machines will be backed up one by one.

- **Limit the number of simultaneously running backups by**

Use this option to manage the parallel backup of virtual machines that are backed up on the hypervisor level (agentless backup).

Protection plans in which this option is selected can run together with other protection plans that are operated by the same agent at the same time. When you select this option, you must specify the number of parallel backups per plan. The total number of machines that are backed up simultaneously by all plans is limited to 10 per agent. To learn how to change the default limit, see "Limiting the total number of simultaneously backed-up virtual machines" (p. 650).

Protection plans in which this option is not selected run the backup operations sequentially, one virtual machine after another.

Sector-by-sector backup

The option is effective only for disk-level backup.

This option defines whether an exact copy of a disk or volume on a physical level is created.

The preset is: **Disabled.**

If this option is enabled, all disk or volume's sectors will be backed up, including unallocated space and those sectors that are free of data. The resulting backup will be equal in size to the disk being

backed up (if the "Compression level" option is set to **None**). The software automatically switches to the sector-by-sector mode when backing up drives with unrecognized or unsupported file systems.

Note

It will be impossible to perform a recovery of application data from the backups which were created in the sector-by-sector mode.

Splitting

This option enables you to select the method of splitting of large backups into smaller files.

Note

Splitting is not available in protection plans that use the cloud storage as a backup location.

The preset is:

- If the backup location is a local or network (SMB) folder, and the backup format is Version 12:
Fixed size - 200 GB
This setting allows the backup software to work with large volumes of data on the NTFS file system, without negative effects caused by file fragmentation.
- Otherwise: **Automatic**

The following settings are available:

- **Automatic**
A backup will be split if it exceeds the maximum file size supported by the file system.
- **Fixed size**
Enter the desired file size or select it from the drop-down list.

Task failure handling

This option determines the program behavior when a scheduled execution of a protection plan fails or your machine restarts while a backup is running. This option is not effective when a protection plan is started manually.

If this option is enabled, the program will try to execute the protection plan again. You can specify the number of attempts and the time interval between the attempts. The program stops trying as soon as an attempt completes successfully or the specified number of attempts is performed, depending on which comes first.

If this option is enabled and your machine restarts while a backup is running, the backup operation will not fail. A few minutes after the restart, the backup operation will continue automatically and complete the backup file with the missing data. In this use case, the option **Interval between attempts** is not relevant.

The preset is: **Enabled**.

Note

This option is not effective in forensic backups.

Task start conditions

This option is effective in Windows and Linux operating systems.

This option determines the program behavior in case a task is about to start (the scheduled time comes or the event specified in the schedule occurs), but the condition (or any of multiple conditions) is not met. For more information about conditions refer to "Start conditions" (p. 404).

The preset is: **Wait until the conditions from the schedule are met.**

Wait until the conditions from the schedule are met

With this setting, the scheduler starts monitoring the conditions and launches the task as soon as the conditions are met. If the conditions are never met, the task will never start.

To handle the situation when the conditions are not met for too long and further delaying the task is becoming risky, you can set the time interval after which the task will run irrespective of the condition. Select the **Run the task anyway after** check box and specify the time interval. The task will start as soon as the conditions are met OR the maximum time delay lapses, depending on which comes first.

Skip the task execution

Delaying a task might be unacceptable, for example, when you need to execute a task strictly at the specified time. Then it makes sense to skip the task rather than wait for the conditions, especially if the tasks occur relatively often.

Volume Shadow Copy Service (VSS)

This option is applicable only to Windows operating systems.

It defines whether a backup can succeed if one or more Volume Shadow Copy Service (VSS) writers fail and which provider has to notify the VSS-aware applications that the backup will start.

Using the Volume Shadow Copy Service ensures the consistent state of all data used by the applications; in particular, completion of all database transactions at the moment of taking the data snapshot by the backup software. Data consistency, in turn, ensures that the application will be recovered in the correct state and become operational immediately after recovery.

The snapshot is used only during the backup operation, and is automatically deleted when the backup operation completes. No temporary files are kept.

You may also use [Pre/Post data capture commands](#) to ensure that the data is backed up in a consistent state. For instance, specify pre-data capture commands that will suspend the database

and flush all caches to ensure that all transactions are completed, and then specify post-data capture commands that will resume the database operations after the snapshot is taken.

Note

Files and folders that are specified in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot** registry key are not backed up. In particular, offline Outlook Data Files (.ost) are not backed up because they are specified in the **OutlookOST** value of this key.

Ignore failed VSS writers

You can select one of the following:

- **Ignore failed VSS writers**

With this option, you can achieve successful backups even when one or more VSS writers fail.

Important

Application-aware backups will always fail if the application-specific writer fails. For example, if you are making application-aware backup of SQL Server data, and **SqlServerWriter** fails, the backup operation will also fail.

When this option is enabled, up to three consecutive attempts for a VSS snapshot will be made. In the first attempt, all VSS writers are required. If this attempt fails, it will be repeated. If the second attempt also fails, the failed VSS writers will be excluded from the scope of the backup operation, and then a third attempt will be made. If the third attempt is successful, the backup will complete with a warning about the failed VSS writers. If the third attempt is not successful, the backup will fail.

Note

If the failed VSS writers are not essential for the consistency of your backups, and you want to remove the warnings, you can permanently exclude these writers from the scope of the backup operation. For more information on how to exclude a VSS writer, refer to [this knowledge base article](#).

- **Require successful processing for all VSS writers**

If any of the VSS writers fails, the backup operation will also fail.

Select the snapshot provider

You can select one of the following:

- **Automatically select snapshot provider**

Automatically select among the hardware snapshot provider, software snapshot providers, and Microsoft Software Shadow Copy provider.

- **Use Microsoft Software Shadow Copy provider**

We recommend that you choose this option when backing up application servers (Microsoft Exchange Server, Microsoft SQL Server, Microsoft SharePoint, or Active Directory).

Enable VSS full backup

If this option is enabled, the logs of Microsoft Exchange Server and of other VSS-aware applications (except for Microsoft SQL Server) will be truncated after each successful full, incremental or differential disk-level backup.

The preset is: **Disabled**.

Leave this option disabled in the following cases:

- If you use Agent for Exchange or third-party software for backing up the Exchange Server data. This is because the log truncation will interfere with the consecutive transaction log backups.
- If you use third-party software for backing up the SQL Server data. The reason for this is that the third-party software will take the resulting disk-level backup for its "own" full backup. As a result, the next differential backup of the SQL Server data will fail. The backups will continue failing until the third-party software creates the next "own" full backup.
- If other VSS-aware applications are running on the machine and you need to keep their logs for any reason.

Important

Enabling this option does not result in the truncation of Microsoft SQL Server logs. To truncate the SQL Server log after a backup, enable the [Log truncation](#) backup option.

Volume Shadow Copy Service (VSS) for virtual machines

This option defines whether quiesced snapshots of virtual machines are taken.

The preset is: **Enabled**.

When this option is disabled, a non-quiesced snapshot is taken. The virtual machine will be backed up in a crash-consistent state.

When this option is enabled, the transactions of all VSS-aware applications running in the virtual machine are completed, and then a quiesced snapshot is taken.

If a quiesced snapshot cannot be taken after the number of re-attempts specified in the "[Error handling](#)" option, and application backup is enabled, the backup fails.

If a quiesced snapshot cannot be taken after the number of re-attempts specified in the "[Error handling](#)" option, and application backup is disabled, a crash-consistent backup is created. To make the backup fail instead of creating a crash-consistent backup, select the **Fail backup if taking a quiesced snapshot is not possible** check box.

The following table summarizes the available settings and their results.

Settings	Quiesced snapshot was taken successfully		Quiesced snapshot was not taken	
	Application backup enabled	Application backup disabled	Application backup enabled	Application backup disabled
Volume Shadow Copy Service (VSS) for virtual machines enabled Fail backup if taking a quiesced snapshot is not possible not selected	Quiesced snapshot is taken. Application-consistent backup is created.	Quiesced snapshot is taken. Application-consistent backup is created.	Backup fails.	Non-quiesced snapshot is taken. Crash-consistent backup is created.
Volume Shadow Copy Service (VSS) for virtual machines enabled Fail backup if taking a quiesced snapshot is not possible selected	Quiesced snapshot is taken. Application-consistent backup is created.	Quiesced snapshot is taken. Application-consistent backup is created.	Backup fails.	Backup fails.
Volume Shadow Copy Service (VSS) for virtual machines disabled	Non-quiesced snapshot is taken. Crash-consistent backup is created.	Non-quiesced snapshot is taken. Crash-consistent backup is created.	Non-quiesced snapshot is taken. Crash-consistent backup is created.	Non-quiesced snapshot is taken. Crash-consistent backup is created.

Enabling **Volume Shadow Copy Service (VSS) for virtual machines** also triggers the pre-freeze and post-thaw scripts that you might have on the backed-up virtual machine. For more information on these scripts, refer to "Running pre-freeze and post-thaw scripts automatically" (p. 644).

To take a quiesced snapshot, the backup software applies VSS inside a virtual machine by using VMware Tools, Hyper-V Integration Services, Virtuozzo Guest Tools, Red Hat Virtualization Guest Tools, or QEMU Guest Tools, respectively.

Note

For Red Hat Virtualization (oVirt) virtual machines, we recommend that you install QEMU Guest Tools instead of Red Hat Virtualization Guest Tools. Some versions of Red Hat Virtualization Guest Tools do not support application-consistent snapshots.

This option does not affect Scale Computing HC3 virtual machines. For them, quiescing depends on whether Scale Tools are installed on the virtual machine.

Weekly backup

This option determines which backups are considered "weekly" in retention rules and backup schemes. A "weekly" backup is the first backup created after a week starts.

The preset is: **Monday**.

Windows event log

This option is effective only in Windows operating systems.

This option defines whether the agents have to log events of the backup operations in the Application Event Log of Windows (to see this log, run eventvwr.exe or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

The preset is: **Disabled**.

Recovery

Recovery cheat sheet

The following table summarizes the available recovery methods. Use the table to choose a recovery method that best fits your need.

Note

You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1012).

What to recover	Recovery method
Physical machine (Windows or Linux)	Using the Cyber Protect console Using bootable media
Physical machine (Mac)	Using bootable media
Virtual machine (VMware, Hyper-V, Red Hat Virtualization (oVirt), or Scale Computing HC3)	Using the Cyber Protect console Using bootable media
Virtual machine or container (Virtuozzo, Virtuozzo Hybrid Server, or Virtuozzo Hybrid Infrastructure)	Using the Cyber Protect console

ESXi configuration	Using bootable media
Files/Folders	Using the Cyber Protect console Downloading files from the cloud storage Using bootable media Extracting files from local backups
System state	Using the Cyber Protect console
SQL databases	Using the Cyber Protect console
Exchange databases	Using the Cyber Protect console
Exchange mailboxes	Using the Cyber Protect console
Websites	Using the Cyber Protect console
Microsoft 365	
Mailboxes (local Agent for Microsoft 365)	Using the Cyber Protect console
Mailboxes (cloud Agent for Microsoft 365)	Using the Cyber Protect console
Public folders	Using the Cyber Protect console
OneDrive files	Using the Cyber Protect console
SharePoint Online data	Using the Cyber Protect console
Google Workspace	
Mailboxes	Using the Cyber Protect console
Google Drive files	Using the Cyber Protect console
Shared drive files	Using the Cyber Protect console

Cross-platform recovery

Cross-platform recovery is available for backups of entire machines and backups of disks that contain an operating system.

A cross-platform recovery is performed in the following cases:

- A backup is created by one type of agent but it is recovered by another type of agent.
- An agent-based backup is recovered on the hypervisor level (agentless recovery), or an agentless backup is recovered by an agent (agent-based recovery).
- A backup is recovered to dissimilar hardware (including virtual hardware).

Note

Some peripheral devices, such as printers, might not be recovered correctly when you perform a cross-platform recovery.

The table below shows a few examples of cross-platform recovery.

Cross-platform recovery	
Agentless backup	Agent-based recovery
Agent-based backup	Agentless recovery
Backup by Agent for Windows	Recovery by Agent for VMware
Backup by Agent for VMware	Recovery by Agent for Hyper-V
Backup by Agent for Windows that is installed on a VMware ESXi virtual machine (agent-based)	Recovery by Agent for VMware (agentless) on the same VMware ESXi host
Backup by Agent for Windows	Recovery by Agent for Windows that is installed on a machine with dissimilar hardware
Backup of a physical machine	Recovery as a virtual machine

Note for Mac users

- Starting with 10.11 El Capitan, certain system files, folders, and processes are flagged for protection with an extended file attribute `com.apple.rootless`. This feature is called System Integrity Protection (SIP). The protected files include preinstalled applications and most of the folders in `/system`, `/bin`, `/sbin`, `/usr`.
The protected files and folders cannot be overwritten during a recovery under the operating system. If you need to overwrite the protected files, perform the recovery under bootable media.
- Starting with macOS Sierra 10.12, rarely used files can be moved to iCloud by the Store in Cloud feature. Small footprints of these files are kept on the file system. These footprints are backed up instead of the original files.
When you recover a footprint to the original location, it is synchronized with iCloud and the original file becomes available. When you recover a footprint to a different location, it cannot be synchronized and the original file will be unavailable.

Safe recovery

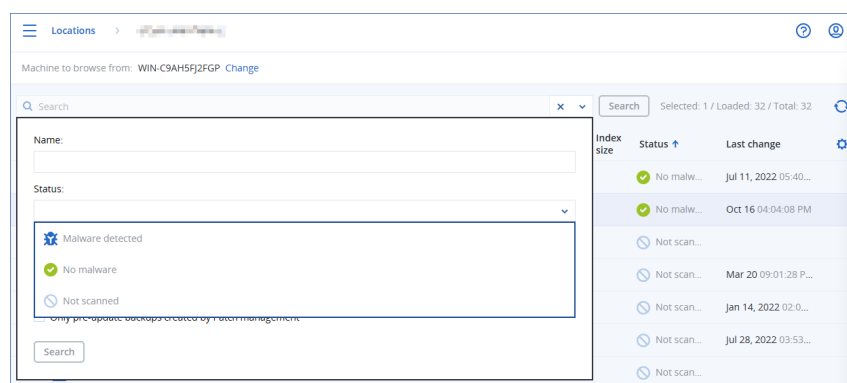
Use safe recovery with **Entire machine** or **Disks/volumes** backups of Windows workloads to ensure that you recover only malware-free data, even if the backup contains infected files.

During a safe recovery operation, the backup is automatically scanned for malware. Then, the protection agent recovers the backup on the target workload and deletes any infected files. As a result, a malware-free backup is recovered.

Additionally, one of the following statuses is assigned to the backup:

- Malware detected
- No malware
- Not scanned

You can use the status to filter the backup archives.



Limitations

- Safe recovery is supported for physical and virtual Windows machines on which a protection agent is installed.
- Safe recovery is supported for **Entire machine** and **Disks/volumes** backups.
- Only NTFS volumes are scanned for malware. Non-NTFS volumes are recovered without antimalware scanning.
- Safe recovery is not supported for the Continuous data protection (CDP) backup in the archive. To recover the data from the CDP backup, run an additional **Files/folders** recovery operation. For more operation about the CDP backups, see "Continuous data protection (CDP)" (p. 386).

Recovering a machine

Recovering physical machines

This section describes recovery of physical machines by using the web interface.

Use bootable media instead of the web interface if you need to recover:

- A machine running macOS
- A machine from a tenant in the Enhanced security mode
- Any operating system to bare metal or to an offline machine
- The structure of logical volumes (volumes created by Logical Volume Manager in Linux). The media enables you to recreate the logical volume structure automatically.

Note

You cannot recover disk-level backups of Intel-based Macs to Macs that use Apple silicon processors, and vice-versa. You can recover files and folders.

Recovery with restart

Recovery of an operating system and recovery of volumes that are encrypted with BitLocker requires a restart. You can choose whether to restart the machine automatically or assign it the **Interaction required** status. The recovered operating system goes online automatically.

Important

Backed-up encrypted volumes are recovered as non-encrypted.

Recovery of BitLocker-encrypted volumes requires that there is a non-encrypted volume on the same machine, and that this volume has at least 1 GB of free space. If either condition is not met, the recovery fails.

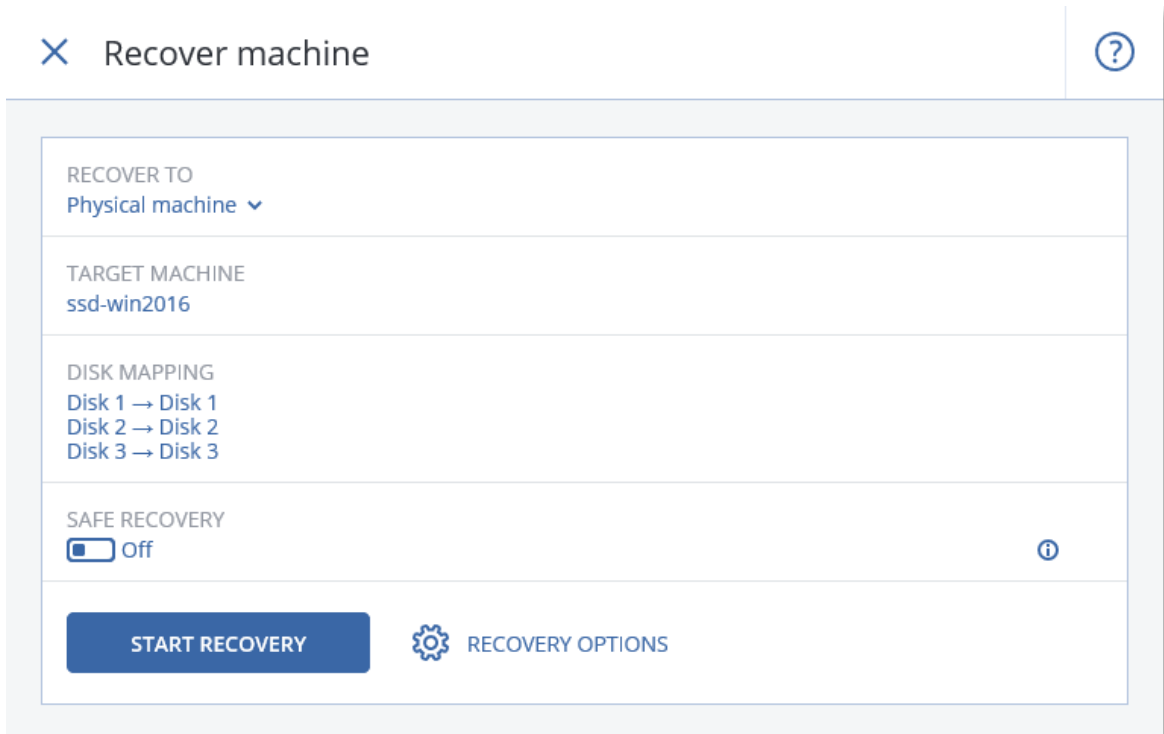
Recovering an encrypted system volume does not require any additional actions. To recover an encrypted non-system volume, you must lock it first, for example, by opening a file that resides on this volume. Otherwise, the recovery will continue without restart and the recovered volume might not be recognized by Windows.

Note

If the recovery fails and your machine restarts with the `Cannot get file from partition error`, try disabling Secure Boot. For more information on how to do it, refer to [Disabling Secure Boot](#) in the Microsoft documentation.

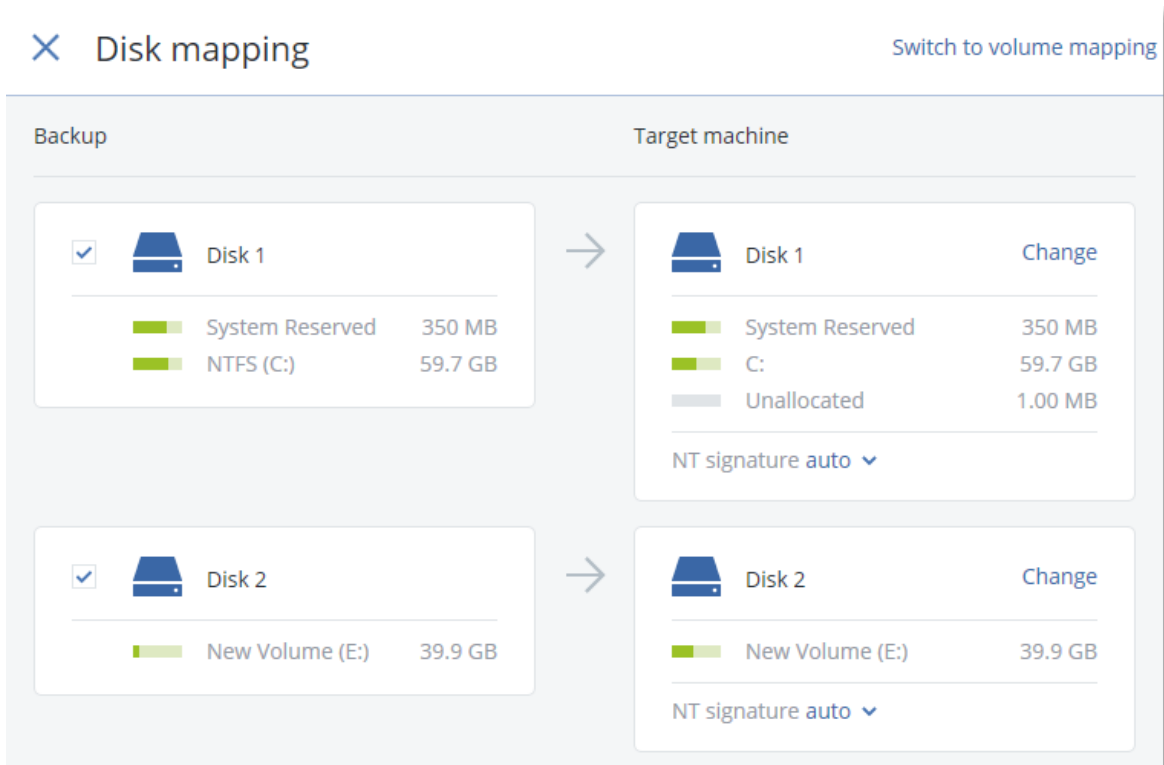
To recover a physical machine

1. Select the backed-up machine.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Do any of the following:
 - If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select a target machine that is online, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).
 - Recover the machine as described in "[Recovering disks by using bootable media](#)".
4. Click **Recover > Entire machine**.
The software automatically maps the disks from the backup to the disks of the target machine. To recover to another physical machine, click **Target machine**, and then select a target machine that is online.



5. If you are unsatisfied with the mapping result or if the disk mapping fails, click **Volume mapping** to re-map the disks manually.

The mapping section also enables you to choose individual disks or volumes for recovery. You can switch between recovering disks and volumes by using the **Switch to...** link in the upper-right corner.



6. [Only available for Windows machines on which a protection agent is installed] Enable the **Safe recovery** switch to ensure that the recovered data is malware-free. For more information about how safe recovery works, see "Safe recovery" (p. 471).
7. Click **Start recovery**.
8. Confirm that you want to overwrite the disks with their backed-up versions. Choose whether to restart the machine automatically.

The recovery progress is shown on the **Activities** tab.

Physical machine to virtual

You can recover a physical machine to a virtual machine on one of the supported hypervisors. This is also a mechanism to migrate a physical machine to a virtual machine. For more information about supported P2V migration paths, refer to "[Machine migration](#)".

This section describes the recovery of a physical machine as a virtual machine by using the web interface. This operation can be performed if at least one agent for the relevant hypervisor is installed and registered in Acronis Management Server. For example, recovery to VMware ESXi requires at least one Agent for VMware, recovery to Hyper-V requires at least one Agent for Hyper-V installed and registered in the environment.

Recovery through the web interface is not available for tenants in the Enhanced security mode.

Note

You cannot recover macOS virtual machines to Hyper-V hosts, because Hyper-V does not support macOS. You can recover macOS virtual machines to a VMware host that is installed on Mac hardware.

Also, you cannot recover backups of macOS physical machines as virtual machines.

To recover a physical machine as a virtual machine

1. Select the backed-up machine.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Do any of the following:
 - If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select a machine that is online, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).
 - Recover the machine as described in "[Recovering disks by using bootable media](#)".
4. Click **Recover > Entire machine**.
5. In **Recover to**, select **Virtual machine**.
6. Click **Target machine**.

- a. Select the hypervisor.

Note

At least one agent for that hypervisor must be installed and registered in Acronis Management Server.


- b. Select whether to recover to a new or existing machine. The new machine option is preferable as it does not require the disk configuration of the target machine to exactly match the disk configuration in the backup.
 - c. Select the host and specify the new machine name, or select an existing target machine.
 - d. Click **OK**.
7. [For Virtuozzo Hybrid Infrastructure] Click **VM settings** to select **Flavor**. Optionally, you can change the memory size, the number of processors, and the network connections of the virtual machine.

Note

Selecting flavor is a required step for Virtuozzo Hybrid Infrastructure.

8. [Optional] Configure additional recovery options:
 - [Not available for Virtuozzo Hybrid Infrastructure] Click **Datastore** for ESXi or **Path** for Hyper-V, and then select the datastore (storage) for the virtual machine.
 - Click **Disk mapping** to select the datastore (storage), interface, and provisioning mode for each virtual disk. The mapping section also enables you to choose individual disks for recovery.

For Virtuozzo Hybrid Infrastructure, you can only select the storage policy for the target disks. To do so, select the desired target disk, and then click Change. In the blade that opens, click the gear icon, select the storage policy, and then click Done.
 - [For VMware ESXi, Hyper-V, and Red Hat Virtualization/oVirt] Click **VM settings** to change the memory size, the number of processors, and the network connections of the virtual machine.

RECOVER TO Virtual machine
TARGET MACHINE New machine on 10.250.22.17 New
DATASTORE datastore1 (1)
DISK MAPPING Disk 1 → datastore1 (1), 50.0 GB Disk 2 → datastore1 (1), 50.0 GB
VM SETTINGS Memory: 2.00 GB Virtual processors: 2 Network adapters: 2
START RECOVERY  RECOVERY OPTIONS

9. [Only available for Windows machines on which a protection agent is installed] Enable the **Safe recovery** switch to ensure that the recovered data is malware-free. For more information about how safe recovery works, see "Safe recovery" (p. 471).
10. Click **Start recovery**.
11. When recovering to an existing virtual machine, confirm that you want to overwrite the disks.

The recovery progress is shown on the **Activities** tab.

Recovering a virtual machine

You can recover virtual machines from their backups.

Note

You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1012).

Prerequisites

- A virtual machine must be stopped during the recovery to this machine. By default, the software stops the machine without a prompt. When the recovery is completed, you have to start the

machine manually. You can change the default behavior by using the VM power management recovery option (click **Recovery options > VM power management**).

Procedure

1. Do one of the following:
 - Select a backed-up machine, click **Recovery**, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).
2. Click **Recover > Entire machine**.
3. If you want to recover to a physical machine, select **Physical machine** in **Recover to**. Otherwise, skip this step.

Recovery to a physical machine is possible only if the disk configuration of the target machine exactly matches the disk configuration in the backup.


If this is the case, continue to step 4 in "[Physical machine](#)". Otherwise, we recommend that you perform the V2P migration by [using bootable media](#).

4. [Optional] By default, the software automatically selects the original machine as the target machine. To recover to another virtual machine, click **Target machine**, and then do the following:
 - a. Select the hypervisor (**VMware ESXi, Hyper-V, Virtuozzo, Virtuozzo Hybrid Infrastructure, Scale Computing HC3, or oVirt**).

Only Virtuozzo virtual machines can be recovered to Virtuozzo. For more information about V2V migration, refer to "[Machine migration](#)".
 - b. Select whether to recover to a new or existing machine.
 - c. Select the host and specify the new machine name, or select an existing target machine.
 - d. Click **OK**.
5. Setup up the additional recovery options that you need.
 - [Optional] [Not available for Virtuozzo Hybrid Infrastructure and Scale Computing HC3] To select the datastore for the virtual machine, click **Datastore** for ESXi, **Path** for Hyper-V and Virtuozzo, or **Storage domain** for Red Hat Virtualization (oVirt), and then select the datastore (storage) for the virtual machine.
 - [Optional] To view the datastore (storage), interface, and the provisioning mode for each virtual disk, click **Disk mapping**. You can change these settings, unless you are recovering a Virtuozzo container or Virtuozzo Hybrid Infrastructure virtual machine.

For Virtuozzo Hybrid Infrastructure, you can only select the storage policy for the target disks. To do so, select the desired target disk, and then click **Change**. In the blade that opens, click the gear icon, select the storage policy, and then click **Done**.

The mapping section also enables you to choose individual disks for recovery.
 - [Optional] [Available for VMware ESXi, Hyper-V, and Virtuozzo] To change the memory size, the number of processors, and the network connections of the virtual machine, click **VM settings**.
 - [For Virtuozzo Hybrid Infrastructure] To change the memory size and the number of processors of the virtual machine, select **Flavor**.

<p>RECOVER TO Virtual machine</p>
<p>TARGET MACHINE New machine on 10.250.22.17 New</p>
<p>DATASTORE datastore1 (1)</p>
<p>DISK MAPPING Disk 1 → datastore1 (1), 50.0 GB Disk 2 → datastore1 (1), 50.0 GB</p>
<p>VM SETTINGS Memory: 2.00 GB Virtual processors: 2 Network adapters: 2</p>
<p>START RECOVERY  RECOVERY OPTIONS</p>

- [Only available for Windows machines on which a protection agent is installed] Enable the **Safe recovery** switch to ensure that the recovered data is malware-free. For more information about how safe recovery works, see "Safe recovery" (p. 471).
- Click **Start recovery**.
- When recovering to an existing virtual machine, confirm that you want to overwrite the disks. The recovery progress is shown on the **Activities** tab.

Recovering disks by using bootable media

For information about how to create bootable media, refer to "Creating physical bootable media" (p. 657).

Note

You cannot recover disk-level backups of Intel-based Macs to Macs that use Apple silicon processors, and vice-versa. You can recover files and folders.

To recover disks by using bootable media

- Boot the target machine by using bootable media.
- [Only when recovering a Mac] If you are recovering APFS-formatted disks/volumes to a non-original machine or to bare metal, re-create the original disk configuration manually:

- a. Click **Disk Utility**.
 - b. Erase and format the target disk into APFS. For instructions, refer to <https://support.apple.com/en-us/HT208496#erasedisk>.
 - c. Re-create the original disk configuration. For instructions, refer to <https://support.apple.com/guide/disk-utility/add-erase-or-delete-apfs-volumes-dskua9e6a110/19.0/mac/10.15>.
 - d. Click **Disk Utility** > **Quit Disk Utility**.
3. Click **Manage this machine locally** or click **Rescue Bootable Media** twice, depending on the media type you are using.
 4. If a proxy server is enabled in your network, click **Tools** > **Proxy server**, and then specify the proxy server host name/IP address, port, and credentials. Otherwise, skip this step.
 5. [Optional] When recovering Windows or Linux, click **Tools** > **Register media in the Cyber Protection service**, and then specify the registration token that you obtained when downloading the media. If you do this, you will not need to enter credentials or a registration code to access the cloud storage, as described in step 8.
 6. On the welcome screen, click **Recover**.
 7. Click **Select data**, and then click **Browse**.
 8. Specify the backup location:
 - To recover from cloud storage, select **Cloud storage**. Enter the credentials of the account to which the backed up machine is assigned.
When recovering Windows or Linux, you have the option to request a registration code and use it instead of the credentials. Click **Use registration code** > **Request the code**. The software shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. The registration code is valid for one hour.
 - To recover from a local or a network folder, browse to the folder under **Local folders** or **Network folders**.Click **OK** to confirm your selection.
 9. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
 10. In **Backup contents**, select the disks that you want to recover. Click **OK** to confirm your selection.
 11. Under **Where to recover**, the software automatically maps the selected disks to the target disks. If the mapping is not successful or if you are unsatisfied with the mapping result, you can re-map disks manually.

Note

Changing disk layout may affect the operating system bootability. Please use the original machine's disk layout unless you feel fully confident of success.

12. [When recovering Linux] If the backed-up machine had logical volumes (LVM) and you want to reproduce the original LVM structure:
 - a. Ensure that the number of the target machine disks and each disk capacity are equal to or exceed those of the original machine, and then click **Apply RAID/LVM**.
 - b. Review the volume structure, and then click **Apply RAID/LVM** to create it.
13. [Optional] Click **Recovery options** to specify additional settings.
14. Click **OK** to start the recovery.

Using Universal Restore

The most recent operating systems remain bootable when recovered to dissimilar hardware, including the VMware or Hyper-V platforms. If a recovered operating system does not boot, use the Universal Restore tool to update the drivers and modules that are critical for the operating system startup.

Universal Restore is applicable to Windows and Linux.

To apply Universal Restore

1. Boot the machine from the bootable media.
2. Click **Apply Universal Restore**.
3. If there are multiple operating systems on the machine, choose the one to apply Universal Restore to.
4. [For Windows only] [Configure the additional settings](#).
5. Click **OK**.

Universal Restore in Windows

Preparation

Prepare drivers

Before applying Universal Restore to a Windows operating system, make sure that you have the drivers for the new HDD controller and the chipset. These drivers are critical to start the operating system. Use the CD or DVD supplied by the hardware vendor or download the drivers from the vendor's website. The driver files should have the *.inf extension. If you download the drivers in the *.exe, *.cab or *.zip format, extract them using a third-party application.

The best practice is to store drivers for all the hardware used in your organization in a single repository sorted by device type or by the hardware configurations. You can keep a copy of the repository on a DVD or a flash drive; pick some drivers and add them to the bootable media; create the custom bootable media with the necessary drivers (and the necessary network configuration) for each of your servers. Or, you can simply specify the path to the repository every time Universal Restore is used.

Check access to the drivers in bootable environment

Make sure you have access to the device with drivers when working under bootable media. Use WinPE-based media if the device is available in Windows but Linux-based media does not detect it.

Universal Restore settings

Automatic driver search

Specify where the program will search for the Hardware Abstraction Layer (HAL), HDD controller driver and network adapter driver(s):

- If the drivers are on a vendor's disc or other removable media, turn on the **Search removable media**.
- If the drivers are located in a networked folder or on the bootable media, specify the path to the folder by clicking **Add folder**.

In addition, Universal Restore will search the Windows default driver storage folder. Its location is determined in the registry value **DevicePath**, which can be found in the registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion**. This storage folder is usually `WINDOWS/inf`.

Universal Restore will perform the recursive search in all the sub-folders of the specified folder, find the most suitable HAL and HDD controller drivers of all those available, and install them into the system. Universal Restore also searches for the network adapter driver; the path to the found driver is then transmitted by Universal Restore to the operating system. If the hardware has multiple network interface cards, Universal Restore will try to configure all the cards' drivers.

Mass storage drivers to install anyway

You need this setting if:

- The hardware has a specific mass storage controller such as RAID (especially NVIDIA RAID) or a fibre channel adapter.
- You migrated a system to a virtual machine that uses a SCSI hard drive controller. Use SCSI drivers bundled with your virtualization software or download the latest drivers versions from the software manufacturer website.
- If the automatic drivers search does not help to boot the system.

Specify the appropriate drivers by clicking **Add driver**. The drivers defined here will be installed, with appropriate warnings, even if the program finds a better driver.

Universal Restore process

After you have specified the required settings, click **OK**.

If Universal Restore cannot find a compatible driver in the specified locations, it will display a prompt about the problem device. Do one of the following:

- Add the driver to any of the previously specified locations and click **Retry**.
- If you do not remember the location, click **Ignore** to continue the process. If the result is not satisfactory, reapply Universal Restore. When configuring the operation, specify the necessary driver.

Once Windows boots, it will initialize the standard procedure for installing new hardware. The network adapter driver will be installed silently if the driver has the Microsoft Windows signature. Otherwise, Windows will ask for confirmation on whether to install the unsigned driver.

After that, you will be able to configure the network connection and specify drivers for the video adapter, USB and other devices.

Universal Restore in Linux

Universal Restore can be applied to Linux operating systems with a kernel version of 2.6.8 or later.

When Universal Restore is applied to a Linux operating system, it updates a temporary file system known as the initial RAM disk (initrd). This ensures that the operating system can boot on the new hardware.

Universal Restore adds modules for the new hardware (including device drivers) to the initial RAM disk. As a rule, it finds the necessary modules in the **/lib/modules** directory. If Universal Restore cannot find a module it needs, it records the module's file name into the log.

Universal Restore may modify the configuration of the GRUB boot loader. This may be required, for example, to ensure the system bootability when the new machine has a different volume layout than the original machine.

Universal Restore never modifies the Linux kernel.

Reverting to the original initial RAM disk

You can revert to the original initial RAM disk if necessary.

The initial RAM disk is stored on the machine in a file. Before updating the initial RAM disk for the first time, Universal Restore saves a copy of it to the same directory. The name of the copy is the name of the file, followed by the **_acronis_backup.img** suffix. This copy will not be overwritten if you run Universal Restore more than once (for example, after you have added missing drivers).

To revert to the original initial RAM disk, do any of the following:

- Rename the copy accordingly. For example, run a command similar to the following:

```
mv initrd-2.6.16.60-0.21-default_acronis_backup.img initrd-2.6.16.60-0.21-default
```

- Specify the copy in the **initrd** line of the GRUB boot loader configuration.

Recovering files

Recovering files in the Cyber Protect console

Note

You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1012).

1. Select the machine that originally contained the data that you want to recover.
2. Click **Recovery**.
3. Select the recovery point. Note that recovery points are filtered by location.
If the selected machine is physical and it is offline, recovery points are not displayed. Do any of the following:
 - [Recommended] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select a target machine that is online, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).
 - [Download the files from the cloud storage](#).
 - [Use bootable media](#).
4. Click **Recover > Files/folders**.
5. Browse to the required folder or use the search bar to obtain the list of the required files and folders.
Search is language-independent.
You can use one or more wildcard characters (* and ?). For more details about using wildcards, refer to "Mask " (p. 437).

Note

Search is not available for disk-level backups that are stored in the cloud storage.

6. Select the files that you want to recover.
7. If you want to save the files as a .zip file, click **Download**, select the location to save the data to, and click **Save**. Otherwise, skip this step.
Downloading is not available if your selection contains folders or the total size of the selected files exceeds 100 MB. To retrieve larger amounts of data from the cloud, use the procedure "Downloading files from the cloud storage" (p. 485).
8. Click **Recover**.
In **Recover to**, click to select the target for the recovery operation, or leave the default target. The default target varies according to the source of the backup.
The following targets are available:

- The source machine (if a protection agent is installed on it).
This is the machine that originally contained the files that you want to recover.
- Other machines on which a protection agent is installed – physical machines, virtual machines, and virtualization hosts on which a protection agent is installed, or virtual appliances.
You can recover files to physical machines, virtual machines, and virtualization hosts on which a protection agent is installed. You cannot recover files to virtual machines on which a protection agent is not installed (except for Virtuozzo virtual machines).
- Virtuozzo containers or virtual machines.
You can recover files to Virtuozzo containers and virtual machines with some limitations. For more information about them, refer to "Limitations for recovering files in the Cyber Protect console" (p. 489).

9. In **Path**, select the recovery destination. You can select one of the following:

- [When recovering to the original machine] The original location.
- A local folder or locally attached storage on the target machine.

Note

Symbolic links are not supported.

- A network folder that is accessible from the target machine.

10. Click **Start recovery**.

11. Select one of the file overwriting options:

- **Overwrite existing files**
- **Overwrite an existing file if it is older**
- **Do not overwrite existing files**

The recovery progress is shown on the **Activities** tab.

Downloading files from the cloud storage

In the Web Restore console, you can browse the cloud storage, view the contents of the backups, and download backed-up files and folders.

You cannot browse backups of system state, SQL databases, and Exchange databases.

You cannot download backed-up disks, volumes, or whole recovery points.

To download files and folders from the cloud storage

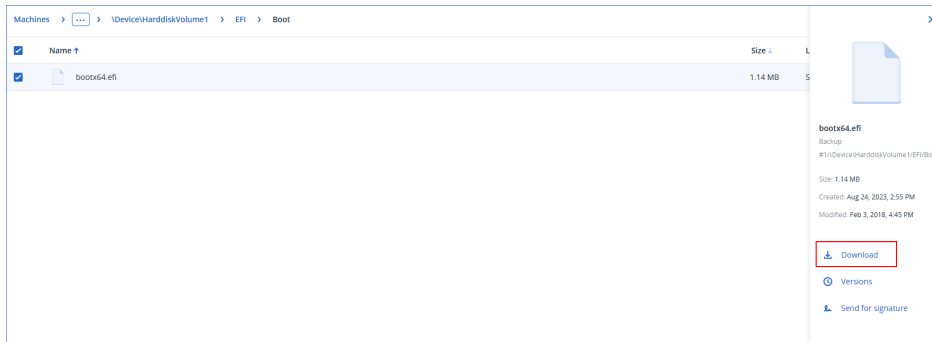
1. In the Cyber Protection console, select the required workload, and then click **Recovery**.
2. [If multiple backup locations are available] Select the backup location, and then click **More ways to recover**.
3. Click **Download files**.
4. Under **Machines**, click the workload name, and then click the backup archive.
A backup archive contains one or more backups (recovery points).
5. Click the backup number (recovery point) from which you want to download files or folders, and then navigate to the required items.

6. Select the check boxes next to the items that you want to download.

Note

If you select multiple items, they will be downloaded as a ZIP file.


7. Click **Download**.



Verifying file authenticity with Notary Service

If notarization [was enabled during backup](#), you can verify the authenticity of a backed-up file.

To verify the file authenticity

1. Select the file as described in steps 1-6 of the "[Recovering files by using the web interface](#)" section, or steps 1-5 of the "[Downloading files from the cloud storage](#)" section.
2. Ensure that the selected file is marked with the following icon: . This means that the file is notarized.
3. Do one of the following:
 - Click **Verify**.
The software checks the file authenticity and displays the result.
 - Click **Get certificate**.
A certificate that confirms the file notarization is opened in a web browser window. The window also contains instructions that allow you to verify the file authenticity manually.

Signing a file with ASign

Note

This feature is available with the Advanced Backup pack.

ASign is a service that allows multiple people to sign a backed-up file electronically. This feature is available only for file-level backups stored in the cloud storage.

Only one file version can be signed at a time. If the file was backed up multiple times, you must choose the version to sign, and only this version will be signed.

For example, ASign can be used for electronic signing of the following files:

- Rental or lease agreements
- Sales contracts
- Asset purchase agreements
- Loan agreements
- Permission slips
- Financial documents
- Insurance documents
- Liability waivers
- Healthcare documents
- Research papers
- Certificates of product authenticity
- Nondisclosure agreements
- Offer letters
- Confidentiality agreements
- Independent contractor agreements

To sign a file version

1. Select the file as described in steps 1-6 of the "[Recovering files by using the web interface](#)" section, or steps 1-5 of the "[Downloading files from the cloud storage](#)" section.
2. Ensure that the correct date and time is selected on the left panel.
3. Click **Sign this file version**.
4. Specify the password for the cloud storage account under which the backup is stored. The login of the account is displayed in the prompt window.
The ASign service interface is opened in a web browser window.
5. Add other signees by specifying their email addresses. It is not possible to add or remove signees after sending invitations, so ensure that the list includes everyone whose signature is required.
6. Click **Invite to sign** to send invitations to the signees.
Each signee receives an email message with the signature request. When all the requested signees sign the file, it is notarized and signed through the notary service.
You will receive notifications when each signee signs the file and when the entire process is complete. You can access the ASign web page by clicking **View details** in any of the email messages that you receive.
7. Once the process is complete, go to the ASign web page and click **Get document** to download a .pdf document that contains:
 - The Signature Certificate page with the collected signatures.
 - The Audit Trail page with history of activities: when the invitation was sent to the signees, when each signee signed the file, and so on.

Recovering files by using bootable media

For information about how to create bootable media, refer to "[Creating bootable media](#)".

To recover files by using bootable media

1. Boot the target machine by using the bootable media.
2. Click **Manage this machine locally** or click **Rescue Bootable Media** twice, depending on the media type you are using.
3. If a proxy server is enabled in your network, click **Tools > Proxy server**, and then specify the proxy server host name/IP address, port, and credentials. Otherwise, skip this step.
4. [Optional] When recovering Windows or Linux, click **Tools > Register media in the Cyber Protection service**, and then specify the registration token that you obtained when downloading the media. If you do this, you will not need to enter credentials or a registration code to access the cloud storage, as described in step 7.
5. On the welcome screen, click **Recover**.
6. Click **Select data**, and then click **Browse**.
7. Specify the backup location:
 - To recover from cloud storage, select **Cloud storage**. Enter the credentials of the account to which the backed up machine is assigned.
When recovering Windows or Linux, you have the option to request a registration code and use it instead of the credentials. Click **Use registration code > Request the code**. The software shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. The registration code is valid for one hour.
 - To recover from a local or a network folder, browse to the folder under **Local folders** or **Network folders**.
Click **OK** to confirm your selection.
8. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
9. In **Backup contents**, select **Folders/files**.
10. Select the data that you want to recover. Click **OK** to confirm your selection.
11. Under **Where to recover**, specify a folder. Optionally, you can prohibit overwriting of newer versions of files or exclude some files from recovery.
12. [Optional] Click **Recovery options** to specify additional settings.
13. Click **OK** to start the recovery.

Extracting files from local backups

You can browse the contents of backups and extract files that you need.

Requirements

- This functionality is available only in Windows by using File Explorer.
- The backed-up file system must be one of the following: FAT16, FAT32, NTFS, ReFS, Ext2, Ext3, Ext4, XFS, or HFS+.

Prerequisites

- A protection agent must be installed on the machine from which you browse a backup.
- The backup must be stored in a local folder or on a network share (SMB/CIFS).

To extract files from a backup

1. Browse to the backup location by using File Explorer.
2. Double-click the backup file. The file names are based on the following template:
<machine name> - <protection plan GUID>
3. If the backup is encrypted, enter the encryption password. Otherwise, skip this step.
File Explorer displays the recovery points.
4. Double-click the recovery point.
File Explorer displays the backed-up data.
5. Browse to the required folder.
6. Copy the required files to any folder on the file system.

Limitations for recovering files in the Cyber Protect console

Tenants in the Enhanced security mode

You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1012).

Recovery to Virtuozzo containers or Virtuozzo virtual machines

- QEMU Guest Agent must be installed on the target virtual machine.
- [Only applicable when recovering to containers] Mount points inside containers cannot be used as target for recovery. For example, you cannot recover files to a second hard disk or an NFS share mounted to a container.
- When recovering files to a Windows virtual machine, and if the "File-level security" (p. 495) recovery option is enabled, the archive bit attribute is set to the recovered files.
- Files with non-ANSI characters in their names are recovered with incorrect names on machines running Windows Server 2012 or older and machines running Windows 7 or older.
- To recover files to CentOS or Red Hat Enterprise Linux virtual machines that run on Virtuozzo Hybrid Server, you must edit the `qemu-ga` file, as follows:
 - On the target virtual machine, navigate to `/etc/sysconfig/`, and then open the `qemu-ga` file for editing.
 - Navigate to the following line, and then delete everything after the equals sign (=):

```
BLACKLIST_RPC=
```

- Restart QEMU Guest Agent by running the following command:

```
systemctl restart qemu-guest-agent
```

Recovering system state

Note

You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1012).

1. Select the machine for which you want to recover the system state.
 2. Click **Recovery**.
 3. Select a system state recovery point. Note that recovery points are filtered by location.
 4. Click **Recover system state**.
 5. Confirm that you want to overwrite the system state with its backed-up version.
- The recovery progress is shown on the **Activities** tab.

Recovering ESXi configuration

To recover an ESXi configuration, you need Linux-based bootable media. For information about how to create bootable media, refer to "Creating physical bootable media" (p. 657).

If you are recovering an ESXi configuration to a non-original host and the original ESXi host is still connected to the vCenter Server, disconnect and remove this host from the vCenter Server to avoid unexpected issues during the recovery. If you want to keep the original host along with the recovered one, you can add it again after the recovery is complete.

The virtual machines running on the host are not included in an ESXi configuration backup. They can be backed up and recovered separately.

To recover an ESXi configuration

1. Boot the target machine by using the bootable media.
2. Click **Manage this machine locally**.
3. On the welcome screen, click **Recover**.
4. Click **Select data**, and then click **Browse**.
5. Specify the backup location:
 - Browse to the folder under **Local folders** or **Network folders**.Click **OK** to confirm your selection.
6. In **Show**, select **ESXi configurations**.
7. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
8. Click **OK**.
9. In **Disks to be used for new datastores**, do the following:

- Under **Recover ESXi to**, select the disk where the host configuration will be recovered. If you are recovering the configuration to the original host, the original disk is selected by default.
 - [Optional] Under **Use for new datastore**, select the disks where new datastores will be created. Be careful because all data on the selected disks will be lost. If you want to preserve the virtual machines in the existing datastores, do not select any disks.
10. If any disks for new datastores are selected, select the datastore creation method in **How to create new datastores: Create one datastore per disk** or **Create one datastore on all selected HDDs**.
 11. [Optional] In **Network mapping**, change the result of automatic mapping of the virtual switches present in the backup to the physical network adapters.
 12. [Optional] Click **Recovery options** to specify additional settings.
 13. Click **OK** to start the recovery.

Recovery options

To modify the recovery options, click **Recovery options** when configuring recovery.

Availability of the recovery options

The set of available recovery options depends on:

- The environment the agent that performs recovery operates in (Windows, Linux, macOS, or bootable media).
- The type of data being recovered (disks, files, virtual machines, application data).

The following table summarizes the availability of the recovery options.

	Disks			Files				Virtual machines	SQL and Exchange
	Windows	Linux	Bootable media	Windows	Linux	macOS	Bootable media	ESXi, Hyper-V, and Virtuozzo	Windows
Backup validation	+	+	+	+	+	+	+	+	+
Boot mode	+	-	-	-	-	-	-	+	-
Date and time for files	-	-	-	+	+	+	+	-	-
Error handling	+	+	+	+	+	+	+	+	+

File exclusions	-	-	-	+	+	+	+	-	-
File-level security	-	-	-	+	-	-	-	-	-
Flashback	+	+	+	-	-	-	-	+	-
Full path recovery	-	-	-	+	+	+	+	-	-
Mount points	-	-	-	+	-	-	-	-	-
Performance	+	+	-	+	+	+	-	+	+
Pre/post commands	+	+	-	+	+	+	-	+	+
SID changing	+	-	-	-	-	-	-	-	-
VM power management	-	-	-	-	-	-	-	+	-
Windows event log	+	-	-	+	-	-	-	Hyper-V only	+

Backup validation

This option defines whether to validate a backup to ensure that the backup is not corrupted, before data is recovered from it. This operation is performed by the protection agent.

The preset is: **Disabled**.

For more information about the validation via checksum verification, refer to "Checksum verification" (p. 200).

Note

Depending on the settings chosen by your service provider, validation might not be available when backing up to the cloud storage.

Boot mode

This option is effective when recovering a physical or a virtual machine from a disk-level backup that contains a Windows operating system.

This option enables you to select the boot mode (BIOS or UEFI) that Windows will use after the recovery. If the boot mode of the original machine is different from the selected boot mode, the software will:

- Initialize the disk to which you are recovering the system volume, according to the selected boot mode (MBR for BIOS, GPT for UEFI).
- Adjust the Windows operating system so that it can start using the selected boot mode.

The preset is: **As on the target machine.**

You can choose one of the following:

- **As on the target machine**

The agent that is running on the target machine detects the boot mode currently used by Windows and makes the adjustments according to the detected boot mode.

This is the safest value that automatically results in bootable system unless the limitations listed below apply. Since the **Boot mode** option is absent under bootable media, the agent on media always behaves as if this value is chosen.

- **As on the backed-up machine**

The agent that is running on the target machine reads the boot mode from the backup and makes the adjustments according to this boot mode. This helps you recover a system on a different machine, even if this machine uses another boot mode, and then replace the disk in the backed-up machine.

- **BIOS**

The agent that is running on the target machine makes the adjustments to use BIOS.

- **UEFI**

The agent that is running on the target machine makes the adjustments to use UEFI.

Once a setting is changed, the disk mapping procedure will be repeated. This will take some time.

Recommendations

If you need to transfer Windows between UEFI and BIOS:

- Recover the entire disk where the system volume is located. If you recover only the system volume on top of an existing volume, the agent will not be able to initialize the target disk properly.
- Remember that BIOS does not allow using more than 2 TB of disk space.

Limitations

- Transferring between UEFI and BIOS is supported for:
 - 64-bit Windows operating systems starting with Windows 7
 - 64-bit Windows Server operating systems starting with Windows Server 2008 SP1
- Transferring between UEFI and BIOS is not supported if the backup is stored on a tape device.

When transferring a system between UEFI and BIOS is not supported, the agent behaves as if the **As on the backed-up machine** setting is chosen. If the target machine supports both UEFI and BIOS, you need to manually enable the boot mode corresponding to the original machine. Otherwise, the system will not boot.

Date and time for files

This option is effective only when recovering files.

This option defines whether to recover the files' date and time from the backup or assign the files the current date and time.

If this option is enabled, the files will be assigned the current date and time.

The preset is: **Enabled**.

Error handling

These options enable you to specify how to handle errors that might occur during recovery.

Re-attempt, if an error occurs

The preset is: **Enabled. Number of attempts: 30. Interval between attempts: 30 seconds.**

When a recoverable error occurs, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts are performed, depending on which comes first.

Do not show messages and dialogs while processing (silent mode)

The preset is: **Disabled**.

With the silent mode enabled, the program will automatically handle situations requiring user interaction where possible. If an operation cannot continue without user interaction, it will fail. Details of the operation, including errors, if any, can be found in the operation log.

Save system information if a recovery with reboot fails

This option is effective for a disk or volume recovery to a physical machine running Windows or Linux.

The preset is: **Disabled**.

When this option is enabled, you can specify a folder on the local disk (including flash or HDD drives attached to the target machine) or on a network share where the log, system information, and crash dump files will be saved. This file will help the technical support personnel to identify the problem.

File exclusions

This option is effective only when recovering files.

The option defines which files and folders to skip during the recovery process and thus exclude from the list of recovered items.

Note

Exclusions override the selection of data items to recover. For example, if you select to recover file MyFile.tmp and to exclude all .tmp files, file MyFile.tmp will not be recovered.

File-level security

This option is effective when recovering files from disk- and file-level backups of NTFS-formatted volumes.

This option defines whether to recover NTFS permissions for files along with the files.

The preset is: **Enabled**.

You can choose whether to recover the permissions or let the files inherit their NTFS permissions from the folder to which they are recovered.

Flashback

This option is effective when recovering disks and volumes on physical and virtual machines, except for Mac.

This option works only if the volume layout of the disk being recovered exactly matches that of the target disk.

If the option is enabled, only the differences between the data in the backup and the target disk data are recovered. This accelerates recovery of physical and virtual machines. The data is compared at the block level.

When recovering a physical machine, the preset is: **Disabled**.

When recovering a virtual machine, the preset is: **Enabled**.

Full path recovery

This option is effective only when recovering data from a file-level backup.

If this option is enabled, the full path to the file will be re-created in the target location.

The preset is: **Disabled**.

Mount points

This option is effective only in Windows for recovering data from a file-level backup.

Enable this option to recover files and folders that were stored on the mounted volumes and were backed up with the enabled [Mount points](#) option.

The preset is: **Disabled**.

This option is effective only when you select for recovery a folder that is higher in the folder hierarchy than the mount point. If you select for recovery folders within the mount point or the mount point itself, the selected items will be recovered regardless of the **Mount points** option value.

Note

Please be aware that if the volume is not mounted at the moment of recovery, the data will be recovered directly to the folder that has been the mount point at the time of backing up.

Performance

This option defines the priority of the recovery process in the operating system.

The available settings are: **Low, Normal, High**.

The preset is: **Normal**.

The priority of a process running in a system determines the amount of CPU and system resources allocated to that process. Decreasing the recovery priority will free more resources for other applications. Increasing the recovery priority might speed up the recovery process by requesting the operating system to allocate more resources to the application that will perform the recovery. However, the resulting effect will depend on the overall CPU usage and other factors like disk I/O speed or network traffic.

Pre/Post commands

The option enables you to define the commands to be automatically executed before and after the data recovery.

Example of how you can use the pre/post commands:

- Launch the **Checkdisk** command in order to find and fix logical file system errors, physical errors or bad sectors to be started before the recovery starts or after the recovery ends.

The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)

A post-recovery command will not be executed if the recovery proceeds with reboot.

Pre-recovery command

To specify a command/batch file to be executed before the recovery process starts

1. Enable the **Execute a command before the recovery** switch.
2. In the **Command...** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.

5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the recovery if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not recover until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the recovery only after the command is successfully executed. Fail the recovery if the command execution failed.	Perform the recovery after the command is executed despite execution failure or success.	N/A	Perform the recovery concurrently with the command execution and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

Post-recovery command

To specify a command/executable file to be executed after the recovery is completed

1. Enable the **Execute a command after the recovery** switch.
2. In the **Command...** field, type a command or browse to a batch file.
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field, specify the command execution arguments, if required.
5. Select the **Fail the recovery if the command execution fails** check box if successful execution of the command is critical for you. The command is considered failed if its exit code is not equal to zero. If the command execution fails, the recovery status will be set to **Error**.

When the check box is not selected, the command execution result does not affect the recovery failure or success. You can track the command execution result by exploring the **Activities** tab.

6. Click **Done**.

Note

A post-recovery command will not be executed if the recovery proceeds with reboot.

SID changing

This option is effective when recovering Windows 8.1/Windows Server 2012 R2 or earlier.

This option is not effective when recovery to a virtual machine is performed by Agent for VMware, Agent for Hyper-V, Agent for Scale Computing HC3, or Agent for oVirt.

The preset is: **Disabled**.

The software can generate a unique security identifier (Computer SID) for the recovered operating system. You only need this option to ensure operability of third-party software that depends on Computer SID.

Microsoft does not officially support changing SID on a deployed or recovered system. So use this option at your own risk.

VM power management

These options are effective when recovery to a virtual machine is performed by Agent for VMware, Agent for Hyper-V, Agent for Virtuozzo, Agent for Scale Computing HC3, or Agent for oVirt.

Power off target virtual machines when starting recovery

The preset is: **Enabled**.

Recovery to an existing virtual machine is not possible if the machine is online, and so the machine is powered off automatically as soon as the recovery starts. Users will be disconnected from the machine and any unsaved data will be lost.

Clear the check box for this option if you prefer to power off virtual machines manually before the recovery.

Power on the target virtual machine when recovery is complete

The preset is: **Disabled**.

After a machine is recovered from a backup to another machine, there is a chance the existing machine's replica will appear on the network. To be on the safe side, power on the recovered virtual machine manually, after you take the necessary precautions.

Windows event log

This option is effective only in Windows operating systems.

This option defines whether the agents have to log events of the recovery operations in the Application Event Log of Windows (to see this log, run eventvwr.exe or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

The preset is: **Disabled**.

Operations with backups

The Backup storage tab

The **Backup storage** tab provides access to all backups, including backups of offline machines, backups of machines that are no longer registered in the Cyber Protection service, backups to public clouds such as Microsoft Azure, and orphaned backups¹.

Backups created via `acrocmbd` are flagged as orphaned. Backups created in the 12.5 version of the product are also identified as orphaned.

Note

Please note that orphaned backups are also charged.

Backups that are stored in a shared location (such as an SMB or NFS share) are visible to all users that have the read permission for the location.

In Windows, backup files inherit the access permissions from their parent folder. Therefore, we recommend that you restrict the read permissions for this folder.

In the cloud storage, users have access only to their own backups.

An administrator can view backups to cloud on behalf of any account that belongs to the given unit or company and its child groups, by selecting the cloud storage for the account. To select the device that you want to use to obtain data from cloud, click **Change** in the **Machine to browse from** row. The **Backup storage** tab shows the backups of all machines ever registered under the selected account.

Backups created by the *cloud* Agent for Microsoft 365 and backups of Google Workspace data are shown not in the **Cloud storage** location, but in a separate section named **Cloud applications backups**.

Backup locations that are used in protection plans are automatically added to the **Backup storage** tab. To add a custom folder (for example, a detachable USB device) to the list of backup locations, click **Browse** and specify the folder path.

If you added or removed some backups by using a file manager, click the gear icon next to the location name, and then click **Refresh**.

Warning!

Do not try editing the backup files manually because this may result in file corruption and make the backups unusable. Also, we recommend that you use the backup replication instead of moving backup files manually.

¹An orphaned backup is a backup that is not associated to a protection plan anymore.

A backup location (except for the cloud storage) disappears from the **Backup storage** tab if all machines that had ever backed up to the location were deleted from the Cyber Protection service. This ensures that you do not have to pay for the backups stored in this location. As soon as a backup to this location occurs, the location is re-added along with all backups that are stored in it.

On the **Backup storage** tab, you can filter backups in the list by using the following criteria:

- **Only with forensic data** – only [backups having forensic data](#) will be shown.
- **Only pre-update backups created by Patch management** – only [backups that were created during patch management run before patch installation](#) will be shown.

To select a recovery point by using the Backup storage tab

1. On the **Backup storage** tab, select the location where the backups are stored.
The software displays all backups that your account is allowed to view in the selected location. The backups are combined in groups. The group names are based on the following template:
<machine name> - <protection plan name>
2. Select a group from which you want to recover the data.
3. [Optional] Click **Change** next to **Machine to browse from**, and then select another machine. Some backups can only be browsed by specific agents. For example, you must select a machine running Agent for SQL to browse the backups of Microsoft SQL Server databases.

Important

Please be aware that the **Machine to browse from** is a default destination for recovery from a physical machine backup. After you select a recovery point and click **Recover**, double check the **Target machine** setting to ensure that you want to recover to this specific machine. To change the recovery destination, specify another machine in **Machine to browse from**.

4. Click **Show backups**.
5. Select the recovery point.

To add a location for a backup

Note

This operation is available only if you have an online agent.

On the **Backup storage** tab, click **Add location**.

Select a location from one of the following locations types, and then click **Done**:

- Local folder
- Network folder
- Secure Zone
- NFS folder
- Public cloud

Mounting volumes from a backup

Mounting volumes from a disk-level backup lets you access the volumes as though they were physical disks.

Mounting volumes in the read/write mode enables you to modify the backup content; that is, save, move, create, delete files or folders, and run executables consisting of one file. In this mode, the software creates an incremental backup that contains the changes you make to the backup content. Note that none of the subsequent backups will contain these changes.

Requirements

- This functionality is available only in Windows by using File Explorer.
- Agent for Windows must be installed on the machine that performs the mount operation.
- The backed-up file system must be supported by the Windows version that the machine is running.
- The backup must be stored in a local folder, on a network share (SMB/CIFS), or in the Secure Zone.

Usage scenarios

- Sharing data
Mounted volumes can be easily shared over the network.
- "Band-aid" database recovery solution
Mount a volume that contains an SQL database from a recently failed machine. This will provide access to the database until the failed machine is recovered. This approach can also be used for granular recovery of Microsoft SharePoint data by using [SharePoint Explorer](#).
- Offline virus removal
If a machine is infected, mount its backup, clean it with an antivirus program (or find the latest backup that is not infected), and then recover the machine from this backup.
- Error check
If a recovery with volume resize has failed, the reason may be an error in the backed-up file system. Mount the backup in the read/write mode. Then, check the mounted volume for errors by using the `chkdsk /r` command. After the errors are fixed and a new incremental backup is created, recover the system from this backup.

To mount a volume from a backup

1. Browse to the backup location by using File Explorer.
2. Double-click the backup file. The file names are based on the following template:
`<machine name> - <protection plan GUID>`
3. If the backup is encrypted, enter the encryption password. Otherwise, skip this step.
File Explorer displays the recovery points.
4. Double-click the recovery point.

File Explorer displays the backed-up volumes.

Note

Double-click a volume to browse its content. You can copy files and folders from the backup to any folder on the file system.

5. Right-click a volume to mount, and then select one of the following options:
 - a. **Mount**

Note

Only the last backup in the archive (backup chain) can be mounted in read-write mode.

- b. **Mount in read-only mode.**
6. If the backup is stored on a network share, provide access credentials. Otherwise, skip this step. The software mounts the selected volume. The first unused letter is assigned to the volume.

To unmount a volume

1. Browse to **Computer (This PC)** in Windows 8.1 and later) by using File Explorer.
2. Right-click the mounted volume.
3. Click **Unmount**.
4. [Optional] If the volume was mounted in the read/write mode, and its content was modified, select whether to create an incremental backup containing the changes. Otherwise, skip this step.

The software unmounts the selected volume.

Validating backups

By validating a backup, you verify that you can recover the data from it. For more information about this operation, refer to "Validation" (p. 196).

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

To validate a backup

1. Select the backed-up workload.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.

If the workload is offline, the recovery points are not displayed. Do any of the following:

 - If the backup location is cloud or shared storage (that is, other agents can access it), click **Select machine**, select a target workload that is online, and then select a recovery point.

- Select a recovery point on the Backup storage tab. For more information about the backups there, refer to "The Backup storage tab" (p. 499).
4. Click the gear icon, and then click **Validate**.
 5. Select the agent that will perform the validation.
 6. Select the validation method.
 7. If the backup is encrypted, provide the encryption password.
 8. Click **Start**.

Exporting backups

The export operation creates a self-sufficient copy of a backup in the location that you specify. The original backup remains untouched. Exporting backups allows you to separate a specific backup from a chain of incremental and differential backups for fast recovery, for writing onto removable or detachable media, or for other purposes.

Note

This functionality is available in customer tenants for which any of the following quota (part of the Advanced Backup pack) is enabled: **Advanced Backup – Workstations**, **Advanced Backup – Servers**, **Advanced Backup – Virtual machines**, or **Advanced Backup – NAS**.

The result of an export operation is always a full backup. If you want to replicate the entire backup chain to a different location and preserve multiple recovery points, use a backup replication plan. For more information about this plan, refer to "Backup replication" (p. 194).

The backup file name of the exported backup is the same as that of the original backup, except for the sequence number. If multiple backups from the same backup chain are exported to the same location, a four-digit sequence number is appended to the file names of all backups except for the first one.

The exported backup inherits the encryption settings and password from the original backup. When exporting an encrypted backup, you must specify the password.

To export a backup

1. Select the backed-up workload.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.
If the workload is offline, the recovery points are not displayed. Do any of the following:
 - If the backup location is cloud or shared storage (that is, other agents can access it), click **Select machine**, select a target workload that is online, and then select a recovery point.
 - Select a recovery point on the Backup storage tab. For more information about the backups there, refer to "The Backup storage tab" (p. 499).
4. Click the gear icon, and then click **Export**.
5. Select the agent that will perform the export.
6. If the backup is encrypted, provide the encryption password. Otherwise, skip this step.

7. Specify the export destination.
8. Click **Start**.

Deleting backups

A backup archive contains one or more backups. You can delete specific backups (recovery points) in an archive or the whole archive.

Deleting the backup archive deletes all backups in it. Deleting all backups of a workload deletes the backup archives that contain these backups.

You can delete backups by using the Cyber Protect console – on the **Devices** tab and on the **Backup storage** tab. Also, you can delete backups from the cloud storage by using the Web Restore console.

Warning!

If immutable storage is disabled, backed-up data is permanently deleted and cannot be recovered.

To delete backups or backup archives

On the Devices tab

This procedure applies only to online workloads.

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the workload backups of which you want to delete.
3. Click **Recovery**.
4. [If more than one backup location is available] Select the backup location.
5. [To delete all backups of the workload] Click **Delete all**.
Deleting all backups also deletes the backup archives that contain these backups.
6. [To delete a specific backup] Select the backup (recovery point) that you want to delete, and then click **Actions > Delete**.
7. [When deleting all backups] Select the check box, and then click **Delete** to confirm your decision.
8. [When deleting a specific backup] Click **Delete** to confirm your decision.

On the Backup storage tab

This procedure applies to online and offline workloads.

1. In the Cyber Protect console, go to **Backup storage**.
2. Select the location from which you want to delete backups.
3. Select the backup archive from which you want to delete backups.
The archive name uses the following template:
 - Non-cloud-to-cloud backup archives: <workload name> - <protection plan name>
 - Cloud-to-cloud backup archives: <user name> or <drive name> or <team name> - <cloud service> - <protection plan name>
4. [To delete the whole backup archive] Click **Delete**.
Deleting a backup archive deletes all backups in that archive.

5. [To delete a specific backup in the backup archive] Click **Show backups**.
 - a. Select the backup (recovery point) that you want to delete.
 - b. Click **Actions > Delete**.
6. [When deleting a backup archive] Select the check box, and then click **Delete** to confirm your decision.
7. [When deleting a specific backup] Click **Delete** to confirm your decision.

In the Web Restore console

This procedure applies only to backup archives in the cloud storage.


1. In the Cyber Protection console, go to **Devices > All devices**.
2. Select the workload backups of which you want to delete, and then click **Recovery**.
3. [If multiple backup locations are available] Select the backup location, and then click **More ways to recover**.
4. Click **Download files**.
You are redirected to the Web Restore console.
5. In the Web Restore console, under **Machines**, click the workload name.
6. Under **Last version**, click the date, and then click **Delete**.
This action is only available on the backup archive level. You cannot drill down the archive and delete specific backups from it.
7. Click **Delete** to confirm your decision.

Deleting backups outside the Cyber Protect console

We recommend that you delete backups by using the Cyber Protect console. If you delete backups from the cloud storage by using the Web Restore console or delete local backups by using a file manager, you must refresh the backup location to sync the changes to the Cyber Protect console.

Prerequisite


- An online agent that can access the backup location must be selected as **Machine to browse from** in the Cyber Protect console.



To refresh a backup location

1. In the Cyber Protect console, go to **Backup storage**.
2. Select the backup location in which the deleted backups were stored.

3. In the **Actions** pane, click **Refresh**.



Understanding the detection of bottlenecks

The bottleneck detection feature helps you understand where you can improve performance by highlighting which component in your system was the slowest during a backup or recovery process.

As bottlenecks *always* occur in any transmission event, it does not necessarily mean they need to be resolved. Your backups may be already fast enough and meet your backup windows perfectly, as well as meet your SLAs, so there is typically nothing you need to actually resolve.

You can easily view and track bottlenecks in the **Activity details** tab. To do this, in the Cyber Protect console, go to **Monitoring > Activities**, and then click the relevant activity. For more information about viewing bottlenecks, see "Viewing bottleneck details" (p. 507) and "On what workloads, agents, and backup locations are bottlenecks shown?" (p. 509).

What is a bottleneck?

Bottlenecks are typically caused due to a slow component in the processing chain, in other words, a component that the other components wait for.

The bottleneck detection feature enables you to track these slow components during the backup and recovery process, helping you understand which of the following component types is the slowest:

- **Source:** At a glance, you can determine if the reading speed from the backup/recovery source is causing a bottleneck.
- **Destination:** Understand if the writing speed to the backup/recovery destination is affecting performance.
- **Agent:** Understand if the agent is processing the data fast enough.

The bottleneck type, whether from the source, destination, or agent, can change at different times during the backup/recovery activity. The percentages shown in the **Bottleneck** section of the **Activity details** tab below (for example, **Read data from source (workload): 63%**), represent the percentage of time when this type of bottleneck was encountered. In this case, for 63% of the recovery activity time, the bottleneck type was reading data, in other words, the slow speed in reading data from the backup archive by the agent.

Similarly, for 30% of the time, the bottleneck was due to the slow speed in writing data to the recovery destination (**Write data to destination: 30%**).

Activity details



15:42 PM — 18:23 PM (2 hrs 41 mins)

Recovering files

Status: Succeeded

Workload: qa-gw3t68hh

Started by: NikolaTesla

Start time: Feb 14, 2020, 15:32:06

Finish time: Feb 14, 2020, 18:23:07

Duration: 2 hrs 41 mins

Backup file name: qa-gw3t68hh-11F95D-412C-9ccF-BCBc8AAF7E9-AFAF230D-D4AB-43242-9ASDQ
13-ASDS7213-DSA7DSA

Backup location: E:/Backups/


What to recover: desktop.ini

Bytes processed: 155 GB

Bytes saved: 177 GB

Speed: 9.8 MB/s

Bottleneck: Read data from source (workload) ⓘ



- Read data from source (workload): 63%
- Write data to destination: 30%
- Data encryption/decryption: 7%

[Hide details](#)

[All properties](#)

Note

It is normal behavior to see bottleneck statistics in the **Activity details** tab. These statistics are only available for tasks more than one minute long.

How to reduce bottlenecks

As mentioned above, the bottleneck detection feature highlights the *read* and *write* data flow between backup components. The *read* statistics refer to the data flow from the data source to the agent which performs the backup/recovery operation, and the *write* statistics refer to the data flow between the agent and the backup archive (the destination).

To reduce bottlenecks and improve the read/write data flow performance, you should analyze the channel between the agent and the data source/backup archive. For example, you can try benchmarking your hard disks if the agent is backing up some local files.

Viewing bottleneck details

You can view detected bottlenecks for any type of backup, backup replication, or recovery process (to any type of destination folder or location), including virtual machine backups, machine backups, and file/folder backups. You can also view bottlenecks for virtual machine replication and failback activities.

For more information on the definition and core concepts of bottleneck types, see "Understanding the detection of bottlenecks" (p. 506).

To view bottleneck details

1. In the Cyber Protect console, go to **Monitoring > Activities**.
2. Click on the relevant activity.

In the **Activity details** tab, the **Bottleneck** section is shown in blue.

Activity details ×

15:42 PM — 18:23 PM (2 hrs 41 mins)

Recovering files

Status: Succeeded
Workload: qa-gw3t68hh
Started by: NikolaTesla

Start time: Feb 14, 2020, 15:32:06
Finish time: Feb 14, 2020, 18:23:07
Duration: 2 hrs 41 mins

Backup file name: qa-gw3t68hh-11F95D-412C-9ccF-BCBc8AAF7E9-AFAF230D-D4AB-43242-9ASDQ
13-ASDS7213-DSA7DSA
Backup location: E:/Backups/
What to recover: desktop.ini

Bytes processed: 155 GB
Bytes saved: 177 GB
Speed: 9.8 MB/s

Bottleneck: Read data from source (workload) ⓘ

[Show details](#)

[All properties](#)

3. Click **Show details** to view the most frequently encountered bottleneck during the backup/recovery operation.

The **Bottleneck** section expands to show a summary of the relevant bottleneck types.

Bottleneck: Read data from source (workload) ⓘ

- Read data from source (workload): 63%
- Write data to destination: 30%
- Data encryption/decryption: 7%

[Hide details](#)

In the example above, the bottleneck that accounted for 63% of the entire operation's time was caused by the *Read* operation (performed by the agent).

Note

The bottleneck values update dynamically every minute while the corresponding activity is running.

On what workloads, agents, and backup locations are bottlenecks shown?

The detection of bottlenecks is available for the following types of workloads, agents, and backup locations:

- Disk/image-level backups performed by:
 - Agent for Azure
 - Agent for Windows
 - Agent for Linux
 - Agent for MAC
 - Agent for VMware (both Virtual Appliance and Windows, including VM replication and failback from replica (restore from replica) activities)
 - Agent for Hyper-V
 - Agent for Scale Computing
 - Agent for oVirt (KVM)
 - Agent for Virtuozzo Infrastructure Platform
 - Agent for Virtuozzo
 - Agent for VMware Cloud Director (vCD-BA)
- File-level backups
 - Agent for Windows
 - Agent for Linux
 - Agent for MAC
- Application-level backups
 - Agent for SQL
 - Agent for Exchange
 - Agent for MySQL/MariaDB
 - Agent for Oracle
 - Agent for SAP HANA
- Backup locations
 - Acronis Cloud storage (including partner hosted storage)
 - Public Cloud storage
 - Network shares (SMB + NFS)
 - Local folders
 - Locations defined by script
 - Acronis Secure Zone

Backing up workloads to public clouds

Note

This feature is part of the Advanced Backup pack, which in turn is part of the Cyber Protection service. Note that when you add this functionality to a protection plan, you may be subject to additional charges.

You can select public cloud services, such as Microsoft Azure, as backup destinations in the Cyber Protect console.

To configure backup locations on Microsoft Azure you must be a Company administrator or Unit administrator, or have one of the following roles defined in the Cyber Protection service: Cyber administrator, Administrator, User.

Defining a backup location in Microsoft Azure

To back up a workload to Microsoft Azure, you need to define the Microsoft Azure backup location in the Cyber Protect console, and connect to the relevant Microsoft Azure subscription. This can be done in the following ways:

- When creating or editing a protection plan.
- When defining and managing backup storage locations.

Important

Both administrators and non-administrator users can back up workloads to Microsoft Azure.

Non-administrator users can add access to a Microsoft Azure subscription (see "Managing access to Microsoft Azure subscriptions" (p. 513)), but can only apply protection plans where the backup location is connected to the Microsoft Azure subscription they added themselves, and for workloads registered in the Cyber Protect console under their name.

Administrators can apply protection plans where the backup location is connected to Microsoft Azure subscriptions they added themselves or to subscriptions added by any other administrator, and for workloads registered in the Cyber Protect console under any user.

To define a backup location in Microsoft Azure

1. In the Cyber Protect console, do one of the following:
 - If you are creating or editing a protection plan, go to **Devices** and select the relevant workload you want to back up to Microsoft Azure. In the **Backup** section of the selected workload's protection plan, click the link in the **Where to back up** row.
For more information about working with protection plans, see "Protection plans and modules" (p. 208).

- If you are managing your backup storage locations and want to add Microsoft Azure as a new location, go to **Backup storage**.

For more information about managing your backup storage locations, see "The Backup storage tab" (p. 499).

2. Click **Add location**.

3. Click **Public cloud**.

By default, **Microsoft Azure** is selected in the **Cloud** field.

4. If the relevant Microsoft Azure subscription is already registered in the Cyber Protect console, select it from the list of subscriptions.

If the relevant subscription is not registered in the Cyber Protect console, click **Add** and in the displayed dialog, click **Sign in**. You are redirected to the Microsoft login page. For more information about adding and defining access to a Microsoft Azure subscription, see "Adding access to a Microsoft Azure subscription" (p. 514).

5. In the **Storage account** field, select the relevant account.

Note





Only Microsoft Azure storage accounts with regular endpoint suffixes that contain `core.windows.net` are currently supported. In addition, the selected storage account must be a StorageV2 account type.

The **Location name** and **Access tier** fields are automatically filled by default, according to the storage account selected. The location name displayed is `microsoft_azure_[storage account]` and the access tier selected is **Default (Hot)**. Both fields can be modified, as required.

Note

When changing the location name, enter a unique location name (the name must be unique to the customer tenant). If the name you add already exists in the storage account, Acronis adds a suffix number to the name. For example, if **Microsoft Azure Storage** already exists, the name is automatically updated to **Microsoft Azure Storage_01**.

✕ Add location

-  Local folder
-  Network folder
-  Defined by a script
-  Public cloud ↑

Public cloud

Cloud
 Microsoft Azure ▼

Microsoft Azure subscription
 Microsoft Azure Enterprise ▼

Storage account
 dktestsa ▼ ⓘ

Location name
 microsoft_azure_dktestsa

Access tier
 Default (Hot) ▼ ⓘ

Add

6. Click **Add**.

If you are creating or editing a protection plan, the Microsoft Azure backup location is set as the location in **Where to back up** row. When the backup is run (either manually or when scheduled), the backup is saved in the defined location.

If you are managing your backup storage locations, you can view and update the location details as required. The Microsoft Azure location is also available when defining a back up location for workloads. For more information, see "Viewing and updating Microsoft Azure backup locations" (p. 512).

Viewing and updating Microsoft Azure backup locations

You can view and update the Microsoft Azure backup locations you define in the **Backup storage** module, or when creating or editing a protection plan.

For information about removing access to a Microsoft Azure subscription from the Cyber Protect console, see "Removing access to a Microsoft Azure subscription" (p. 516).

Note

You cannot refresh or delete a Microsoft Azure backup location in the **Backup storage** module.

To view Microsoft Azure backup locations

1. In the Cyber Protect console, go to **Backup storage**.
A list of backup locations is displayed, including Microsoft Azure storage locations, with details of the storage capacity and number of backups assigned to each location.
For more information about working with the listed backup locations, see "The Backup storage tab" (p. 499).
2. Select the relevant location.
Any current backups for the selected location are listed.
3. (Optional) Click on a backup to view more details for the backup.

To update a Microsoft Azure backup location in a protection plan

1. Go to the relevant protection plan, and select **Edit**.
2. Click the link in the **Where to back up** row.
3. Select from the list of existing backup locations, or click **Add location** to add a new location.
If the relevant Microsoft Azure subscription is already registered in the Cyber Protect console, select it from the list of subscriptions. For more information, see "Defining a backup location in Microsoft Azure" (p. 510).
If you are adding a new Microsoft Azure subscription, you will be prompted to authenticate your Microsoft account details (see "Adding access to a Microsoft Azure subscription" (p. 514). For more information about the required permissions when connecting to Microsoft Azure, refer to article [Microsoft Azure connection security and audit \(72684\)](#).

Managing public cloud account access

To enable Acronis Cyber Protection services in public cloud platforms, access to the relevant public cloud accounts needs to be configured.

For example, when working with Microsoft Azure, access to your Microsoft Azure subscription is required. Once added in the Cyber Protect console, the subscription can be selected when you configure a direct backup to Microsoft Azure.

Access to public clouds is managed through the **Infrastructure** menu in the Cyber Protect console.

Managing access to Microsoft Azure subscriptions

By connecting to the relevant Microsoft Azure subscriptions in the Cyber Protect console, you can directly back up the relevant workloads to Microsoft Azure.

Connection to a subscription can be configured when creating a backup location via the **Devices** or **Backup storage** menu, as described in "Defining a backup location in Microsoft Azure" (p. 510).

Alternatively, these Microsoft Azure subscriptions can be configured in the **Public clouds** screen (go to **Infrastructure > Public clouds**). Here you can also manage your subscriptions, including renewing access to the subscription, viewing subscription properties and activities, or removing the subscription.

Depending on your assigned administrator role, you may be able to manage Microsoft Azure subscriptions added by other users within your organization. For example, if you are a Company administrator or Unit administrator, or have been assigned the Cyber administrator or Administrator role in the Cyber Protection service, you can view and manage Microsoft Azure subscriptions added by other administrators, as well as subscriptions added by non-administrator users. Non-administrator users can only view and access Microsoft Azure subscriptions they added to the Cyber Protect console.

Note

Partners can manage the Microsoft Azure subscriptions of customers below their level in the hierarchy. However, when a partner selects **All customers**, the **Infrastructure** menu in the Cyber Protect console is not available.

Important

When connecting to a Microsoft Azure subscription, Acronis requires the minimum permissions to connect to the subscription. For more information about the required permissions, refer to article [Microsoft Azure connection security and audit \(72684\)](#).

Adding access to a Microsoft Azure subscription

By adding a Microsoft Azure subscription in the Cyber Protect console, Acronis can securely access your subscription and directly back up the relevant workloads to Microsoft Azure.

To add access to a Microsoft Azure subscription

1. In the Cyber Protect console, go to **Infrastructure > Public clouds**.
2. Click **Add**, and in the displayed dialog, click **Sign in**. You are redirected to the Microsoft login page.

Note

You must be assigned with one of the following roles in Microsoft Azure AD in order to complete the connection to the subscription: Cloud Application Administrator, Application Administrator, or Global Administrator. You must also be assigned the Owner role for each selected subscription.

3. In the Microsoft login screen, enter your login credentials and accept the requested permissions. The connection process starts, and may take several minutes.
For more information about securely accessing your Microsoft Azure and subscription, refer to article [Microsoft Azure connection security and audit \(72684\)](#).
4. When the connection is complete, select the relevant subscription from the drop-down list in the displayed dialog, and click **Add subscription**.

Add subscription



✓ Authenticated with your Azure account

Select a subscription from the list.

Microsoft Azure subscription
Microsoft Azure Enterprise - 656771587-8174-4688-6867-67478478527

Cancel Add subscription

The subscription is added to the list of public clouds.

To renew the annual access certificate for the subscription, see "Renewing access to a Microsoft Azure subscription" (p. 515).

To remove access to the subscription, see "Removing access to a Microsoft Azure subscription" (p. 516).

Note

If the Microsoft Azure account you are logged into includes access to multiple Microsoft Azure ADs, including ADs in which you were invited as a guest user, only the default user directory is selected. If you want to use a directory in which you are a guest user, you need to create a new user in that specific Microsoft Azure AD. You can then log in to that account and connect to the relevant subscription.

Renewing access to a Microsoft Azure subscription

Once registered in the Cyber Protect console, access to a Microsoft Azure subscription is automatically set for one year by Acronis using a free and unique access certificate. When the certificate nears its expiry date, you can quickly and easily renew it.

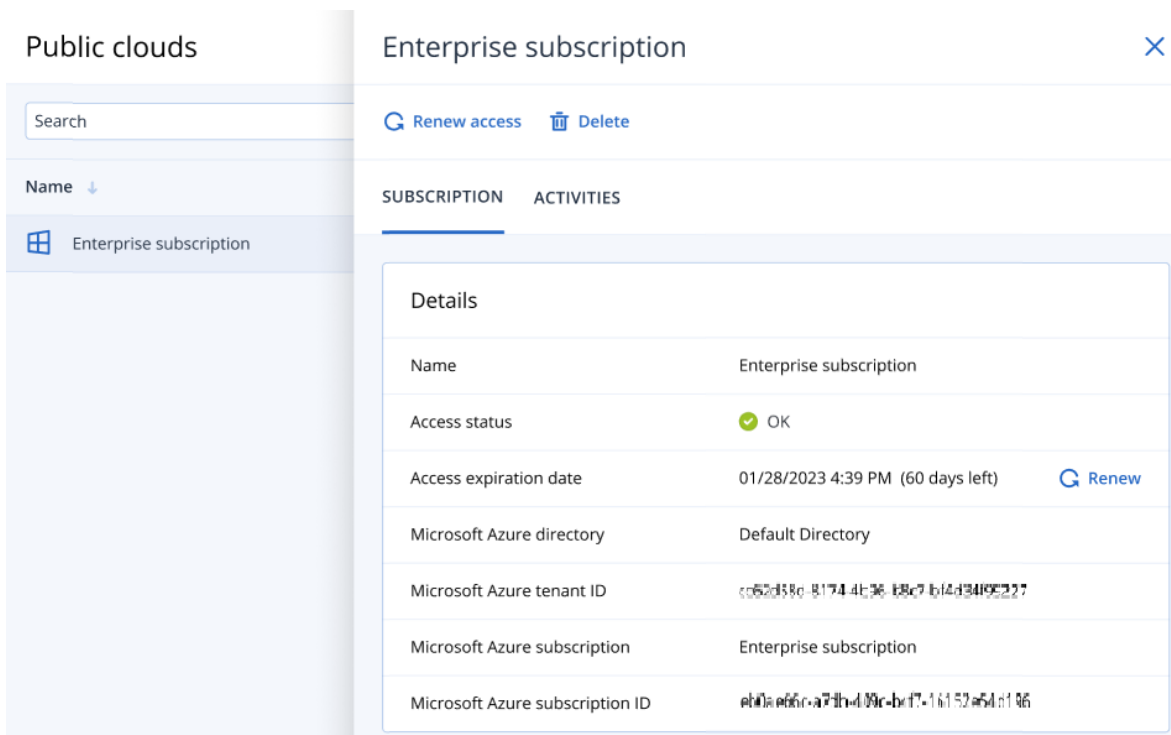
To renew the access certificate for your Microsoft Azure subscription

1. In the Cyber Protect console, go to **Infrastructure > Public clouds**.
2. Select the relevant subscription from the displayed list.

Note

The **Access status** column indicates the current status of the access certificate for each subscription and shows one of two statuses: **OK** or **Expired**.

3. In the right pane, click **Renew access**.
Alternatively, click the **Subscription** tab, and then click **Renew** in the **Access expiration date** field.



- In the Microsoft login screen, enter your login credentials and accept the requested permissions. The connection process starts, and may take several minutes. When the authentication is successful, access is automatically renewed for one year. For more information about the required permissions, refer to article [Microsoft Azure connection security and audit \(72684\)](#).

Removing access to a Microsoft Azure subscription

You should remove access to the Microsoft Azure subscription if you are not backing up workloads to Microsoft Azure.

To remove access to a Microsoft Azure subscription

Important

You cannot remove a subscription if it is currently being used to backup to Microsoft Azure.

- In the Cyber Protect console, go to **Infrastructure > Public clouds**.
- Select the relevant subscription from the displayed list.
- In the right pane, click **Delete**.

Note

You can only remove a subscription you added. You can also remove a subscription if you are a Company administrator or Unit administrator, or were assigned the role of Cyber administrator or Administrator in the Cyber Protection service.

- In the displayed confirmation message, click **Remove**.

Protecting Microsoft applications

Protecting Microsoft SQL Server and Microsoft Exchange Server

There are two methods of protecting these applications:

- **Database backup**

This is a file-level backup of the databases and the metadata associated with them. The databases can be recovered to a live application or as files.

- **Application-aware backup**

This is a disk-level backup that also collects the applications' metadata. This metadata enables browsing and recovery of the application data without recovering the entire disk or volume. The disk or volume can also be recovered as a whole. This means that a single solution and a single protection plan can be used for both disaster recovery and data protection purposes.

For Microsoft Exchange Server, you can opt for **Mailbox backup**. This is a backup of individual mailboxes via the Exchange Web Services protocol. The mailboxes or mailbox items can be recovered to a live Exchange Server or to Microsoft 365. Mailbox backup is supported for Microsoft Exchange Server 2010 Service Pack 1 (SP1) and later.

Protecting Microsoft SharePoint

A Microsoft SharePoint farm consists of front-end servers that run SharePoint services, database servers that run Microsoft SQL Server, and (optionally) application servers that offload some SharePoint services from the front-end servers. Some front-end and application servers may be identical to each other.

To protect an entire SharePoint farm:

- Back up all of the database servers with application-aware backup.
- Back up all of the unique front-end servers and application servers with usual disk-level backup.

The backups of all servers should be done on the same schedule.

To protect only the content, you can back up the content databases separately.

Protecting a domain controller

A machine running Active Directory Domain Services can be protected by application-aware backup. If a domain contains more than one domain controller, and you recover one of them, a nonauthoritative restore is performed and a USN rollback will not occur after the recovery.

Recovering applications

The following table summarizes the available application recovery methods.

	From a database backup	From an application-aware backup	From a disk backup
Microsoft SQL Server	Databases to a live SQL Server instance Databases as files	Entire machine Databases to a live SQL Server instance Databases as files	Entire machine
Microsoft Exchange Server	Databases to a live Exchange Databases as files Granular recovery to a live Exchange or to Microsoft 365*	Entire machine Databases to a live Exchange Databases as files Granular recovery to a live Exchange or to Microsoft 365*	Entire machine
Microsoft SharePoint database servers	Databases to a live SQL Server instance Databases as files Granular recovery by using SharePoint Explorer	Entire machine Databases to a live SQL Server instance Databases as files Granular recovery by using SharePoint Explorer	Entire machine
Microsoft SharePoint front-end web servers	-	-	Entire machine
Active Directory Domain Services	-	Entire machine	-

* Granular recovery is also available from a mailbox backup. Recovery of Exchange data items to Microsoft 365, and vice versa, is supported on the condition that Agent for Microsoft 365 is installed locally.

Prerequisites

Before configuring the application backup, ensure that the requirements listed below are met.

To check the VSS writers state, use the `vssadmin list writers` command.

Common requirements

For Microsoft SQL Server, ensure that:

- At least one Microsoft SQL Server instance is started.
- The SQL writer for VSS is turned on.

For Microsoft Exchange Server, ensure that:

- The Microsoft Exchange Information Store service is started.
- Windows PowerShell is installed. For Exchange 2010 or later, the Windows PowerShell version must be at least 2.0.
- Microsoft .NET Framework is installed.
For Exchange 2007, the Microsoft .NET Framework version must be at least 2.0.
For Exchange 2010 or later, the Microsoft .NET Framework version must be at least 3.5.
- The Exchange writer for VSS is turned on.

Note

Agent for Exchange needs a temporary storage to operate. By default, the temporary files are located in %ProgramData%\Acronis\Temp. Ensure that you have at least as much free space on the volume where the %ProgramData% folder is located as 15 percent of an Exchange database size. Alternatively, you can change the location of the temporary files before creating Exchange backups as described in [Changing Temp Files and Folder Location \(40040\)](#).

On a domain controller, ensure that:

- The Active Directory writer for VSS is turned on.

When creating a protection plan, ensure that:

- For physical machines and machines with the agent installed inside, the [Volume Shadow Copy Service \(VSS\)](#) backup option is enabled.
- For virtual machines, the [Volume Shadow Copy Service \(VSS\) for virtual machines](#) backup option is enabled.

Additional requirements for application-aware backups

When creating a protection plan, ensure that **Entire machine** is selected for backup. The **Sector-by-sector** backup option must be disabled in a protection plan, otherwise it will be impossible to perform a recovery of application data from such backups. If the plan is executed in the **Sector-by-sector** mode due to an automatic switch to this mode, then recovery of application data will also be impossible.

Requirements for ESXi virtual machines

If the application runs on a virtual machine that is backed up by Agent for VMware, ensure that:

- The virtual machine being backed up meets the requirements for application-consistent backup and restore listed in the article "Windows Backup Implementations" in the VMware documentation: <https://code.vmware.com/docs/1674/virtual-disk-programming-guide/doc/vddkBkupVadp.9.6.html>.
- VMware Tools is installed and up-to-date on the machine.
- User Account Control (UAC) is disabled on the machine. If you do not want to disable UAC, you must provide the credentials of the built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

If you do not want to disable UAC, you must provide the credentials of the built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

Note

Use the built-in domain administrator account that was configured as part of the creation of the domain. Accounts created later are not supported.

Requirements for Hyper-V virtual machines

If the application runs on a virtual machine that is backed up by Agent for Hyper-V, ensure that:

- The guest operating system is Windows Server 2008 or later.
- For Hyper-V 2008 R2: the guest operating system is Windows Server 2008/2008 R2/2012.
- The virtual machine has no dynamic disks.
- The network connection exists between the Hyper-V host and the guest operating system. This is required to execute remote WMI queries inside the virtual machine.
- User Account Control (UAC) is disabled on the machine. If you do not want to disable UAC, you must provide the credentials of the built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

If you do not want to disable UAC, you must provide the credentials of the built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

Note

Use the built-in domain administrator account that was configured as part of the creation of the domain. Accounts created later are not supported.

- The virtual machine configuration matches the following criteria:
 - Hyper-V Integration Services is installed and up-to-date. The critical update is <https://support.microsoft.com/en-us/help/3063109/hyper-v-integration-components-update-for-windows-virtual-machines>
 - In the virtual machine settings, the **Management > Integration Services > Backup (volume checkpoint)** option is enabled.
 - For Hyper-V 2012 and later: the virtual machine has no checkpoints.
 - For Hyper-V 2012 R2 and later: the virtual machine has a SCSI controller (check **Settings > Hardware**).

Database backup

Before backing up databases, ensure that the requirements listed in "[Prerequisites](#)" are met.

Select the databases as described below, and then specify other settings of the protection plan [as appropriate](#).

Selecting SQL databases

A backup of an SQL database contains the database files (.mdf, .ndf), log files (.ldf), and other associated files. The files are backed with the help of the SQL Writer service. The service must be running at the time that the Volume Shadow Copy Service (VSS) requests a backup or recovery.

The SQL transaction logs are truncated after each successful backup. SQL log truncation can be disabled in the [protection plan options](#).

To select SQL databases

1. Click **Devices > Microsoft SQL**.

The software shows the tree of SQL Server Always On Availability Groups (AAG), machines running Microsoft SQL Server, SQL Server instances, and databases.

2. Browse to the data that you want to back up.

Expand the tree nodes or double-click items in the list to the right of the tree.

3. Select the data that you want to back up. You can select AAGs, machines running SQL Server, SQL Server instances, or individual databases.

- If you select an AAG, all databases that are included into the selected AAG will be backed up. For more information about backing up AAGs or individual AAG databases, refer to ["Protecting Always On Availability Groups \(AAG\)"](#).
- If you select a machine running an SQL Server, all databases that are attached to all SQL Server instances running on the selected machine will be backed up.
- If you select a SQL Server instance, all databases that are attached to the selected instance will be backed up.
- If you select databases directly, only the selected databases will be backed up.

4. Click **Protect**. If prompted, provide credentials to access the SQL Server data.

If you use Windows authentication, the account must be a member of the **Backup Operators** or **Administrators** group on the machine and a member of the **sysadmin** role on each of the instances that you are going to back up.

If you use SQL Server authentication, the account must be a member of the **sysadmin** role on each of the instances that you are going to back up.

Selecting Exchange Server data

The following table summarizes the Microsoft Exchange Server data that you can select for backup and the minimal user rights required to back up the data.

Exchange version	Data items	User rights
2007	Storage groups	Membership in the Exchange Organization Administrators role group
2010/2013/2016/2019	Databases, Database Availability Groups (DAG)	Membership in the Server Management role group.

A full backup contains all of the selected Exchange Server data.

An incremental backup contains the changed blocks of the database files, the checkpoint files, and a small number of the log files that are more recent than the corresponding database checkpoint. Because changes to the database files are included in the backup, there is no need to back up all the transaction log records since the previous backup. Only the log that is more recent than the checkpoint needs to be replayed after a recovery. This makes for faster recovery and ensures successful database backup, even with circular logging enabled.

The transaction log files are truncated after each successful backup.

To select Exchange Server data

1. Click **Devices > Microsoft Exchange**.

The software shows the tree of Exchange Server Database Availability Groups (DAG), machines running Microsoft Exchange Server, and Exchange Server databases. If you configured Agent for Exchange as described in "Mailbox backup" (p. 528), mailboxes are also shown in this tree.

2. Browse to the data that you want to back up.

Expand the tree nodes or double-click items in the list to the right of the tree.

3. Select the data that you want to back up.

- If you select a DAG, one copy of each clustered database will be backed up. For more information about backing up DAGs, refer to "Protecting Database Availability Groups (DAG)" (p. 524).
- If you select a machine running Microsoft Exchange Server, all databases that are mounted to the Exchange Server running on the selected machine will be backed up.
- If you select databases directly, only the selected databases will be backed up.
- If you configured Agent for Exchange as described in "Mailbox backup" (p. 528), you can select mailboxes for backup.

If your selection includes multiple databases, they are processed two at a time. When the backup of the first group finishes, the backup of the next group will begin.

4. If prompted, provide the credentials to access the data.

5. Click **Protect**.

Protecting Always On Availability Groups (AAG)

Note

This feature is available with the Advanced Backup pack.

SQL Server high-availability solutions overview

The Windows Server Failover Clustering (WSFC) functionality enables you to configure a highly available SQL Server through redundancy at the instance level (Failover Cluster Instance, FCI) or at the database level (AlwaysOn Availability Group, AAG). You can also combine both methods.

In a Failover Cluster Instance, SQL databases are located on a shared storage. This storage can only be accessed from the active cluster node. If the active node fails, a failover occurs and a different node becomes active.

In an availability group, each database replica resides on a different node. If the primary replica becomes not available, a secondary replica residing on a different node is assigned the primary role.

Thus, the clusters are already serving as a disaster recovery solution themselves. However, there might be cases when the clusters cannot provide data protection: for example, in case of a database logical corruption, or when the entire cluster is down. Also cluster solutions do not protect from harmful content changes, as they usually immediately replicate to all cluster nodes.

Supported cluster configurations

This backup software supports *only* the Always On Availability Group (AAG) for SQL Server 2012 or later. Other cluster configurations, such as Failover Cluster Instances, database mirroring, and log shipping are *not* supported.

How many agents are required for cluster data backup and recovery?

For successful data backup and recovery of a cluster Agent for SQL has to be installed on each node of the WSFC cluster.

Backing up databases included in an AAG

1. Install Agent for SQL on each node of the WSFC cluster.
2. Select the AAG to backup as described in "Selecting SQL databases".
You must select the AAG itself to backup all databases of the AAG. To backup a set of databases, define this set of databases in all nodes of the AAG.

Warning!

The database set must be exactly the same in all nodes. If even one set is different, or not defined on all nodes, the cluster backup will not work correctly.

3. Configure the "[Cluster backup mode](#)" backup option.

Recovery of databases included in an AAG

1. Select the databases that you want to recover, and then select the recovery point from which you want to recover the databases.

When you select a clustered database under **Devices > Microsoft SQL > Databases**, and then click **Recover**, the software shows only the recovery points that correspond to the times when the selected copy of the database was backed up.

The easiest way to view all recovery points of a clustered database is to select the backup of the entire AAG [on the Backup storage tab](#). The names of AAG backups are based on the following template: <AAG name> - <protection plan name> and have a special icon.

2. To configure recovery, follow the steps described in "[Recovering SQL databases](#)", starting from step 5.

The software automatically defines a cluster node to which the data will be recovered. The node's name is displayed in the **Recover to** field. You can manually change the target node.

Important

A database that is included in an Always On Availability Group cannot be overwritten during a recovery because Microsoft SQL Server prohibits this. You need to exclude the target database from the AAG before the recovery. Or, just recover the database as a new non-AAG one. When the recovery is completed, you can reconstruct the original AAG configuration.

Protecting Database Availability Groups (DAG)

Note

This feature is available with the Advanced Backup pack.

Exchange Server clusters overview

The main idea of Exchange clusters is to provide high database availability with fast failover and no data loss. Usually, it is achieved by having one or more copies of databases or storage groups on the members of the cluster (cluster nodes). If the cluster node hosting the active database copy or the active database copy itself fails, the other node hosting the passive copy automatically takes over the operations of the failed node and provides access to Exchange services with minimal downtime. Thus, the clusters are already serving as a disaster recovery solution themselves.

However, there might be cases when failover cluster solutions cannot provide data protection: for example, in case of a database logical corruption, or when a particular database in a cluster has no copy (replica), or when the entire cluster is down. Also cluster solutions do not protect from harmful content changes, as they usually immediately replicate to all cluster nodes.

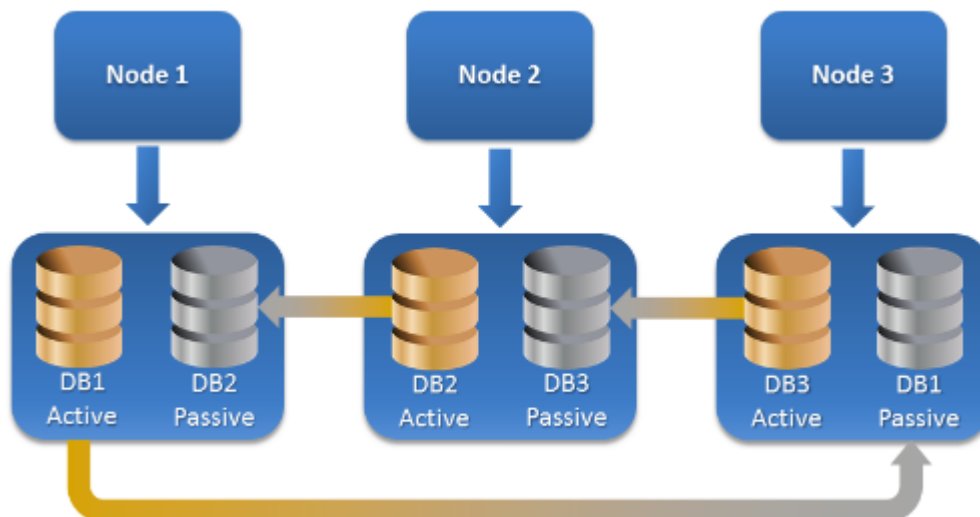
Cluster-aware backup

With cluster-aware backup, you back up only one copy of the clustered data. If the data changes its location within the cluster (due to a switchover or a failover), the software will track all relocations of this data and safely back it up.

Supported cluster configurations

Cluster-aware backup is supported *only* for Database Availability Group (DAG) in Exchange Server 2010 or later. Other cluster configurations, such as Single Copy Cluster (SCC) and Cluster Continuous Replication (CCR) for Exchange 2007, are *not* supported.

DAG is a group of up to 16 Exchange Mailbox servers. Any node can host a copy of mailbox database from any other node. Each node can host passive and active database copies. Up to 16 copies of each database can be created.



How many agents are required for cluster-aware backup and recovery?

For successful backup and recovery of clustered databases, Agent for Exchange has to be installed on each node of the Exchange cluster.

Note

After you install the agent on one of the nodes, the Cyber Protect console displays the DAG and its nodes under **Devices > Microsoft Exchange > Databases**. To install Agents for Exchange on the rest of the nodes, select the DAG, click **Details**, and then click **Install agent** next to each of the nodes.

Backing up the Exchange cluster data

1. When creating a protection plan, select the DAG as described in "Selecting Exchange Server data" (p. 521).
2. Configure the "Cluster backup mode" (p. 432) backup option.
3. Specify other settings of the protection plan [as appropriate](#).

Important

For cluster-aware backup, ensure to select the DAG itself. If you select individual nodes or databases inside the DAG, only the selected items will be backed up and the **Cluster backup mode** option will be ignored.

Recovering the Exchange cluster data

1. Select the recovery point for the database that you want to recover. Selecting an entire cluster for recovery is not possible.

When you select a copy of a clustered database under **Devices > Microsoft Exchange > Databases > <cluster name> > <node name>** and click **Recover**, the software shows only the recovery points that correspond to the times when this copy was backed up.


The easiest way to view all recovery points of a clustered database is to select its backup [on the Backup storage tab](#).

2. Follow the steps described in "Recovering Exchange databases" (p. 538), starting from step 5. The software automatically defines a cluster node to which the data will be recovered. The node's name is displayed in the **Recover to** field. You can manually change the target node.

Application-aware backup

Application-aware disk-level backup is available for physical machines, ESXi virtual machines, and Hyper-V virtual machines.

When you back up a machine running Microsoft SQL Server, Microsoft Exchange Server, or Active Directory Domain Services, enable **Application backup** for additional protection of these applications' data.



Why use application-aware backup?

By using application-aware backup, you ensure that:

- The applications are backed up in a consistent state and thus will be available immediately after the machine is recovered.
- You can recover the SQL and Exchange databases, mailboxes, and mailbox items without recovering the entire machine.
- The SQL transaction logs are truncated after each successful backup. SQL log truncation can be disabled in the [protection plan options](#). The Exchange transaction logs are truncated on virtual machines only. You can enable the [VSS full backup option](#) if you want to truncate Exchange transaction logs on a physical machine.
- If a domain contains more than one domain controller, and you recover one of them, a nonauthoritative restore is performed and a USN rollback will not occur after the recovery.

What do I need to use application-aware backup?

On a physical machine, Agent for SQL and/or Agent for Exchange must be installed, in addition to Agent for Windows.

On a virtual machine, no agent installation is required; it is presumed that the machine is backed up by Agent for VMware (Windows) or Agent for Hyper-V.

Note

For Hyper-V and VMware ESXi virtual machines that are running Windows Server 2022, application-aware backup is not supported in the agentless mode – that is, when the backup is performed by Agent for Hyper-V or Agent for VMware, respectively. To protect Microsoft applications on these machines, install Agent for Windows inside the guest operating system.

Agent for VMware (Virtual Appliance) can create application-aware backups, but cannot recover application data from them. To recover application data from backups created by this agent, you need Agent for VMware (Windows), Agent for SQL, or Agent for Exchange on a machine that has access to the location where the backups are stored. When configuring recovery of application data, select the recovery point on the **Backup storage** tab, and then select this machine in **Machine to browse from**.

Other requirements are listed in the "[Prerequisites](#)" and "[Required user rights](#)" sections.

Note

Application-aware backups of Hyper-V virtual machines may fail with the error "WMI 'ExecQuery' failed executing query." or "Failed to create a new process via WMI" if the backups are performed on a host under high load, due to no or delayed response from Windows Management Instrumentation. Retry these backups in a time slot when the load on the host is lower.

Required user rights for application-aware backups

An application-aware backup contains metadata of VSS-aware applications that are present on the disk. To access this metadata, the agent needs an account with the appropriate rights, which are listed below. You are prompted to specify this account when enabling application backup.

- For SQL Server:
The account must be a member of the **Backup Operators** or **Administrators** group on the machine and a member of the **sysadmin** role on each of the instances that you are going to back up.

Note

Only Windows authentication is supported.

- For Exchange Server:
Exchange 2007: The account must be a member of the **Administrators** group on the machine, and a member of the **Exchange Organization Administrators** role group.
Exchange 2010 and later: The account must be a member of the **Administrators** group on the machine, and a member of the **Organization Management** role group.
- For Active Directory:
The account must be a domain administrator.

Additional requirement for virtual machines

If the application runs on a virtual machine that is backed up by Agent for VMware or Agent for Hyper-V, ensure that User Account Control (UAC) is disabled on the machine.

If you do not want to disable UAC, you must provide the credentials of the built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

Note

Use the built-in domain administrator account that was configured as part of the creation of the domain. Accounts created later are not supported.

Additional requirements for machines running Windows

For all Windows versions, you must disable the User Account Control (UAC) policies to allow application-aware backups.

If you do not want to disable UAC, you must provide the credentials of the built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

Note

Use the built-in domain administrator account that was configured as part of the creation of the domain. Accounts created later are not supported.

To disable the UAC policies in Windows

1. In the Registry Editor, locate the following registry key:
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System
2. Change the **EnableLUA** value to **0**.
3. Restart the machine.

Mailbox backup

Mailbox backup is supported for Microsoft Exchange Server 2010 Service Pack 1 (SP1) and later.

Mailbox backup is available if at least one Agent for Exchange is registered on the management server. The agent must be installed on a machine that belongs to the same Active Directory forest as Microsoft Exchange Server.

Before backing up mailboxes, you must connect Agent for Exchange to the machine running the **Client Access** server role (CAS) of Microsoft Exchange Server. In Exchange 2016 and later, the CAS role is not available as a separate installation option. It is automatically installed as part of the Mailbox server role. Thus, you can connect the agent to any server running the **Mailbox role**.

Note

You can recover mailboxes and mailbox items also from database backups and application-aware backups. For more information, see "Recovering Exchange mailboxes and mailbox items" (p. 540). With database backups and application-aware backups you cannot create protection plans for individual mailboxes.

To connect Agent for Exchange to CAS

1. Click **Devices > Add**.
2. Click **Microsoft Exchange Server**.
3. Click **Exchange mailboxes**.
If no Agent for Exchange is registered on the management server, the software suggests that you install the agent. After the installation, repeat this procedure from step 1.
4. [Optional] If multiple Agents for Exchange are registered on the management server, click **Agent**, and then change the agent that will perform the backup.
5. In **Client Access server**, specify the fully qualified domain name (FQDN) of the machine where the **Client Access** role of Microsoft Exchange Server is enabled.
In Exchange 2016 and later, the Client Access services are automatically installed as part of the Mailbox server role. Thus, you can specify any server running the **Mailbox role**. We refer to this server as CAS later in this section.
6. In **Authentication type**, select the authentication type that is used by the CAS. You can select **Kerberos** (default) or **Basic**.
7. [Only for basic authentication] Select which protocol will be used. You can select **HTTPS** (default) or **HTTP**.
8. [Only for basic authentication with the HTTPS protocol] If the CAS uses an SSL certificate that was obtained from a certification authority, and you want the software to check the certificate when connecting to the CAS, select the **Check SSL certificate** check box. Otherwise, skip this step.
9. Provide the credentials of an account that will be used to access the CAS. The requirements for this account are listed in "[Required user rights](#)".
10. Click **Add**.

As a result, the mailboxes appear under **Devices > Microsoft Exchange > Mailboxes**.

Selecting Exchange Server mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan [as appropriate](#).

To select Exchange mailboxes

1. Click **Devices > Microsoft Exchange**.
The software shows the tree of Exchange databases and mailboxes.
2. Click **Mailboxes**, and then select the mailboxes that you want to back up.
3. Click **Protect**.

Required user rights

To access mailboxes, Agent for Exchange needs an account with the appropriate rights. You are prompted to specify this account when configuring various operations with mailboxes.

Membership of the account in the **Organization Management** role group enables access to any mailbox, including mailboxes that will be created in the future.

The minimum required user rights are as follows:

- The account must be a member of the **Server Management** and **Recipient Management** role groups.
- The account must have the **ApplicationImpersonation** management role enabled for all users or groups of users whose mailboxes the agent will access.

For information about configuring the **ApplicationImpersonation** management role, refer to the following Microsoft knowledge base article: <https://msdn.microsoft.com/en-us/library/office/dn722376.aspx>.

Recovering SQL databases

You can recover SQL databases from database backups and application-aware backups. For more information about the difference between the two backup types, refer to "Protecting Microsoft SQL Server and Microsoft Exchange Server" (p. 517).

You can recover SQL databases to the original instance, to a different instance on the original machine, or to an instance on a non-original machine. When you perform recovery to a non-original machine, Agent for SQL must be installed on the target machine.

Also, you can recover databases as files.

If you use Windows authentication for the SQL instance, you must provide credentials for an account that is a member of the **Backup Operators** or **Administrators** group on the machine and a member of the **sysadmin** role on the target instance. If you use SQL Server authentication, you must provide credentials for an account that is a member of the **sysadmin** role on the target instance.

System databases are recovered as user databases, with some distinctions. To learn more about these distinctions, refer to "Recovering system databases" (p. 537).

During a recovery, you can check the progress of the operation in the Cyber Protect console, on the **Monitoring > Activities** tab.

Recovering SQL databases to the original machine

You can recover SQL databases to their original instance, to a different instance on the original machine, or to an instance on a non-original target machine.

To recover SQL databases to the original machine

From a database backup

1. In the Cyber Protect console, go to **Devices > Microsoft SQL**.
2. Select the SQL Server instance or click the instance name to select specific databases that you want to recover, and then click **Recovery**.
If the machine is offline, the recovery points are not displayed. To recover data to a non-original machine, refer to "Recovering SQL databases to a non-original machine" (p. 532).
3. Select a recovery point.
The recovery points are filtered by location.
4. Click **Recover > Databases to an instance**.
By default, the instance and the databases are recovered to the original ones. You can also recover an original database as a new database.
5. [When recovering to a non-original instance on the same machine] Click **Target SQL Server instance**, select the target instance, and then click **Done**.
6. [When recovering a database as a new database] Click the database name, and then in **Recover to**, select **New database**.
7. [Optional] [Not available when recovering a database as a new database] To change the database state after recovery, click the database name, choose one of the following states, and then click **Done**.

- **Ready to use (RESTORE WITH RECOVERY)** (default)

After the recovery completes, the database will be ready for use. Users will have full access to it. The software will roll back all uncommitted transactions of the recovered database that are stored in the transaction logs. You will not be able to recover additional transaction logs from the native Microsoft SQL backups.

- **Non-operational (RESTORE WITH NORECOVERY)**

After the recovery completes, the database will be non-operational. Users will have no access to it. The software will keep all uncommitted transactions of the recovered database. You will be able to recover additional transaction logs from the native Microsoft SQL backups and thus reach the necessary recovery point.

- **Read-only (RESTORE WITH STANDBY)**

After the recovery completes, users will have read-only access to the database. The software will undo any uncommitted transactions. However, it will save the undo actions in a temporary standby file so that the recovery effects can be reverted.

This value is primarily used to detect the point in time when a SQL Server error occurred.

8. Click **Start recovery**.

From an application-aware backup

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the machine that originally contained the data that you want to recover, and then click **Recovery**.

If the machine is offline, the recovery points are not displayed. To recover data to a non-original machine, refer to "Recovering SQL databases to a non-original machine" (p. 532).

3. Select a recovery point.

The recovery points are filtered by location.

4. Click **Recover > SQL databases**.

5. Select the SQL Server instance or click the instance name to select specific databases that you want to recover, and then click **Recover**.

By default, the instance and the databases are recovered to the original ones. You can also recover an original database as a new database.

6. [When recovering to a non-original instance on the same machine] Click **Target SQL Server instance**, select the target instance, and then click **Done**.

7. [When recovering a database as a new database] Click the database name, and then in **Recover to**, select **New database**.

- Specify the new database name.
- Specify the new database path.
- Specify the log path.

8. [Optional] [Not available when recovering a database as a new database] To change the database state after recovery, click the database name, choose one of the following states, and then click **Done**.

- **Ready to use (RESTORE WITH RECOVERY)** (default)

After the recovery completes, the database will be ready for use. Users will have full access to it. The software will roll back all uncommitted transactions of the recovered database that are stored in the transaction logs. You will not be able to recover additional transaction logs from the native Microsoft SQL backups.

- **Non-operational (RESTORE WITH NORECOVERY)**

After the recovery completes, the database will be non-operational. Users will have no access to it. The software will keep all uncommitted transactions of the recovered database. You will be able to recover additional transaction logs from the native Microsoft SQL backups and thus reach the necessary recovery point.

- **Read-only (RESTORE WITH STANDBY)**

After the recovery completes, users will have read-only access to the database. The software will undo any uncommitted transactions. However, it will save the undo actions in a temporary standby file so that the recovery effects can be reverted.

This value is primarily used to detect the point in time when a SQL Server error occurred.

9. Click **Start recovery**.

Recovering SQL databases to a non-original machine

You can recover both application-aware backups and database backups to SQL Server instances on non-original target machines on which Agent for SQL is installed. The backups must be located on the cloud storage or on a shared storage that the target machine can access.

The SQL Server version on the target machine must be the same as the version on the source machine, or newer.


To recover SQL databases to a non-original machine

From Backup storage

This procedure applies to application-aware backups and database backups.


1. In the Cyber Protect console, go to **Backup storage**.
2. Select the location of the backup set from which you want to recover data.
3. In **Machine to browse from**, select the target machine.

This is the machine to which you will recover data. The target machine must be online.




4. Select the backup set, and then in the **Actions** pane, click **Show backups**.

Application-aware backup sets and database backup sets have different icons.



5. Select the recovery point from which you want to recover data.
6. [For database backups] Click **Recover SQL databases**.
7. [For application-aware backups] Click **Recover > SQL databases**.
8. Select the SQL Server instance or click the instance name to select specific databases that you want to recover, and then click **Recover**.
9. [If there is more than one SQL instance on the target machine] Click **Target SQL Server instance**, select the target instance, and then click **Done**.
10. Click the database name, specify the new database path and log path, and then click **Done**.
You can specify the same path in both fields, for example:

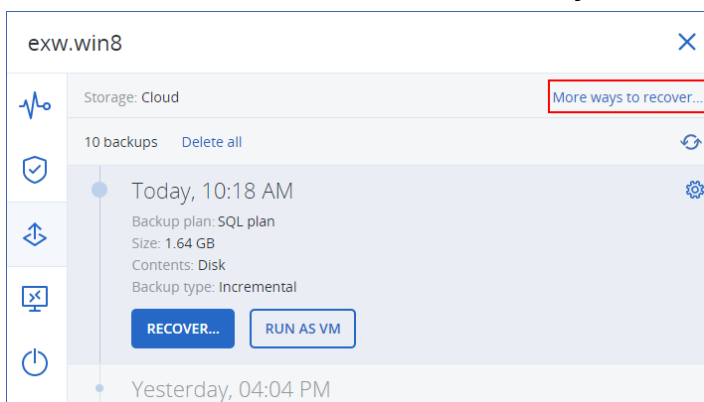


11. Click **Start recovery**.

From Devices

This procedure only applies to application-aware backups.

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the machine that originally contained the data that you want to recover, and then click **Recovery**.
3. [If the source machine is online] Click **More ways to recover**.



4. Click **Select machine** to select the target machine, and then click **OK**.
This is the machine to which you will recover data. The target machine must be online.
5. Select a recovery point.
The recovery points are filtered by location.
6. Click **Recover > SQL databases**.
7. Select the SQL Server instance or click the instance name to select specific databases that you want to recover, and then click **Recover**.
8. [If there is more than one SQL instance on the target machine] Click **Target SQL Server instance**, select the target instance, and then click **Done**.
9. Click the database name, specify the new database path and log path, and then click **Done**.
You can specify the same path in both fields, for example:

```
C:\Program Files\Microsoft SQL Server\MSSQL16.MSSQLSERVER\MSSQL\DATA\
```

10. Click **Start recovery**.

Recovering SQL databases as files

You can recover databases as files. This option might be useful if you need to extract data for data mining, audit, or further processing by third-party tools. To learn how to attach the SQL database files to a SQL Server instance, refer to "Attaching SQL Server databases" (p. 537).

You can recover databases as files to the original machine or to non-original target machines, on which Agent for SQL is installed. When you recover data to non-original machines, the backups must be located on the cloud storage or on a shared storage that the target machine can access.

Note

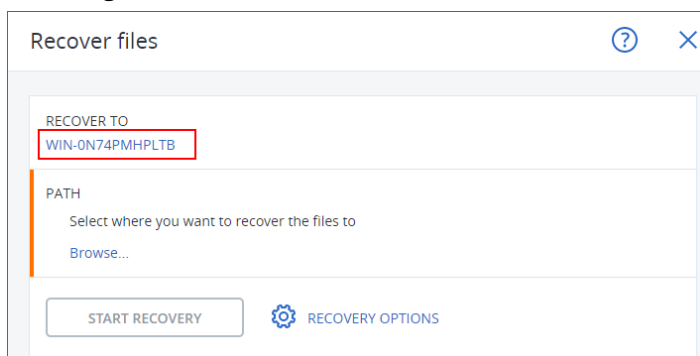
Recovering databases as files is the only recovery method if you use Agent for VMware (Windows). Recovering databases by using Agent for VMware (Virtual Appliance) is not possible.

To recover SQL databases as files

From a database backup

This procedure applies to online source machines.

1. In the Cyber Protect console, go to **Devices > Microsoft SQL**.
2. Select the databases that you want to recover, and then click **Recovery**.
3. Select a recovery point.
The recovery points are filtered by location.
4. Click **Recover > Databases as files**.
5. [When recovering to a non-original machine] In **Recover to**, select the target machine.
This is the machine to which you will recover data. The target machine must be online.
To change the selection, click the machine name, select another machine, and then click **OK**.




6. In **Path**, click **Browse**, select a local or network folder to save the files to, and then click **Done**.
7. Click **Start recovery**.

From an application-aware backup

This procedure applies to online source machines.

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the machine that originally contained the data that you want to recover, and then click **Recovery**.
3. Select a recovery point.
The recovery points are filtered by location.
4. Click **Recover > SQL databases**, select the databases that you want to recover, and then click **Recover as files**.
5. [When recovering to a non-original machine] In **Recover to**, select the target machine.
This is the machine to which you will recover data. The target machine must be online.
To change the selection, click the machine name, select another machine, and then click **OK**.




6. In **Path**, click **Browse**, select a local or network folder to save the files to, and then click **Done**.
7. Click **Start recovery**.

From a backup on an offline machine

This procedure applies to application-aware backups and database backups on source machines that are offline.


1. In the Cyber Protect console, go to **Backup storage**.
2. Select the location of the backup set from which you want to recover data.
3. In **Machine to browse from**, select the target machine.

This is the machine to which you will recover data. The target machine must be online.



4. Select the backup set, and then in the **Actions** pane, click **Show backups**.

Application-aware backup sets and database backup sets have different icons.



5. Select the recovery point from which you want to recover data.
6. [For database backups] Click **Recover SQL databases**.
7. [For application-aware backups] Click **Recover > SQL databases**.
8. Select the SQL Server instance or click the instance name to select specific databases that you want to recover, and then click **Recover as files**.
9. In **Path**, click **Browse**, select a local or a network folder to save the files to, and then click **Done**.
10. Click **Start recovery**.

Recovering system databases

All system databases of an instance are recovered at once. When recovering system databases, the software automatically restarts the destination instance in the single-user mode. After the recovery completes, the software restarts the instance and recovers other databases (if any).

Other things to consider when recovering system databases:

- System databases can only be recovered to an instance of the same version as the original instance.
- System databases are always recovered in the "ready to use" state.

Recovering the master database

System databases include the **master** database. The **master** database records information about all databases of the instance. Hence, the **master** database in a backup contains information about databases which existed in the instance at the time of the backup. After recovering the **master** database, you may need to do the following:

- Databases that have appeared in the instance after the backup was done are not visible by the instance. To bring these databases back to production, attach them to the instance manually by using SQL Server Management Studio.
- Databases that have been deleted after the backup was done are displayed as offline in the instance. Delete these databases by using SQL Server Management Studio.

Attaching SQL Server databases

This section describes how to attach a database in SQL Server by using SQL Server Management Studio. Only one database can be attached at a time.

Attaching a database requires any of the following permissions: **CREATE DATABASE**, **CREATE ANY DATABASE**, or **ALTER ANY DATABASE**. Normally, these permissions are granted to the **sysadmin** role of the instance.

To attach a database

1. Run Microsoft SQL Server Management Studio.
2. Connect to the required SQL Server instance, and then expand the instance.
3. Right-click **Databases** and click **Attach**.
4. Click **Add**.
5. In the **Locate Database Files** dialog box, find and select the .mdf file of the database.
6. In the **Database Details** section, make sure that the rest of database files (.ndf and .ldf files) are found.

Details. SQL Server database files may not be found automatically, if:

- They are not in the default location, or they are not in the same folder as the primary database file (.mdf). Solution: Specify the path to the required files manually in the **Current**

File Path column.

- You have recovered an incomplete set of files that make up the database. Solution: Recover the missing SQL Server database files from the backup.

7. When all of the files are found, click **OK**.

Recovering Exchange databases

This section describes recovery from both database backups and application-aware backups.

You can recover Exchange Server data to a live Exchange Server. This may be the original Exchange Server or an Exchange Server of the same version running on the machine with the same fully qualified domain name (FQDN). Agent for Exchange must be installed on the target machine.

The following table summarizes the Exchange Server data that you can select for recovery and the minimal user rights required to recover the data.

Exchange version	Data items	User rights
2007	Storage groups	Membership in the Exchange Organization Administrators role group.
2010/2013/2016/2019	Databases	Membership in the Server Management role group.

Alternatively, you can recover the databases (storage groups) as files. The database files, along with transaction log files, will be extracted from the backup to a folder that you specify. This can be useful if you need to extract data for an audit or further processing by third-party tools, or when the recovery fails for some reason and you are looking for a workaround to [mount the databases manually](#).

If you use only Agent for VMware (Windows), recovering databases as files is the only available recovery method. Recovering databases by using Agent for VMware (Virtual Appliance) is not possible.

We will refer to both databases and storage groups as "databases" throughout the below procedures.

To recover Exchange databases to a live Exchange Server

1. Do one of the following:
 - When recovering from an application-aware backup, under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the databases that you want to recover.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Do one of the following:

 - [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine

that has Agent for Exchange, and then select a recovery point.

- Select a recovery point on [the Backup storage tab](#).

The machine chosen for browsing in either of the above actions becomes a target machine for the Exchange data recovery.

4. Do one of the following:
 - When recovering from an application-aware backup, click **Recover > Exchange databases**, select the databases that you want to recover, and then click **Recover**.
 - When recovering from a database backup, click **Recover > Databases to an Exchange server**.
5. By default, the databases are recovered to the original ones. If the original database does not exist, it will be recreated.
To recover a database as a different one:
 - a. Click the database name.
 - b. In **Recover to**, select **New database**.
 - c. Specify the new database name.
 - d. Specify the new database path and log path. The folder you specify must not contain the original database and log files.
6. Click **Start recovery**.

The recovery progress is shown on the Activities tab.

To recover Exchange databases as files

1. Do one of the following:
 - When recovering from an application-aware backup, under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the databases that you want to recover.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Do one of the following:
 - [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine that has Agent for Exchange or Agent for VMware, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).The machine chosen for browsing in either of the above actions becomes a target machine for the Exchange data recovery.
4. Do one of the following:
 - When recovering from an application-aware backup, click **Recover > Exchange databases**, select the databases that you want to recover, and then click **Recover as files**.
 - When recovering from a database backup, click **Recover > Databases as files**.

5. Click **Browse**, and then select a local or a network folder to save the files to.
6. Click **Start recovery**.

The recovery progress is shown on the Activities tab.

Mounting Exchange Server databases

After recovering the database files, you can bring the databases online by mounting them. Mounting is performed by using Exchange Management Console, Exchange System Manager, or Exchange Management Shell.

The recovered databases will be in a Dirty Shutdown state. A database that is in a Dirty Shutdown state can be mounted by the system if it is recovered to its original location (that is, information about the original database is present in Active Directory). When recovering a database to an alternate location (such as a new database or as the recovery database), the database cannot be mounted until you bring it to a Clean Shutdown state by using the `Eseutil /r <Enn>` command. `<Enn>` specifies the log file prefix for the database (or storage group that contains the database) into which you need to apply the transaction log files.

The account you use to attach a database must be delegated an Exchange Server Administrator role and a local Administrators group for the target server.

For details about how to mount databases, see the following articles:

- Exchange 2010 or later: <http://technet.microsoft.com/en-us/library/aa998871.aspx>
- Exchange 2007: [http://technet.microsoft.com/en-us/library/aa998871\(v=EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa998871(v=EXCHG.80).aspx)

Recovering Exchange mailboxes and mailbox items

You can recover Exchange mailboxes and mailbox items from the following backups:

- Database backups
- Application-aware backups
- Mailbox backups

You can recover the following items:

- Mailboxes (except for archive mailboxes)
- Public folders

Note

Available only from database backups. See "Selecting Exchange Server data" (p. 521).

- Public folder items
- Email folders
- Email messages
- Calendar events
- Tasks

- Contacts
- Journal entries
- Notes

You can use search to locate the items.

The mailboxes or mailbox items can be recovered to a live Exchange Server or to Microsoft 365.

Recovery to an Exchange Server

Granular recovery can be performed to Microsoft Exchange Server 2010 Service Pack 1 (SP1) and later. The source backup may contain databases or mailboxes of any supported Exchange version.

Granular recovery can be performed by Agent for Exchange or Agent for VMware (Windows). The target Exchange Server and the machine running the agent must belong to the same Active Directory forest.

When a mailbox is recovered to an existing mailbox, the existing items with matching IDs are overwritten.

Recovery of mailbox items does not overwrite anything. Instead, the full path to a mailbox item is recreated in the target folder.

Requirements on user accounts

A mailbox being recovered from a backup must have an associated user account in Active Directory.

User mailboxes and their contents can be recovered only if their associated user accounts are *enabled*. Shared, room, and equipment mailboxes can be recovered only if their associated user accounts are *disabled*.

A mailbox that does not meet the above conditions is skipped during recovery.

If some mailboxes are skipped, the recovery will succeed with warnings. If all mailboxes are skipped, the recovery will fail.

Recovery to Microsoft 365

Recovery of Exchange data items to Microsoft 365, and vice versa, is supported on the condition that Agent for Microsoft 365 is installed locally.

Recovery can be performed from backups of Microsoft Exchange Server 2010 and later.

When a mailbox is recovered to an existing Microsoft 365 mailbox, the existing items are kept intact, and the recovered items are placed next to them.

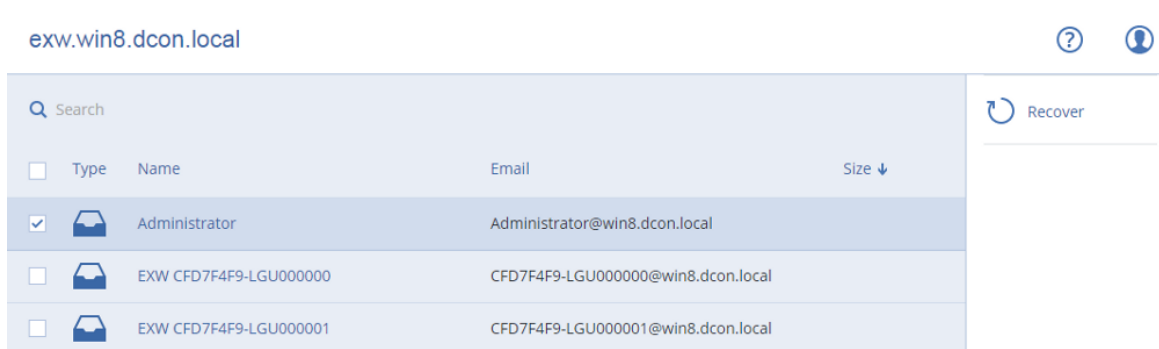
When recovering a single mailbox, you need to select the target Microsoft 365 mailbox. When recovering several mailboxes within one recovery operation, the software will try to recover each mailbox to the mailbox of the user with the same name. If the user is not found, the mailbox is skipped. If some mailboxes are skipped, the recovery will succeed with warnings. If all mailboxes are skipped, the recovery will fail.

For more information about recovery to Microsoft 365, refer to "Protecting Microsoft 365 data" (p. 553).

Recovering mailboxes

To recover mailboxes from an application-aware backup or a database backup

1. [Only when recovering from a database backup to Microsoft 365] If Agent for Microsoft 365 is not installed on the machine running Exchange Server that was backed up, do one of the following:
 - If there is not Agent for Microsoft 365 in your organization, install Agent for Microsoft 365 on the machine that was backed up (or on another machine with the same Microsoft Exchange Server version).
 - If you already have Agent for Microsoft 365 in your organization, copy libraries from the machine that was backed up (or from another machine with the same Microsoft Exchange Server version) to the machine with Agent for Microsoft 365, as described in "[Copying Microsoft Exchange libraries](#)".
2. Do one of the following:
 - When recovering from an application-aware backup: under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the database that originally contained the data that you want to recover.
3. Click **Recovery**.
4. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Use other ways to recover:
 - [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine that has Agent for Exchange or Agent for VMware, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).The machine chosen for browsing in either of the above actions will perform the recovery instead of the original machine that is offline.
5. Click **Recover > Exchange mailboxes**.
6. Select the mailboxes that you want to recover.
You can search mailboxes by name. Wildcards are not supported.



7. Click **Recover**.
8. [Only when recovering to Microsoft 365]:
 - a. In **Recover to**, select **Microsoft 365**.
 - b. [If you selected only one mailbox in step 6] In **Target mailbox**, specify the target mailbox.
 - c. Click **Start recovery**.

Further steps of this procedure are not required.

Click **Target machine with Microsoft Exchange Server** to select or change the target machine. This step allows recovery to a machine that is not running Agent for Exchange.

Specify the fully qualified domain name (FQDN) of a machine where the **Client Access** role (in Microsoft Exchange Server 2010/2013) or **Mailbox role** (in Microsoft Exchange Server 2016 or later) is enabled. The machine must belong to the same Active Directory forest as the machine that performs the recovery.

9. If prompted, provide the credentials of an account that will be used to access the machine. The requirements for this account are listed in "[Required user rights](#)".
10. [Optional] Click **Database to re-create any missing mailboxes** to change the automatically selected database.
11. Click **Start recovery**.

The recovery progress is shown on the **Activities** tab.

To recover a mailbox from a mailbox backup

1. Click **Devices > Microsoft Exchange > Mailboxes**.
2. Select the mailbox to recover, and then click **Recovery**.

You can search mailboxes by name. Wildcards are not supported.

If the mailbox was deleted, select it on [the Backup storage tab](#), and then click **Show backups**.
3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Mailbox**.
5. Perform steps 8-11 of the above procedure.

Recovering mailbox items

To recover mailbox items from an application-aware backup or a database backup

1. [Only when recovering from a database backup to Microsoft 365] If Agent for Microsoft 365 is not installed on the machine running Exchange Server that was backed up, do one of the following:
 - If there is not Agent for Microsoft 365 in your organization, install Agent for Microsoft 365 on the machine that was backed up (or on another machine with the same Microsoft Exchange Server version).
 - If you already have Agent for Microsoft 365 in your organization, copy libraries from the machine that was backed up (or from another machine with the same Microsoft Exchange Server version) to the machine with Agent for Microsoft 365, as described in "[Copying Microsoft Exchange libraries](#)".
2. Do one of the following:

- When recovering from an application-aware backup: under **Devices**, select the machine that originally contained the data that you want to recover.
- When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the database that originally contained the data that you want to recover.

3. Click **Recovery**.

4. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Use other ways to recover:

- [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine that has Agent for Exchange or Agent for VMware, and then select a recovery point.
- Select a recovery point on [the Backup storage tab](#).

The machine chosen for browsing in either of the above actions will perform the recovery instead of the original machine that is offline.

5. Click **Recover > Exchange mailboxes**.

6. Click the mailbox that originally contained the items that you want to recover.

7. Select the items that you want to recover.

The following search options are available. Wildcards are not supported.

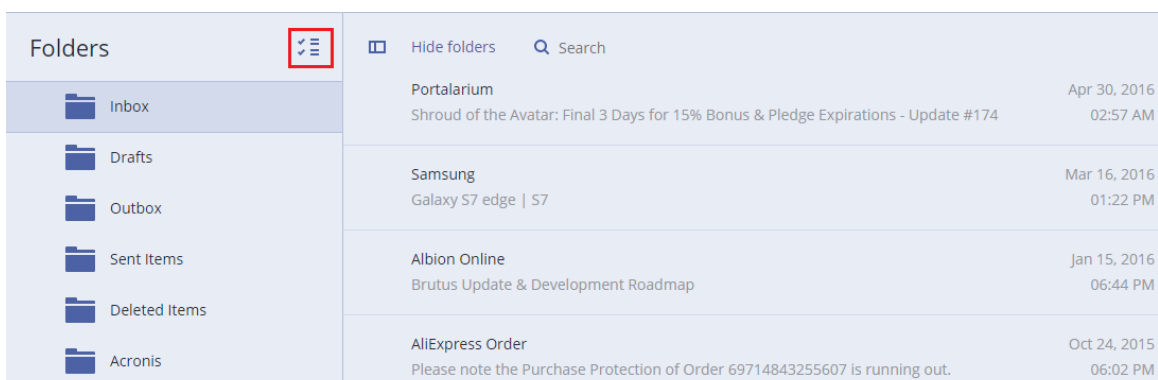
- For email messages: search by subject, sender, recipient, and date.
- For events: search by title and date.
- For tasks: search by subject and date.
- For contacts: search by name, email address, and phone number.

When an email message is selected, you can click **Show content** to view its contents, including attachments.

Note

Click the name of an attached file to download it.

To be able to select folders, click the recover folders icon.



8. Click **Recover**.

9. To recover to Microsoft 365, select **Microsoft 365** in **Recover to**.

To recover to an Exchange Server, keep the default **Microsoft Exchange** value in **Recover to**.

[Only when recovering to an Exchange Server] Click **Target machine with Microsoft Exchange Server** to select or change the target machine. This step allows recovery to a machine that is not running Agent for Exchange.

Specify the fully qualified domain name (FQDN) of a machine where the **Client Access** role (in Microsoft Exchange Server 2010/2013) or **Mailbox role** (in Microsoft Exchange Server 2016 or later) is enabled. The machine must belong to the same Active Directory forest as the machine that performs the recovery.

10. If prompted, provide the credentials of an account that will be used to access the machine. The requirements for this account are listed in "[Required user rights](#)".
11. In **Target mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist or a non-original target machine is selected, you must specify the target mailbox.
12. [Only when recovering email messages] In **Target folder**, view or change the target folder in the target mailbox. By default, the **Recovered items** folder is selected. Due to Microsoft Exchange limitations, events, tasks, notes, and contacts are restored to their original location regardless of any different **Target folder** specified.
13. Click **Start recovery**.

The recovery progress is shown on the **Activities** tab.

To recover a mailbox item from a mailbox backup

1. Click **Devices > Microsoft Exchange > Mailboxes**.
2. Select the mailbox that originally contained the items that you want to recover, and then click **Recovery**.

You can search mailboxes by name. Wildcards are not supported.

If the mailbox was deleted, select it on [the Backup storage tab](#), and then click **Show backups**.

3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Email messages**.
5. Select the items that you want to recover.

The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, and date.
- For events: search by title and date.
- For tasks: search by subject and date.
- For contacts: search by name, email address, and phone number.

When an email message is selected, you can click **Show content** to view its contents, including attachments.

Note

Click the name of an attached file to download it.

When an email message is selected, you can click **Send as email** to send the message to an email address. The message is sent from your administrator account's email address.

To be able to select folders, click the recover folders icon: 

6. Click **Recover**.
7. Perform steps 9-13 of the above procedure.

Copying Microsoft Exchange Server libraries

When [recovering Exchange mailboxes or mailbox items to Microsoft 365](#), you may need to copy the following libraries from the machine that was backed up (or from another machine with the same Microsoft Exchange Server version) to the machine with Agent for Microsoft 365.

Copy the following files, according to the Microsoft Exchange Server version that was backed up.

Microsoft Exchange Server version	Libraries	Default location
Microsoft Exchange Server 2010	ese.dll esecli2.dll store.exe	%ProgramFiles%\Microsoft\Exchange Server\V14\bin
Microsoft Exchange Server 2013	ese.dll	%ProgramFiles%\Microsoft\Exchange Server\V15\bin
	msvcr110.dll	%WINDIR%\system32
Microsoft Exchange Server 2016, 2019	ese.dll	%ProgramFiles%\Microsoft\Exchange Server\V15\bin
	msvcr110.dll	%WINDIR%\system32
	msvcp110.dll	

The libraries should be placed in the folder %ProgramData%\Acronis\ese. If this folder does not exist, create it manually.

Changing the SQL Server or Exchange Server access credentials

You can change access credentials for SQL Server or Exchange Server without re-installing the agent.

To change the SQL Server or Exchange Server access credentials

1. Click **Devices**, and then click **Microsoft SQL** or **Microsoft Exchange**.
2. Select the Always On Availability Group, Database Availability Group, SQL Server instance, or Exchange Server for which you want to change the access credentials.
3. Click **Specify credentials**.
4. Specify the new access credentials, and then click **OK**.

To change the Exchange Server access credentials for mailbox backup

1. Click **Devices > Microsoft Exchange**, and then expand **Mailboxes**.
2. Select the Exchange Server for which you want to change the access credentials.
3. Click **Settings**.

4. Under **Exchange administrator account**, specify the new access credentials, and then click **Save**.

Protecting mobile devices

The Cyber Protect app allows you to back up your mobile data to the Cloud storage and then recover it in case of loss or corruption. Note that backup to the cloud storage requires an account and the Cloud subscription.

Supported mobile devices

You can install the Cyber Protect app on a mobile device that runs one of the following operating systems:

- iOS 14 to iOS 16 (iPhone, iPod, iPad)
- Android 9 to Android 13

What you can back up

- Contacts (name, phone number, and email)
- Photos (the original size and format of your photos are preserved)
- Videos
- Calendars
- Reminders (only on iOS devices)

What you need to know

- You can back up the data only to the cloud storage.
- Any time you open the app, you will see the summary of data changes and can start a backup manually.
- The **Continuous backup** functionality is enabled by default. If this setting is turned on, the Cyber Protect app automatically detects new data on the fly and uploads it to the Cloud.
- The **Use Wi-Fi only** option is enabled by default in the app settings. If this setting is turned on, the Cyber Protect app will back up your data only when a Wi-Fi connection is available. If the Wi-Fi connection is lost, a backup process does not start. For the app to use cellular connection as well, turn this option off.
- The battery optimization on your device might prevent the Cyber Protect app from proper operation. To run backups on time, you should stop the battery optimization for the app.
- You have two ways to save energy:
 - The **Back up while charging** functionality which is disabled by default. If this setting is turned on, the Cyber Protect app will back up your data only when your device is connected to a power source. When the device is disconnected from a power source during a continuous backup process, the backup is paused.

- The **Save power mode** which is enabled by default. If this setting is turned on, the Cyber Protect app will back up your data only when your device battery is not low. When the device battery gets low, the continuous backup is paused.
- You can access the backed-up data from any mobile device registered under your account. This helps you transfer the data from an old mobile device to a new one. Contacts and photos from an Android device can be recovered to an iOS device and vice versa. You can also download a photo, video, or contact to any device by using the Cyber Protect console.
- The data backed up from mobile devices registered under your account is available only under this account. Nobody else can view or recover your data.
- In the Cyber Protect app, you can recover only the latest data versions. If you need to recover from a specific backup version, use the Cyber Protect console on either a tablet or a computer.
- Retention rules are not applied to backups of mobile devices.
- [Only for Android devices] If an SD card is present during a backup, the data stored on this card is also backed up. The data will be recovered to an SD card, to the folder **Recovered by Backup** if it is present during recovery, or the app will ask for a different location to recover the data to.

Where to get the Cyber Protect app

Depending on your mobile device, install the app from the App Store or Google Play.

How to start backing up your data

1. Open the app.
2. Sign in with your account.
3. Tap **Set up** to create your backup. Note that this button occurs only when you have no backup of your mobile device.
4. Select the data categories that you want to back up. By default, all categories are selected.
5. [optional step] Enable **Encrypt Backup** to protect your backup by encryption. In this case, you will need to also:
 - a. Enter an encryption password twice.

Note

Make sure you remember the password, because a forgotten password can never be restored or changed.

- b. Tap **Encrypt**.
6. Tap **Back up**.
7. Allow the app access to your personal data. If you deny access to some data categories, they will not be backed up.

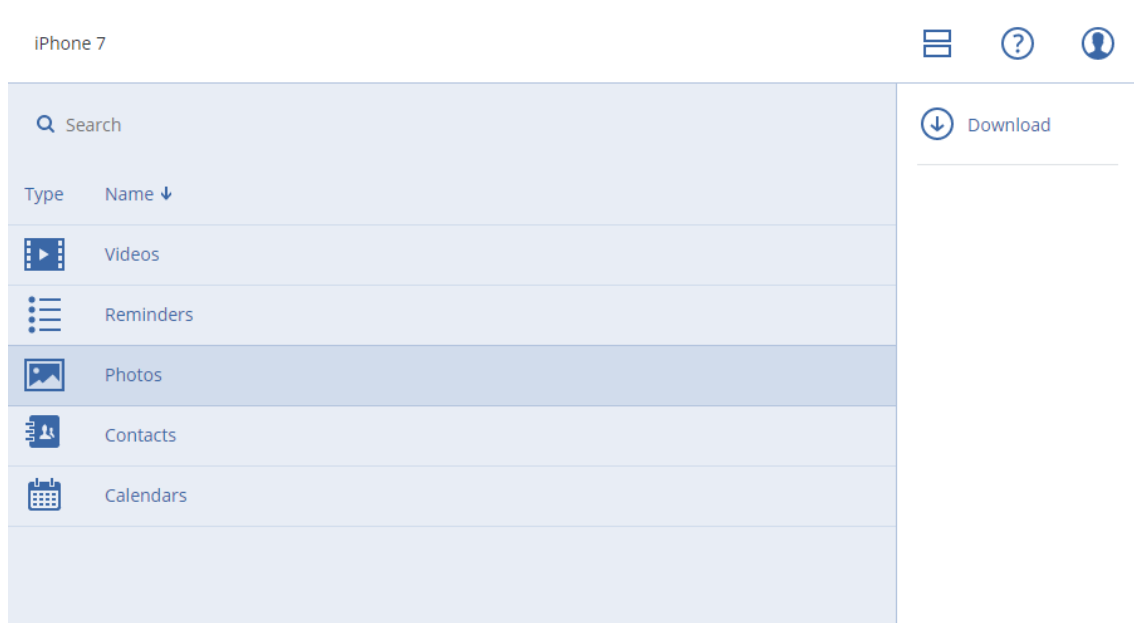
The backup starts.

How to recover data to a mobile device


1. Open the Cyber Protect app.
2. Tap **Browse**.
3. Tap the device name.
4. Do one of the following:
 - To recover all of the backed-up data, tap **Recover all**. No more actions are required.
 - To recover one or more data categories, tap **Select**, and then tap the check boxes for the required data categories. Tap **Recover**. No more actions are required.
 - To recover one or more data items belonging to the same data category, tap the data category. Proceed to further steps.
5. Do one of the following:
 - To recover a single data item, tap it.
 - To recover several data items, tap **Select**, and then tap the check boxes for the required data items.
6. Tap **Recover**.

How to review data via the Cyber Protect console

1. On a computer, open a browser and type the Cyber Protect console URL.
2. Sign in with your account.
3. In **All devices**, click **Recover** under your mobile device name.
4. Do any of the following:
 - To download all photos, videos, contacts, calendars, or reminders, select the respective data category. Click **Download**.



- To download individual photos, videos, contacts, calendars, or reminders, click the respective data category name, and then select the check boxes for the required data items. Click **Download**.



- To preview a photo, or a contact, click the respective data category name, and then click the required data item.

Protecting Hosted Exchange data

What items can be backed up?

You can back up user mailboxes, shared mailboxes, and group mailboxes. Optionally, you can choose to back up the archive mailboxes (**In-Place Archive**) of the selected mailboxes.

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders
- Email messages
- Calendar events
- Tasks
- Contacts
- Journal entries
- Notes

You can use search to locate the items.

When recovering mailboxes, mailbox items, public folders, and public folder items, you can select whether to overwrite the items in the target location.

When a mailbox is recovered to an existing mailbox, the existing items with matching IDs are overwritten.

Recovery of mailbox items does not overwrite anything. Instead, the full path to a mailbox item is recreated in the target folder.

Selecting Exchange Online mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan [as appropriate](#).

To select Exchange Online mailboxes

1. Click **Devices > Hosted Exchange**.
2. If multiple Hosted Exchange organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the mailboxes of all users and all shared mailboxes (including mailboxes that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up individual user or shared mailboxes, expand the **Users** node, select **All users**, select the users whose mailboxes you want to back up, and then click **Backup**.
 - To back up all group mailboxes (including mailboxes of groups that will be created in the future), expand the **Groups** node, select **All groups**, and then click **Group backup**.
 - To back up individual group mailboxes, expand the **Groups** node, select **All groups**, select the groups whose mailboxes you want to back up, and then click **Backup**.

Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **Devices > Hosted Exchange**.
2. If multiple Hosted Exchange organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox you want to recover, and then click **Recovery**.
 - To recover a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that you want to recover, and then click **Recovery**.
 - To recover a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox you want to recover, and then click **Recovery**.



- If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Entire mailbox**.
6. If multiple Hosted Exchange organizations are added to the Cyber Protection service, click **Hosted Exchange organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.
8. Click **Start recovery**.
9. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
10. Click **Proceed** to confirm your decision.

Recovering mailbox items

1. Click **Devices > Hosted Exchange**.
2. If multiple Hosted Exchange organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover items from a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.
You can search users and groups by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Email messages**.
6. Browse to the required folder or use search to obtain the list of the required items.
The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, attachment name, and date.
 - For events: search by title and date.
 - For tasks: search by subject and date.
 - For contacts: search by name, email address, and phone number.
7. Select the items that you want to recover. To be able to select folders, click the "recover folders" icon:  icon: 
- Additionally, you can do any of the following:
- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
 - When an email message or a calendar item is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
 - Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. Click **Recover**.
 9. If multiple Hosted Exchange organizations were added to the Cyber Protection service, click **Hosted Exchange organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
 10. In **Recover to mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.
 11. [Only when recovering to a user or a shared mailbox] In **Path**, view or change the target folder in the target mailbox. By default, the **Recovered items** folder is selected.
Group mailbox items are always recovered to the **Inbox** folder.
 12. Click **Start recovery**.
 13. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
 14. Click **Proceed** to confirm your decision.

Protecting Microsoft 365 data

Why back up Microsoft 365 data?

Even though Microsoft 365 is a set of cloud services, regular backups provide an additional layer of protection from user errors and intentional malicious actions. You can recover deleted items from a backup even after the Microsoft 365 retention period has expired. Also, you can keep a local copy of the Exchange Online mailboxes if it is required for regulatory compliance.

Backed-up data is automatically compressed and it uses less space on the backup location than on its original location. The compression level for cloud-to-cloud backups is fixed and corresponds to the **Normal** level of non-cloud-to-cloud backups. For more information about these levels, refer to "Compression level" (p. 433).

Cloud agent and local agent

For Microsoft 365 workloads, two agents are available:

- **Cloud agent**
The cloud agent provides extended backup functionality, which is directly accessible in the Cyber Protect console. No installation is required. For more information, see "Using the cloud Agent for Microsoft 365" (p. 562).
- **Local agent**
The local agent only provides backup of Exchange online mailboxes. This agent must be installed on a Windows machine that is connected to the Internet. For more information, see "Using the locally installed Agent for Office 365" (p. 558).

Azure Information Protection (AIP) is supported with both agents.

Note

For tenants in the Enhanced security mode, only the local agent is available. These tenants can back up only Microsoft 365 mailboxes. They cannot use the extended functionality provided by the cloud agent.

The following table summarizes the functionality of the agents.

	Local agent	Cloud agent
Data items that can be backed up	Exchange Online: user mailboxes and shared mailboxes (including mailboxes of users on a Kiosk plan and mailboxes on litigation hold)	<ul style="list-style-type: none"> • Exchange Online: <ul style="list-style-type: none"> ◦ user mailboxes and shared mailboxes (including mailboxes of users on a Kiosk plan and mailboxes on litigation hold) ◦ group mailboxes ◦ public folders • OneDrive: user files and folders • SharePoint Online: <ul style="list-style-type: none"> ◦ classic site collections ◦ group (team) sites ◦ communication sites ◦ individual data items • Microsoft 365 Teams: <ul style="list-style-type: none"> ◦ entire teams

	Local agent	Cloud agent
		<ul style="list-style-type: none"> ◦ team channels ◦ channel files ◦ team mailboxes ◦ files and email messages in team mailboxes ◦ meetings ◦ team sites • OneNote notebooks: as part of OneDrive, SharePoint Online, and Microsoft 365 Teams backups
Backup of archive mailboxes (In-Place Archive)	No	Yes
Backup schedule	User-defined	Up to six times per day*
Backup locations	Cloud storage, local folder, network folder	Cloud storage only (including partner-hosted storage)
Automatic protection of new Microsoft 365 users, groups, sites, and teams	No	Yes, by applying a protection plan to the All users, All groups, All sites, All teams groups
Protecting more than one Microsoft 365 organization	No	Yes
Granular recovery	Yes	Yes
Recovery to another user within one organization	Yes	Yes
Recovery to another organization	No	Yes
Recovery to an on-premises Microsoft Exchange Server	No	No
Maximum number of items that can be backed up without performance degradation	<p>When backing up to the cloud storage: 5000 mailboxes per company</p> <p>When backing up to other destinations: 2000 mailboxes per protection plan (no limitation for number of mailboxes per company)</p>	10 000 protected items (mailboxes, OneDrives, or sites) per company**

	Local agent	Cloud agent
Maximum number of manual backup runs	No	10 manual runs during an hour
Maximum number of simultaneous recovery operations	No	10 operations, including Google Workspace recovery operations

* The default option is **Once a day**. With the Advanced Backup pack, you can schedule up to six backups per day. The backups start at approximate intervals that depend on the current load of the cloud agent, which serves multiple customers in a data center. This ensures even load during the day and equal quality of service for all customers.

Note

The protection schedule might be affected by the operation of third-party services, for example, the accessibility of Microsoft 365 servers, throttling settings on the Microsoft servers, and others. See also <https://docs.microsoft.com/en-us/graph/throttling>.

** We recommend that you back up your protected items gradually and in this order:

1. Mailboxes.
2. After all mailboxes are backed up, proceed with OneDrives.
3. After OneDrive backup is completed, proceed with the SharePoint Online sites.

The first full backup may take several days, depending on the number of protected items and their size.

Required user rights

In Cyber Protection

The local agent must be registered under a company administrator account and used on the customer tenant level. Company administrators acting on the unit level, unit administrators, and users cannot back up or recover Microsoft 365 data.

The cloud agent can be used both on a customer tenant level and on a unit level. For more information about these levels and their respective administrators, see "Administering Microsoft 365 organizations added on different levels" (p. 563).

In Microsoft 365

Your account must be assigned the global administrator role in Microsoft 365.

To discover, back up, and recover Microsoft 365 public folders, at least one of your Microsoft 365 administrator accounts must have a mailbox and read/write rights to the public folders that you want to back up.

- The local agent will log in to Microsoft 365 by using this account. To enable the agent to access the contents of all mailboxes, this account will be assigned the **ApplicationImpersonation** management role. If you change the account password, update the password in the Cyber Protect console, as described in "Changing the Microsoft 365 access credentials" (p. 560).
- The cloud agent does not log in to Microsoft 365. You need to log into Microsoft 365 as a global administrator once, in order to grant the cloud agent the permissions required for its operation. The following permissions in Microsoft 365 are required:
 - Sign in and read user profiles
 - Read and write files in all site collections
 - Read and write all users' full profiles
 - Read and write all groups
 - Read directory data
 - Read all channel messages
 - Read and write managed metadata
 - Read and write items and lists in all site collections
 - Have full control of all site collection
 - Read and write items in all site collections
 - Use Exchange Web Services with full access to all mailboxes
- The cloud agent does not store your account credentials and does not use them to perform backup and recovery. Changing the credentials, disabling the account, or deleting the account does not affect the operation of the cloud agent.

Limitations

- All users with a mailbox or OneDrive are shown in the Cyber Protect console, including users without a Microsoft 365 license and users who are blocked from signing in to the Microsoft 365 services.
- A mailbox backup includes only folders visible to users. The **Recoverable items** folder and its subfolders (**Deletions, Versions, Purges, Audits, DiscoveryHold, Calendar Logging**) are not included in a mailbox backup.
- Automatic creation of users, public folders, groups, or sites during a recovery is not possible. For example, if you want to recover a deleted SharePoint Online site, first create a new site manually, and then specify it as the target site during a recovery.
- You cannot simultaneously recover items from different recovering points, even though you can select such items from the search results.
- During a backup, any sensitivity labels that are applied to the content will be preserved. Therefore, sensitive content might not be shown if it is recovered to a non-original location and its user has different access permissions.
- You cannot apply more than one individual backup plan to the same workload.
- When an individual backup plan and a group backup plan are applied to the same workload, the settings in the individual plan take precedence.

Microsoft 365 seats licensing report

Company administrators can download a report about the protected Microsoft 365 seats and their licensing. The report is in the CSV format and includes information about the licensing status of a seat and the reason why a license is used. The report includes also the protected seat name, associated email, group, Microsoft 365 organization, name and type of the protected workload.

This report is only available for tenants in which a Microsoft 365 Organization is registered.

To download the Microsoft 365 seats licensing report

1. Log in to the Cyber Protect console as a company administrator.
2. Click the account icon in the upper-right corner.
3. Click **Microsoft 365 seats licensing report**.

Logging

Actions with cloud-to-cloud resources, such as viewing the content of backed-up emails, downloading attachments or files, recovering emails to non-original mailboxes, or sending them as emails may violate user privacy. These actions are logged in **Monitoring > Audit log** in the Management Portal.

Using the locally installed Agent for Office 365

Adding a Microsoft 365 organization

To add a Microsoft 365 organization

1. Log in to the Cyber Protect console as a company administrator.
2. Click the account icon in the upper-right corner, and then click **Downloads > Agent for Office 365**.
3. Download the agent and install it on a Windows machine that is connected to the Internet.
4. In the Cyber Protect console, go to **Devices > Microsoft Office 365 (Local agent)**.
5. In the window that opens, enter your application ID, application secret, and Microsoft 365 tenant ID. For more information on how to find these, refer to "Obtaining application ID and application secret" (p. 559).
6. Click **OK**.

As a result, your organization data items appear in the Cyber Protect console, on the **Microsoft Office 365 (Local agent)** tab.

Important

There must be only one locally installed Agent for Office 365 in an organization (company group).

Obtaining application ID and application secret

To use the modern authentication for Office 365, you need to create a custom application in the Azure Active Directory and grant it specific API permissions. Thus, you will obtain the **application ID**, **application secret**, and **directory (tenant) ID** that you need to [enter in the Cyber Protect console](#).

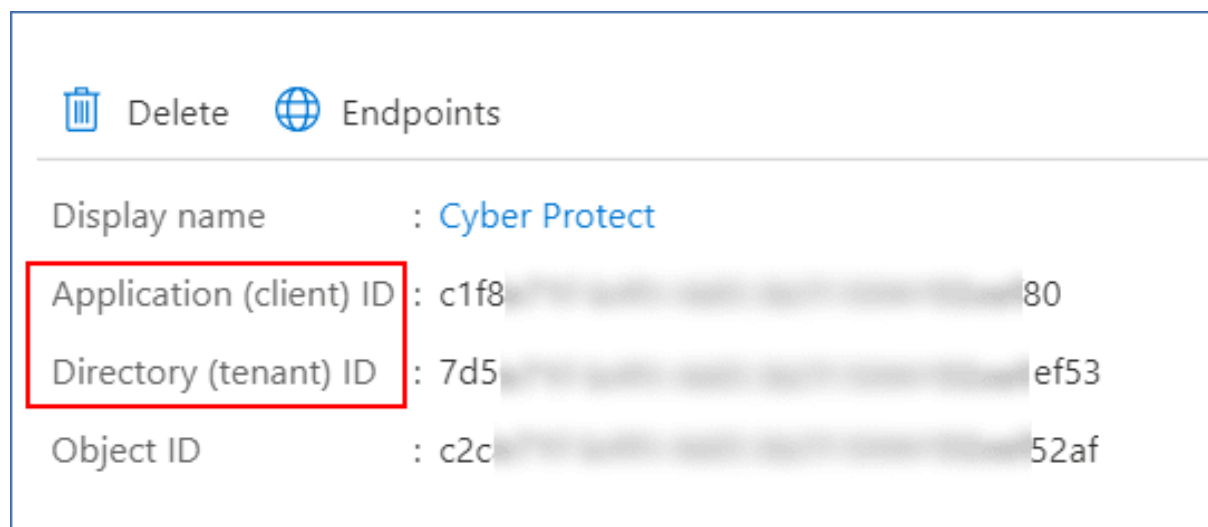
Note

On the machine where Agent for Office 365 is installed, ensure that you allow access to graph.microsoft.com through port 443.

To create an application in Azure Active Directory

1. Log in to the [Azure portal](#) as an administrator.
2. Navigate to **Azure Active Directory** > **App registrations**, and then click **New registration**.
3. Specify a name for your custom application, for example, Cyber Protection.
4. In **Supported Account types**, select **Accounts in this organizational directory only**.
5. Click **Register**.

Your application is now created. In the Azure portal, navigate to the application's **Overview** page and check your application (client) ID and directory (tenant) ID.



The screenshot shows the 'Overview' page for an application named 'Cyber Protect' in the Azure portal. At the top, there are 'Delete' and 'Endpoints' options. Below, the application details are listed:

Display name	: Cyber Protect
Application (client) ID	: c1f8 [redacted] 80
Directory (tenant) ID	: 7d5 [redacted] ef53
Object ID	: c2c [redacted] 52af

The 'Application (client) ID' and 'Directory (tenant) ID' rows are highlighted with a red box.

For more information on how to create an application in the Azure portal, refer to the [Microsoft documentation](#).


To grant your application the necessary API permissions

1. In the Azure portal, navigate to the application's **API permissions**, and then click **Add a permission**.
2. Select the **APIs my organization uses** tab, and then search for **Office 365 Exchange Online**.
3. Click **Office 365 Exchange Online**, and then click **Application permissions**.
4. Select the **full_access_as_app** check box, and then click **Add permissions**.

5. In **API permissions**, click **Add a permission**.
6. Select **Microsoft Graph**.
7. Select **Application permissions**.
8. Expand the **Directory** tab, and then select the **Directory.Read.All** check box. Click **Add permissions**.
9. Check all permissions, and then click **Grant admin consent for <your application's name>**.
10. Confirm your choice by clicking **Yes**.

To create an application secret

1. In the Azure portal, navigate to your application's **Certificates & secrets > New client secret**.
2. In the dialog box that opens, select Expires: **Never**, and then click **Add**.
3. Check your application secret in the **Value** field and make sure that you remember it.



For more information on the application secret, refer to the [Microsoft documentation](#).

Changing the Microsoft 365 access credentials

You can change access credentials for Microsoft 365 without re-installing the agent.

To change the Microsoft 365 access credentials

1. Click **Devices > Microsoft Office 365 (Local agent)**.
2. Select the Microsoft 365 organization.
3. Click **Specify credentials**.
4. Enter your application ID, application secret, and Microsoft 365 tenant ID. For more information on how to find these, refer to "Obtaining application ID and application secret" (p. 559).
5. Click **OK**.

Protecting Exchange Online mailboxes

What items can be backed up?

You can back up user mailboxes and shared mailboxes. Group mailboxes and archive mailboxes (**In-Place Archive**) cannot be backed up.

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders
- Email messages
- Calendar events
- Tasks
- Contacts
- Journal entries
- Notes

You can use search to locate the items.

When a mailbox is recovered to an existing mailbox, the existing items with matching IDs are overwritten.

Recovery of mailbox items does not overwrite anything. Instead, the full path to a mailbox item is recreated in the target folder.

Selecting Microsoft 365 mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan [as appropriate](#).

To select mailboxes

1. Click **Microsoft Office 365 (Local agent)**.
2. Select the mailboxes that you want to back up.
3. Click **Backup**.

Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **Microsoft Office 365 (Local agent)**.
2. Select the mailbox to recover, and then click **Recovery**.
You can search mailboxes by name. Wildcards are not supported.
If the mailbox was deleted, select it on [the Backup storage tab](#), and then click **Show backups**.
3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Mailbox**.
5. In **Target mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist, you must specify the target mailbox.
6. Click **Start recovery**.

Recovering mailbox items

1. Click **Microsoft Office 365 (Local agent)**.
2. Select the mailbox that originally contained the items that you want to recover, and then click **Recovery**.

You can search mailboxes by name. Wildcards are not supported.

If the mailbox was deleted, select it on [the Backup storage tab](#), and then click **Show backups**.

3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Email messages**.
5. Select the items that you want to recover.

The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, attachment name, and date.
- For events: search by title and date.
- For tasks: search by subject and date.
- For contacts: search by name, email address, and phone number.

When an email message is selected, you can click **Show content** to view its contents, including attachments.

Note

Click the name of an attached file to download it.

When an email message is selected, you can click **Send as email** to send the message to an email address. The message is sent from your administrator account's email address.

To be able to select folders, click the "recover folders" icon: 

6. Click **Recover**.
7. In **Target mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist, you must specify the target mailbox.
8. Click **Start recovery**.
9. Confirm your decision.

The mailbox items are always recovered to the **Recovered items** folder of the target mailbox.

Using the cloud Agent for Microsoft 365

Adding a Microsoft 365 organization

An administrator can add one or more Microsoft 365 organizations to a customer tenant or to a unit.

Company administrators add organizations to customer tenants. Unit administrators and customer administrators acting on the unit level add organizations to units.

To add a Microsoft 365 organization

1. Depending on where you need to add the organization, log in to the Cyber Protect console as a company administrator or unit administrator.
2. [For company administrators acting on the unit level] In the management portal, navigate to the desired unit.
3. Click **Devices > Add > Microsoft 365 Business**.
The software redirects you to the Microsoft 365 login page.
4. Sign in with the Microsoft 365 global administrator credentials.
Microsoft 365 displays a list of permissions that are necessary to back up and recover your organization's data.
5. Confirm that you grant the Cyber Protection service these permissions.

As a result, your Microsoft 365 organization appears under the **Devices** tab in the Cyber Protect console.

Useful tips

- The cloud agent synchronizes with Microsoft 365 every 24 hours, starting from the moment when the organization is added to the Cyber Protection service. If you add or remove a user, group, or site, you will not see this change in the Cyber Protect console immediately. To synchronize the change immediately, select the organization on the **Microsoft 365** page, and then click **Refresh**. For more information about synchronizing the resources of a Microsoft 365 organization and the Cyber Protect console, refer to "Discovering Microsoft 365 resources" (p. 565).
- If you applied a protection plan to the **All users, All groups, or All sites** group, the newly added items will be included in the backup only after synchronization.
- According to Microsoft policy, when a user, group, or site is removed from the Microsoft 365 graphical user interface, it remains available via an API for a few days. During this period, the removed item is inactive (grayed out) in the Cyber Protect console and is not backed up. When the removed item becomes unavailable via the API, it disappears from the Cyber Protect console. Its backups (if any) can be found at **Backup Storage > Cloud applications backups**.

Administering Microsoft 365 organizations added on different levels

Company administrators have full access to the Microsoft 365 organizations that are added to the customer tenant level.

Company administrators have limited access to the organizations that are added to a unit. In these organizations, shown with the unit name in brackets, company administrators can do the following:

- Recover data from backups.
Company administrators can recover data to all organizations in the tenant, regardless of the level on which these organizations are added.
- Browse backups and recovery points in backups.

- Delete backups and recovery points in backups.
- View alerts and activities.

Company administrators, when acting on the customer tenant level, cannot do the following:

- Add Microsoft 365 organizations to units.
- Delete Microsoft 365 organizations from units.
- Synchronize Microsoft 365 organizations that were added to a unit.
- View, create, edit, delete, apply, run, or revoke protection plans for data items in the Microsoft 365 organizations that are added to a unit.

Unit administrators and company administrators acting on the unit level have full access to the organizations that are added to a unit. However, they do not have access to any resources from the parent customer tenant, including the protection plans that are created in it.

Deleting a Microsoft 365 organization

Deleting a Microsoft 365 organization does not affect the existing backups of this organization's data. If you do not need these backups anymore, delete them first, and then delete the Microsoft 365 organization. Otherwise, the backups will still use cloud storage space that might be billed.

For more information about how to delete backups, see "To delete backups or backup archives" (p. 504).

To delete a Microsoft 365 organization

1. Depending on where the organization is added, sign in to the Cyber Protect console as a company administrator or unit administrator.
2. [For company administrators acting on the unit level] In the management portal, navigate to the desired unit.
3. Go to **Devices > Microsoft 365**.
4. Select the organization, and then click **Delete group**.

As a result, the backup plans applied to this group will be revoked.

However, you should additionally revoke access rights of the Backup Service application to Microsoft 365 organization data manually.

To revoke access rights

1. Log in to Microsoft 365 under a global administrator.
2. Go to **Admin Center > Azure Active Directory > Enterprise applications > All applications**.
3. Select the **Backup Service** application and drill down to it.
4. Go to the **Properties** tab, and then, on the action panel, click **Delete**.
5. Confirm the deletion operation.

As a result, access rights to the Microsoft 365 organization data will be revoked from the Backup Service application.

Discovering Microsoft 365 resources

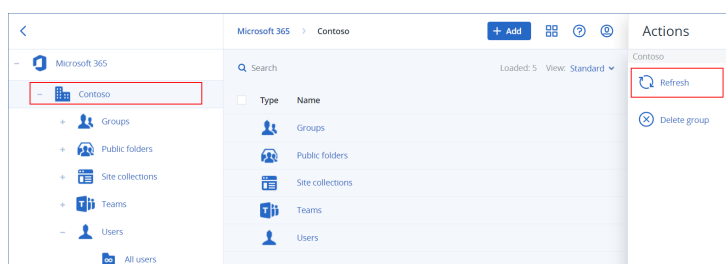
When you add a Microsoft 365 organization to the Cyber Protection service, the resources in this organization, such as mailboxes, OneDrive storages, Microsoft Teams, and SharePoint sites, are synchronized to the Cyber Protect console. This operation is called discovery and it is logged in **Monitoring > Activities**.

After the discovery operation completes, you can see the resources of the Microsoft 365 organization on the **Devices > Microsoft 365** tab in the Cyber Protect console, and you can apply backup plans to them.

An automatic discovery operation runs once a day to keep the list of resources in the Cyber Protect console up to date. You can also synchronize this list on demand, by re-running a discovery operation manually.

To re-run a discovery operation manually

1. In the Cyber Protect console, go to **Devices > Microsoft 365**.
2. Select your Microsoft 365 organization, and then, in the **Actions** pane, click **Refresh**.



Note

You can manually run a discovery operation up to 10 times per hour. When this number is reached, the allowed runs are reset to one per hour, and then every hour an additional run becomes available, until a total of 10 runs per hour is reached again.

Setting the frequency of Microsoft 365 backups

By default, Microsoft 365 backups run once a day and no additional scheduling options are available.

If the Advanced Backup pack is enabled in your tenant, you can configure more frequent backups. You can select the number of backups per day, but you cannot configure the backup start time. The backups start automatically at approximate intervals that depend on the current load of the cloud agent, which serves multiple customers in a data center. This ensures even load during the day, and equal quality of service for all customers.

The following options are available.

Scheduling options	Approximate interval between each backup
Once a day	24 hours
Twice a day (default)	12 hours
3 times a day	8 hours
6 times a day	4 hours

Note

Depending on the load on the cloud agent and possible throttling on the Microsoft 365 side, a backup might start later than scheduled or take longer to complete. If a backup takes longer than the average interval between two backups, the next backup will be rescheduled, which might result in fewer backups per day than selected. For example, only two backups per day might be able to complete, even though you selected six per day.

Backups of group mailboxes can only run once a day.

Protecting Exchange Online data

What items can be backed up?

You can back up user mailboxes, shared mailboxes, and group mailboxes. Optionally, you can choose to back up the online archive mailboxes (**In-Place Archive**) of the selected mailboxes.

Starting from version 8.0 of the Cyber Protection service, you can back up public folders. If your organization was added to the Cyber Protection service before the version 8.0 release, you need to re-add the organization to obtain this functionality. Do not delete the organization, simply repeat the steps described in "Adding a Microsoft 365 organization" (p. 562). As a result, the Cyber Protection service obtains the permission to use the corresponding API.

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders
- Email messages
- Calendar events
- Tasks
- Contacts
- Journal entries
- Notes

The following items can be recovered from a public folder backup:

- Subfolders
- Posts
- Email messages

You can use search to locate the items.

When recovering mailboxes, mailbox items, public folders, and public folder items, you can select whether to overwrite the items in the target location.

Selecting mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan [as appropriate](#).

To select Exchange Online mailboxes

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the mailboxes of all users and all shared mailboxes (including mailboxes that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up individual user or shared mailboxes, expand the **Users** node, select **All users**, select the users whose mailboxes you want to back up, and then click **Backup**.
 - To back up all group mailboxes (including mailboxes of groups that will be created in the future), expand the **Groups** node, select **All groups**, and then click **Group backup**.
 - To back up individual group mailboxes, expand the **Groups** node, select **All groups**, select the groups whose mailboxes you want to back up, and then click **Backup**.

Note

The cloud Agent for Microsoft 365 uses an account with the appropriate rights to access a group mailbox. Thus, to back up a group mailbox, at least one of the group owners must be licensed Microsoft 365 user with a mailbox. If the group is private or with hidden membership, the owner must also be a member of the group.

4. On the protection plan panel:
 - Ensure that the **Microsoft 365 mailboxes** item is selected in **What to back up**.
If some of the individually selected users do not have the Exchange service included in their Microsoft 365 plan, you will not be able to select this option.
If some of the selected users for group backup do not have the Exchange service included in their Microsoft 365 plan, you will be able to select this option, but the protection plan will not be applied to those users.
 - If you do not want to backup the archive mailboxes, disable the **Archive mailbox** switch.

Selecting public folders

Select the public folders as described below, and then specify other settings of the protection plan [as appropriate](#).

Note

Public folders consume licenses from your backup quota for Microsoft 365 seats.

To select Exchange Online public folders

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, expand the organization whose data you want to back up. Otherwise, skip this step.
3. Expand the **Public folders** node, and then select **All public folders**.
4. Do one of the following:
 - To back up all public folders (including public folders that will be created in the future), click **Group backup**.
 - To back up individual public folders, select the public folders that you want to back up, and then click **Backup**.
5. On the protection plan panel, ensure that the **Microsoft 365 mailboxes** item is selected in **What to back up**.

Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox you want to recover, and then click **Recovery**.
 - To recover a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that you want to recover, and then click **Recovery**.
 - To recover a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.
You can search users and groups by name. Wildcards are not supported.
4. Select a recovery point.

Note

To see only the recovery points that contain mailboxes, select **Mailboxes** in **Filter by content**.

5. Click **Recover > Entire mailbox**.
6. If multiple Microsoft 365 organizations are added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.
You cannot create a new target mailbox during recovery. To recover a mailbox to a new one, first you need to create the target mailbox in the desired Microsoft 365 organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Microsoft 365 every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Microsoft 365** page, and then click **Refresh**.
8. Click **Start recovery**.
9. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
10. Click **Proceed** to confirm your decision.

Recovering mailbox items

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover items from a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.
4. Select a recovery point.

Note

To see only the recovery points that contain mailboxes, select **Mailboxes** in **Filter by content**.

5. Click **Recover > Email messages**.

6. Browse to the required folder or use search to obtain the list of the required items.

The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, attachment name, and date. You can select a start date or an end date (both inclusive), or both dates to search within a time range.
- For events: search by title and date.
- For tasks: search by subject and date.
- For contacts: search by name, email address, and phone number.

7. Select the items that you want to recover. To be able to select folders, click the "recover folders"

icon: 

You cannot create a new target mailbox during recovery. To recover a new mailbox item to a new mailbox, first you need to create the target new mailbox item in Microsoft 365 organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Microsoft 365 every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Microsoft 365** page, and then click **Refresh**.

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
- When an email message or a calendar item is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
- Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. Click **Recover**.

9. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

10. In **Recover to mailbox**, view, change, or specify the target mailbox.

By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.

11. [Only when recovering to a user or a shared mailbox] In **Path**, view or change the target folder in the target mailbox. By default, the **Recovered items** folder is selected.

Group mailbox items are always recovered to the **Inbox** folder.

12. Click **Start recovery**.

13. Select one of the overwriting options:

- **Overwrite existing items**
- **Do not overwrite existing items**

14. Click **Proceed** to confirm your decision.

Recovering entire mailboxes to PST data files

To recover mailbox

Note

Mailbox recovery to PST files can be time-consuming, as it involves not only data transfer, but also data transformation using complex algorithms.

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover a user mailbox to PST data file, expand the **Users** node, select **All users**, select the mailbox you want to recover, and then click **Recovery**.
 - To recover a shared mailbox to PST data file, expand the **Users** node, select **All users**, select the mailbox that you want to recover, and then click **Recovery**.
 - To recover a group mailbox to PST data file, expand the **Groups** node, select **All groups**, select the group whose mailbox you want to recover, and then click **Recovery**.

You can search users and groups by name. Wildcards are not supported.

If the user, group, or the shared Outlook data file was deleted, select the item in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

4. Click **Recover > As PST files**.
5. Set the password to encrypt the archive with the PST files.
The password must contain at least one symbol.
6. Confirm the password and click **Done**.
7. The selected mailbox items will be recovered as PST data files and archived in ZIP format. The maximum size of one PST file is limited to 2 GB, so if the data you are recovering exceeds 2 GB, it will be split into several PST files. The ZIP archive will be protected with the password you set.
8. You will receive an email with a link to a ZIP archive containing the created PST files.
9. The administrator will receive an email notification that you have performed the recovery procedure.

To download the archive with PST files and complete recovery

1. Do one of the following:
 - To download the archive from the email, follow the **Download files** link.
The archive is available for download within 24 hours. If the link expires, repeat the recovery procedure.
 - To download the archive from the Cyber Protect console:
 - a. Go to **Backup Storage > PST files**.
 - b. Select the latest highlighted archive.
 - c. Click **Download** in the right pane.

The archive will be downloaded to the default download directory on your computer.

2. Extract the PST files from the archive using the password you set to encrypt the archive.
3. Open the PST files with Microsoft Outlook.

Important

Do not import these files to Microsoft Outlook by using the **Import and Export Wizard**.

Open the files by double-clicking them or right-clicking them and selecting **Open with... >**

Microsoft Outlook in the context menu.

Recovering mailbox items to PST files

To recover mailbox items

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover items from a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Click **Recover > Email messages**.
5. Browse to the required folder or use search to obtain the list of the required items.

The following search options are available. Wildcards are not supported.

 - For email messages: search by subject, sender, recipient, attachment name, and date.
 - For events: search by title and date.
 - For tasks: search by subject and date.
 - For contacts: search by name, email address, and phone number.
6. Select the items that you want to recover. To be able to select folders, click the "recover folders"

icon: 

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.

- When an email message or a calendar item is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
- Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

7. Click **Recover as PST files**.

8. Set the password to encrypt the archive with the PST files.

The password should contain at least one symbol.

9. Confirm the password and click **DONE**.

The selected mailbox items will be recovered as PST data files and archived in ZIP format. The maximum size of one PST file is limited to 2 GB, so if the data you are recovering exceeds 2 GB, it will be split into several PST files. The ZIP archive will be protected with the password you set.

You will receive an email with a link to a ZIP archive containing the created PST files.

The administrator will receive an email notification that you have performed the recovery procedure.

To download the archive with PST files and complete recovery

1. Do one of the following:

- To download the archive from the email, follow the **Download files** link.
The archive is available for download within 24 hours. If the link expires, repeat the recovery procedure.
- To download the archive from the Cyber Protect console:
 - a. Go to **Backup Storage > PST files**.
 - b. Select the latest highlighted archive.
 - c. Click **Download** in the right pane.

The archive will be downloaded to the default download directory on your computer.

2. Extract the PST files from the archive using the password you set to encrypt the archive.

3. Open the PST files with Microsoft Outlook.

Important

Do not import these files to Microsoft Outlook by using the **Import and Export Wizard**.

Open the files by double-clicking them or right-clicking them and selecting **Open with... >**

Microsoft Outlook in the context menu.


Recovering public folders and folder items

In order to recover a public folder or public folder items, at least one administrator of the target Microsoft 365 organization must have the **Owner's** rights for the target public folder. If the recovery fails with an error about denied access, assign these rights in the target folder properties, select the target organization in the Cyber Protect console, click **Refresh**, and then repeat the recovery.

To recover a public folder or folder items

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations are added to the Cyber Protection service, expand the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - Expand the **Public folders** node, select **All public folders**, select the public folder that you want to recover or that originally contained the items that you want to recover, and then click **Recovery**.
 - If the public folder was deleted, select it in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search public folders by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover data**.
6. Browse to the required folder or use search to obtain the list of the required items.
You can search email messages and posts by subject, sender, recipient, and date. Wildcards are not supported.
7. Select the items that you want to recover. To be able to select folders, click the "recover folders" icon: 

Additionally, you can do any of the following:

- When an email message or a post is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
- When an email message or a post is selected, click **Send as email** to send the item to specified email addresses. You can select the sender and write a text to be added to the forwarded item.
- Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. Click **Recover**.
9. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
10. In **Recover to public folder**, view, change, or specify the target public folder.
By default, the original folder is selected. If this folder does not exist or a non-original organization is selected, you must specify the target folder.
You cannot create a new public folder during recovery. To recover a public folder to a new one, first you need to create the target public folder in the desired Microsoft 365 organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Microsoft 365 every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Microsoft 365** page, and then click **Refresh**.

11. In **Path**, view or change the target subfolder in the target public folder. By default, the original path will be recreated.
12. Click **Start recovery**.
13. Select one of the overwriting options:

Option	Description
Overwrite existing items	All existing files in the destination location are overwritten.
Do not overwrite existing items	If the destination location contains a file of the same name, that file is not overwritten and the source file is not saved to the destination location.

14. Click **Proceed** to confirm your decision.

Protecting OneDrive files

What items can be backed up?

You can back up an entire OneDrive, or individual files and folders.

A separate option in the backup plan enables the backup of OneNote notebooks.

Files are backed up together with their sharing permissions. Advanced permission levels (**Design, Full, Contribute**) are not backed up.

Some files may contain sensitive information and the access to them may be blocked by a data loss prevention (DLP) rule in Microsoft 365. These files are not backed up, and no warnings are displayed after the backup operation completes.

Limitations

Backing up OneDrive content is not supported for shared mailboxes. To back up this content, convert the shared mailbox to a regular user account and ensure that OneDrive is enabled for that account.

What items can be recovered?

You can recover an entire OneDrive or any file or folder that was backed up.

You can use search to locate the items.

You can choose whether to recover the sharing permissions or let the files inherit the permissions from the folder to which they are recovered.

Sharing links for files and folders are not recovered.

Selecting OneDrive files

Select the files as described below, and then specify other settings of the protection plan [as appropriate](#).

To select OneDrive files

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the files of all users (including users that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up the files of individual users, expand the **Users** node, select **All users**, select the users whose files you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **OneDrive** item is selected in **What to back up**.
If some of the individually selected users do not have the OneDrive service included in their Microsoft 365 plan, you will not be able to select this option.
If some of the selected users for group backup do not have the OneDrive service included in their Microsoft 365 plan, you will be able to select this option, but the protection plan will not be applied to those users.
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all files).
 - Specify the files and folders to back up by adding their names or paths.
You can use wildcard characters (*, **, and ?). For more details about specifying paths and using wildcards, refer to "[File filters](#)".
 - Specify the files and folders to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single user.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the files and folders to skip during the backup.
File exclusions override the file selection; i.e. if you specify the same file in both fields, this file will be skipped during a backup.
 - [Optional] To back up the OneNote notebooks, enable the **Include OneNote** switch.

Recovering OneDrive and OneDrive files

Recovering an entire OneDrive

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.

- Expand the **Users** node, select **All users**, select the user whose OneDrive you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of the [Backup storage tab](#), and then click **Show backups**.

You can search users by name. Wildcards are not supported.

- Select a recovery point.

Note

To see only the recovery points that contain OneDrive files, select **OneDrive** in **Filter by content**.

- Click **Recover > Entire OneDrive**.
- If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
You cannot create a new OneDrive target during recovery. To recover a OneDrive to a new one, first you need to create the target OneDrive in Microsoft 365 organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Microsoft 365 every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Microsoft 365** page, and then click **Refresh**.
- In **Recover to drive**, view, change, or specify the target user.
By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user.
- Select whether to recover the sharing permissions for the files.
- Click **Start recovery**.
- Select one of the overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite existing files	All existing files in the destination location are overwritten, regardless of their last modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

Note

When you recover OneNote notebooks, both **Overwrite an existing file if it is older** and **Overwrite existing files** will result in overwriting the existing OneNote notebooks.

11. Click **Proceed** to confirm your decision.

Recovering OneDrive files

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose OneDrive files you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud Applications Backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search users by name. Wildcards are not supported.

4. Select a recovery point.

Note

To see only the recovery points that contain OneDrive files, select **OneDrive** in **Filter by content**.

5. Click **Recover > Files/folders**.
6. Browse to the required folder or use search to obtain the list of the required files and folders.
7. Select the files that you want to recover.

If the backup is not encrypted and you selected a single file, you can click **Show versions** to select the file version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. If you want to download a file, select the file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
9. Click **Recover**.
10. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

You cannot create a new OneDrive during recovery. To recover a file to a new OneDrive, first you need to create the target OneDrive in the desired Microsoft 365 organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Microsoft 365 every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Microsoft 365** page, and then click **Refresh**.
11. In **Recover to drive**, view, change, or specify the target user.

By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user.
12. In **Path**, view or change the target folder in the target user's OneDrive. By default, the original location is selected.
13. Select whether to recover the sharing permissions for the files.
14. Click **Start recovery**.

15. Select one of the file overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite existing files	All existing files in the destination location are overwritten, regardless of their last modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

Note

When you recover OneNote notebooks, both **Overwrite an existing file if it is older** and **Overwrite existing files** will result in overwriting the existing OneNote notebooks.

16. Click **Proceed** to confirm your decision.

Protecting SharePoint Online sites

What items can be backed up?

You can back up SharePoint classic site collections, group (modern team) sites, and communication sites. Also, you can select individual subsites, lists, and libraries for backup.

A separate option in the backup plan enables the backup of OneNote notebooks.

The following items are *skipped* during a backup:

- The **Look and Feel** site settings (except for **Title, description, and logo**).
- Site page comments and page comments settings (comments **On/Off**).
- The **Site features** site settings.
- Web part pages and web parts embedded in the wiki pages (due to SharePoint Online API limitations).
- Checked out files—files that are manually checked out for editing and all files that are created or uploaded in libraries, for which the option **Require Check Out** was enabled. To backup these files, first check them in.
- External data and Managed Metadata types of columns.
- The default site collection "domain-my.sharepoint.com". This is a collection where all of the organization users' OneDrive files reside.
- The contents of the recycle bin.

Limitations

- Titles and descriptions of sites/subsites/lists/columns are truncated during a backup if the title/description size is greater than 10000 bytes.
- You cannot back up previous versions of files created in SharePoint Online. Only the latest versions of the files are protected.
- You cannot back up the Preservation Hold library.
- You cannot back up sites created in the Business Productivity Online Suite (BPOS), the predecessor of Microsoft 365.
- You cannot back up the settings for sites that use the managed path /portals (for example, <https://<tenant>.sharepoint.com/portals/...>).
- Information Rights Management (IRM) settings of a list or a library can be recovered only if IRM is enabled in the target Microsoft 365 organization.

What items can be recovered?

The following items can be recovered from a site backup:

- Entire site
- Subsites
- Lists
- List items
- Document libraries
- Documents
- List item attachments
- Site pages and wiki pages

You can use search to locate the items.

Items can be recovered to the original or a non-original site. The path to a recovered item is the same as the original one. If the path does not exist, it is created.

You can choose whether to recover the sharing permissions or let the items inherit the permissions from the parent object after the recovery.

What items cannot be recovered?

- Subsites based on the **Visio Process Repository** template.
- Lists of the following types: **Survey list, Task list, Picture library, Links, Calendar, Discussion Board, External, and Import Spreadsheet.**
- Lists for which multiple content types are enabled.

Selecting SharePoint Online data

Select the data as described below, and then specify other settings of the protection plan [as appropriate](#).

To select SharePoint Online data

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up all classic SharePoint sites in the organization, including sites that will be created in the future, expand the **Site collections** node, select **All site collections**, and then click **Group backup**.
 - To back up individual classic sites, expand the **Site collections** node, select **All site collections**, select the sites that you want to back up, and then click **Backup**.
 - To back up all group (modern team) sites, including sites that will be created in the future, expand the **Groups** node, select **All groups**, and then click **Group backup**.
 - To back up individual group (modern team) sites, expand the **Groups** node, select **All groups**, select the groups whose sites you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **SharePoint sites** item is selected in **What to back up**.
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all items of the selected sites).
 - Specify the subsites, lists, and libraries to back up by adding their names or paths.
To back up a subsite or a top-level site list/library, specify its display name in the following format: /display name/**
To back up a subsite list/library, specify its display name in the following format: /subsite display name/list display name/**
The display names of subsites, lists, and libraries are shown on the **Site contents** page of a SharePoint site or subsite.
 - Specify the subsites to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single site.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the subsites, lists, and libraries to skip during the backup.
Item exclusions override the item selection; i.e. if you specify the same subsite in both fields, this subsite will be skipped during a backup.
 - [Optional] To back up the OneNote notebooks, enable the **Include OneNote** switch.

Recovering SharePoint Online data

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover data from a group (modern team) site, expand the **Groups** node, select **All groups**, select the group whose site originally contained the items that you want to recover, and then

click **Recovery**.

- To recover data from a classic site, expand the **Site Collections** node, select **All site collections**, select the site that originally contained the items that you want to recover, and then click **Recovery**.
- If the site was deleted, select it in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search groups and sites by name. Wildcards are not supported.

4. Select a recovery point.

Note

To see only the recovery points that contain SharePoint sites, select **SharePoint sites** in **Filter by content**.

5. Click **Recover SharePoint files**.
6. Browse to the required folder or use search to obtain the list of the required data items.
7. Select the items that you want to recover.
If the backup is not encrypted, you used search, and selected a single item in the search results, you can click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. [Optional] To download an item, select the item, click **Download**, select the location in which you want to save the item, and then click **Save**.
9. Click **Recover**.
10. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
11. In **Recover to site**, view, change, or specify the target site.
You cannot create a new SharePoint site during recovery. To recover a SharePoint site to a new one, first you need to create the target site in the desired Microsoft 365 organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Microsoft 365 every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Microsoft 365** page, and then click **Refresh**.
12. Select whether to recover the sharing permissions of the recovered items.
13. Click **Start recovery**.
14. Select one of the overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite	All existing files in the destination location are overwritten, regardless of their last

Option	Description
existing files	modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

Note

When you recover OneNote notebooks, both **Overwrite an existing file if it is older** and **Overwrite existing files** will result in overwriting the existing OneNote notebooks.

15. Click **Proceed** to confirm your decision.

Protecting Microsoft 365 Teams

What items can be backed up?

You can back up entire teams. This includes team name, team members list, team channels and their content, team mailbox and meetings, and team site.

A separate option in the backup plan enables the backup of OneNote notebooks.

What items can be recovered?

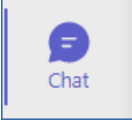
- Entire team
- Team channels
- Channel files
- Team mailbox
- Email folders in the team mailbox
- Email messages in the team mailbox
- Meetings
- Team site

You cannot recover conversations in team channels, but you can download them as a single html file.

Limitations

The following items are not backed up:

- The settings of the general channel (moderation preferences) – due to a [Microsoft Teams beta API](#) limitation.
- The settings of the custom channels (moderation preferences) – due to a [Microsoft Teams beta API](#) limitation.
- Meeting notes.

Messages in the chat section . These are private one-on-one or group chats.

-
- Stickers and praises.

Backup and recovery are supported for the following channel tabs:

- Word
- Excel
- PowerPoint
- PDF
- Document Library

Files that are shared in private channels are backed up, but not restored due to an API limitation.

Note

These files are stored in specific locations, separately from the files that are shared in public channels.

Selecting teams

Select teams as described below, and then specify other settings of the protection plan [as appropriate](#).

To select teams

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose teams you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up all the teams in the organization (including teams that will be created in the future), expand the **Teams** node, select **All teams**, and then click **Group backup**.
 - To back up individual teams, expand the **Teams** node, select **All teams**, select the teams that you want to back up, and then click **Backup**.

You can search teams by name. Wildcards are not supported.

4. On the protection plan panel:
 - Ensure that the **Microsoft Teams** item is selected in **What to back up**.
 - [Optional] In **How long to keep**, set the cleanup options.
 - [Optional] If you want to encrypt your backup, enable the **Encryption** switch, and then set your password and select the encryption algorithm.
 - [Optional] To back up the OneNote notebooks, enable the **Include OneNote** switch.

Recovering an entire team

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team that you want to recover, and then click **Recovery**.

You can search teams by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Entire Team**.

If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

6. In **Recover to team**, view the target team or select another.
By default, the original team is selected. If this team does not exist (for example, it was deleted) or you selected an organization that does not contain the original team, you must select a target team from the drop-down list.

You can recover a team only into an existing team. You cannot create teams during recovery operations.

7. Click **Start recovery**.
8. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**

Note

When you recover OneNote notebooks, both of the options **Overwrite existing content if it is older** and **Overwrite existing content** will result in overwriting the existing OneNote notebooks.

9. Click **Proceed** to confirm your decision.

When you delete a channel in Microsoft Teams' graphic interface, it is not immediately removed from the system. Thus, when you recover the whole team, this channel's name cannot be used and a postfix will be added to it.

Conversations are recovered as a single html file in the **Files** tab of the channel. You can find this file in a folder named according to the following pattern: <Team name>_<Channel name>_conversations_backup_<date of recovery>T<time of recovery>Z.

Note

After recovering a team or team channels, go to Microsoft Teams, select the channels that were recovered, and then click their **Files** tab. Otherwise, the subsequent backups of these channels will not include this tab's content – due to a [Microsoft Teams beta API](#) limitation.

Recovering team channels or files in team channels

To recover team channels

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose channels you want to recover, and then click **Recovery**.
4. Select a recovery point.
5. Click **Recover > Channels**.
6. Select the channels that you want to recover, and then click **Recover**. To select a channel in the main pane, select the check box in front of its name.

The following search options are available:

- For **Conversations**: sender, subject, content, language, attachment name, date or date range.
- For **Files**: file name or folder name, file type, size, date or date range of the last change.

Note

You can also download the files locally, instead of recovering them.

7. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
8. In **Recover to team**, view, change, or specify the target team.
By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target team.
9. In **Recover to channel**, view, change, or specify the target channel.
10. Click **Start recovery**.
11. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**

Note

When you recover OneNote notebooks, both of the options **Overwrite existing content if it is**

older and **Overwrite existing content** will result in overwriting the existing OneNote notebooks.

12. Click **Proceed** to confirm your decision.

Conversations are recovered as a single html file in the **Files** tab of the channel. You can find this file in a folder named according to the following pattern: <Team name>_<Channel name>_conversations_backup_<date of recovery>T<time of recovery>Z.

Note

After recovering a team or team channels, go to Microsoft Teams, select the channels that were recovered, and then click their **Files** tab. Otherwise, the subsequent backups of these channels will not include this tab's content – due to a [Microsoft Teams beta API](#) limitation.

To recover files in a team channel

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose channels you want to recover, and then click **Recovery**.
4. Select a recovery point.
5. Click **Recover > Channels**.
6. Select the desired channel, and then open the **Files** folder.

Browse to the required items or use search to obtain the list of the required items. The following search options are available: file name or folder name, file type, size, date or date range of the last change.
7. [Optional] To download an item, select the item, click **Download**, select the location in which you want to save the item, and then click **Save**.
8. Select the items that you want to recover, and then click **Recover**
9. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click Microsoft 365 organization to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
10. In **Recover to team**, view, change, or specify the target team.


By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target team.
11. In **Recover to channel**, view, change, or specify the target channel.
12. Select whether to recover the sharing permissions of the recovered items.
13. Click **Start recovery**.
14. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**

- **Do not overwrite existing content**

Note

When you recover OneNote notebooks, both of the options **Overwrite existing content if it is older** and **Overwrite existing content** will result in overwriting the existing OneNote notebooks.

15. Click **Proceed** to confirm your decision.

You cannot recover individual conversations. In the main pane, you can only browse the **Conversation** folder or download its content as a single html file. To do so, click the "recover folders" icon , select the desired **Conversations** folder, and then click **Download**.


You can search the messages in the **Conversation** folder by:

- Sender
- Content
- Attachment name
- Date

Recovering a team mailbox

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose mailbox you want to recover, and then click **Recovery**.

You can search teams by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Email messages**.
6. Click the "recover folders" icon , select the root mailbox folder, and then click **Recover**.

Note


You can also recover individual folders from the selected mailbox.

7. Click **Recover**.
8. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
9. In **Recover to mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.
10. Click **Start recovery**.
11. Select one of the overwriting options:

- **Overwrite existing items**
 - **Do not overwrite existing items**
12. Click **Proceed** to confirm your decision.

Recovering team mailbox items to PST files

To recover team mailbox items

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. You can search users and groups by name. Wildcards are not supported.
4. Expand the **Teams** node, select **All teams**, select a team whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
5. Click **Recover > Email messages**.
6. Browse to the required folder or use search to obtain the list of the required items.
The following search options are available. Wildcards are not supported.
 - For email messages: search by subject, sender, recipient, attachment name, and date.
 - For events: search by title and date.
 - For tasks: search by subject and date.
 - For contacts: search by name, email address, and phone number.
7. Select the items that you want to recover. To be able to select folders, click the "recover folders" icon: 
Additionally, you can do any of the following:
 - When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
 - When an email message or a calendar item is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
 - When the backup is not encrypted, you used search, and selected a single item from the search results: click **Show versions** to view the item version. You can select any backed-up version, no matter if it is earlier or later than the selected recovery point.
8. Click **Recover as PST files**.
9. Set the password to encrypt the archive with the PST files.
The password should contain at least one symbol.
10. Confirm the password and click **DONE**.

The selected mailbox items will be recovered as PST data files and archived in ZIP format. The maximum size of one PST file is limited to 2 GB, so if the data you are recovering exceeds 2 GB, it will be split into several PST files. The ZIP archive will be protected with the password you set.

You will receive an email with a link to a ZIP archive containing the created PST files.

The administrator will receive an email notification that you have performed the recovery procedure.

To download the archive with PST files and complete recovery

1. Do one of the following:
 - To download the archive from the email, follow the **Download files** link.
The archive is available for download within 24 hours. If the link expires, repeat the recovery procedure.
 - To download the archive from the Cyber Protect console:
 - a. Go to **Backup Storage > PST files**.
 - b. Select the latest highlighted archive.
 - c. Click **Download** in the right pane.The archive will be downloaded to the default download directory on your computer.
2. Extract the PST files from the archive using the password you set to encrypt the archive.
3. In Microsoft Outlook open or import the PST files. To learn how to do it, refer to Microsoft documentation.

Recovering email messages and meetings

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose email messages or meetings you want to recover, and then click **Recovery**.
You can search teams by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Email messages**.
6. Browse to the required item or use search to obtain the list of the required items.
The following search options are available:
 - For email messages: search by subject, sender, recipient, and date.
 - For meetings: search by event name and date.
7. Select the items that you want to recover, and then click **Recover**.

Note

You can find the meetings in the **Calendar** folder.

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
- When an email message or a meeting is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.

8. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
9. In **Recover to mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.
10. Click **Start recovery**.
11. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
12. Click **Proceed** to confirm your decision.

Recovering a team site or specific items of a site

1. Click **Microsoft 365**.
2. If multiple Microsoft 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose site you want to recover, and then click **Recovery**.
You can search teams by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Team site**.
6. Browse to the required item or use search to obtain the list of the required items.
7. [Optional] To download an item, select the item, click **Download**, select the location in which you want to save the item, and then click **Save**.
8. Select the items that you want to recover, and then click **Recover**.
9. If multiple Microsoft 365 organizations were added to the Cyber Protection service, click **Microsoft 365 organization** to view, change, or specify the target organization.
By default, the original organization and team are selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
10. In **Recover to team**, view, change, or specify the target team.
By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target site.
11. Select whether to recover the sharing permissions of the recovered items.
12. Click **Start recovery**.
13. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**

Note

When you recover OneNote notebooks, both of the options **Overwrite existing content if it is older** and **Overwrite existing content** will result in overwriting the exiting OneNote notebooks.

14. Click **Proceed** to confirm your decision.

Protecting OneNote notebooks

By default, OneNote notebooks are included in the backups of OneDrive files, Microsoft Teams, and SharePoint sites.

To exclude the OneNote notebooks from these backups, disable the **Include OneNote** switch in the respective backup plan.

Recovering backed-up OneNote notebooks

To learn how to recover a backed-up OneNote notebook, refer to the respective topic:

- For OneDrive backups, see "Recovering an entire OneDrive" (p. 576) or "Recovering OneDrive files" (p. 578).
- For Teams backups, see "Recovering an entire team" (p. 585), "Recovering team channels or files in team channels" (p. 586) or "Recovering a team site or specific items of a site" (p. 591).
- For SharePoint site backups, see "Recovering SharePoint Online data" (p. 581).

Supported versions

- OneNote (OneNote 2016 and later)
- OneNote for Windows 10

Limitations and known issues

- OneNote notebooks saved in OneDrive or SharePoint are limited to 2 GB. You cannot recover larger OneNote notebooks to OneDrive or SharePoint targets.
- OneNote notebooks with section groups are not supported.
- In backed-up OneNote notebooks that contain sections with non-default names, the first section is shown with the default name (such as *New section* or *Untitled section*). This might affect the section order in notebooks with multiple sections.
- When you recover OneNote notebooks, both of the options **Overwrite existing content if it is older** and **Overwrite existing content** will result in overwriting the exiting OneNote notebooks.
- When you recover an entire team, a team site, or the *Site Assets* folder of a team site, and you selected the **Overwrite existing content if it is older** or the **Overwrite existing content** option, the default OneNote notebook of that team will not be overwritten. The recovery succeeds with the warning *Failed to update the properties of file '/sites/<Team name>/SiteAssets/<OneNote notebook name>'.*

Protecting Microsoft 365 collaboration app seats

You can use the Advanced Email Security pack, that provides real-time protection for your Microsoft 365, Google Workspace, or Open-Xchange mailboxes:

- Antimalware and anti-spam
- URL scan in emails
- DMARC analysis
- Anti-phishing
- Impersonation protection
- Attachments scan
- Content disarm and reconstruction
- Graph of trust

You can also enable Microsoft 365 collaboration app seats, which allows the protection of Microsoft 365 cloud collaboration applications from content-borne security threats. These applications include OneDrive, SharePoint, and Teams.

Advanced Email Security can be enabled per workload or per gigabyte and will impact your licensing model.

To get to Advanced Email Security onboarding from Cyber Protect Cloud console

1. Click **Devices** > **Microsoft 365**.
2. Click the **Users** node and then click the **Go to Email Security** link at the top right.

Learn more about Advanced Email Security in the [Advanced Email Security data sheet](#).

For configuration instructions, see [Advanced Email Security with Perception Point](#).

Protecting Google Workspace data

Note

This feature is not available for tenants in the Enhanced security mode. For more information, refer to "Enhanced security mode" (p. 1011).

What does Google Workspace protection mean?

- Cloud-to-cloud backup and recovery of Google Workspace user data (Gmail mailboxes, Calendars, Contacts, Google Drives) and Google Workspace Shared drives.
- Granular recovery of emails, files, contacts, and other items.
- Support for several Google Workspace organizations and cross-organization recovery.
- Optional notarization of the backed-up files by means of the Ethereum blockchain database. When enabled, you can prove that a file is authentic and unchanged since it was backed up.
- Optional full-text search. When enabled, you can search emails by their content.

- Up to 5000 items (mailboxes, Google Drives, and Shared drives) per company can be protected without performance degradation.
- Backed-up data is automatically compressed and it uses less space on the backup location than on its original location. The compression level for cloud-to-cloud backups is fixed and corresponds to the **Normal** level of non-cloud-to-cloud backups. For more information about these levels, refer to "Compression level" (p. 433).

Required user rights

In Cyber Protection

In Cyber Protection, you need to be a company administrator acting on a customer tenant level. Company administrators acting on a unit level, unit administrators, and users cannot back up or recover Google Workspace data.

In Google Workspace

To add your Google Workspace organization to the Cyber Protection service, you must be signed in as a Super Admin with enabled API access (**Security > API reference > Enable API access** in the Google Admin console).

The Super Admin password is not stored anywhere and is not used to perform backup and recovery. Changing this password in Google Workspace does not affect Cyber Protection service operation.

If the Super Admin who added the Google Workspace organization is deleted from Google Workspace or assigned a role with less privileges, the backups will fail with an error like 'Access denied'. In this case, repeat the procedure described in "Adding a Google Workspace organization" (p. 595), and specify valid Super Admin credentials. To avoid this case, we recommend that you create a dedicated Super Admin user for backup and recovery purposes.

About the backup schedule

Because the cloud agent serves multiple customers, it determines the start time for each protection plan on its own, to ensure an even load during a day and an equal quality of service for all of the customers.

Each protection plan runs daily at the same time of day.

The default option is **Once a day**. With the Advanced Backup pack, you can schedule up to six backups per day. The backups start at approximate intervals that depend on the current load of the cloud agent, which serves multiple customers in a data center. This ensures even load during the day and equal quality of service for all customers.

Limitations

- The console shows only users that have an assigned Google Workspace license and a mailbox or Google Drive.
- Documents in the native Google formats are backed up as generic office documents and are shown with a different extension in the Cyber Protect console – such as .docx or .pptx, for example. The documents are converted back to their original format during recovery.
- No more than [10 manual backup runs during an hour](#).
- No more than 10 simultaneous recovery operations (this number includes both Microsoft 365 and Google Workspace recovery).
- You cannot simultaneously recover items from different recovering points, even though you can select such items from the search results.
- The backups of deleted Google Workspace user accounts are not automatically deleted from the cloud storage. These backups are billed for the storage space that they use.
- You cannot apply more than one individual backup plan to the same workload.
- When an individual backup plan and a group backup plan are applied to the same workload, the settings in the individual plan take precedence.

Logging

Actions with cloud-to-cloud resources, such as viewing the content of backed-up emails, downloading attachments or files, recovering emails to non-original mailboxes, or sending them as emails may violate user privacy. These actions are logged in **Monitoring > Audit log** in the Management Portal.

Adding a Google Workspace organization

To add a Google Workspace organization to the Cyber Protection service, you need a dedicated personal Google Cloud project. For more information about how to create and configure such a project, refer to "Creating a personal Google Cloud project" (p. 596).

To add a Google Workspace organization by using a dedicated personal Google Cloud project

1. Log in to the Cyber Protect console as a company administrator.
2. Click **Devices > Add > Google Workspace**.
3. Enter the email address of a Super Administrator of your Google Workspace account.
For this procedure, it is irrelevant whether 2-Step Verification is enabled for the Super Administrator email account.
4. Browse for the JSON file that contains the private key of the service account that you created in your Google Cloud project.
You can also paste the file content as text.
5. Click **Confirm**.

As a result, your Google Workspace organization appears under the **Devices** tab in the Cyber Protect console.

Useful tips

- After adding a Google Workspace organization, the user data and Shared drives in both the primary domain and all the secondary domains, if there are any, will be backed up. The backed-up resources will be displayed in one list, and will not be grouped by their domain.
- The cloud agent synchronizes with Google Workspace every 24 hours, starting from the moment when the organization is added to the Cyber Protection service. If you add or remove a user or Shared drive, you will not see this change in the Cyber Protect console immediately. To synchronize the change immediately, select the organization on the **Google Workspace** page, and then click **Refresh**.

For more information about synchronizing the resources of a Google Workspace organization and the Cyber Protect console, refer to "Discovering Google Workspace resources" (p. 599).

- If you applied a protection plan to the **All users** or **All Shared drives** group, the newly added items will be included in the backup only after the synchronization.
- According to Google policy, when a user or Shared drive is removed from the Google Workspace graphical user interface, it remains available via an API for a few days. During this period, the removed item is inactive (grayed out) in the Cyber Protect console and is not backed up. When the removed item becomes unavailable via the API, it disappears from the Cyber Protect console. Its backups (if any) can be found at **Backup storage > Cloud applications backups**.

Creating a personal Google Cloud project

To add your Google Workspace organization to the Cyber Protection service by using a dedicated Google Cloud project, you need to do the following:


1. Create a new Google Cloud project.
2. Enable the required APIs for this project.
3. Configure the credentials for this project:
 - a. Configure the OAuth consent screen.
 - b. Create and configure the service account for the Cyber Protection service.
4. Grant the new project access to your Google Workspace account.

Note


This topic contains a description of third-party user interface that might be subject to change without prior notice.

To create a new Google Cloud project

1. Sign in to the Google Cloud Platform (console.cloud.google.com) as a Super Administrator.
2. In the Google Cloud Platform console, click the project picker in the upper-left corner.



3. In the screen that opens, select an organization, and then click **New project**.



4. Specify a name for your new project.
5. Click **Create**.

As a result, your new Google Cloud project is created.

To enable the required APIs for this project

1. In the Google Cloud Platform console, select your new project.
2. From the navigation menu, select **APIs and services > Enabled APIs and services**.
3. Disable all the APIs that are enabled by default in this project, one by one:
 - a. Scroll down the **Enabled APIs and services** page, and then click the name of an enabled API. The **API/Service details** page of the selected API opens.
 - b. Click **Disable API**, and then confirm your choice by clicking **Disable**.
 - c. [If prompted] Confirm your choice by clicking **Confirm**.
 - d. Go back to **APIs and services > Enabled APIs and services**, and disable the next API.
4. From the navigation menu, select **APIs and services > Library**.
5. In the API library, enable the following APIs, one by one:
 - Admin SDK API
 - Gmail API
 - Google Calendar API
 - Google Drive API
 - Google People API

Use the search bar to find the required APIs. To enable an API, click its name, and then click **Enable**. To search for the next API, go back to the API library, by selecting **APIs and services > Library** from the navigation menu.

To configure the OAuth consent screen

1. From the navigation menu in the Google Cloud Platform, select **APIs and services > OAuth consent screen**.
2. In the window that opens, select **Internal** for user type, and then click **Create**.
3. In the **App name** field, specify a name for your application.
4. In the **User support email** field, enter the Super Administrator email.
5. In the **Developer contact information** field, enter the Super Administrator email.
6. Leave all other fields blank, and then click **Save and continue**.

7. On the **Scopes** page, click **Save and continue**, without changing anything.
8. On the **Summary** page, verify your settings, and then click **Back to dashboard**.

To create and configure the service account for the Cyber Protection service

1. From the navigation menu in the Google Cloud Platform, select **IAM & Admin > Service accounts**.
2. Click **Create service account**.
3. Specify a name for the service account.
4. [Optional] Specify a description for the service account.
5. Click **Create and continue**.
6. Do not change anything in the **Grant this service account access to the project** and **Grant users access to this service account** steps.
7. Click **Done**.
The **Service accounts** page opens.
8. On the **Service accounts** page, select the new service account, and then under **Actions**, click **Manage keys**.
9. Under **Keys**, click **Add key > Create new key**, and then select the **JSON** key type.
10. Click **Create**.

As a result, a JSON file with the private key of the service account is automatically downloaded to your machine. Store this file securely because you need it to add your Google Workspace organization to the Cyber Protection service.

To grant the new project access to your Google Workspace account

1. From the navigation menu in the Google Cloud Platform, select **IAM & Admin > Service Accounts**.
2. In the list, find the service account that you created, and then copy the client ID that is shown in the **OAuth 2.0 Client ID** column.
3. Sign in to the Google Admin console (admin.google.com) as a Super Administrator.
4. From the navigation menu, select **Security > Access and data control > API controls**.
5. Scroll down the **API controls** page, and then under **Domain-wide delegation**, click **Manage domain-wide delegation**.
The **Domain-wide delegation** page opens.
6. On the **Domain-wide delegation** page, click **Add new**.
The **Add a new client ID** window opens.
7. In the **Client ID** field, enter the client ID of your service account client.
8. In the **OAuth scopes** field, copy and paste the following comma-delimited list of scopes:

```
https://mail.google.com,https://www.googleapis.com/auth/contacts,https://www.googleapis.com/auth/calendar,https://www.googleapis.com/auth/admin.directory.user.readonly,https://www.googleapis.com/auth/admin.directory.domain.readonly,https://www.googleapis.com/auth/drive,https://www.googleapis.com/auth/gmail.modify
```

Alternatively, you can add the scopes one per line:

- <https://mail.google.com>
- <https://www.googleapis.com/auth/contacts>
- <https://www.googleapis.com/auth/calendar>
- <https://www.googleapis.com/auth/admin.directory.user.readonly>
- <https://www.googleapis.com/auth/admin.directory.domain.readonly>
- <https://www.googleapis.com/auth/drive>
- <https://www.googleapis.com/auth/gmail.modify>

9. Click **Authorise**.

As a result, your new Google Cloud project can access the data in your Google Workspace account. To back up the data, you need to link this project to the Cyber Protection service. For more information on how to do this, refer to "To add a Google Workspace organization by using a dedicated personal Google Cloud project" (p. 595).

If you need to revoke the access of your Google Cloud project to your Google Workspace account, and respectively the access of the Cyber Protection service, delete the API client that your project uses.

To revoke access to your Google Workspace account

1. In the Google Admin console (admin.google.com), sign in as a Super Administrator.
2. From the navigation menu, select **Security > Access and data control > API controls**.
3. Scroll down the **API controls** page, and then under **Domain-wide delegation**, click **Manage domain-wide delegation**.

The **Domain-wide delegation** page opens.

4. On the **Domain-wide delegation** page, select the API client that your project uses, and then click **Delete**.

As a result, your Google Cloud project and the Cyber Protection service will not be able to access your Google Workspace account and back up the data in it.

Discovering Google Workspace resources


When you add a Google Workspace organization to the Cyber Protection service, the resources in this organization, such as mailboxes and Google Drives, are synchronized to the Cyber Protect console. This operation is called discovery and it is logged in **Monitoring > Activities**.

After the discovery operation completes, you can see the resources of the Google Workspace organization on the **Devices > Google Workspace** tab in the Cyber Protect console, and you can apply backup plans to them.

An automatic discovery operation runs once a day to keep the list of resources in the Cyber Protect console up to date. You can also synchronize this list on demand, by re-running a discovery operation manually.

To re-run a discovery operation manually

1. In the Cyber Protect console, go to **Devices > Google Workspace**.
2. Select your Google Workspace organization, and then, in the **Actions** pane, click **Refresh**.



Note

You can manually run a discovery operation up to 10 times per hour. When this number is reached, the allowed runs are reset to one per hour, and then every hour an additional run becomes available, until a total of 10 runs per hour is reached again.

Setting the frequency of Google Workspace backups

By default, Google Workspace backups run once a day and no additional scheduling options are available.

If the Advanced Backup pack is enabled in your tenant, you can configure more frequent backups. You can select the number of backups per day, but you cannot configure the backup start time. The backups start automatically at approximate intervals that depend on the current load of the cloud agent, which serves multiple customers in a data center. This ensures even load during the day, and equal quality of service for all customers.

The following options are available.

Scheduling options	Approximate interval between each backup
Once a day	24 hours
Twice a day (default)	12 hours
3 times a day	8 hours
6 times a day	4 hours

Note

Depending on the load on the cloud agent and possible throttling on the Google Workspace side, a backup might start later than scheduled or take longer to complete. If a backup takes longer than the average interval between two backups, the next backup will be rescheduled, which might result in fewer backups per day than selected. For example, only two backups per day might be able to complete, even though you selected six per day.

Protecting Gmail data

What items can be backed up?

You can back up Gmail users' mailboxes. A mailbox backup also includes the Calendar and Contacts data. Optionally, you can choose to back up the shared calendars.

The following items are *skipped* during a backup:

- The **Birthdays, Reminders, Tasks** calendars
- Folders attached to calendar events
- The **Directory** folder in Contacts

The following Calendar items are *skipped*, due to Google Calendar API limitations:

- Appointment slots
- The conferencing field of an event
- The calendar setting **All-day event notifications**
- The calendar setting **Auto-accept invitations** (in calendars for rooms or shared spaces)

The following Contacts items are *skipped*, due to Google People API limitations:

- The **Other contacts** folder
- The external profiles of a contact (**Directory profile, Google profile**)
- The contact field **File as**

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders (According to Google terminology, "labels". **Labels** are presented in the backup software as folders, for consistency with other data presentation.)
- Email messages
- Calendar events
- Contacts

You can use search to locate items in a backup.

When recovering mailboxes and mailbox items, you can select whether to overwrite the items in the target location.

Limitations

- Contact photos cannot be recovered
- The **Out of office** calendar item is recovered as a regular calendar event, due to Google Calendar API limitations

Selecting Gmail mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan [as appropriate](#).

To select Gmail mailboxes

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the mailboxes of all users (including mailboxes that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up individual user mailboxes, expand the **Users** node, select **All users**, select the users whose mailboxes you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **Gmail** item is selected in **What to back up**.
 - If you want to back up calendars that are shared with the selected users, enable the **Include shared calendars** switch.
 - Decide whether you need [full-text search](#) through the backed-up email messages. To access this option, click the gear icon > **Backup options** > **Full-text search**.

Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose mailbox you want to recover, and then click **Recovery**.
If the user was deleted, select the user in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.
You can search users and groups by name. Wildcards are not supported.
4. Select a recovery point.

Note

To see only the recovery points that contain mailboxes, select **Gmail** in **Filter by content**.

5. Click **Recover** > **Entire mailbox**.
6. If multiple Google Workspace organizations are added to the Cyber Protection service, click **Google Workspace organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must select a new target organization from the available registered organizations.

7. In **Recover to mailbox**, view, change, or specify the target mailbox.

By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.

You cannot create a new target mailbox during recovery. To recover a mailbox to a new one, first you need to create the target mailbox in the desired Google Workspace organization, and then let the cloud agent synchronize the change. The cloud agent automatically synchronizes with Google Workspace every 24 hours. To synchronize the change immediately, in the Cyber Protect console, select the organization on the **Google Workspace** page, and then click **Refresh**.

8. Click **Start recovery**.
9. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
10. Click **Proceed** to confirm your decision.

Recovering mailbox items

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.

Note

To see only the recovery points that contain mailboxes, select **Gmail** in **Filter by content**.

5. Click **Recover > Email messages**.
6. Browse to the required folder. If the backup is not encrypted, you can use search to obtain the list of the required items.

The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, date, attachment name, and message content.

When searching by date, you can select a start date or an end date (both inclusive), or both dates to search within a time range.

Searching by attachment name or in the message content gives results only if the **Full-text search** option was enabled during backup. You can specify the language of the message fragment that will be searched as an additional parameter.

- For events: search by title and date.
- For contacts: search by name, email address, and phone number.

7. Select the items that you want to recover. To be able to select folders, click the "recover folders"

icon: 

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
- Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. Click **Recover**.

9. If multiple Google Workspace organizations were added to the Cyber Protection service, click **Google Workspace organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must select a new target organization from the available registered organizations.

10. In **Recover to mailbox**, view, change, or specify the target mailbox.

By default, the original mailbox is selected. If this mailbox does not exist or a non-original organization is selected, you must specify the target mailbox.

11. In **Path**, view or change the target folder in the target mailbox. By default, the original folder is selected.

12. Click **Start recovery**.

13. Select one of the overwriting options:

- **Overwrite existing items**
- **Do not overwrite existing items**

14. Click **Proceed** to confirm your decision.

Protecting Google Drive files

What items can be backed up?

You can back up an entire Google Drive, or individual files and folders. Files are backed up together with their sharing permissions.

Important

The following items are not backed up:

- The **Shared with me** folder
 - The **Computers** folder (created by the Backup and Sync client)
-

Limitations

Out of the Google-specific file formats, Google Docs, Google Sheets, and Google Slides are fully supported for backup and recovery. Other Google-specific formats might not be fully supported or might not be supported at all – for example, Google Drawings files are recovered as .svg files, Google Sites files are recovered as .txt files, Google Jamboard files are recovered as .pdf files, and Google My Maps files are skipped during a backup.

Note

File formats that are not Google-specific – for example, .txt, .docx, .pptx, .pdf, .jpg, .png, .zip, are fully supported for backup and recovery.

What items can be recovered?

You can recover an entire Google Drive, or any file or folder that was backed up.

You can choose whether to recover the sharing permissions or let the files inherit the permissions from the folder to which they are recovered.

Limitations

- Comments in files are not recovered.
- Sharing links for files and folders are not recovered.
- The read-only **Owner settings** for shared files (**Prevent editors from changing access and adding new people** and **Disable options to download, print and copy for commenters and viewers**) cannot be changed during a recovery.
- Ownership of a shared folder cannot be changed during a recovery if the **Prevent editors from changing access and adding new people** option is enabled for this folder. This setting prevents the Google Drive API from listing the folder permissions. Ownership of the files in the folder is recovered correctly.

Selecting Google Drive files

Select the files as described below, and then specify other settings of the protection plan [as appropriate](#).

To select Google Drive files

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.

3. Do one of the following:
 - To back up the files of all users (including users that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up the files of individual users, expand the **Users** node, select **All users**, select the users whose files you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **Google Drive** item is selected in **What to back up**.
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all files).
 - Specify the files and folders to back up by adding their names or paths.
You can use wildcard characters (*, **, and ?). For more details about specifying paths and using wildcards, refer to "[File filters](#)".
 - Specify the files and folders to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single user.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the files and folders to skip during the backup.
File exclusions override the file selection; i.e. if you specify the same file in both fields, this file will be skipped during a backup.
 - If you want to enable notarization of all files selected for backup, enable the **Notarization** switch. For more information about notarization, refer to "[Notarization](#)".

Recovering Google Drive and Google Drive files

Recovering an entire Google Drive

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose Google Drive you want to recover, and then click **Recovery**.
If the user was deleted, select the user in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.
You can search users by name. Wildcards are not supported.
4. Select a recovery point.

Note

To see only the recovery points that contain Google Drive files, select **Google Drive** in **Filter by content**.

5. Click **Recover > Entire Drive**.
6. If multiple Google Workspace organizations were added to the Cyber Protection service, click **Google Workspace organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must select a new target organization from the available registered organizations.

7. In **Recover to drive**, view, change, or specify the target user or the target Shared drive.

By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user or the target Shared drive.

If the backup contains shared files, the files will be recovered to the root folder of the target drive.

8. Select whether to recover the sharing permissions for the files.

9. Click **Start recovery**.

10. Select one of the overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite existing files	All existing files in the destination location are overwritten, regardless of their last modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

11. Click **Proceed** to confirm your decision.

Recovering Google Drive files

1. Click **Google Workspace**.

2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.

3. Expand the **Users** node, select **All users**, select the user whose Google Drive files you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.

You can search users by name. Wildcards are not supported.

4. Select a recovery point.

Note

To see only the recovery points that contain Google Drive files, select **Google Drive** in **Filter by content**.

5. Click **Recover > Files/folders**.

6. Browse to the required folder or use search to obtain the list of the required files and folders.

7. Select the files that you want to recover.

If the backup is not encrypted and you selected a single file, you can click **Show versions** to select the file version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. If you want to download a file, select the file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
9. Click **Recover**.
10. If multiple Google Workspace organizations were added to the Cyber Protection service, click **Google Workspace organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must select a new target organization from the available registered organizations.
11. In **Recover to drive**, view, change, or specify the target user or the target Shared drive.
By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user or the target Shared drive.
12. In **Path**, view or change the target folder in the target user's Google Drive or in the target Shared drive. By default, the original location is selected.
13. Select whether to recover the sharing permissions for the files.
14. Click **Start recovery**.
15. Select one of the file overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite existing files	All existing files in the destination location are overwritten, regardless of their last modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

16. Click **Proceed** to confirm your decision.

Protecting Shared drive files

What items can be backed up?

You can back up an entire Shared drive, or individual files and folders. Files are backed up together with their sharing permissions.

Important

The **Shared with me** folder is not backed up.

Limitations

- A Shared drive without members cannot be backed up, due to Google Drive API limitations.
- Out of the Google-specific file formats, Google Docs, Google Sheets, and Google Slides are fully supported for backup and recovery. Other Google-specific formats might not be fully supported or might not be supported at all – for example, Google Drawings files are recovered as .svg files, Google Sites files are recovered as .txt files, Google Jamboard files are recovered as .pdf files, and Google My Maps files are skipped during a backup.

Note

File formats that are not Google-specific – for example, .txt, .docx, .pptx, .pdf, .jpg, .png, .zip, are fully supported for backup and recovery.

What items can be recovered?

You can recover an entire Shared drive, or any file or folder that was backed up.

You can choose whether to recover the sharing permissions or let the files inherit the permissions from the folder to which they are recovered.

The following items are not recovered:

- Sharing permissions for a file that was shared with a user outside the organization are not recovered if sharing outside the organization is disabled in the target Shared drive.
- Sharing permissions for a file that was shared with a user who is not a member of the target Shared drive are not recovered if **Sharing with non-members** is disabled in the target Shared drive.

Limitations

- Comments in files are not recovered.
- Sharing links for files and folders are not recovered.

Selecting Shared drive files

Select the files as described below, and then specify other settings of the protection plan [as appropriate](#).

To select Shared drive files

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the files of all Shared drive (including Shared drive that will be created in the future), expand the **Shared drives** node, select **All Shared drives**, and then click **Group backup**.

- To back up the files of individual Shared drives, expand the **Shared drives** node, select **All Shared drives**, select the Shared drives to back up, and then click **Backup**.
4. On the protection plan panel:
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all files).
 - Specify the files and folders to back up by adding their names or paths.
You can use wildcard characters (*, **, and ?). For more details about specifying paths and using wildcards, refer to "[File filters](#)".
 - Specify the files and folders to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single Shared drive.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the files and folders to skip during the backup.
File exclusions override the file selection; i.e. if you specify the same file in both fields, this file will be skipped during a backup.
 - If you want to enable notarization of all files selected for backup, enable the **Notarization** switch. For more information about notarization, refer to "[Notarization](#)".

Recovering Shared drive and Shared drive files

Recovering an entire Shared drive

1. Click **Google Workspace**.
2. If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Shared drives** node, select **All Shared drives**, select the Shared drive that you want to recover, and then click **Recovery**.
If the Shared drive was deleted, select it in the **Cloud applications backups** section of the [Backup storage tab](#), and then click **Show backups**.
You can search Shared drives by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Entire Shared drive**.
6. If multiple Google Workspace organizations were added to the Cyber Protection service, click **Google Workspace organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must select a new target organization from the available registered organizations.
7. In **Recover to drive**, view, change, or specify the target Shared drive or the target user. If you specify a user, the data will be recovered to this user's Google Drive.
By default, the original Shared drive is selected. If this Shared drive does not exist or a non-original organization is selected, you must specify the target Shared drive or the target user.
8. Select whether to recover the sharing permissions for the files.
9. Click **Start recovery**.

- Select one of the overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite existing files	All existing files in the destination location are overwritten, regardless of their last modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

- Click **Proceed** to confirm your decision.

Recovering Shared drive files

- Click **Google Workspace**.
- If multiple Google Workspace organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
- Expand the **Shared drives** node, select **All Shared drives**, select the Shared drive that originally contained the files you want to recover, and then click **Recovery**.
If the Shared drive was deleted, select it in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.
You can search Shared drives by name. Wildcards are not supported.
- Select a recovery point.
- Click **Recover > Files/folders**.
- Browse to the required folder or use search to obtain the list of the required files and folders.
- Select the files that you want to recover.
If the backup is not encrypted and you selected a single file, you can click **Show versions** to select the file version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
- If you want to download a file, select the file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
- Click **Recover**.
- If multiple Google Workspace organizations were added to the Cyber Protection service, click **Google Workspace organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must select a new target organization from the available registered organizations.
- In **Recover to drive**, view, change, or specify the target Shared drive or the target user. If you specify a user, the data will be recovered to this user's Google Drive.

By default, the original Shared drive is selected. If this Shared drive does not exist or a non-original organization is selected, you must specify the target Shared drive or the target user.

12. In **Path**, view or change the target folder in the target Shared drive or the target user's Google Drive. By default, the original location is selected.
13. Select whether to recover the sharing permissions for the files.
14. Click **Start recovery**.
15. Select one of the file overwriting options:

Option	Description
Overwrite an existing file if it is older	If there is a file with the same name in the destination location, and it is older than the source file, the source file will be saved in the destination location, replacing the older version.
Overwrite existing files	All existing files in the destination location are overwritten, regardless of their last modified date.
Do not overwrite existing files	If there is a file with the same name in the destination location, no changes are applied to it, and the source file is not saved to the destination location.

16. Click **Proceed** to confirm your decision.

Notarization

Notarization enables you to prove that a file is authentic and unchanged since it was backed up. We recommend that you enable notarization when backing up your legal document files or other files that require proved authenticity.

Notarization is available only for backups of Google Drive files and Google Workspace Shared drive files.

How to use notarization

To enable notarization of all files selected for backup, enable the **Notarization** switch when creating a protection plan.

When configuring recovery, the notarized files will be marked with a special icon, and you can [verify the file authenticity](#).

How it works

During a backup, the agent calculates the hash codes of the backed-up files, builds a hash tree (based on the folder structure), saves the tree in the backup, and then sends the hash tree root to the notary service. The notary service saves the hash tree root in the Ethereum blockchain database to ensure that this value does not change.


When verifying the file authenticity, the agent calculates the hash of the file, and then compares it with the hash that is stored in the hash tree inside the backup. If these hashes do not match, the file is considered not authentic. Otherwise, the file authenticity is guaranteed by the hash tree.

To verify that the hash tree itself was not compromised, the agent sends the hash tree root to the notary service. The notary service compares it with the one stored in the blockchain database. If the hashes match, the selected file is guaranteed to be authentic. Otherwise, the software displays a message that the file is not authentic.

Verifying file authenticity with Notary Service

If notarization was enabled during backup, you can verify the authenticity of a backed-up file.

To verify the file authenticity

1. Do one of the following:
 - To verify the authenticity of a Google Drive file, select the file as described in steps 1-7 of the ["Recovering Google Drive files"](#) section.
 - To verify the authenticity of a Google Workspace Shared drive file, select the file as described in steps 1-7 of the ["Recovering Shared drive files"](#) section.
2. Ensure that the selected file is marked with the following icon: . This means that the file is notarized.
3. Do one of the following:
 - Click **Verify**.
The software checks the file authenticity and displays the result.
 - Click **Get certificate**.
A certificate that confirms the file notarization is opened in a web browser window. The window also contains instructions that allow you to verify the file authenticity manually.

Search in cloud-to-cloud backups

When recovering data, you can search for specific backed-up items instead of browsing the backup archive.

In non-encrypted backups, search is always available. Only enhanced (index-based) search is supported.

The index-based search is faster and provides additional options, such as showing versions of the backed-up items, searching in attachment names, and full-text search in Gmail backups.

In encrypted backups, you can also enable enhanced (index-based) search. If you do not enable the enhanced search, basic search will be available for backups of Microsoft 365 mailboxes. For all other workloads, search will not be available.

The table below summarizes the available options for encrypted backups.

Workload type	What to recover	Enhanced search is disabled	Enhanced search is enabled
Microsoft 365 workloads			
Mailbox	Email messages	Basic (non-index based) search is available	Enhanced (index-based) search is available
OneDrive	Files/folders	Search is not available	Enhanced (index-based) search is available
SharePoint site	SharePoint files	Search is not available	Enhanced (index-based) search is available
Teams	Channels	Search is not available	Enhanced (index-based) search is available
	Email messages	Basic (non-index based) search is available	Enhanced (index-based) search is available
	Team site	Search is not available	Enhanced (index-based) search is available
Google Workspace workloads			
Mailbox	Email messages	Search is not available	Enhanced (index-based) search is available
Google Drive	Files/folders	Search is not available	Enhanced (index-based) search is available
Shared Drives	Files/folders	Search is not available	Enhanced (index-based) search is available

Full-text search

Full-text search is available only for Gmail backups, and it is enabled by default. With it, you can search in the body text of the backed-up emails. If this option is disabled, you can search only by subject, sender, recipient, and date.

A full-text search index takes between 10 and 30 percent of the storage space occupied by the Gmail backup. An index without full-text search data is significantly smaller. To save storage space, you can disable the full-text search and clear the portion of the index that contains the full-text search data.

Search indexes

Search indexes provide enhanced search capabilities in cloud-to-cloud backup archives.

The archives are automatically indexed after each backup operation. The indexing process does not affect the backup performance because indexing and backing up are done by different software components.

Showing search results becomes available after the indexing operation completes, which might take up to 24 hours. Indexing the first backup, which is full, usually takes longer than indexing the successive incremental backups.

All indexes contain metadata that supports the main searching functionality— search by subject, sender, recipient, or date. The indexes for Gmail backups contain additional data if full-text search is enabled.

Checking the size of a search index

Search indexes grow bigger with time. The indexes for backup archives in which full-text search is enabled, might take up to 30 percent of the archive size.

To check the size of a search index

1. Log in to the Cyber Protect console as administrator.
2. On the **Backup storage** tab, click **Cloud applications backup**.
3. Check the value in the **Index size** column.

Updating, rebuilding, or deleting indexes

To troubleshoot search-related issues in cloud-to-cloud backups, you can update, rebuild, or delete search indexes.

Note

We recommend that you contact the Support team before updating, rebuilding, or deleting an index.

To update, rebuild, or delete an index

1. Log in to the Cyber Protect console as an administrator.
2. On the **Backup storage** tab, click **Cloud applications backup**.
Select the archive the index of which you want to update, rebuild, or delete.

The availability of these actions depend on the administrator level and role, as follows:

Account level	Role	Can update index	Can rebuild index	Can delete index
Partner tenant	Company administrator	+	+	+
	Protection cyber administrator	+	-	-
	Protection administrator	+	-	-
	Protection read-only administrator	-	-	-

Account level	Role	Can update index	Can rebuild index	Can delete index
Customer tenant	Company administrator	+	-	-
	Protection administrator	+	-	-
	Protection read-only administrator	-	-	-
Unit	Unit administrator	+	-	-
	Protection administrator	+	-	-
	Protection read-only administrator	-	-	-

- In the **Actions** pane, select the action that you want to perform:
 - **Update index**—the recovery points in the archive are checked, and the missing indexes are added.
 - **Rebuild index**—the indexes for all recovery points in the archive are deleted, and then the indexes are created again.
 - **Delete index**—the indexes for all recovery points in the archive are deleted.
- [For encrypted archives] Specify the encryption password, and then click **OK**.
- Select the scope of the action, and then click **OK**.
Depending on the archive and the selected action, one or more of the following options are available:
 - **Metadata only**
 - **Content only**
 - **Metadata and content search**

Enabling enhanced search in encrypted backups

When creating a backup plan for encrypted cloud-to-cloud backup, you can enable enhanced (index-based) search.

If you do not enable enhanced search, basic search will be available for backups of Microsoft 365 mailboxes. For all other workloads, search will not be available. For more information about the available options, see "Search in cloud-to-cloud backups" (p. 613).

Note

This functionality is available in selected data centers and might not be accessible to all customers.

To enable search in encrypted backups

- When creating a backup plan, enable the **Encryption** switch.
- Specify and confirm the encryption password.
- Select the **Allow enhanced search in encrypted backups** check box.
- Click **Done**.

Note

You cannot disable encryption or change the encryption password later. To create a non-encrypted backup or change the encryption password, create a new backup plan.

Enabling or disabling enhanced search in existing plans

You can edit an existing plan for encrypted backup to enable or disable enhanced (index-based) search.

If you do not enable enhanced search, basic search will be available for backups of Microsoft 365 mailboxes. For all other workloads, search will not be available. For more information about the available options, see "Search in cloud-to-cloud backups" (p. 613).

In non-encrypted backups, enhanced search is always available. This option cannot be disabled.

To enable or disable enhanced search in encrypted backups

1. When editing a backup plan in which encryption is enabled, click the gear icon in the upper right corner.
2. On the **Search options** tab, toggle the switch as required.
3. Click **Done**.
4. Click **Save settings**.

Note

If you re-enable enhanced search, all archives created by this backup plan will be indexed again. This is a time-consuming operation.

Disabling full-text search for Gmail backups

Full-text search is available only for Gmail backups, and it is enabled by default. With it, you can search in the body text of the backed-up emails. If this option is disabled, you can search only by subject, sender, recipient, and date.

You might want to disable full-text search if you need to keep the size of the search index minimal.

To disable full-text search

1. When creating or editing a backup plan, click the gear icon in the upper right corner.
2. On the **Full-text search** tab, disable the switch.
3. Click **Done**.
4. [When creating a plan] Click **Apply**.
5. [When editing a plan] Click **Save settings**.

Note

If you re-enable full-text search, all archives created by this backup plan will be indexed again. This is a time-consuming operation.

Protecting Oracle Database

Note

This feature is available with the Advanced Backup pack.

Protection of Oracle Database is described in a separate document available at https://dl.managed-protection.com/u/pdf/OracleBackup_whitepaper_en-US.pdf

Protecting SAP HANA

Note

This feature is available with the Advanced Backup pack.

Protection of SAP HANA is described in a separate document available at https://dl.managed-protection.com/u/pdf/SAP_HANA_backup_whitepaper_en-US.pdf

Protecting MySQL and MariaDB data

You can protect MySQL or MariaDB data with application-aware backup. It collects application metadata and allows granular recovery on the instance, database, or table level.

Note

Application-aware backup of MySQL or MariaDB data is available with the Advanced Backup pack.

To protect a physical or virtual machine that runs MySQL or MariaDB instances with application-aware backup, you need to install Agent for MySQL/MariaDB on this machine. Agent for MySQL/MariaDB is bundled with Agent for Linux (64-bit) and therefore can be installed only on 64-bit Linux-based operating systems. See "Supported operating systems and environments" (p. 27).

To download the Agent for Linux (64-bit) installation file

1. Log in to the Cyber Protect console.
2. Click the account icon in the upper-right corner, and then select **Downloads**.
3. Click **Agent for Linux (64-bit)**.

The installation file is downloaded to your machine. To install the agent, proceed as described in "Installing protection agents in Linux" (p. 80) or "Unattended installation or uninstallation in Linux" (p. 104). Ensure that you select Agent for MySQL/MariaDB, which is an optional component.

To recover databases and tables to a live instance, Agent for MySQL/MariaDB needs a temporary storage to operate. By default, the /tmp directory is used. You can change this directory by setting the ACRONIS_MYSQL_RESTORE_DIR environment variable.

Limitations

- MySQL or MariaDB clusters are not supported.
- MySQL or MariaDB instances running in Docker containers are not supported.
- MySQL or MariaDB instances running on operating systems that use BTRFS file system are not supported.
- System databases (`sys`, `mysql`, `information-schema`, and `performance_schema`) and databases that do not contain any tables cannot be recovered to live instances. However, these databases can be recovered as files, when recovering the whole instance.
- Recovery is supported only to target instances of the same version as the backed-up instance or later, with the following restrictions:
 - Recovery from MySQL 5.x instances to MySQL 8.x instances is not supported.
 - Recovery to a later MySQL 5.x version (including the minor versions) is supported only via recovery of the whole instance as files. Before attempting recovery, consult the official MySQL upgrade guide for the target version, for example, the [MySQL 5.7 upgrade guide](#).
- Recovery from backups stored on Secure Zone is not supported.
- Databases and tables cannot be recovered by Agent for MySQL/MariaDB that is running on a machine on which AppArmor is installed. You can still recover an instance as files, or the entire machine.
- Recovery to target databases that are configured with symbolic links is not supported. You can recover the backed-up databases as new databases, by changing their name.

Known issues

If you encounter issues while recovering data from password protected Samba shares, log out from the Cyber Protect console, and then log in back to it. Select the desired recovery point, and then click **MySQL/MariaDB databases**. Do not click **Entire machine** or **Files/folders**.

Configuring an application-aware backup

Prerequisites

- At least one MySQL or MariaDB instance must be running on the selected machine.
- On the machine where the MySQL or MariaDB instance is running, the protection agent must be started under the root user.
- Application-aware backup is available only when the **Entire machine** is selected as a backup source in the protection plan.
- The **Sector-by-sector** backup option must be disabled in the protection plan. Otherwise, it is impossible to recover application data.

To configure an application-aware backup

1. In the Cyber Protect console, select one or more machines on which MySQL or MariaDB instances are running.
You can have one or more instances on each machine.
2. Create a protection plan with the backup module enabled.
3. In **What to back up**, select **Entire machine**.
4. Click **Application backup**, and then enable the switch next to **MySQL/MariaDB Server**.
5. Select how to specify the MySQL or MariaDB instances:
 - **For all workloads**
Use this option if you run instances with identical configurations on multiple servers. The same connection parameters and access credentials will be used for all instances.
 - **For specific workloads**
Use this option to specify the connection parameters and access credentials for each instance.
6. Click **Add instance** to configure the connection parameters and access credentials.
 - a. Select the connection type, and then specify the following:
 - [For TCP socket] IP address and port.
 - [For Unix socket] Socket path.
 - b. Specify the credentials of a user account that has the following privileges for the instance:
 - FLUSH_TABLES or RELOAD for all databases and tables (*.*)
 - SELECT for the information_schema.tables
 - c. Click **OK**.
7. Click **Done**.

Recovering data from an application-aware backup

From an application-aware backup, you can recover MySQL or MariaDB instances, databases, and tables. You can also recover the entire server on which the instances are running, or files and folders from this server.

The table below summarizes all recovery options.

What to recover	Recover as	Recover to
MySQL Server MariaDB Server	Entire machine	Machine* on which Agent for Linux is installed
MySQL Server MariaDB Server	Files or folders	Machine* on which Agent for Linux is installed

What to recover	Recover as	Recover to
Instance	Files	Machine* on which Agent for MySQL/MariaDB is installed
Database	The same database New database	Machine* on which Agent for MySQL/MariaDB is installed <ul style="list-style-type: none"> • Original instance • Another instance • Original database • New database
Table	The same table New table	Machine* on which Agent for MySQL/MariaDB is installed <ul style="list-style-type: none"> • Original instance • Another instance • Original database • Original table • New table

* A virtual machine with an agent inside is treated as a physical machine from the backup standpoint.

Recovering the entire server

To learn how to recover the entire server on which MySQL or MariaDB instances are running, refer to "Recovering a machine" (p. 472).

Recovering instances

From an application-aware backup, you can recover MySQL or MariaDB instances as files.

To recover an instance

1. In the Cyber Protect console, select the machine that originally contained the data that you want to recover.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Do one of the following:

- If the backup location is cloud or shared storage (that is, other agents can access it), click **Select machine**, select an online machine that has Agent for MySQL/MariaDB, and then select a recovery point.
- Select a recovery point on the **Backup storage** tab.

The machine chosen for browsing in either of the above actions becomes a target machine for the recovery.

4. Click **Recover > MySQL/MariaDB databases**.
5. Select the instance that you want to recover, and then click **Recover as files**.

6. Under **Path**, select the directory to which the files will be recovered.
7. Click **Start recovery**.

Recovering databases

From an application-aware backup, you can recover databases to live MySQL or MariaDB instances.

1. In the Cyber Protect console, select the machine that originally contained the data that you want to recover.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Do one of the following:
 - If the backup location is cloud or shared storage (that is, other agents can access it), click **Select machine**, select an online machine that has Agent for MySQL/MariaDB, and then select a recovery point.
 - Select a recovery point on the **Backup storage** tab.The machine chosen for browsing in either of the above actions becomes a target machine for the recovery.
4. Click **Recover > MySQL/MariaDB databases**.
5. Click the name of the desired instance to drill down to its databases.
6. Select one or more databases that you want to recover.
7. Click **Recover**.
8. Click **Target MySQL/MariaDB instance** to specify the connection parameters and access credentials for the target instance.
 - Verify the instance to which you want to recover data. By default, the original instance is selected.
 - Specify the credentials of a user account that can access the target instance. This user account must have the following privileges assigned for all databases and tables (*.*):
 - INSERT
 - CREATE
 - DROP
 - LOCK_TABLES
 - ALTER
 - SELECT
 - Click **OK**.
9. Verify the target database.
By default, the original database is selected.
To recover a database as a new one, click the name of the target database and change it. This action is only available when you recover a single database.
10. Under **Overwrite existing databases**, select the overwriting mode.
By default, overwriting is enabled and the backed-up database will replace the target database that has the same name.

If overwriting is disabled, the backed-up database will be skipped during the recovery operation and will not replace the target database that has the same name.

11. Click **Start recovery**.

Recovering tables

From an application-aware backup, you can recover tables to live MySQL or MariaDB instances.

1. In the Cyber Protect console, select the machine that originally contained the data that you want to recover.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Do one of the following:

- If the backup location is cloud or shared storage (that is, other agents can access it), click **Select machine**, select an online machine that has Agent for MySQL/MariaDB, and then select a recovery point.
- Select a recovery point on the **Backup storage** tab.

The machine chosen for browsing in either of the above actions becomes a target machine for the recovery.

4. Click **Recover** > **MySQL/MariaDB databases**.
5. Click the name of the desired instance to drill down to its databases.
6. Click the name of the desired database to drill down to its tables.
7. Select one or more tables that you want to recover.
8. Click **Recover**.
9. Click **Target MySQL/MariaDB instance** to specify the connection parameters and access credentials for the target instance.
 - Verify the instance to which you want to recover data. By default, the original instance is selected.
 - Specify the credentials of a user account that can access the target instance. This user account must have the following privileges assigned for all databases and tables (*.*):
 - INSERT
 - CREATE
 - DROP
 - LOCK_TABLES
 - ALTER
 - SELECT
 - Click **OK**.
10. Verify the target table.

By default, the original table is selected.

To recover a table as a new one, click the name of the target table and change it. This action is only available when you recover a single table.
11. Under **Overwrite existing tables**, select the overwriting mode.

By default, overwriting is enabled and the backed-up table will replace the target table that has the same name.

If overwriting is disabled, the backed-up table will be skipped during the recovery operation and will not replace the target table that has the same name.

12. Click **Start recovery**.

Recovering stored routines

When you recover a whole MySQL instance, the stored routines are automatically recovered.

When you recover an individual database to a non-original instance or recover it as a new database, the stored routines are not automatically recovered. You can recover them manually, by exporting them in an SQL file, and then adding them to the recovered database.

To export the stored routines and add them to a recovered database

1. On the machine with the original MySQL instance, open Terminal.
2. Run the following command to export the stored routines.

3.

```
mysqldump -p [source_database_name] --routines --no-create-info --no-data > [exported_db_routines.sql]
```

4. On the machine where the database is recovered, open the MySQL command line client.
5. Run the following commands to add the routines to the recovered database.

```
mysql> use [recovered_database_name];
```

```
mysql> source [path_to_exported_db_routines.sql];
```

Protecting websites and hosting servers

Protecting websites

A website can be corrupted as a result of unauthorized access or a malware attack. Back up your website if you want to easily revert it to a healthy state, in case of corruption.

What do I need to back up a website?

The website must be accessible via the SFTP or SSH protocol. You do not need to install an agent, just add a website as described later in this section.

What items can be backed up?

You can back up the following items:

- **Website content files**

All files accessible to the account you specify for the SFTP or SSH connection.

- **Linked databases (if any) hosted on MySQL servers.**

All databases accessible to the MySQL account you specify.

If your website employs databases, we recommend that you back up both the files and the databases, to be able to recover them to a consistent state.

Limitations

- The only backup location available for website backup is the cloud storage.
- It is possible to apply several protection plans to a website, but only one of them can run on a schedule. Other plans need to be started manually.
- The only available backup option is "[Backup file name](#)".
- The website protection plans are not shown on the **Management > Protection plan** tab.

Backing up a website

To add a website

1. Click **Devices > Add**.
2. Click **Website**.
3. Configure the following access settings for the website:
 - In **Website name**, create and type a name for your website. This name will be displayed in the Cyber Protect console.
 - In **Host**, specify the host name or IP address that will be used to access the website via SFTP or SSH. For example, `my.server.com` or `10.250.100.100`.
 - In **Port**, specify the port number.
 - In **User name** and **Password**, specify the credentials of the account that can be used to access the website via SFTP or SSH.

Important

Only the files that are accessible to the specified account will be backed up.

Instead of a password, you can specify your private SSH key. To do this, select the **Use SSH private key instead of password** check box, and then specify the key.

4. Click **Next**.
5. If your website uses MySQL databases, configure the access settings for the databases. Otherwise, click **Skip**.
 - a. In **Connection type**, select how to access the databases from the cloud:
 - **Via SSH from host**—The databases will be accessed via the host specified in step 3.
 - **Direct connection**—The databases will be accessed directly. Choose this setting only if the databases are accessible from the Internet.
 - b. In **Host**, specify the name or IP address of the host where the MySQL server is running.
 - c. In **Port**, specify the port number for the TCP/IP connection to the server. The default port number is 3306.

- d. In **User name** and **Password**, specify the MySQL account credentials.

Important

Only the databases that are accessible to the specified account will be backed up.

- e. Click **Create**.

The website appears in the Cyber Protect console under **Devices > Websites**.

To change the connection settings

1. Select the website under **Devices > Websites**.
2. Click **Details**.
3. Click the pencil icon next to the website or the database connection settings.
4. Do the necessary changes, and then click **Save**.

To create a protection plan for websites

1. Select a website or several websites under **Devices > Websites**.
2. Click **Protect**.
3. [Optional] Enable backup of databases.
If several websites are selected, backup of databases is disabled by default.
4. [Optional] Change the [retention rules](#).
5. [Optional] Enable [encryption of backups](#).
6. [Optional] Click the gear icon to edit the **Backup file name** option. This makes sense in two cases:
 - If you backed up this website earlier and want to continue the existing sequence of backups
 - If you want to see the custom name on the **Backup storage** tab
7. Click **Apply**.

You can edit, revoke, and delete protection plans for websites in the same way as for machines. These operations are described in "Operations with protection plans".

Recovering a website

To recover a website

1. Do one of the following:
 - Under **Devices > Websites**, select the website that you want to recover, and then click **Recovery**.
You can search websites by name. Wildcards are not supported.
 - If the website was deleted, select it in the **Cloud applications backups** section of [the Backup storage tab](#), and then click **Show backups**.
To recover a deleted website, you need to add the target site as a device.
2. Select the recovery point.

3. Click **Recover**, and then select what you want to recover: **Entire website**, **Databases** (if any), or **Files/folders**.

To ensure that your website is in a consistent state, we recommend that you recover both files and databases, in any order.

4. Depending on your choice, follow one of the procedures described below.

To recover the entire website

1. In **Recover to website**, view or change the target website.
By default, the original website is selected. If it does not exist, you must select the target website.
2. Select whether to recover the sharing permissions of the recovered items.
3. Click **Start recovery**, and then confirm the action.

To recover the databases

1. Select the databases that you want to recover.
2. If you want to download a database as a file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
3. Click **Recover**.
4. In **Recover to website**, view or change the target website.
By default, the original website is selected. If it does not exist, you must select the target website.
5. Click **Start recovery**, and then confirm the action.

To recover the website files/folders

1. Select the files/folders that you want to recover.
2. If you want to save a file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
3. Click **Recover**.
4. In **Recover to website**, view or change the target website.
By default, the original website is selected. If it does not exist, you must select the target website.
5. Select whether to recover the sharing permissions of the recovered items.
6. Click **Start recovery**, and then confirm the action.

Protecting web hosting servers

You can protect Linux-based web hosting servers that run Plesk, cPanel, DirectAdmin, VirtualMin , or ISPManager control panels. Servers that run web hosting control panels from other vendors are protected as regular workloads.

Quotas

Servers that run Plesk, cPanel, DirectAdmin, VirtualMin , or ISPManager control panels are considered web hosting servers. Each backed-up web hosting server consumes the **Web hosting servers** quota. If this quota is disabled or the overage for this quota is exceeded, a quota will be assigned as follows or the backups will fail:

- If the server is physical, the **Servers** quota will be used. If this quota is disabled or the overage for this quota is exceeded, the backup will fail.
- If the server is virtual, the **Virtual machines** quota will be used. If this quota is disabled or the overage for this quota is exceeded, the backup will fail.

Integrations for DirectAdmin, cPanel, and Plesk

Web hosting administrators that use DirectAdmin, Plesk or cPanel, can integrate these control panels with the Cyber Protection service to gain several powerful capabilities, including:

- Backing up entire web hosting server to the cloud storage with disk-level backup
- Recovering the entire server, including all websites and accounts
- Performing granular recovery and downloading of accounts, websites, individual files, mailboxes, or databases
- Enabling resellers and customers to perform self-service recovery of their own data

To perform the integration, you need to use a Cyber Protection service extension. For detailed information, please refer to the corresponding integration guides:

- [DirectAdmin Integration Guide](#)
- [WHM and cPanel Integration Guide](#)
- [Plesk Integration Guide](#)

Special operations with virtual machines

Running a virtual machine from a backup (Instant Restore)

You can run a virtual machine from a disk-level backup that contains an operating system. This operation, also known as instant restore, enables you to spin up a virtual server in seconds. The virtual disks are emulated directly from the backup and thus do not consume space on the datastore (storage). The storage space is required only to keep changes to the virtual disks.

We recommend that you leave this temporary virtual machine working for up to three days. Then, you can completely remove it or convert it to a regular virtual machine (finalize) without downtime.

As long as the temporary virtual machine exists, retention rules cannot be applied to the backup being used by that machine. Backups of the original machine can continue to run.

Usage examples

- **Disaster recovery**
Instantly bring a copy of a failed machine online.
- **Testing a backup**
Run the machine from the backup and ensure that the guest OS and applications are functioning properly.
- **Accessing application data**

While the machine is running, use application's native management tools to access and extract the required data.


Prerequisites

- At least one Agent for VMware or Agent for Hyper-V must be registered in the Cyber Protection service.
- The backup can be stored in a network folder or in a local folder of the machine where Agent for VMware or Agent for Hyper-V is installed. If you select a network folder, it must be accessible from that machine. A virtual machine can also be run from a backup stored in the cloud storage, but it works slower because this operation requires intense random-access reading from the backup.
- The backup must contain an entire machine or all of the volumes that are required for the operating system to start.
- Backups of both physical and virtual machines can be used. Backups of Virtuozzo *containers* cannot be used.
- Backups that contain Linux logical volumes (LVM) must be created by Agent for VMware or Agent for Hyper-V. The virtual machine must be of the same type as the original machine (ESXi or Hyper-V).

Running the machine

1. Do one of the following:
 - Select a backed-up machine, click **Recovery**, and then select a recovery point.
 - Select a recovery point on [the Backup storage tab](#).
2. Click **Run as VM**.

The software automatically selects the host and other required parameters.



3. [Optional] Click **Target machine**, and then change the virtual machine type (ESXi or Hyper-V), the host, or the virtual machine name.


4. [Optional] Click **Datastore** for ESXi or **Path** for Hyper-V, and then select the datastore for the virtual machine.


Changes to the virtual disks accumulate while the machine is running. Ensure that the selected datastore has enough free space. If you are planning to preserve these changes by [making the virtual machine permanent](#), select a datastore that is suitable for running the machine in production.

5. [Optional] Click **VM settings** to change the memory size and network connections of the virtual machine.

6. [Optional] Select the VM power state (**On/Off**).

7. Click **Run now**.

As a result, the machine appears in the web interface with one of the following icons:  or

. Such virtual machines cannot be selected for backup.

Note

You can perform the Run as virtual machine (Instant Restore) operation with backups in Microsoft Azure. However, this operation results in significant egress traffic, which will be added to your Microsoft Azure subscription bill. Typical egress traffic for a Windows machine running from a Microsoft Azure backup would be approximately 5 GB from virtual machine power up until login.

Deleting the machine

We recommend that you do not delete a temporary virtual machine directly in vSphere/Hyper-V. This might lead to artifacts in the web interface. Also, the backup from which the machine was running may remain locked for a while (it cannot be deleted by retention rules).

To delete a virtual machine that is running from a backup

1. On the **All devices** tab, select a machine that is running from a backup.
2. Click **Delete**.

The machine is removed from the web interface. It is also removed from the vSphere or Hyper-V inventory and datastore (storage). All changes that occurred to the data while the machine was running are lost.

Finalizing the machine

While a virtual machine is running from a backup, the virtual disks' content is taken directly from that backup. Therefore, the machine will become inaccessible or even corrupted if the connection is lost to the backup location or to the protection agent.

You have the option to make this machine permanent, i.e. recover all of its virtual disks, along with the changes that occurred while the machine was running, to the datastore that stores these changes. This process is named finalization.

Finalization is performed without downtime. The virtual machine will *not* be powered off during finalization.

The location of the final virtual disks is defined in the parameters of the **Run as VM** operation (**Datastore** for ESXi or **Path** for Hyper-V). Prior to starting the finalization, ensure that free space, sharing capabilities, and performance of this datastore are suitable for running the machine in production.

Note

Finalization is not supported for Hyper-V running in Windows Server 2008/2008 R2 and Microsoft Hyper-V Server 2008/2008 R2 because the necessary API is missing in these Hyper-V versions.

To finalize a machine that is running from a backup

1. On the **All devices** tab, select a machine that is running from a backup.
2. Click **Finalize**.

3. [Optional] Specify a new name for the machine.
4. [Optional] Change the disk provisioning mode. The default setting is **Thin**.
5. Click **Finalize**.

The machine name changes immediately. The recovery progress is shown on the **Activities** tab. Once the recovery is completed, the machine icon changes to that of a regular virtual machine.

What you need to know about finalization

Finalization vs. regular recovery

The finalization process is slower than a regular recovery for the following reasons:

- During a finalization, the agent performs random access to different parts of the backup. When an entire machine is being recovered, the agent reads data from the backup sequentially.
- If the virtual machine is running during the finalization, the agent reads data from the backup more often, to maintain both processes simultaneously. During a regular recovery, the virtual machine is stopped.

Finalization of machines running from cloud backups

Because of intensive access to the backed-up data, the finalization speed highly depends on the connection bandwidth between the backup location and the agent. The finalization will be slower for backups located in the cloud as compared to local backups. If the Internet connection is very slow or unstable, the finalization of a machine running from a cloud backup may fail. We recommend that you run virtual machines from local backups if you are planning to perform finalization, and have the choice.

Note

Finalization speed depends on whether the agent is connected to a VMware ESXi host or vCenter, as described in step 3 of "Configuring the virtual appliance" (p. 135). Connection to a VMware vCenter can slow down the finalization operation due to the specifics of VMware APIs. To speed up the finalization operation, use a separate Agent for VMware for performing the **Run as VM** operation followed by finalization, where this Agent will be connected to an ESXi host instead of a vCenter.

Working in VMware vSphere

This section describes operations that are specific for VMware vSphere environments.

Replication of virtual machines

Replication is available only for VMware ESXi virtual machines.

Replication is the process of creating an exact copy (replica) of a virtual machine, and then maintaining the replica in sync with the original machine. By replicating a critical virtual machine, you will always have a copy of this machine in a ready-to-start state.

The replication can be started manually or on the schedule you specify. The first replication is full (copies the entire machine). All subsequent replications are incremental and are performed with [Changed Block Tracking](#), unless this option is disabled.

Replication vs. backing up

Unlike scheduled backups, a replica keeps only the latest state of the virtual machine. A replica consumes datastore space, while backups can be kept on a cheaper storage.

However, powering on a replica is much faster than a recovery and faster than running a virtual machine from a backup. When powered on, a replica works faster than a VM running from a backup and does not load the Agent for VMware.

Usage examples

- **Replicate virtual machines to a remote site.**

Replication enables you to withstand partial or complete datacenter failures, by cloning the virtual machines from a primary site to a secondary site. The secondary site is usually located in a remote facility that is unlikely to be affected by environmental, infrastructure, or other factors that might cause the primary site failure.

- **Replicate virtual machines within a single site (from one host/datastore to another).**

Onsite replication can be used for high availability and disaster recovery scenarios.

What you can do with a replica

- **Test a replica**

The replica will be powered on for testing. Use vSphere Client or other tools to check if the replica works correctly. Replication is suspended while testing is in progress.

- **Failover to a replica**

Failover is a transition of the workload from the original virtual machine to its replica. Replication is suspended while a failover is in progress.

- **Back up the replica**

Both backup and replication require access to virtual disks, and thus impact the performance of the host where the virtual machine is running. If you want to have both a replica and backups of a virtual machine, but don't want to put additional load on the production host, replicate the machine to a different host, and set up backups of the replica.

Restrictions

The following types of virtual machines cannot be replicated:

- Fault-tolerant machines running on ESXi 5.5 and lower.
- Machines running from backups.
- Replicas of virtual machines.


Creating a replication plan

A replication plan must be created for each machine individually. It is not possible to apply an existing plan to other machines.

To create a replication plan

1. Select a virtual machine to replicate.
2. Click **Replication**.
The software displays a new replication plan template.
3. [Optional] To modify the replication plan name, click the default name.
4. Click **Target machine**, and then do the following:
 - a. Select whether to create a new replica or use an existing replica of the original machine.
 - b. Select the ESXi host and specify the new replica name, or select an existing replica.
The default name of a new replica is **[Original Machine Name]_replica**.
 - c. Click **OK**.
5. [Only when replicating to a new machine] Click **Datastore**, and then select the datastore for the virtual machine.
6. [Optional] Click **Schedule** to change the replication schedule.
By default, replication is performed on a daily basis, Monday to Friday. You can select the time to run the replication.
If you want to change the replication frequency, move the slider, and then specify the schedule.
You can also do the following:
 - Set a date range for when the schedule is effective. Select the **Run the plan within a date range** check box, and then specify the date range.
 - Disable the schedule. In this case, replication can be started manually.
7. [Optional] Click the gear icon to modify the [replication options](#).
8. Click **Apply**.
9. [Optional] To run the plan manually, click **Run now** on the plan panel.

As a result of running a replication plan, the virtual machine replica appears in the **All devices** list

with the following icon: 

Testing a replica

To prepare a replica for testing

1. Select a replica to test.
2. Click **Test replica**.
3. Click **Start testing**.
4. Select whether to connect the powered-on replica to a network. By default, the replica will not be connected to a network.

5. [Optional] If you chose to connect the replica to the network, select the **Stop original virtual machine** check box to stop the original machine before powering on the replica.
6. Click **Start**.

To stop testing a replica

1. Select a replica for which testing is in progress.
2. Click **Test replica**.
3. Click **Stop testing**.
4. Confirm your decision.

Failing over to a replica

To failover a machine to a replica

1. Select a replica to failover to.
2. Click **Replica actions**.
3. Click **Failover**.
4. Select whether to connect the powered-on replica to a network. By default, the replica will be connected to the same network as the original machine.
5. [Optional] If you chose to connect the replica to the network, clear the **Stop original virtual machine** check box to keep the original machine online.
6. Click **Start**.

While the replica is in a failover state, you can choose one of the following actions:

- **Stop failover**
Stop failover if the original machine was fixed. The replica will be powered off. Replication will be resumed.
- **Perform permanent failover to the replica**
This instant operation removes the 'replica' flag from the virtual machine, so that replication to it is no longer possible. If you want to resume replication, edit the replication plan to select this machine as a source.
- **Failback**
Perform failback if you failed over to the site that is not intended for continuous operations. The replica will be recovered to the original or a new virtual machine. Once the recovery to the original machine is complete, it is powered on and replication is resumed. If you choose to recover to a new machine, edit the replication plan to select this machine as a source.

Stopping failover

To stop a failover

1. Select a replica that is in the failover state.
2. Click **Replica actions**.

3. Click **Stop failover**.
4. Confirm your decision.

Performing a permanent failover

To perform a permanent failover

1. Select a replica that is in the failover state.
2. Click **Replica actions**.
3. Click **Permanent failover**.
4. [Optional] Change the name of the virtual machine.
5. [Optional] Select the **Stop original virtual machine** check box.
6. Click **Start**.

Failing back

To failback from a replica

1. Select a replica that is in the failover state.
2. Click **Replica actions**.
3. Click **Failback from replica**.

The software automatically selects the original machine as the target machine.
4. [Optional] Click **Target machine**, and then do the following:
 - a. Select whether to failback to a new or existing machine.
 - b. Select the ESXi host and specify the new machine name, or select an existing machine.
 - c. Click **OK**.
5. [Optional] When failing back to a new machine, you can also do the following:
 - Click **Datastore** to select the datastore for the virtual machine.
 - Click **VM settings** to change the memory size, the number of processors, and the network connections of the virtual machine.
6. [Optional] Click **Recovery options** to modify the [failback options](#).
7. Click **Start recovery**.
8. Confirm your decision.

Replication options

To modify the replication options, click the gear icon next to the replication plan name, and then click **Replication options**.

Changed Block Tracking (CBT)

This option is similar to the backup option "[Changed Block Tracking \(CBT\)](#)".

Disk provisioning

This option defines the disk provisioning settings for the replica.

The preset is: **Thin provisioning**.

The following values are available: **Thin provisioning**, **Thick provisioning**, **Keep the original setting**.

Error handling

This option is similar to the backup option "[Error handling](#)".

Pre/Post commands

This option is similar to the backup option "[Pre/Post commands](#)".

Volume Shadow Copy Service VSS for virtual machines

This option is similar to the backup option "[Volume Shadow Copy Service VSS for virtual machines](#)".

Failback options

To modify the failback options, click **Recovery options** when configuring failback.

Error handling

This option is similar to the recovery option "[Error handling](#)".

Performance

This option is similar to the recovery option "[Performance](#)".

Pre/Post commands

This option is similar to the recovery option "[Pre/Post commands](#)".

VM power management

This option is similar to the recovery option "[VM power management](#)".

Seeding an initial replica

To speed up replication to a remote location and save network bandwidth, you can perform replica seeding.

Important

To perform replica seeding, Agent for VMware (Virtual Appliance) must be running on the target ESXi.

To seed an initial replica

1. Do one of the following:
 - If the original virtual machine can be powered off, power it off, and then skip to step 4.
 - If the original virtual machine cannot be powered off, continue to the next step.

2. Create a replication plan.

When creating the plan, in **Target machine**, select **New replica** and the ESXi that hosts the original machine.

3. Run the plan once.

A replica is created on the original ESXi.

4. Export the virtual machine (or the replica) files to an external hard drive.

- a. Connect the external hard drive to the machine where vSphere Client is running.
- b. Connect vSphere Client to the original vCenter\ESXi.
- c. Select the newly created replica in the inventory.
- d. Click **File > Export > Export OVF template**.
- e. In **Directory**, specify the folder on the external hard drive.
- f. Click **OK**.

5. Transfer the hard drive to the remote location.

6. Import the replica to the target ESXi.

- a. Connect the external hard drive to the machine where vSphere Client is running.
- b. Connect vSphere Client to the target vCenter\ESXi.
- c. Click **File > Deploy OVF template**.
- d. In **Deploy from a file or URL**, specify the template that you exported in step 4.
- e. Complete the import procedure.


7. Edit the replication plan that you created in step 2. In **Target machine**, select **Existing replica**, and then select the imported replica.

As a result, the software will continue updating the replica. All replications will be incremental.

Agent for VMware - LAN-free backup

If your ESXi uses a SAN attached storage, install the agent on a machine connected to the same SAN. The agent will back up the virtual machines directly from the storage rather than via the ESXi host and LAN. This capability is called a LAN-free backup.

The diagram below illustrates a LAN-based and a LAN-free backup. LAN-free access to virtual machines is available if you have a fibre channel (FC) or iSCSI Storage Area Network. To completely eliminate transferring the backed-up data via LAN, store the backups on a local disk of the agent's machine or on a SAN attached storage.



To enable the agent to access a datastore directly

1. Install Agent for VMware on a Windows machine that has network access to the vCenter Server.
2. Connect the logical unit number (LUN) that hosts the datastore to the machine. Consider the following:
 - Use the same protocol (i.e. iSCSI or FC) that is used for the datastore connection to the ESXi.
 - The LUN *must not* be initialized and must appear as an "offline" disk in **Disk Management**. If Windows initializes the LUN, it may become corrupted and unreadable by VMware vSphere.

As a result, the agent will use the SAN transport mode to access the virtual disks, i.e. it will read raw LUN sectors over iSCSI/FC without recognizing the VMFS file system (which Windows is not aware of).

Limitations

- In vSphere 6.0 and later, the agent cannot use the SAN transport mode if some of the VM disks are located on a VMware Virtual Volume (VVol) and some are not. Backups of such virtual machines will fail.
- Encrypted virtual machines, introduced in VMware vSphere 6.5, will be backed up via LAN, even if you configure the SAN transport mode for the agent. The agent will fall back on the NBD transport because VMware does not support SAN transport for backing up encrypted virtual disks.

Example

If you are using an iSCSI SAN, configure the iSCSI initiator on the machine running Windows where Agent for VMware is installed.

To configure the SAN policy

1. Log on as an administrator, open the command prompt, type `diskpart`, and then press **Enter**.
2. Type `san`, and then press **Enter**. Ensure that **SAN Policy : Offline All** is displayed.
3. If another value for SAN Policy is set:
 - a. Type `san policy=offlineall`.
 - b. Press **Enter**.
 - c. To check that the setting has been applied correctly, perform step 2.
 - d. Restart the machine.

To configure an iSCSI initiator

1. Go to **Control Panel > Administrative Tools > iSCSI Initiator**.


Note

To find the **Administrative Tools** applet, you may need to change the **Control Panel** view to something other than **Home** or **Category**, or use search.

2. If this is the first time that Microsoft iSCSI Initiator is launched, confirm that you want to start the Microsoft iSCSI Initiator service.
3. On the **Targets** tab, type the fully qualified domain name (FQDN) name or the IP address of the target SAN device, and then click **Quick Connect**.
4. Select the LUN that hosts the datastore, and then click **Connect**.

If the LUN is not displayed, ensure that the zoning on the iSCSI target enables the machine running the agent to access the LUN. The machine must be added to the list of allowed iSCSI initiators on this target.
5. Click **OK**.

The ready SAN LUN should appear in **Disk Management** as shown in the screenshot below.



Using a locally attached storage

You can attach an additional disk to Agent for VMware (Virtual Appliance) so the agent can back up to this locally attached storage. This approach eliminates the network traffic between the agent and the backup location.

A virtual appliance that is running on the same host or cluster with the backed-up virtual machines has direct access to the datastore(s) where the machines reside. This means the appliance can attach the backed-up disks by using the HotAdd transport, and therefore the backup traffic is directed from one local disk to another. If the datastore is connected as **Disk/LUN** rather than **NFS**, the backup will be completely LAN-free. In the case of NFS datastore, there will be network traffic between the datastore and the host.

Using a locally attached storage presumes that the agent always backs up the same machines. If multiple agents work within the vSphere, and one or more of them use locally attached storages, you need to [manually bind](#) each agent to all machines it has to back up. Otherwise, if the machines are redistributed among the agents by the management server, a machine's backups may be dispersed over multiple storages.

You can add the storage to an already working agent or when deploying the agent [from an OVF template](#).

To attach a storage to an already working agent

1. In VMware vSphere inventory, right click the Agent for VMware (Virtual Appliance).
2. Add the disk by editing the settings of the virtual machine. The disk size must be at least 10 GB.

Warning!

Be careful when adding an already existing disk. Once the storage is created, all data previously contained on this disk will be lost.

3. Go to the virtual appliance console. The **Create storage** link is available at the bottom of the screen. If it is not, click **Refresh**.
4. Click the **Create storage** link, select the disk and specify a label for it. The label length is limited to 16 characters, due to file system restrictions.

To select a locally attached storage as a backup destination

When [creating a protection plan](#), in **Where to back up**, select **Local folders**, and then type the letter corresponding to the locally attached storage, for example, **D:**.

Virtual machine binding

This section gives you an overview of how the Cyber Protection service organizes the operation of multiple agents within VMware vCenter.

The below distribution algorithm works for both virtual appliances and agents installed in Windows.

Distribution algorithm

The virtual machines are automatically evenly distributed between Agents for VMware. By evenly, we mean that each agent manages an equal number of machines. The amount of storage space occupied by a virtual machine is not counted.

However, when choosing an agent for a machine, the software tries to optimize the overall system performance. In particular, the software considers the agent and the virtual machine location. An agent hosted on the same host is preferred. If there is no agent on the same host, an agent from the same cluster is preferred.

Once a virtual machine is assigned to an agent, all backups of this machine are delegated to this agent.

Redistribution

Redistribution takes place each time the established balance breaks, or, more precisely, when a load imbalance among the agents reaches 20 percent. This may happen when a machine or an agent is added or removed, or a machine migrates to a different host or cluster, or if you manually bind a machine to an agent. If this happens, the Cyber Protection service redistributes the machines using the same algorithm.

For example, you realize that you need more agents to help with throughput and deploy an additional virtual appliance to the cluster. The Cyber Protection service will assign the most appropriate machines to the new agent. The old agents' load will reduce.

When you remove an agent from the Cyber Protection service, the machines assigned to the agent are distributed among the remaining agents. However, this will not happen if an agent gets corrupted or is deleted manually from vSphere. Redistribution will start only after you remove such agent from the web interface.

Viewing the distribution result

You can view the result of the automatic distribution:

- in the **Agent** column for each virtual machine on the **All devices** section
- in the **Assigned virtual machines** section of the **Details** panel when an agent is selected in the **Settings > Agents** section

Manual binding

The Agent for VMware binding lets you exclude a virtual machine from this distribution process by specifying the agent that must always back up this machine. The overall balance will be maintained, but this particular machine can be passed to a different agent only if the original agent is removed.

To bind a machine with an agent

1. Select the machine.
2. Click **Details**.
In the **Assigned agent** section, the software shows the agent that currently manages the selected machine.
3. Click **Change**.
4. Select **Manual**.
5. Select the agent to which you want to bind the machine.
6. Click **Save**.

To unbind a machine from an agent

1. Select the machine.
2. Click **Details**.
In the **Assigned agent** section, the software shows the agent that currently manages the selected machine.
3. Click **Change**.
4. Select **Automatic**.
5. Click **Save**.

Disabling automatic assignment for an agent

You can disable the automatic assignment for Agent for VMware to exclude it from the distribution process by specifying the list of machines that this agent must back up. The overall balance will be maintained between other agents.

Automatic assignment cannot be disabled for an agent if there are no other registered agents, or if automatic assignment is disabled for all other agents.

To disable automatic assignment for an agent

1. Click **Settings > Agents**.
2. Select Agent for VMware for which you want to disable the automatic assignment.
3. Click **Details**.
4. Disable the **Automatic assignment** switch.

Usage examples

- Manual binding comes in handy if you want a particular (very large) machine to be backed up by Agent for VMware (Windows) via a fibre channel while other machines are backed up by virtual appliances.
- It is necessary to bind VMs to an agent if the agent has a locally attached storage.
- Disabling the automatic assignment enables you to ensure that a particular machine is predictably backed up on the schedule you specify. The agent that only backs up one VM cannot be busy backing up other VMs when the scheduled time comes.
- Disabling the automatic assignment is useful if you have multiple ESXi hosts that are separated geographically. If you disable the automatic assignment, and then bind the VMs on each host to the agent running on the same host, you can ensure that the agent will never back up any machines running on the remote ESXi hosts, thus saving network traffic.

Running pre-freeze and post-thaw scripts automatically

With VMware Tools, you can automatically run custom pre-freeze and post-thaw scripts on virtual machines that you back up in the agentless mode. Thus, for example, you can run custom quiescing scripts and create application-consistent backups for virtual machines running applications that are not VSS-aware.

Prerequisites

The pre-freeze and post-thaw scripts must be located in a specific folder on the virtual machine.

- For Windows virtual machines, the location of this folder depends on the ESXi version of the host. For example, for virtual machines running on an ESXi 6.5 host, this folder is `C:\Program Files\VMware\VMware Tools\backupScripts.d`. You must create the `backupScripts.d` folder manually. Do not store other types of files in this folder because this may cause VMware Tools to become unstable.
For more information about the location of the pre-freeze and post-thaw scripts for other ESXi versions, refer to the VMware documentation.
- For Linux virtual machines, copy your scripts to the `/usr/sbin/pre-freeze-script` and `/usr/sbin/post-thaw-script` directories, respectively. The scripts in `/usr/sbin/pre-freeze-script` are run when you create a snapshot and those in `/usr/sbin/post-thaw-script` are run when the snapshot is finalized. The scripts must be executable by the VMware Tools user.

To run pre-freeze and post-thaw scripts automatically

1. Ensure that VMware Tools are installed on the virtual machine.
2. On the virtual machine, put your custom scripts in the required folder.
3. In the protection plan for this machine, enable the **Volume Shadow Copy Service (VSS) for virtual machines** option.

This creates a VMware snapshot with the **Quiesce guest file system** option enabled, which in turn triggers the pre-freeze and post-thaw scripts inside the virtual machine.

You do not need to run custom quiescing scripts on virtual machines running VSS-aware applications, such as Microsoft SQL Server or Microsoft Exchange. To create an application-consistent backup for such machines, enable the **Volume Shadow Copy Service (VSS) for virtual machines** option in the protection plan.

Support for virtual machine migration

This section contains information about migration of virtual machines within a vSphere environment, including migration between ESXi hosts that are part of a vSphere cluster.

vMotion allows moving the state and configuration of a virtual machine to another host, while the machine's disks remain in the same location on a shared storage. Storage vMotion allows moving the disks of a virtual machine from one datastore to another.

- Migration with vMotion, including Storage vMotion, is not supported for a virtual machine that runs Agent for VMware (Virtual Appliance), and is disabled automatically. This virtual machine is added to the **VM overrides** list in the vSphere cluster configuration.
- When a backup of a virtual machine starts, migration with vMotion, including Storage vMotion, is automatically disabled. This virtual machine is temporarily added to the **VM overrides** list in the vSphere cluster configuration. After the backup finishes, the **VM overrides** settings are automatically reverted to their previous state.
- A backup cannot start for a virtual machine while its migration with vMotion, including Storage vMotion, is in progress. The backup for this machine will start when its migration finishes.

Managing virtualization environments

You can view the vSphere, Hyper-V, and Virtuozzo environments in their native presentation. Once the corresponding agent is installed and registered, the **VMware**, **Hyper-V**, or **Virtuozzo** tab appears under **Devices**.

In the **VMware** tab, you can back up the following vSphere infrastructure objects:

- Data center
- Folder
- Cluster

This information appears in the virtual machine summary (**Summary > Custom attributes/Annotations/Notes**, depending on the client type and vSphere version). You can also enable the **Last backup** and **Backup status** columns on the **Virtual Machines** tab for any host, datacenter, folder, resource pool, or the entire vCenter Server.

To provide these attributes, Agent for VMware must have the following privileges in addition to those described in "[Agent for VMware - necessary privileges](#)":

- **Global > Manage custom attributes**
- **Global > Set custom attribute**


Agent for VMware – necessary privileges

To perform any operations with vCenter objects, such as virtual machines, ESXi hosts, clusters, vCenter, and more, Agent for VMware authenticates on vCenter or ESXi host by using the vSphere credentials provided by a user. The vSphere account, used for connection to vSphere by Agent for VMware, must have the required privileges on all levels of vSphere infrastructure starting from the vCenter level.

Specify the vSphere account with the necessary privileges during Agent for VMware installation or configuration. If you need to change the account later, refer to "Managing virtualization environments" (p. 645).

To assign the permissions to a vSphere user on the vCenter level

1. Log in to vSphere web client.
2. Right-click on vCenter and then click **Add permission**.
3. Select or add a new user with the required role (the role must include all the required permissions from the table below).
4. Select the **Propagate to children** option.



Object	Privilege	Operation			
		Back up a VM	Recover to a new VM	Recover to an existing VM	Run VM from backup
Cryptographic operations (starting with vSphere 6.5)	Add disk	+			
	Direct Access	+			
Datastore	Allocate space		+	+	+
	Browse datastore				+
	Configure datastore	+	+	+	+
	Low level file operations				+
Global	Licenses	+	+	+	+
	Disable methods	+	+	+	
	Enable methods	+	+	+	
	Manage custom attributes	+	+	+	
	Set custom attribute	+	+	+	
Host > Configuration	Storage partition configuration				+
Host > Local operations	Create VM				+
	Delete VM				+
	Reconfigure VM				+
Network	Assign network		+	+	+
Resource	Assign VM to resource pool		+	+	+
Virtual machine > Configuration	Add existing disk	+	+		+
	Add new disk		+	+	+
	Add or remove device		+		+
	Advanced	+	+	+	

	Change CPU count		+		
	Disk change tracking	+		+	
	Disk lease	+		+	
	Memory		+		
	Remove disk	+	+	+	+
	Rename		+		
	Set annotation				+
	Settings		+	+	+
Virtual machine > Guest Operations	Guest Operation Program Execution	+++			
	Guest Operation Queries	+++			
	Guest Operation Modifications	+++			
Virtual machine > Interaction	Acquire guest control ticket (in vSphere 4.1 and 5.0)				+
	Configure CD media		+	+	
	Guest operating system management by VIX API (in vSphere 5.1 and later)				+
	Power off			+	+
	Power on		+	+	+
Virtual machine > Inventory	Create from existing		+	+	+
	Create new		+	+	+
	Register				+
	Remove		+	+	+
	Unregister				+
Virtual machine > Provisioning	Allow disk access		+	+	+
	Allow read-only disk access	+		+	
	Allow virtual machine	+	+	+	+

	download				
Virtual machine > State	Create snapshot	+		+	+
Virtual machine > Snapshot management (vSphere 6.5 and later)					
	Remove snapshot	+		+	+
vApp	Add virtual machine				+

* This privilege is required for backing up encrypted machines only.

** This privilege is required for application-aware backups only.

Backing up clustered Hyper-V machines

In a Hyper-V cluster, virtual machines may migrate between cluster nodes. Follow these recommendations to set up a correct backup of clustered Hyper-V machines:

1. A machine must be available for backup no matter what node it migrates to. To ensure that Agent for Hyper-V can access a machine on any node, the agent service must run under a domain user account that has administrative privileges on each of the cluster nodes. We recommend that you specify such an account for the agent service during the Agent for Hyper-V installation.
2. Install Agent for Hyper-V on each node of the cluster.
3. Register all of the agents in the Cyber Protection service.

High Availability of a recovered machine

When you recover backed-up disks to an *existing* Hyper-V virtual machine, the machine's High Availability property remains as is.

When you recover backed-up disks to a *new* Hyper-V virtual machine, the resulting machine is not highly available. It is considered as a spare machine and is normally powered off. If you need to use the machine in the production environment, you can configure it for High Availability from the **Failover Cluster Management** snap-in.

Limiting the total number of simultaneously backed-up virtual machines

In the **Scheduling** backup option, you can limit the number of simultaneously backed-up virtual machines per protection plan.

When an agent runs multiple plans at the same time, the number of simultaneously backed-up machines adds up. This might affect the backup performance and overload the host and the virtual machine storage. You can avoid such issues by configuring a limitation on the agent level.

To limit the number of simultaneous backups on the agent level

Agent for VMware (Windows)

1. On the machine with the agent, create a new text document, and then open it in a text editor.
2. Copy and paste the following lines into the file.

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\MMS\Configuration\ManagedMachine\SimultaneousBackupsLimits]
"MaxNumberOfSimultaneousBackups"=dword:00000001
```

3. Replace 00000001 with the hexadecimal value of the limit that you want to set.
For example, 00000001 is 1 and 0000000A is 10.
4. Save the document as **limit.reg**.
5. Run the file as an administrator.
6. Confirm that you want to edit the Windows registry.
7. Restart the agent.
 - a. In the **Start** menu, click **Run**.
 - b. Type **cmd**, and then click **OK**.
 - c. On the command line, run the following commands:

```
net stop mms
net start mms
```

Agent for Hyper-V

1. On the machine with the agent, create a new text document, and then open it in a text editor.
2. Copy and paste the following lines into the file.

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\MMS\Configuration\ManagedMachine\SimultaneousBackupsLimits]
"MaxNumberOfSimultaneousBackups"=dword:00000001
```

3. Replace 00000001 with the hexadecimal value of the limit that you want to set.
For example, 00000001 is 1 and 0000000A is 10.
4. Save the document as **limit.reg**.
5. Run the file as an administrator.
6. Confirm that you want to edit the Windows registry.
7. Restart the agent.
 - a. In the **Start** menu, click **Run**.
 - b. Type **cmd**, and then click **OK**.
 - c. On the command line, run the following commands:

```
net stop mms
net start mms
```

Virtual appliances

This procedure applies to Agent for VMware (Virtual Appliance), Agent for Scale Computing, Agent for Virtuozzo Hybrid Infrastructure, and Agent for oVirt.

1. In the console of the virtual appliance, press CTRL+SHIFT+F2 to open the command-line interface.
2. Open the `/etc/Acronis/MMS.config` file in a text editor.
3. Locate the following section:

```
<key name="SimultaneousBackupsLimits">
  <value name="MaxNumberOfSimultaneousBackups" type="Tdwor d">"10"</value>
</key>
```

4. Replace 10 with the maximum number of simultaneous backups that you want to set.
5. Save the file.
6. Restart the agent by running the `reboot` command.

Machine migration

You can perform machine migration by recovering its backup to a non-original machine.

The following table summarizes the available migration options.

Backed-up machine type	Available recovery destinations							
	Physical machine	ESXi virtual machine	Hyper-V virtual machine	Virtuozzo		Virtuozzo Hybrid Infrastructure virtual machine	Scale Computing HC3 virtual machine	RHV/oVirt virtual machine
				Virtual machine	Container			
Physical machine	+	+	+	-	-	+	+*	+
VMware ESXi virtual machine	+	+	+	-	-	+	+*	+
Hyper-V virtual machine	+	+	+	-	-	+	+*	+
Virtuozzo	+	+	+	+	-	+	+*	+

virtual machine								
Virtuozzo container	-	-	-	-	+	-	-	-
Virtuozzo Hybrid Infrastructure virtual machine	+	+	+	-	-	+	+*	+
Scale Computing HC3 virtual machine	+	+	+	-	-	+	+	+
Red Hat Virtualization/oVirt virtual machine	+	+	+	-	-	+	+*	+

*If Secure Boot is enabled on the source machine, the recovered VM will not be able to start up unless you disable Secure Boot in the VM console after the recovery.

Note

You cannot recover macOS virtual machines to Hyper-V hosts, because Hyper-V does not support macOS. You can recover macOS virtual machines to a VMware host that is installed on Mac hardware.

For more information on how to perform the migration operations, see the following topics:

- For physical-to-virtual (P2V) migration, see "Physical machine to virtual" (p. 475).
- For virtual-to-virtual (V2V) migration, see "Recovering a virtual machine". You can recover virtual machines from their backups. You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1). Prerequisites A virtual machine must be stopped during the recovery to this machine. By default, the software stops the machine without a prompt. When the recovery is completed, you have to start the machine manually. You can change the default behavior by using the VM power management recovery option (click Recovery options > VM power management). Procedure Do one of the following: Select a backed-up machine, click Recovery, and then select a recovery point. Select a recovery point on the Backup storage tab. Click Recover > Entire machine. If you want to recover to a physical machine, select Physical machine in Recover to. Otherwise, skip this step. Recovery to a physical machine is possible only if the disk configuration of the target machine exactly matches the disk configuration in the backup. If this is the case, continue to step 4 in "Physical machine".

Otherwise, we recommend that you perform the V2P migration by using bootable media.

[Optional] By default, the software automatically selects the original machine as the target machine. To recover to another virtual machine, click Target machine, and then do the following: Select the hypervisor (VMware ESXi, Hyper-V, Virtuozzo, Virtuozzo Hybrid Infrastructure, Scale Computing HC3, or oVirt). Only Virtuozzo virtual machines can be recovered to Virtuozzo. For more information about V2V migration, refer to "Machine migration". Select whether to recover to a new or existing machine. Select the host and specify the new machine name, or select an existing target machine. Click OK. Setup up the additional recovery options that you need. [Not available for Virtuozzo Hybrid Infrastructure and Scale Computing HC3] To select the datastore for the virtual machine, click Datastore for ESXi, Path for Hyper-V and Virtuozzo, or Storage domain for Red Hat Virtualization (oVirt), and then select the datastore (storage) for the virtual machine. To view the datastore (storage), interface, and the provisioning mode for each virtual disk, click Disk mapping. You can change these settings, unless you are recovering a Virtuozzo container or Virtuozzo Hybrid Infrastructure virtual machine. For Virtuozzo Hybrid Infrastructure, you can only select the storage policy for the target disks. To do so, select the desired target disk, and then click Change. In the blade that opens, click the gear icon, select the storage policy, and then click Done. The mapping section also enables you to choose individual disks for recovery. [Available for VMware ESXi, Hyper-V, and Virtuozzo] To change the memory size, the number of processors, and the network connections of the virtual machine, click VM settings. [For Virtuozzo Hybrid Infrastructure] To change the memory size and the number of processors of the virtual machine, select Flavor. [Only available for Windows machines on which a protection agent is installed] Enable the Safe recovery switch to ensure that the recovered data is malware-free. For more information about how safe recovery works, see "Safe recovery" (p. 1). Click Start recovery. When recovering to an existing virtual machine, confirm that you want to overwrite the disks. The recovery progress is shown on the Activities tab." (p. 1).

- For virtual-to-physical (V2P) migration, see "Recovering a virtual machine You can recover virtual machines from their backups. You cannot recover backups in the Cyber Protect console for tenants in the Enhanced security mode. For more information on how to recover such backups, refer to "Recovering backups for tenants in the Enhanced security mode" (p. 1). Prerequisites A virtual machine must be stopped during the recovery to this machine. By default, the software stops the machine without a prompt. When the recovery is completed, you have to start the machine manually. You can change the default behavior by using the VM power management recovery option (click Recovery options > VM power management). Procedure Do one of the following: Select a backed-up machine, click Recovery, and then select a recovery point. Select a recovery point on the Backup storage tab. Click Recover > Entire machine. If you want to recover to a physical machine, select Physical machine in Recover to. Otherwise, skip this step. Recovery to a physical machine is possible only if the disk configuration of the target machine exactly matches the disk configuration in the backup. If this is the case, continue to step 4 in "Physical machine". Otherwise, we recommend that you perform the V2P migration by using bootable media. [Optional] By default, the software automatically selects the original machine as the target machine. To recover to another virtual machine, click Target machine, and then do the following: Select the hypervisor (VMware ESXi, Hyper-V, Virtuozzo, Virtuozzo Hybrid Infrastructure, Scale Computing HC3, or oVirt). Only Virtuozzo virtual machines can be recovered to Virtuozzo.

For more information about V2V migration, refer to "Machine migration". Select whether to recover to a new or existing machine. Select the host and specify the new machine name, or select an existing target machine. Click OK. Setup up the additional recovery options that you need. [Not available for Virtuozzo Hybrid Infrastructure and Scale Computing HC3] To select the datastore for the virtual machine, click Datastore for ESXi, Path for Hyper-V and Virtuozzo, or Storage domain for Red Hat Virtualization (oVirt), and then select the datastore (storage) for the virtual machine. To view the datastore (storage), interface, and the provisioning mode for each virtual disk, click Disk mapping. You can change these settings, unless you are recovering a Virtuozzo container or Virtuozzo Hybrid Infrastructure virtual machine. For Virtuozzo Hybrid Infrastructure, you can only select the storage policy for the target disks. To do so, select the desired target disk, and then click Change. In the blade that opens, click the gear icon, select the storage policy, and then click Done. The mapping section also enables you to choose individual disks for recovery. [Available for VMware ESXi, Hyper-V, and Virtuozzo] To change the memory size, the number of processors, and the network connections of the virtual machine, click VM settings. [For Virtuozzo Hybrid Infrastructure] To change the memory size and the number of processors of the virtual machine, select Flavor. [Only available for Windows machines on which a protection agent is installed] Enable the Safe recovery switch to ensure that the recovered data is malware-free. For more information about how safe recovery works, see "Safe recovery" (p. 1). Click Start recovery. When recovering to an existing virtual machine, confirm that you want to overwrite the disks. The recovery progress is shown on the Activities tab." (p. 1) and "Recovering disks by using bootable media" (p. 479).

Migration via a bootable media

As an alternative to the machine migration that you perform in the Cyber Protect console, you can recover a machine by using a bootable media.

We recommend that you use a bootable media in the following cases:

- Performing a migration that is not natively supported.
For example, use a bootable media to recover a physical machine or a non-Virtuozzo virtual machine as a Virtuozzo virtual machine on a Virtuozzo host.
- Performing migration of a Linux machine that contains logical volumes (LVM).
Use Agent for Linux or bootable media to create the backup, and then use a bootable media to recover the backup.
- Providing drivers for specific hardware that is critical for the system bootability.
Build a bootable media that can use the required drivers. For more information, see "Bootable Media Builder" (p. 658).

Microsoft Azure and Amazon EC2 virtual machines

To back up a Microsoft Azure or Amazon EC2 virtual machine, install a protection agent on the machine. The backup and recovery operations are the same as with a physical machine. Nevertheless, the machine is counted as virtual when you set quotas for the number of machines.

The difference from a physical machine is that Microsoft Azure and Amazon EC2 virtual machines cannot be booted from bootable media. If you need to recover to a new Microsoft Azure or Amazon EC2 virtual machine, follow the procedure below.

Note

The following recovery procedure applies only for backups of machines that contain all necessary drivers to run in Microsoft Azure natively (backups created of an Azure VM, a local Hyper-V machine, or the source machine being a Windows Server 2016 and up). For cross-platform recovery, please see [this knowledge base article](#).

To recover a machine as a Microsoft Azure or Amazon EC2 virtual machine

1. Create a new virtual machine from an image/template in Microsoft Azure or Amazon EC2. The new machine must have the same disk configuration as the machine that you want to recover.
2. Install Agent for Windows or Agent for Linux on the new machine.
3. Recover the backed-up machine as described in "[Physical machine](#)". When configuring the recovery, select the new machine as the target machine.

Creating bootable media to recover operating systems

Bootable media is a CD, DVD, USB flash drive, or other removable media that allows you to run the protection agent either in a Linux-based environment or a Windows Preinstallation Environment/Windows Recovery Environment (WinPE/WinRE), without the help of an operating system. The main purpose of the bootable media is to recover an operating system that cannot start.

Note

Bootable media does not support hybrid drives.

Custom or ready-made bootable media?

By using Bootable Media Builder, you can create custom bootable media (Linux-based or WinPE-based) for Windows, Linux, or macOS computers. In the both Linux-based and WinPE/WinRE-based custom bootable media, you can configure additional settings, such as automatic registration, network settings, or proxy server settings. In the WinPE/WinRE-based custom bootable media, you can also add additional drivers.

Alternatively, you can download a ready-made bootable media (Linux-based only). You can use the ready-made bootable media for recovery operations and access to the Universal Restore feature.

Linux-based or WinPE/WinRE-based bootable media?

Linux-based

Linux-based bootable media contains a protection agent based on a Linux kernel. The agent can boot and perform operations on any PC-compatible hardware, including bare metal, and machines with corrupted or non-supported file systems.

WinPE/WinRE-based

WinPE-based bootable media contains a minimal Windows system called Windows Preinstallation Environment (WinPE) and a Cyber Protection plugin for WinPE, that is, a modification of the protection agent that can run in the preinstallation environment. WinRE-based bootable media uses Windows Recovery Environment and does not require installation of additional Windows packages.

WinPE proved to be the most convenient bootable solution in large environments with heterogeneous hardware.

Advantages:

- Using Cyber Protection in Windows Preinstallation Environment provides more functionality than using Linux-based bootable media. Having booted PC-compatible hardware into WinPE, you can use not only the protection agent, but also PE commands and scripts, and other plugins that you have added to the PE.
- PE-based bootable media helps overcome some Linux-related bootable media issues, such as support for certain RAID controllers or certain levels of RAID arrays only. Media based on WinPE 2.x and later allows dynamic loading of the necessary device drivers.

Limitations:

- Bootable media based on WinPE versions earlier than 4.0 cannot boot on machines that use Unified Extensible Firmware Interface (UEFI).

Creating physical bootable media

We highly recommend that you create and test the bootable media as soon as you start using disk-level backup. Also, it is a good practice to re-create the media after each major update of the protection agent.

You can recover either Windows or Linux by using the same media. To recover macOS, create a separate media on a machine running macOS.

To create physical bootable media in Windows or Linux

1. Create a custom bootable media ISO file or download the ready-made ISO file.
To create a custom ISO file, use "Bootable Media Builder" (p. 658).

To download the ready-made ISO file, in the Cyber Protect console, select a machine, and then click **Recover > More ways to recover... > Download ISO image**.

2. [Optional] In the Cyber Protect console, generate a registration token. The registration token is displayed automatically when you download a ready-made ISO file.

This token allows the bootable media to access the cloud storage, without prompting you to enter a login and password.

3. Create physical bootable media in one of the following ways:

- Burn the ISO file to a CD/DVD.
- Create a bootable USB flash drive by using the ISO file and one of the free tools available online.

Use ISO to USB or RUFUS if you need to boot an UEFI machine, and Win32DiskImager for a BIOS machine. In Linux, using the dd utility makes sense.

For virtual machines, you can connect the ISO file as a CD/DVD drive to the machine that you want to recover.

To create physical bootable media in macOS

1. On a machine where Agent for Mac is installed, click **Applications > Rescue Media Builder**.
2. The software displays the connected removable media. Select the one that you want to make bootable.

Warning!

All data on the disk will be erased.

3. Click **Create**.
4. Wait while the software creates the bootable media.

Bootable Media Builder

Bootable Media Builder is a dedicated tool for creating bootable media. It is installed as an optional component on the machine where the protection agent is installed.

Why use Bootable Media Builder?

The ready-made bootable media that is available for download in the Cyber Protect console is based on a Linux kernel. Unlike Windows PE, it does not allow injecting custom drivers on the fly.

Bootable Media Builder allows you to create customized Linux-based and WinPE-based bootable media images.

32-bit or 64-bit?

Bootable Media Builder creates bootable media with both 32-bit and 64-bit components. In most cases, you will need a 64-bit media to boot a machine that uses Unified Extensible Firmware Interface (UEFI).

Linux-based bootable media

To create a Linux-based bootable media

1. Start **Bootable Media Builder**.
2. In **Bootable media type**, select **Default (Linux-based media)**.
3. Select how volumes and network resources will be represented:
 - Bootable media with a Linux-like volume representation displays the volumes as, for example, hda1 and sdb2. It tries to reconstruct MD devices and logical volumes (LVM) before starting a recovery.
 - Bootable media with Windows-like volume representation displays the volumes as, for example, C: and D:. It provides access to dynamic volumes (LDM).
4. [Optional] Specify the parameters of the Linux kernel. Separate multiple parameters with spaces. For example, to be able to select a display mode for the bootable agent each time the media starts, type: **vga=ask**. For more information about the available parameters, refer to "Kernel parameters" (p. 659).
5. [Optional] Select the language for the bootable media.
6. [Optional] Select the boot mode (BIOS or UEFI) that Windows will use after the recovery.
7. Select the component to be placed on the media – the Cyber Protection bootable agent.
8. [Optional] Specify the timeout interval for the boot menu. If this setting is not configured, the loader will wait for you to select whether to boot the operating system (if present) or the component.
9. [Optional] If you want to automate the bootable agent operations, select the **Use the following script** check box. Then, select one of the scripts and specify the script parameters. For more information about the scripts, refer to "Scripts in bootable media" (p. 662).
10. [Optional] Select how to register the bootable media in the Cyber Protection service on booting up. For more information about the registration settings, refer to "Registering the bootable media" (p. 670).
11. Specify the network settings for the network adapters of the booted machine or keep the automatic DHCP configuration.
12. [Optional] If a proxy server is enabled in your network, specify its host name/IP address and port.
13. Select the file type of the created bootable media:
 - ISO image
 - ZIP file
14. Specify a file name for the bootable media file.
15. Check your settings in the summary screen, and then click **Proceed**.

Kernel parameters

You can specify one or more parameters of the Linux kernel that will be automatically applied when the bootable media starts. These parameters are typically used when you experience problems

while working with the bootable media. Normally, you can leave this field empty.

You can also specify any of these parameters by pressing F11 while you are in the boot menu.

Parameters

When specifying multiple parameters, separate them with spaces.

- **acpi=off**
Disables Advanced Configuration and Power Interface (ACPI). You may want to use this parameter when experiencing problems with a particular hardware configuration.
- **noapic**
Disables Advanced Programmable Interrupt Controller (APIC). You may want to use this parameter when experiencing problems with a particular hardware configuration.
- **vga=ask**
Prompts for the video mode to be used by the bootable media's graphical user interface. Without the **vga** parameter, the video mode is detected automatically.
- **vga= *mode_number***
Specifies the video mode to be used by the bootable media's graphical user interface. The mode number is given by *mode_number* in the hexadecimal format—for example: **vga=0x318**
The screen resolution and the number of colors corresponding to a mode number may be different on different machines. We recommend that you use the **vga=ask** parameter first to choose a value for *mode_number*.
- **quiet**
Disables displaying of startup messages when the Linux kernel is loading, and starts the management console after the kernel is loaded.
This parameter is implicitly specified when creating the bootable media, but you can remove this parameter while you are in the boot menu.
If this parameter is removed, all startup messages will be displayed, followed by a command prompt. To start the management console from the command prompt, run the command:
/bin/product
- **nousb**
Disables loading of the USB (Universal Serial Bus) subsystem.
- **nousb2**
Disables USB 2.0 support. USB 1.1 devices still work with this parameter. This parameter allows you to use some USB drives in the USB 1.1 mode if they do not work in the USB 2.0 mode.
- **nodma**
Disables direct memory access (DMA) for all IDE hard disk drives. Prevents the kernel from freezing on some hardware.
- **nofw**
Disables the FireWire (IEEE1394) interface support.
- **nopcmcia**
Disables the detection of PCMCIA hardware.

- **nomouse**
Disables the mouse support.
- **module_name =off**
Disables the module whose name is given by *module_name*. For example, to disable the use of the SATA module, specify: **sata_sis=off**
- **pci=bios**
Forces the use of PCI BIOS instead of accessing the hardware device directly. You may want to use this parameter if the machine has a non-standard PCI host bridge.
- **pci=nobios**
Disables the use of PCI BIOS; only direct hardware access methods will be allowed. You may want to use this parameter when the bootable media fails to start, which may be caused by the BIOS.
- **pci=biosirq**
Uses PCI BIOS calls to get the interrupt routing table. You may want to use this parameter if the kernel is unable to allocate interrupt requests (IRQs) or discover secondary PCI buses on the motherboard.
These calls might not work properly on some machines. But this may be the only way to get the interrupt routing table.
- **LAYOUTS=en-US, de-DE, fr-FR, ...**
Specifies the keyboard layouts that can be used in the bootable media's graphical user interface. Without this parameter, only two layouts can be used: English (USA) and the layout that corresponds to the language selected in the media's boot menu.
You can specify any of the following layouts:
Belgian: **be-BE**
Czech: **cz-CZ**
English: **en-GB**
English (USA): **en-US**
French: **fr-FR**
French (Swiss): **fr-CH**
German: **de-DE**
German (Swiss): **de-CH**
Italian: **it-IT**
Polish: **pl-PL**
Portuguese: **pt-PT**
Portuguese (Brazilian): **pt-BR**
Russian: **ru-RU**
Serbian (Cyrillic): **sr-CR**
Serbian (Latin): **sr-LT**
Spanish: **es-ES**
When working under a bootable media, use CTRL + SHIFT to cycle through the available layouts.

Scripts in bootable media

If you want the bootable media to perform a predefined set of operations, you can specify a script while creating the media with Bootable Media Builder. Thus, every time a machine is booted from the media, the specified script will run and the user interface will not be shown.

You can select one of the predefined scripts or create a custom script by following the scripting conventions.

Predefined scripts

Bootable Media Builder provides the following predefined scripts:

- Recovery from the cloud storage (**entire_pc_cloud**)
- Recovery from a network share (**entire_pc_share**)

The scripts are located in the following folders on the machine where Bootable Media Builder is installed:

- In Windows: **%ProgramData%\Acronis\MediaBuilder\scripts**
- In Linux: **/var/lib/Acronis/MediaBuilder/scripts/**

Recovery from the cloud storage

In Bootable Media Builder, specify the following script parameters:

1. The backup file name.
2. [Optional] A password that the script will use to access encrypted backups.

Recovery from a network share

In Bootable Media Builder, specify the following script parameters:

- The path to the network share.
- The user name and password for the network share.
- The backup file name. To find out the backup file name:
 - a. In the Cyber Protect console, go to **Backup storage > Locations**.
 - b. Select the network share (click **Add location** if the share is not listed).
 - c. Select the backup.
 - d. Click **Details**. The file name is displayed under **Backup file name**.
- [Optional] A password that the script will use to access encrypted backups.

Custom scripts

Important

Creating custom scripts requires the knowledge of the Bash command language and JavaScript Object Notation (JSON). If you are not familiar with Bash, a good place to learn it is <http://www.tldp.org/LDP/abs/html>. The JSON specification is available at <http://www.json.org>.

Files of a script

Your script must be located in the following directories on the machine where Bootable Media Builder is installed:

- In Windows: **%ProgramData%\Acronis\MediaBuilder\scripts**
- In Linux: **/var/lib/Acronis/MediaBuilder/scripts/**

The script must consist of at least three files:

- **<script_file>.sh** - a file with your Bash script. When creating the script, use only a limited set of shell commands, which you can find at <https://busybox.net/downloads/BusyBox.html>. Also, the following commands can be used:

- **acrocmd** - the command-line utility for backup and recovery
- **product** - the command that starts the bootable media user interface

This file and any additional files that the script includes (for example, by using the dot command) must be located in the **bin** subfolder. In the script, specify the additional file paths as **/ConfigurationFiles/bin/<some_file>**.

- **autostart** - a file for starting **<script_file>.sh**. The file contents must be as follows:

```
#!/bin/sh
. /ConfigurationFiles/bin/variables.sh
. /ConfigurationFiles/bin/<script_file>.sh
. /ConfigurationFiles/bin/post_actions.sh
```

- **autostart.json** - a JSON file that contains the following:
 - The script name and description to be displayed in Bootable Media Builder.
 - The names of the script variables to be configured via Bootable Media Builder.
 - The parameters of controls that will be displayed in Bootable Media Builder for each variable.

Structure of autostart.json

Top-level object

Pair		Required	Description
Name	Value type		
displayName	string	Yes	The script name to be displayed in Bootable Media Builder.
description	string	No	The script description to be displayed in Bootable Media Builder.
timeout	number	No	A timeout (in seconds) for the boot menu before starting the script. If the pair is not specified, the timeout will be ten seconds.
variables	object	No	Any variables for <script_file>.sh that you want to configure via Bootable Media Builder. The value should be a set of the following pairs: the string identifier of a variable and the object of the variable (see the table below).

Variable object

Pair		Required	Description
Name	Value type		
displayName	string	Yes	The variable name used in <script_file>.sh .
type	string	Yes	The type of a control that is displayed in Bootable Media Builder. This control is used to configure the variable value. For all supported types, see the table below.
description	string	Yes	The control label that is displayed above the control in Bootable Media Builder.
default	string if type is string, multiString, password, or enum number if	No	The default value for the control. If the pair is not specified, the default value will be an empty string or a zero, based on the control type. The default value for a check box can be 0 (the cleared state) or 1 (the selected state).

	type is number, spinner, or checkbox		
order	number (non-negative)	Yes	The control order in Bootable Media Builder. The higher the value, the lower the control is placed relative to other controls defined in autostart.json . The initial value must be 0.
min (for spinner only)	number	No	The minimum value of the spin control in a spin box. If the pair is not specified, the value will be 0.
max (for spinner only)	number	No	The maximum value of the spin control in a spin box. If the pair is not specified, the value will be 100.
step (for spinner only)	number	No	The step value of the spin control in a spin box. If the pair is not specified, the value will be 1.
items (for enum only)	array of strings	Yes	The values for a drop-down list.
required (for string, multiString, password, and enum)	number	No	Specifies if the control value can be empty (0) or not (1). If the pair is not specified, the control value can be empty.

Control type

Name	Description
string	A single-line, unconstrained text box used to enter or edit short strings.
multiString	A multi-line, unconstrained text box used to enter or edit long strings.
password	A single-line, unconstrained text box used to enter passwords securely.
number	A single-line, numeric-only text box used to enter or edit numbers.
spinner	A single-line, numeric-only text box used to enter or edit numbers, with a spin control. Also, called a spin box.
enum	A standard drop-down list, with a fixed set of predetermined values.
checkbox	A check box with two states - the cleared state or the selected state.

The sample **autostart.json** below contains all possible types of controls that can be used to configure variables for **<script_file>.sh**.

```
{
  "displayName": "Autostart script name",
  "description": "This is an autostart script description.",
  "variables": {
    "var_string": {
      "displayName": "VAR_STRING",
      "type": "string", "order": 1,
      "description": "This is a 'string' control:", "default": "Hello,
world!"
    },
    "var_multistring": {
      "displayName": "VAR_MULTISTRING",
      "type": "multiString", "order": 2,
      "description": "This is a 'multiString' control:",
      "default": "Lorem ipsum dolor sit amet,\nconsectetur adipiscing elit."
    },
    "var_number": {
      "displayName": "VAR_NUMBER",
      "type": "number", "order": 3,
      "description": "This is a 'number' control:", "default": 10
    },
    "var_spinner": {
      "displayName": "VAR_SPINNER",
      "type": "spinner", "order": 4,
      "description": "This is a 'spinner' control:",
      "min": 1, "max": 10, "step": 1, "default": 5
    },
    "var_enum": {
      "displayName": "VAR_ENUM",
```

```

        "type": "enum", "order": 5,
        "description": "This is an 'enum' control:",
        "items": ["first", "second", "third"], "default": "second"
    },
    "var_password": {
        "displayName": "VAR_PASSWORD",
        "type": "password", "order": 6,
        "description": "This is a 'password' control:", "default": "qwe"
    },
    "var_checkbox": {
        "displayName": "VAR_CHECKBOX",
        "type": "checkbox", "order": 7,
        "description": "This is a 'checkbox' control", "default": 1
    }
}
}
}

```

WinPE-based and WinRE-based bootable media

You can create WinRE images without any additional preparation, or create WinPE images after installing [Windows Automated Installation Kit \(AIK\)](#) or [Windows Assessment and Deployment Kit \(ADK\)](#).

WinRE images

Creating WinRE images is supported for the following operation systems:

- Windows 7 (64-bit)
- Windows 8, 8.1, 10 (32-bit and 64-bit)
- Windows Server 2012, 2016, 2019 (64-bit)

WinPE images

After installing Windows Automated Installation Kit (AIK), or Windows Assessment and Deployment Kit (ADK), Bootable Media Builder supports WinPE distributions that are based on any the following kernels:

- Windows Vista (PE 2.0)
- Windows Vista SP1 and Windows Server 2008 (PE 2.1)

- Windows 7 (PE 3.0) with or without the supplement for Windows 7 SP1 (PE 3.1)
- Windows 8 (PE 4.0)
- Windows 8.1 (PE 5.0)
- Windows 10 (PE for Windows 10)

Bootable Media Builder supports both 32-bit and 64-bit WinPE distributions. The 32-bit WinPE distributions can also work on 64-bit hardware. However, you need a 64-bit distribution to boot a machine that uses Unified Extensible Firmware Interface (UEFI).

Note

PE images based on WinPE 4 and later require approximately 1 GB of RAM to work.

Creating WinPE or WinRE bootable media

Bootable Media Builder provides two methods of integrating Cyber Protection with WinPE and WinRE:

- Creating an ISO file with the Cyber Protection plugin from scratch.
- Adding the Cyber Protection plugin to a WIM file for any future purpose (manual ISO building, adding other tools to the image and so on).

To create WinPE or WinRE bootable media

1. On the machine where the protection agent is installed, run Bootable Media Builder.
2. In **Bootable media type**, select **Windows PE** or **Windows PE (64-bit)**. A 64-bit media is required to boot a machine that uses Unified Extensible Firmware Interface (UEFI).
3. Select the subtype of the bootable media: **WinRE** or **WinPE**.

Creating WinRE bootable media does not require installation of any additional packages.

To create a 64-bit WinPE media, you must download Windows Automated Installation Kit (AIK) or Windows Assessment and Deployment Kit (ADK). To create 32-bit WinPE media, in addition to downloading the AIK or ADK, you need to do the following:

- a. Click **Download the Plug-in for WinPE (32-bit)**.
 - b. Save the plugin to **%PROGRAM_FILES%\BackupClient\BootableComponents\WinPE32**.
4. [Optional] Select the language for the bootable media.
 5. [Optional] Select the boot mode (BIOS or UEFI) that Windows will use after the recovery.
 6. Specify the network settings for the network adapters of the booted machine or keep the automatic DHCP configuration.
 7. [Optional] Select how to register the bootable media in the Cyber Protection service on booting up. For more information about the registration settings, refer to "Registering the bootable media" (p. 670).
 8. [Optional] Specify the Windows drivers to be added to the bootable media.

After you boot a machine into Windows PE or Windows RE, the drivers can help you access the device where the backup is located. Add 32-bit drivers if you use a 32-bit WinPE or WinRE distribution or 64-bit drivers if you use a 64-bit WinPE or WinRE distribution.

To add the drivers:

- Click **Add**, and then specify the path to the necessary .inf file for a corresponding SCSI, RAID, SATA controller, network adapter, tape drive, or other device.
- Repeat this procedure for each driver that you want to include in the resulting WinPE or WinRE media.

9. Select the file type of the created bootable media:

- ISO image
- WIM image

10. Specify the full path to the resulting image file, including the file name.

11. Check your settings in the summary screen, and then click **Proceed**.

To create a PE image (ISO file) from the resulting WIM file

- Replace the default boot.wim file in your Windows PE folder with the newly created WIM file. For the above example, type:

```
copy c:\RecoveryWIMMedia.wim c:\winpe_x86\ISO\sources\boot.wim
```

- Use the **Oscdimg** tool. For the above example, type:

```
oscdimg -n -bc:\winpe_x86\etfsboot.com c:\winpe_x86\ISO c:\winpe_x86\winpe_x86.iso
```

Warning!

Do not copy and paste this example. Type the command manually, otherwise it will fail.

Preparation: WinPE 2.x and 3.x

To be able to create or modify PE 2.x or 3.x images, install Bootable Media Builder on a machine where Windows Automated Installation Kit (AIK) is installed. If you do not have a machine with AIK, prepare it as follows.

To prepare a machine with AIK

1. Download and install Windows Automated Installation Kit.

Automated Installation Kit (AIK) for Windows Vista (PE 2.0):

<http://www.microsoft.com/Downloads/details.aspx?familyid=C7D4BC6D-15F3-4284-9123-679830D629F2&displaylang=en>

Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008 (PE 2.1):

<http://www.microsoft.com/downloads/details.aspx?FamilyID=94bb6e34-d890-4932-81a5-5b50c657de08&DisplayLang=en>

Automated Installation Kit (AIK) for Windows 7 (PE 3.0):

<http://www.microsoft.com/downloads/details.aspx?familyid=696DD665-9F76-4177-A811-39C26D3B3B34&displaylang=en>

Automated Installation Kit (AIK) Supplement for Windows 7 SP1 (PE 3.1):

<http://www.microsoft.com/download/en/details.aspx?id=5188>

You can find system requirements for installation by following the above links.

2. [Optional] Burn the WAIK to DVD or copy it to a flash drive.
3. Install the Microsoft .NET Framework from this kit (NETFXx86 or NETFXx64, depending on your hardware).
4. Install Microsoft Core XML (MSXML) 5.0 or 6.0 Parser from this kit.
5. Install Windows AIK from this kit.
6. Install Bootable Media Builder on the same machine.

Preparation: WinPE 4.0 and later

To be able to create or modify PE 4 or later images, install Bootable Media Builder on a machine where Windows Assessment and Deployment Kit (ADK) is installed. If you do not have a machine with ADK, prepare it as follows.

To prepare a machine with ADK

1. Download the setup program of Assessment and Deployment Kit.
Assessment and Deployment Kit (ADK) for Windows 8 (PE 4.0): <http://www.microsoft.com/en-us/download/details.aspx?id=30652>.
Assessment and Deployment Kit (ADK) for Windows 8.1 (PE 5.0): <http://www.microsoft.com/en-US/download/details.aspx?id=39982>.
Assessment and Deployment Kit (ADK) for Windows 10 (PE for Windows 10):
<https://msdn.microsoft.com/en-us/windows/hardware/dn913721%28v=vs.8.5%29.aspx>.
You can find system requirements for installation by following the above links.
2. Install Assessment and Deployment Kit on the machine.
3. Install Bootable Media Builder on the same machine.

Registering the bootable media

Registering the bootable media in the Cyber Protection service allows accessing the cloud storage for your backups. You can preconfigure the registration while creating the bootable media. If the registration is not preconfigured, you can register the media after booting a machine with it.

To preconfigure the registration in the Cyber Protection service

1. In Bootable Media Builder, navigate to **Bootable media registration**.
2. In **Service URL**, specify the Cyber Protection service address.
3. [Optional] In **Display name**, specify a name for the booted machine.
4. To set the automatic registration in the Cyber Protection service, select the **Register the bootable media automatically** check box, and then select the level of automatic registration:
 - **Ask for registration token at booting up**
The token has to be provided every time when a machine is booted from this bootable media.
 - **Use the following token**
The machine will be registered automatically when it is booted from this bootable media.

To register the bootable media after booting a machine from it

1. Boot the machine from the bootable media.
2. In the startup window, click **Register media**.
3. In **Server**, specify the Cyber Protection service address.
4. In **Registration token**, enter the registration token.
5. Click **Register**.

Network settings

While creating bootable media, you can preconfigure the network connections that will be used by the bootable agent. The following parameters can be preconfigured:

- IP address
- Subnet mask
- Gateway
- DNS server
- WINS server

After the bootable agent starts on a machine, the configuration is applied to the machine's network interface card (NIC). If the settings have not been preconfigured, the agent uses DHCP auto configuration.

You can also configure the network settings manually when the bootable agent is running on the machine.

Preconfiguring multiple network connections

You can preconfigure TCP/IP settings for up to ten network interface cards (NICs). To ensure that each NIC will be assigned the appropriate settings, create the media on the server for which the media is customized. When you select an existing NIC in the wizard window, its settings are selected and saved on the media. The MAC address of each existing NIC is also saved on the media.

You can change the settings, except for the MAC address, or configure the settings for a non-existent NIC.

After the bootable agent starts on the server, it retrieves the list of available NICs. This list is sorted by the slots that the NICs occupy, the closest to the processor is on top.

The bootable agent assigns each known NIC the appropriate settings, and identifies the NICs by their MAC addresses. After the NICs with known MAC addresses are configured, the remaining NICs are assigned the settings that you made for non-existent NICs, starting from the upper non-assigned NIC.

You can customize the bootable media for any machine, and not only for the machine where the media is created. To do so, configure the NICs according to their slot order on that machine: NIC1 occupies the slot closest to the processor, NIC2 is in the next slot, and so on. When the bootable agent starts on that machine, it will not find the NICs with known MAC addresses and will configure the NICs in the same order as you did.

Example

The bootable agent can use one of the network adapters for communication with the management console through the production network. Automatic configuration can be done for this connection. Sizeable data for recovery can be transferred through the second NIC, included in the dedicated backup network by means of static TCP/IP settings.

Connecting to a machine booted from bootable media

Local connection

To operate directly on the machine booted from bootable media, click **Manage this machine locally** in the startup window.

After a machine boots from bootable media, the machine terminal displays a startup window with the IP addresses obtained from DHCP or set according to the preconfigured values.

Configuring network settings

To change the network settings for a current session, in the startup window, click **Configure network**. The **Network Settings** window that appears allows you to configure the network settings for each network interface card (NIC) of the machine.

The changes that are made during a session will be lost after the machine reboots.

Adding VLANs

In the **Network Settings** window, you can add virtual local area networks (VLANs). Use this functionality if you need access to a backup location that is included in a specific VLAN.

VLANs are mainly used to divide a local area network into segments. A NIC that is connected to an *access* port of the switch always has access to the VLAN specified in the port configuration. A NIC connected to a *trunk* port of the switch can access the VLANs allowed in the port configuration only if you specify the VLANs in the network settings.

To enable access to a VLAN via a trunk port

1. Click **Add VLAN**.
2. Select the NIC that provides access to the local area network that includes the required VLAN.
3. Specify the VLAN identifier.

After you click **OK**, a new entry appears in the list of network adapters.

If you need to remove a VLAN, click the required VLAN entry, and then click **Remove VLAN**.

Local operations with bootable media

Operations with bootable media are similar to the recovery operations that are performed under a running operating system. The differences are as follows:

1. Under bootable media with a Windows-like volume representation, a volume has the same drive letter as in Windows. Volumes that do not have drive letters in Windows (such as the System Reserved volume) are assigned free letters in order of their sequence on the disk.

If the bootable media cannot detect Windows on the machine or detects more than one, all volumes, including those without drive letters, are assigned letters in order of their sequence on the disk. Thus, the volume letters may differ from those seen in Windows. For example, the D: drive under the bootable media might correspond to the E: drive in Windows.

Note

It is advisable to assign unique names to the volumes.

2. The bootable media with a Linux-like volume representation shows local disks and volumes as unmounted (sda1, sda2...).
3. Tasks cannot be scheduled. If you need to repeat an operation, configure it from scratch.
4. The log lifetime is limited to the current session. You can save the entire log or the filtered log entries to a file.

Setting up a display mode

When you boot a machine via Linux-based bootable media, a display video mode is detected automatically based on the hardware configuration (monitor and graphics card specifications). If the video mode is detected incorrectly, do the following:

1. In the boot menu, press F11.
2. On the command line, enter **vga=ask**, and then proceed with booting.
3. From the list of supported video modes, choose the appropriate mode by typing its number (for example, **318**), and then press **Enter**.

If you do not want to follow this procedure every time you boot a given hardware configuration, recreate the bootable media with the appropriate mode number (in the example above, **vga=0x318**) specified in the **Kernel parameters** field.

Recovery with bootable media on-premises

1. Boot your machine from the bootable media.
2. Click **Manage this machine locally**.
3. Click **Recover**.
4. In **What to recover**, click **Select data**.
5. Select the backup file that you want to recover from.
6. In the lower left pane, select the drives/volumes or files/folders that you want to recover, and then click **OK**.
7. Configure the overwriting rules.
8. Configure the recovery exclusions.

9. Configure the recovery options.
10. Check that your settings are correct, and then click **OK**.

Remote operations with bootable media

Note

This feature is available with the Advanced Backup pack.

To see the bootable media in the Cyber Protect console, first you need to register it as described in "Registering the bootable media" (p. 670).

After you register the media in the Cyber Protect console, it appears on the **Devices > Bootable media** tab. A bootable media disappears from this tab when it has been offline for more than 30 days.

You can manage the bootable media remotely in the Cyber Protect console. For example, you can recover data, restart the or shut down the machine booted with the media, or view information, activities, and alerts about the media.

Important

You cannot update the bootable media remotely, on the **Settings > Agents** tab in the Cyber Protect console.

To update the bootable media, create a new one, as described in the "Bootable Media Builder" (p. 658) section. Alternatively, download the ready-made media, by clicking your account icon **> Downloads > Bootable media** in the Cyber Protect console.

To recover files or folders with bootable media remotely

1. In the Cyber Protect console, go to **Devices > Bootable media**.
1. Select the media that you want to use for data recovery.
2. Click **Recovery**.
3. Select the location, and then select the backup that you need. Note that backups are filtered by location.
4. Select the recovery point, and then click **Recover files/folders**.
5. Browse to the required folder or use the search bar to obtain the list of the required files and folders.
Search is language-independent.
You can use one or more wildcard characters (* and ?). For more details about using wildcards, refer to "File filters (Inclusions/Exclusions)" (p. 435).
6. Click to select the files that you want to recover, and then click **Recover**.
7. In **Path**, select the recovery destination.
8. [Optional] For advanced recovery configuration, click **Recovery options**. For more information, refer to "Recovery options" (p. 491).
9. Click **Start recovery**.

10. Select one of the file overwriting options:

- **Overwrite existing files**
- **Overwrite an existing file if it is older**
- **Do not overwrite existing files**

Choose whether to restart the machine automatically.

11. Click **Proceed** to start the recovery. The recovery progress is shown on the **Activities** tab.

To recover disks, volumes, or entire machines with bootable media remotely

1. On the **Devices** tab, go to the **Bootable media** group, and then select the media that you want to use for data recovery.
2. Click **Recovery**.
3. Select the location, and then select the backup that you need. Note that backups are filtered by location.
4. Select the recovery point, and then click **Recover > Entire machine**.

If necessary, configure the target machine and volume mapping as described in "Recovering physical machines". This section describes recovery of physical machines by using the web interface. Use bootable media instead of the web interface if you need to recover:

- A machine running macOS
- A machine from a tenant in the Enhanced security mode
- Any operating system to bare metal or to an offline machine

The structure of logical volumes (volumes created by Logical Volume Manager in Linux). The media enables you to recreate the logical volume structure automatically. You cannot recover disk-level backups of Intel-based Macs to Macs that use Apple silicon processors, and vice-versa. You can recover files and folders.

Recovery with restart
Recovery of an operating system and recovery of volumes that are encrypted with BitLocker requires a restart. You can choose whether to restart the machine automatically or assign it the Interaction required status. The recovered operating system goes online automatically. Backed-up encrypted volumes are recovered as non-encrypted. Recovery of BitLocker-encrypted volumes requires that there is a non-encrypted volume on the same machine, and that this volume has at least 1 GB of free space. If either condition is not met, the recovery fails.

Recovering an encrypted system volume does not require any additional actions. To recover an encrypted non-system volume, you must lock it first, for example, by opening a file that resides on this volume. Otherwise, the recovery will continue without restart and the recovered volume might not be recognized by Windows.

If the recovery fails and your machine restarts with the Cannot get file from partition error, try disabling Secure Boot. For more information on how to do it, refer to Disabling Secure Boot in the Microsoft documentation.

To recover a physical machine
Select the backed-up machine. Click Recovery. Select a recovery point. Note that recovery points are filtered by location. If the machine is offline, the recovery points are not displayed. Do any of the following:
If the backup location is cloud or shared storage (i.e. other agents can access it), click Select machine, select a target machine that is online, and then select a recovery point. Select a recovery point on the Backup storage tab. Recover the machine as described in "Recovering disks by using bootable media". Click Recover > Entire machine. The software automatically maps the disks from the backup to the disks of the target machine. To recover to another physical machine, click Target machine, and then select a target machine that

is online. If you are unsatisfied with the mapping result or if the disk mapping fails, click Volume mapping to re-map the disks manually. The mapping section also enables you to choose individual disks or volumes for recovery. You can switch between recovering disks and volumes by using the Switch to... link in the upper-right corner. [Only available for Windows machines on which a protection agent is installed] Enable the Safe recovery switch to ensure that the recovered data is malware-free. For more information about how safe recovery works, see "Safe recovery" (p. 1). Click Start recovery. Confirm that you want to overwrite the disks with their backed-up versions. Choose whether to restart the machine automatically. The recovery progress is shown on the Activities tab." (p. 1).

5. For advanced recovery configuration, click **Recovery options**. For more information, refer to "Recovery options" (p. 491).
6. Click **Start recovery**.
7. Confirm that you want to overwrite the disks with their backed-up versions. Choose whether to restart the machine automatically.
8. The recovery progress is shown on the **Activities** tab.

To restart the booted machine remotely

1. On the **Devices** tab, go to the **Bootable media** group, and then select the media that you want to use for data recovery.
2. Click **Reboot**.
3. Confirm that you want to restart the machine booted with the media.

To shut down the booted machine remotely

1. On the **Devices** tab, go to the **Bootable media** group, and then select the media that you want to use for data recovery.
2. Click **Shut down**.
3. Confirm that you want to shut down the machine booted with the media.

To view information about the bootable media

1. On the **Devices** tab, go to the **Bootable media** group, and then select the media that you want to use for data recovery.
2. Click **Details**, **Activities**, or **Alerts** to see the corresponding information.

To delete bootable media remotely

1. On the **Devices** tab, go to the **Bootable media** group, and then select the media that you want to use for data recovery.
2. Click **Delete** to delete the bootable media from the Cyber Protect console.
3. Confirm that you want to delete the bootable media.

Startup Recovery Manager

Startup Recovery Manager is a bootable component that resides on the hard drive. With Startup Recovery Manager, you can start the bootable rescue utility without using a separate bootable

media.

If a failure occurs, restart the machine, wait for the prompt **Press F11 for Acronis Startup Recovery Manager** to appear, and then press F11 or select the Startup Recovery Manager from the boot menu (if you use the GRUB boot loader). Startup Recovery Manager starts and you can perform a recovery.

Limitations

- [Not applicable to GRUB that is installed to the master boot record] Activating Startup Recovery Manager overwrites the master boot record (MBR) with its own boot code. As a result, you might need to reactivate any third-party boot loaders after the activation.
- [Not applicable to GRUB] Before activating Startup Recovery Manager in Linux, we recommend that you install the boot loader to the root partition's boot record or to the /boot partitions' boot record instead of installing it to the master boot record. Otherwise, manually reconfigure the boot loader after the activation.

Activating Startup Recovery Manager

To enable the boot-time prompt **Press F11 for Acronis Startup Recovery Manager** (or add the **Startup Recovery Manager** item to GRUB menu), you must activate Startup Recovery Manager.

Note

Activating Startup Recovery Manager on a machine with non-encrypted system volume requires at least 100 MB of free space on this machine. Recovery with restart requires additional 100 MB.

To activate Startup Recovery Manager on a machine that has a BitLocker-encrypted volume, this machine must have at least one non-encrypted volume on which there are at least 500 MB of free space. Recovery with restart requires additional 500 MB of free space.

Backup operations that create One-click recovery backups will fail if Startup Recovery Manager is not activated.

To activate Startup Recovery Manager

On a Windows or Linux machine with an agent

1. In the Cyber Protect console, select the machine on which you want to activate Startup Recovery Manager.
2. Click **Details**.
3. Enable the **Startup Recovery Manager** switch.

On a machine without an agent

1. Boot the machine by using a bootable media.
2. In the bootable media graphical interface, click **Tools > Activate Startup Recovery Manager**.
3. Select **Activate**.
4. Click **OK**.

5. On the **Details** tab, check the **Result** row to verify that the activation succeeded, and then click **Close**.

Deactivating Startup Recovery Manager

Deactivation disables the boot-time prompt **Press F11 for Acronis Startup Recovery Manager** (or removes the **Startup Recovery Manager** item from the GRUB menu).

If Startup Recovery Manager is not activated, you can still recover a machine that fails to boot by using a separate bootable media.

Note

Backup operations that create One-click recovery backups will fail if Startup Recovery Manager is not activated.

To deactivate Startup Recovery Manager

On a Windows or Linux machine with an agent

1. In the Cyber Protect console, select the machine on which you want to deactivate Startup Recovery Manager.
2. Click **Details**.
3. Disable the **Startup Recovery Manager** switch.

On a machine without an agent

1. Boot the machine by using a bootable media.
2. In the bootable media graphical interface, click **Tools > Deactivate Startup Recovery Manager**.
3. Select **Deactivate**.
4. Click **OK**.
5. On the **Details** tab, check the **Result** row to verify that the deactivation succeeded, and then click **Close**.

Implementing disaster recovery

Note

- This functionality does not support Microsoft Azure backup locations.
-

About Cyber Disaster Recovery Cloud

Cyber Disaster Recovery Cloud (DR) – a part of Cyber Protection that provides disaster recovery as a service (DRaaS). Cyber Disaster Recovery Cloud provides you with a fast and stable solution to launch the exact copies of your machines on the cloud site and switch the workload from the corrupted original machines to the recovery servers in the cloud in case of a man-made or a natural disaster.

You can set up and configure disaster recovery in the following ways:

- Create a protection plan that includes the disaster recovery module and apply it to your devices. This will automatically set up default disaster recovery infrastructure. See [Create a disaster recovery protection plan](#).
- Set up the disaster recovery cloud infrastructure manually and control each step. See "Setting up recovery servers" (p. 722).

The key functionality

Note

Some features might require additional licensing, depending on the applied licensing model.

- Manage the Cyber Disaster Recovery Cloud service from a single console
- Extend up to 23 local networks to the cloud, by using a secure VPN tunnel
- Establish the connection to the cloud site without any VPN appliance¹ deployment (the cloud-only mode)
- Establish the point-to-site connection to your local and cloud sites
- Protect your machines by using recovery servers in the cloud
- Protect applications and appliances by using primary servers in the cloud
- Perform automatic disaster recovery operations for encrypted backups
- Perform a test failover in the isolated network
- Use runbooks to spin up the production environment in the cloud

¹[Disaster Recovery] A special virtual machine that enables connection between the local network and the cloud site via a secure VPN tunnel. The VPN appliance is deployed on the local site.

Software requirements

Supported operating systems

Protection with a recovery server has been tested for the following operating systems:

- CentOS 6.6, 7.x, 8.x
- Debian 9.x, 10.x, 11.x
- Red Hat Enterprise Linux 6.6, 7.x, 8.x
- Ubuntu 16.04, 18.04, 20.x, 21.x
- Oracle Linux 7.3 and 7.9 with Unbreakable Enterprise Kernel
- Windows Server 2008 R2
- Windows Server 2012/2012 R2
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows Server 2022 – all installation options, except for Nano Server

Windows desktop operating systems are not supported due to Microsoft product terms.

The software may work with other Windows operating systems and Linux distributions, but this is not guaranteed.

Note

Protection with a recovery server has been tested for Microsoft Azure VM with the following operating systems.

- Windows Server 2008 R2
- Windows Server 2012/2012 R2
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows Server 2022 – all installation options, except for Nano Server
- Ubuntu Server 20.04 LTS - Gen2 (Canonical). For more information about accessing the recovery server console, see <https://kb.acronis.com/content/71616>.

Supported virtualization platforms

Protection of virtual machines with a recovery server has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7, 7.0
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server

- Windows Server 2022 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016
- Kernel-based Virtual Machines (KVM) — fully virtualized guests (HVM) only. Paravirtualized guests (PV) are not supported.
- Red Hat Enterprise Virtualization (RHEV) 3.6
- Red Hat Virtualization (RHV) 4.0
- Citrix XenServer: 6.5, 7.0, 7.1, 7.2

The VPN appliance has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2022 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016

The software may work with other virtualization platforms and versions, but this is not guaranteed.

Limitations

The following platforms and configurations are not supported in Cyber Disaster Recovery Cloud:

1. Unsupported platforms:

- Agents for Virtuozzo
- macOS

2. Unsupported configurations:

Microsoft Windows

- Dynamic disks are not supported
- Windows desktop operating systems are not supported (due to Microsoft product terms)
- Active Directory service with FRS replication is not supported
- Removable media without either GPT or MBR formatting (so-called "superfloppy") are not supported

Linux

- File systems without a partition table
- Linux workloads that are backed up with an agent from a guest OS and have volumes with the following advanced Logical Volume Manager (LVM) configurations: Striped volumes, Mirrored volumes, RAID 0, RAID 4, RAID 5, RAID 6, or RAID 10 volumes.

Note

Workloads with multiple operating systems installed are not supported.

3. Unsupported backup types:

- Continuous data protection (CDP) recovery points are incompatible.

Important

If you create a recovery server from a backup having a CDP recovery point, then during the failback or creating backup of a recovery server, you will lose the data contained in the CDP recovery point.

- Forensic backups cannot be used for creating recovery servers.

A recovery server has one network interface. If the original machine has several network interfaces, only one is emulated.

Cloud servers are not encrypted.

Cyber Disaster Recovery Cloud trial version

You can use a trial version of Acronis Cyber Disaster Recovery Cloud for a period of 30 days. In this case, Disaster Recovery has the following limitations for partner tenants:

- No access to public internet for recovery and primary servers. You cannot assign public IP addresses to the servers.
- IPsec Multi-site VPN is not available.

Limitations when using Geo-redundant Cloud Storage

Geo-redundant Cloud Storage provides a secondary location for your backup data. The secondary location is in a region that is geographically distinct from the primary storage location. Geographical separation of regions ensures that - if there is a disaster that affects one of the regions and makes the backup data unrecoverable - the other region will not be affected, and operations will continue.

Important

The Disaster Recovery service is not supported if the backup storage location is switched from the primary location to the geo-redundant secondary location.

Disaster Recovery compatibility with encryption software

Disaster recovery is compatible with the following disk-level encryption software:

- Microsoft BitLocker Drive Encryption
- McAfee Endpoint Encryption

- PGP Whole Disk Encryption

Note

- For workloads with disk-level encryption, we recommend that you install the protection agent in the guest operating system of the workload, and perform agent-based backups.
 - Failover and failback will not be supported for agentless backups of encrypted workloads.
-

For more information about the Cyber Protection compatibility with encryption software, see "Compatibility with encryption software" (p. 47).

Compute points

In Disaster Recovery, compute points are used for primary servers and recovery servers during test failover and production failover. Compute points reflect the compute resources used for running the servers (virtual machines) in the cloud.

The consumption of compute points during disaster recovery depends on the server's parameters, and the duration of the time period in which the server is in failover state. The more powerful the server and the longer the time period, the more compute points will be consumed. And the more compute points are consumed, the higher the price that you will be charged.

In the table below you can see eight different flavors for servers in the cloud. You can change the flavors of the servers in the **Details** tab.

Type	CPU	RAM	Compute points
F1	1 vCPU	2 GB	1
F2	1 vCPU	4 GB	2
F3	2 vCPU	8 GB	4
F4	4 vCPU	16 GB	8
F5	8 vCPU	32 GB	16
F6	16 vCPU	64 GB	32
F7	16 vCPU	128 GB	64
F8	16 vCPU	256 GB	128

Using the information in the table, you can easily estimate how much compute points a server (virtual machine) will consume.

For example, if you want to protect with Disaster Recovery one virtual machine with 4 vCPU* of 16 GB RAM, and one virtual machine with 2 vCPU with 8 GB of RAM, the first virtual machine will consume 8 compute points per hour, and the second virtual machine – 4 compute points per hour. If both virtual machines are in failover, the total consumption will be 12 compute points per hour, or 288 compute points for the whole day (12 compute points x 24 hours = 288 compute points).

*vCPU refers to a physical central processing unit (CPU) that is assigned to a virtual machine and is a time dependent entity.

Setting up the disaster recovery functionality

Note

Some features might require additional licensing, depending on the applied licensing model.

To set up the disaster recovery functionality

1. Configure the connectivity type to the cloud site:
 - [Point-to-site connection](#)
 - [Site-to-site OpenVPN connection](#)
 - [Multi-site IPsec VPN connection](#)
 - [Cloud-only mode](#)
2. Create a protection plan with the backup module enabled and select the entire machine or system plus boot volumes for backing up. At least one protection plan is required for creating a recovery server.
3. Apply the protection plan to the local servers to be protected.
4. [Create the recovery servers](#) for each of your local servers that you want to protect.
5. [Perform a test failover](#) to check how it works.
6. [Optional] [Create the primary servers](#) for application replication.

As a result, you have set up the disaster recovery functionality to protect your local servers from a disaster.

If a disaster occurs, you can [fail over the workload](#) to the recovery servers in the cloud. At least one recovery point must be created before failing over to recovery servers. When your local site is recovered from a disaster, you can switch the workload back to your local site by performing failback. For more information about the failback process, see "Prerequisites" (p. 735) and "Prerequisites" (p. 740).

Create a disaster recovery protection plan

Create a protection plan that includes the Disaster Recovery module and apply it to your devices.

By default, when creating a new protection plan, the Disaster Recovery module is disabled. After you enable the disaster recovery functionality and apply the plan to your devices, the cloud network infrastructure is created, including a *recovery server* for each protected device. The *recovery server* is a virtual machine in the cloud that is a copy of the selected device. For each of the selected devices a recovery server with default settings is created in a standby state (virtual machine not running). The recovery server is sized automatically depending on the CPU and RAM of the protected device. Default cloud network infrastructure is also created automatically: VPN gateway and networks on the cloud site, to which the recovery servers are connected.

If you revoke, delete, or switch off the Disaster Recovery module of a protection plan, the recovery servers and cloud networks are not deleted automatically. You can remove the disaster recovery infrastructure manually, if needed.

Note

- After you configure disaster recovery, you will be able to perform a test or production failover from any of the recovery points generated after the recovery server was created for the device. Recovery points that were generated before the device was protected with disaster recovery (e.g. before the recovery server was created) cannot be used for failover.
- A disaster recovery protection plan cannot be enabled if the IP address of a device cannot be detected. For example, when virtual machines are backed up agentless and are not assigned an IP address.
- When you apply a protection plan, the same networks and IP addresses are assigned in the cloud site. The IPsec VPN connectivity requires that network segments of the cloud and local sites do not overlap. If a Multi-site IPsec VPN connectivity is configured, and you apply a protection plan to one or several devices later, you must additionally update the cloud networks and reassign the IP addresses of the cloud servers. For more information, see "Reassigning IP addresses" (p. 712).

To create a disaster recovery protection plan

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the machines that you want to protect.
3. Click **Protect**, and then click **Create plan**.
The protection plan default settings open.
4. Configure the backup options.
To use the disaster recovery functionality, the plan must back up the entire machine, or only the disks, required for booting up and providing the necessary services, to a cloud storage.
5. Enable the Disaster recovery module by clicking the switch next to the module name.
6. Click **Create**.
The plan is created and applied to the selected machines.

What to do next

- You can edit the default configuration of the recovery server. For more information, see "Setting up recovery servers" (p. 722).
- You can edit the default networking configuration. For more information, see "Setting up connectivity" (p. 687).
- You can learn more about the recovery server default parameters and the cloud network infrastructure. For more information, see "Editing the Recovery server default parameters" (p. 686) and "Cloud network infrastructure" (p. 687).

Editing the Recovery server default parameters

When you create and apply a disaster recovery protection plan, a recovery server with default parameters is created. You can edit these default parameters later.

Note

A recovery server is created only if it does not exist. Existing recovery servers are not changed or recreated.

To edit the recovery server default parameters

1. Go to **Devices > All devices**.
2. Select a device, and click **Disaster recovery**.
3. Edit the recovery server default parameters.

The recovery server parameters are described in the following table.

Recovery server parameter	Default value	Description
CPU and RAM	auto	The number of virtual CPUs and the amount of RAM for the recovery server. The default settings will be automatically determined based on the original device CPU and RAM configuration.
Cloud network	auto	Cloud network to which the server will be connected. For details on how cloud networks are configured, see Cloud network infrastructure .
IP address in production network	auto	The IP address that the server will have in the production network. By default, the IP address of the original machine is set.
Test IP address	disabled	Test IP address gives you the capability to test a failover in the isolated test network and to connect to the recovery server via RDP or SSH during a test failover. In the test failover mode, the VPN gateway will replace the test IP address with the production IP address by using the NAT protocol. If a test IP address is not specified, the console will be the only way to access the server during a test failover.
Internet Access	enabled	Enable the recovery server to access the Internet during a real or test failover. By default, TCP port 25 is denied for outbound connections.
Use Public address	disabled	Having a public IP address makes the recovery

		server available from the Internet during a failover or test failover. If you do not use a public IP address, the server will be available only in your production network. To use a public IP address, you must enable internet access. The public IP address will be shown after you complete the configuration. By default, TCP port 443 is open for inbound connections.
Set RPO threshold	disabled	RPO threshold defines the maximum allowable time interval between the last recovery point and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

Cloud network infrastructure

The cloud network infrastructure consists of the VPN gateway on the cloud site and the cloud networks to which the recovery servers will be connected.

Note

Applying a disaster recovery protection plan creates recovery cloud network infrastructure only if it does not exist. Existing cloud networks are not changed or recreated.

The system checks devices IP addresses and if there are no existing cloud networks where an IP address fits, it automatically creates suitable cloud networks. If you already have existing cloud networks where the recovery servers IP addresses fit, the existing cloud networks will not be changed or recreated.

- If you do not have existing cloud networks or you setup disaster recovery configuration for the first time, the cloud networks will be created with maximum ranges recommended by IANA for private use (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16) based on your devices IP address range. You can narrow your network by editing the network mask.
- If you have devices on multiple local networks, the network on the cloud site may become a superset of the local networks. You may reconfigure networks in the **Connectivity** section. See "Managing networks" (p. 706).
- If you need to set up Site-to-site Open VPN connectivity, download the VPN appliance and set up it. See "Configuring Site-to-site Open VPN" (p. 698). Make sure your cloud networks ranges match your local network ranges connected to the VPN appliance.
- To change the default network configuration, click the **Go to connectivity** link on the Disaster Recovery module of the Protection plan, or navigate to **Disaster Recovery > Connectivity**.

Setting up connectivity

This section explains the network concepts necessary for you to understand how it all works in Cyber Disaster Recovery Cloud. You will learn how to configure different types of connectivity to the

cloud site, depending on your needs. Finally, you will learn how to manage your networks in the cloud and manage the settings of the VPN appliance and VPN gateway.

Networking concepts

Note

Some features might require additional licensing, depending on the applied licensing model.

With Cyber Disaster Recovery Cloud you can define the following connectivity types to the cloud site:

- **Cloud-only mode**

This type of connection does not require a VPN appliance deployment on the local site.

The local and cloud networks are independent networks. This type of connection implies either the failover of all the local site's protected servers or partial failover of independent servers that do not need to communicate with the local site.

Cloud servers on the cloud site are accessible through the point-to-site VPN, and public IP addresses (if assigned).

- **Site-to-site Open VPN connection**

This type of connection requires a VPN appliance deployment on the local site.

The Site-to-site Open VPN connection allows to extend your networks to the cloud and retain the IP addresses.

Your local site is connected to the cloud site by means of a secure VPN tunnel. This type of connection is suitable in case you have tightly dependent servers on the local site, such as a web server and a database server. In case of partial failover, when one of these servers is recreated on the cloud site while the other stays on the local site, they will still be able to communicate with each other via a VPN tunnel.

Cloud servers on the cloud site are accessible through the local network, point-to-site VPN, and public IP addresses (if assigned).

- **Multi-site IPsec VPN connection**

This type of connection requires a local VPN device that supports IPsec IKE v2.

When you start configuring the Multi-site IPsec VPN connection, Cyber Disaster Recovery Cloud automatically creates a cloud VPN gateway with a public IP address.

With Multi-site IPsec VPN your local sites are connected to the cloud site by means of a secure IPsec VPN tunnel.

This type of connection is suitable for Disaster Recovery scenarios when you have one or several local sites hosting critical workloads or tightly dependent services.

In case of partial failover of one of the servers, the server is recreated on the cloud site while the others remain on the local site, and they are still able to communicate with each other through an IPsec VPN tunnel.

In case of partial failover of one of the local sites, the rest of the local sites remain operational, and will still be able to communicate with each other through an IPsec VPN tunnel.

- **Point-to-site remote VPN access**

A secure Point-to-site remote VPN access to your cloud and local site workloads from outside by using your endpoint device.

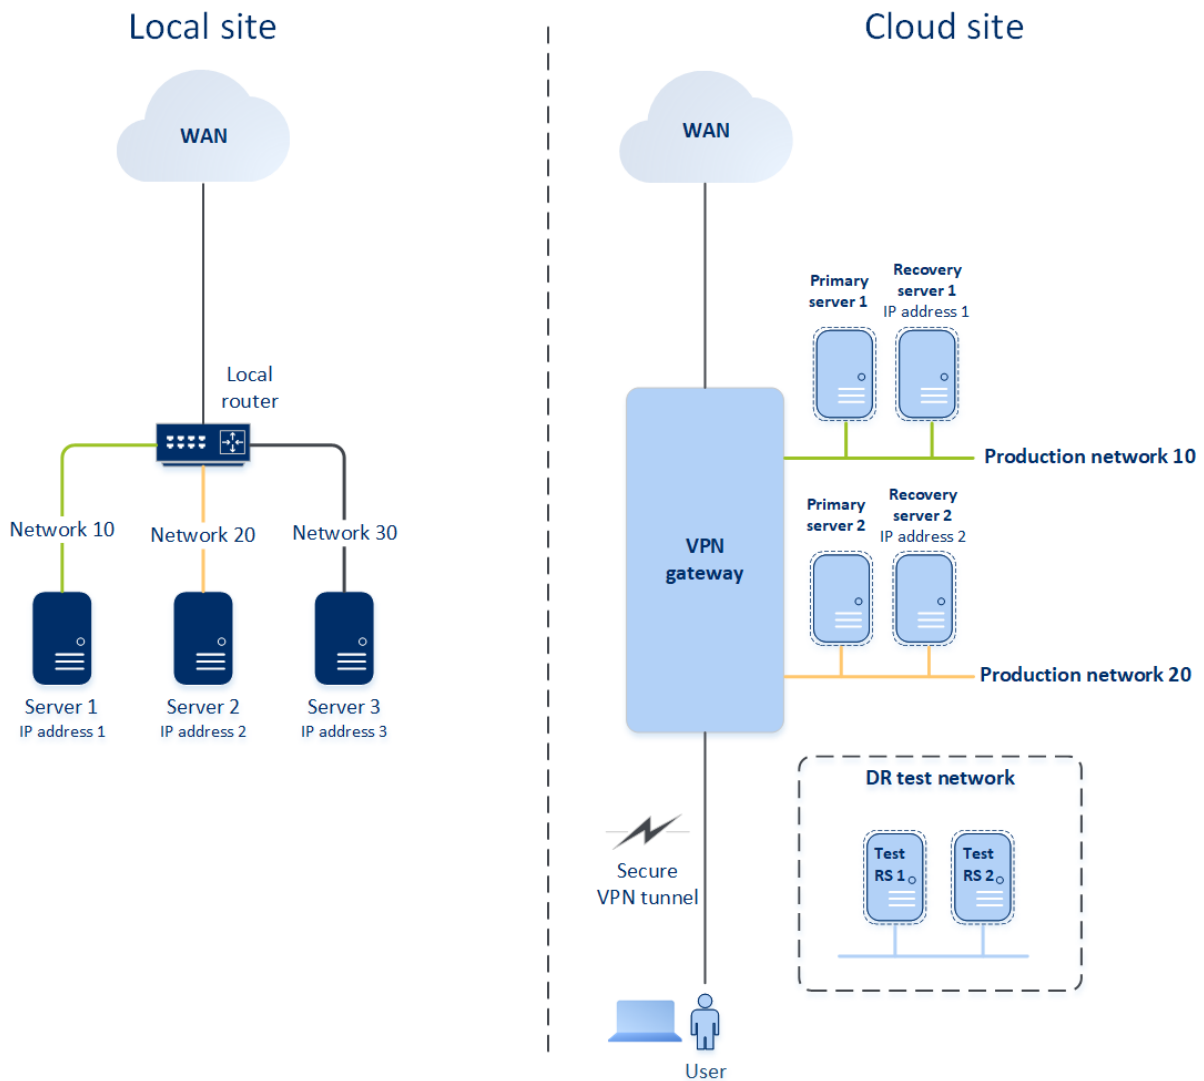
For a local site access, this type of connection requires a VPN appliance deployment on the local site.

Cloud-only mode

The cloud-only mode does not require a VPN appliance deployment on the local site. It implies that you have two independent networks: one on the local site, another on the cloud site. Routing is performed with the router on the cloud site.

How routing works

In case the cloud-only mode is established, routing is performed with the router on the cloud site so that servers from different cloud networks can communicate with each other.



Site-to-site Open VPN connection

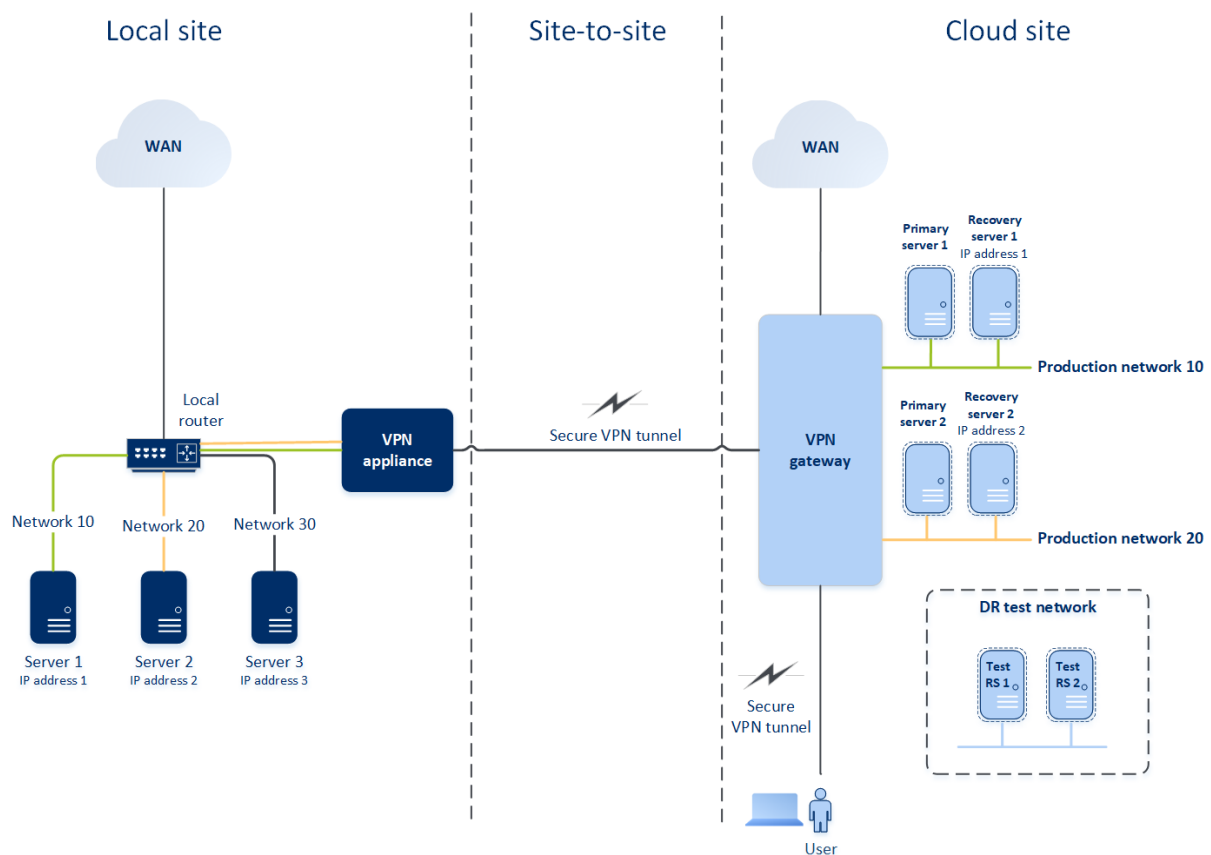
Note

The availability of this feature depends on the service quotas that are enabled for your account.

To understand how networking works in Cyber Disaster Recovery Cloud, we will consider a case when you have three networks with one machine each in the local site. You are going to configure the protection from a disaster for the two networks – Network 10 and Network 20.

On the diagram below, you can see the local site where your machines are hosted, and the cloud site where the cloud servers are launched in case of a disaster.

With the Cyber Disaster Recovery Cloud solution you can fail over all the workload from the corrupted machines in the local site to the cloud servers in the cloud. You can protect up to 23 networks with Cyber Disaster Recovery Cloud.



To establish a Site-to-site Open VPN communication between the local and cloud sites, a **VPN appliance** and a **VPN gateway** are used. When you start configuring the Site-to-site Open VPN connection in the Cyber Protect console, the VPN gateway is automatically deployed in the cloud site. Then, you must deploy the VPN appliance on your local site, add the networks to be protected, and register the appliance in the cloud. Cyber Disaster Recovery Cloud creates a replica of your local network in the cloud. A secure VPN tunnel is established between the VPN appliance and the VPN

gateway. It provides your local network extension to the cloud. The production networks in the cloud are bridged with your local networks. The local and cloud servers can communicate through this VPN tunnel as if they are all in the same Ethernet segment. Routing is performed with your local router.

For each source machine to be protected, you must create a recovery server on the cloud site. It stays in the **Standby** state until a failover event happens. If a disaster happens and you start a failover process (in the **production mode**), the recovery server representing the exact copy of your protected machine is launched in the cloud. It may be assigned the same IP address as the source machine and it can be launched in the same Ethernet segment. Your clients can continue working with the server, without noticing any background changes.

You can also start a failover process in the **test mode**. This means that the source machine is still working and at the same time the respective recovery server with the same IP address is launched in the cloud. To prevent IP address conflicts, a special virtual network is created in the cloud – **test network**. The test network is isolated to prevent duplication of the source machine IP address in one Ethernet segment. To access the recovery server in the test failover mode, when you create a recovery server, you must assign a **Test IP address** to it. There are other parameters for the recovery server that can be specified, they will be considered in the respective sections below.

How routing works

When a Site-to-site connection is established, routing between cloud networks is performed with your local router. The VPN server does not perform routing between cloud servers located in different cloud networks. If a cloud server from one network wants to communicate to a server from another cloud network, the traffic goes through the VPN tunnel to the local router on the local site, then the local router routes it to another network, and it goes back through the tunnel to the destination server on the cloud site.

VPN gateway

The major component that allows communication between the local and cloud sites is the **VPN gateway**. It is a virtual machine in the cloud on which special software is installed, and network is specifically configured. The VPN gateway has the following functions:

- Connects the Ethernet segments of your local network and production network in the cloud in the L2 mode.
- Provides iptables and ebtables rules.
- Works as a default router and NAT for the machines in the test and production networks.
- Works as a DHCP server. All machines in the production and test networks get the network configuration (IP addresses, DNS settings) via DHCP. Every time a cloud server will get the same IP address from the DHCP server. If you need to set up the custom DNS configuration, you should contact the support team.
- Works as a caching DNS.

VPN gateway network configuration

The VPN gateway has several network interfaces:

- External interface, connected to the Internet
- Production interfaces, connected to the production networks
- Test interface, connected to the test network

In addition, two virtual interfaces are added for Point-to-site and Site-to-site connections.

When the VPN gateway is deployed and initialized, the bridges are created – one for the external interface, and one for the client and production interfaces. Though the client-production bridge and the test interface use the same IP addresses, the VPN gateway can route packages correctly by using a specific technique.

VPN appliance

The **VPN appliance** is a virtual machine on the local site with Linux that has special software installed, and a special network configuration. It allows communication between the local and cloud sites.

Recovery servers

A **recovery server** – a replica of the original machine based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers in case of a disaster.

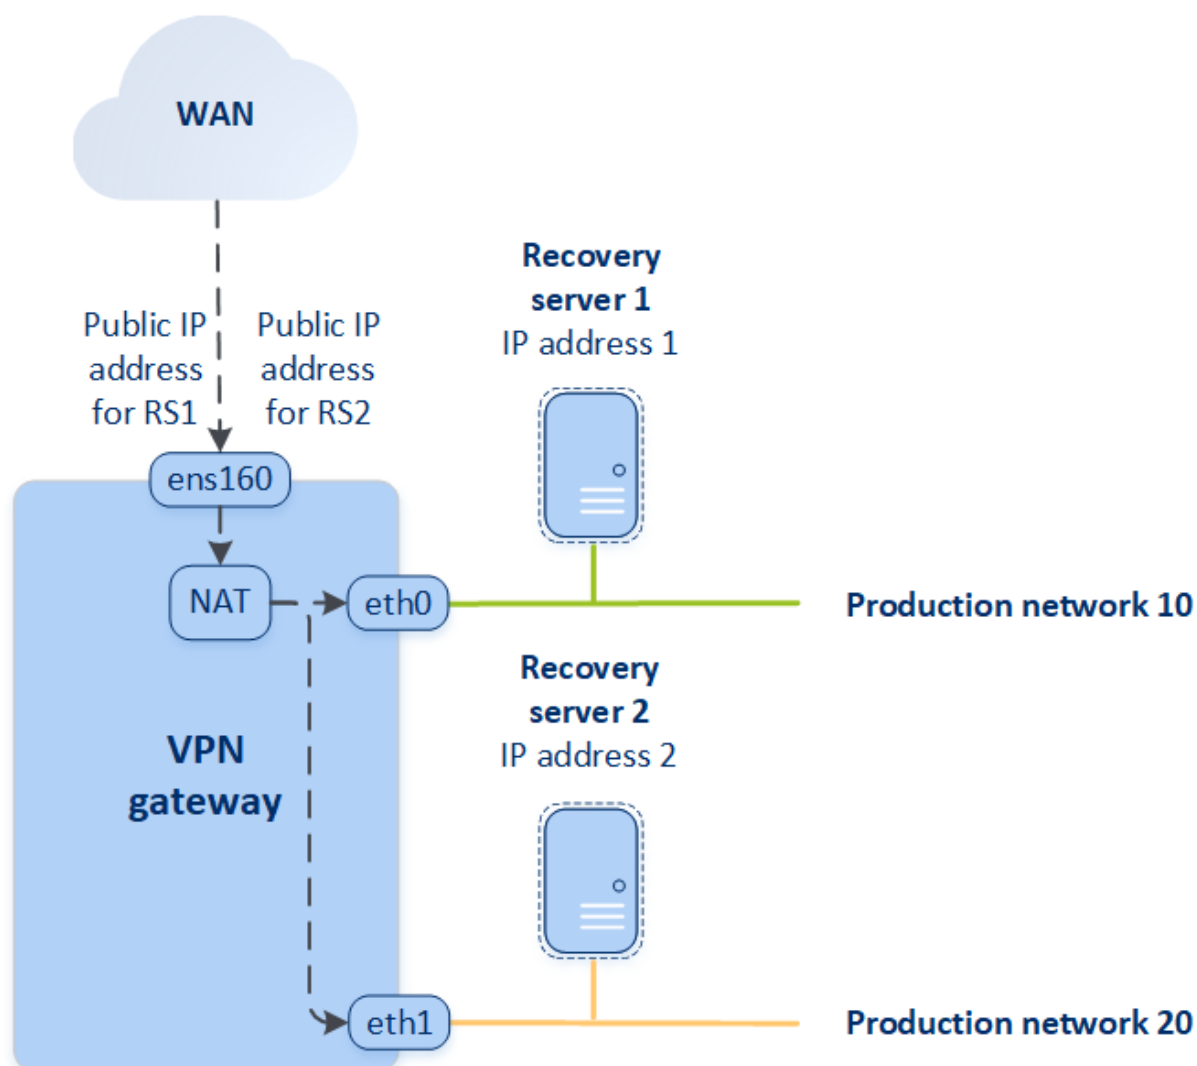
When creating a recovery server, you must specify the following network parameters:

- **Cloud network** (required): a cloud network to which a recovery server will be connected.
- **IP address in production network** (required): an IP address with which a virtual machine for a recovery server will be launched. This address is used in both the production and test networks. Before launching, the virtual machine is configured for getting the IP address via DHCP.
- **Test IP address** (optional): an IP address to access a recovery server from the client-production network during the test failover, to prevent the production IP address from being duplicated in the same network. This IP address is different from the IP address in the production network. Servers in the local site can reach the recovery server during the test failover via the test IP address, while access in the reverse direction is not available. Internet access from the recovery server in the test network is available if the **Internet access** option was selected during the recovery server creation.
- **Public IP address** (optional): an IP address to access a recovery server from the Internet. If a server has no public IP address, it can be reached only from the local network.
- **Internet access** (optional): it allows a recovery server to access the Internet (in both the production and test failover cases).

Public and test IP address


If you assign the public IP address when creating a recovery server, the recovery server becomes available from the Internet through this IP address. When a packet comes from the Internet with the destination public IP address, the VPN gateway remaps it to the respective production IP address by using NAT, and then sends it to the corresponding recovery server.

Cloud site



If you assign the test IP address when creating a recovery server, the recovery server becomes available in the test network through this IP address. When you perform the test failover, the original machine is still running while the recovery server with the same IP address is launched in the test network in the cloud. There is no IP address conflict as the test network is isolated. The recovery servers in the test network are reachable by their test IP addresses, which are remapped to the production IP addresses through NAT.

Cloud site



For more information about Site-to-site Open VPN, see "Site-to-site Open VPN - Additional information" (p. 183).

Primary servers

A **primary server** – a virtual machine that does not have a linked machine on the local site if compared to a recovery server. Primary servers are used for protecting an application by replication, or running various auxiliary services (such as a web server).

Typically, a primary server is used for real-time data replication across servers running crucial applications. You set up the replication by yourself, using the application's native tools. For example, Active Directory replication, or SQL replication, can be configured among the local servers and the primary server.

Alternatively, a primary server can be included in an AlwaysOn Availability Group (AAG) or Database Availability Group (DAG).

Both methods require a deep knowledge of the application and the administrator rights. A primary server constantly consumes computing resources and space on the fast disaster recovery storage. It needs maintenance on your side: monitoring the replication, installing software updates, and backing up. The benefits are the minimal RPO and RTO with a minimal load on the production environment (as compared to backing up entire servers to the cloud).

Primary servers are always launched only in the production network and have the following network parameters:

- **Cloud network** (required): a cloud network to which a primary server will be connected.
- **IP address in production network** (required): an IP address that the primary server will have in the production network. By default, the first free IP address from your production network is set.
- **Public IP address** (optional): an IP address to access a primary server from the Internet. If a server has no public IP address, it can be reached only from the local network, not through the Internet.
- **Internet access** (optional): allows a primary server to access the Internet.

Multi-site IPsec VPN connection

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can use the Multi-site IPsec VPN connectivity to connect a single local site, or multiple local sites to the Cyber Disaster Recovery Cloud through a secure L3 IPsec VPN connection.

This connectivity type is useful for Disaster Recovery scenarios if you have one of the following use cases:

- you have one local site hosting critical workloads.
- you have multiple local sites hosting critical workloads, for example offices in different locations.
- you use third-party software sites, or managed service providers sites and are connected to them through an IPsec VPN tunnel.

To establish a Multi-site IPsec VPN communication between the local sites and the cloud site, a **VPN gateway** is used. When you start configuring the Multi-site IPsec VPN connection in the Cyber Protect console, the VPN gateway is automatically deployed in the cloud site. You should configure the cloud network segments and make sure that they do not overlap with the local network segments. A secure VPN tunnel is established between local sites and the cloud site. The local and cloud servers can communicate through this VPN tunnel as if they are all in the same Ethernet segment.

For each source machine to be protected, you must create a recovery server on the cloud site. It stays in the **Standby** state until a failover event happens. If a disaster happens and you start a failover process (in the **production mode**), the recovery server representing the exact copy of your protected machine is launched in the cloud. Your clients can continue working with the server, without noticing any background changes.

You can also launch a failover process in the **test mode**. This means that the source machine is still working and at the same time the respective recovery server is launched in the cloud in a special virtual network that is created in the cloud – **test network**. The test network is isolated to prevent duplication of IP addresses in the other cloud network segments.

VPN gateway

The major component that allows communication between the local sites and the cloud site is the **VPN gateway**. It is a virtual machine in the cloud on which the special software is installed, and the network is specifically configured. The VPN gateway serves the following functions:

- Connects the Ethernet segments of your local network and production network in the cloud in the L3 IPsec mode.
- Works as a default router and NAT for the machines in the test and production networks.
- Works as a DHCP server. All machines in the production and test networks get the network configuration (IP addresses, DNS settings) via DHCP. Every time a cloud server will get the same IP address from the DHCP server.

If you prefer, you can set up a custom DNS configuration. For more information, see "Configuring custom DNS servers" (p. 713).

- Works as a caching DNS.

How routing works

Routing between the cloud networks is performed with the router on the cloud site so that servers from different cloud networks can communicate with each other.

Point-to-site remote VPN access

Note


The availability of this feature depends on the service quotas that are enabled for your account.

The Point-to-site connection is a secure connection from the outside by using your endpoint devices (such as computer or laptop) to the cloud and local sites through a VPN. It is available after you establish a Site-to-site Open VPN connection to the Cyber Disaster Recovery Cloud site. This type of connection is useful in the following cases:

- In many companies, the corporate services and web resources are available only from the corporate network. You can use the Point-to-site connection to securely connect to the local site.
- In case of a disaster, when a workload is switched to the cloud site and your local network is down, you may need direct access to your cloud servers. This is possible through the Point-to-site connection to the cloud site.

For the Point-to-site connection to the local site, you need to install the VPN appliance on the local site, configure the Site-to-site connection, and then the Point-to-site connection to the local site. Thus, your remote employees will have access to the corporate network through L2 VPN.

The scheme below shows the local site, cloud site, and communications between servers highlighted in green. The L2 VPN tunnel connects your local and cloud sites. When a user establishes a Point-to-site connection, the communications to the local site are performed through the cloud site.



The Point-to-site configuration uses certificates to authenticate to the VPN client. Additionally user credentials are used for authentication. Note the following about the Point-to-site connection to the local site:

- Users should use their Cyber Protect Cloud credentials to authenticate in the VPN client. They must have either a "Company Administrator" or a "Cyber Protection" user role.
- If you [re-generated the OpenVPN configuration](#), you need to provide the updated configuration to all of the users using the Point-to-site connection to the cloud site.

Automatic deletion of unused customer environments on the cloud site

The Disaster Recovery service tracks the usage of the customer environments created for disaster recovery purposes and automatically deletes them if they are unused.

The following criteria are used to define if the customer tenant is active:

- Currently, there is at least one cloud server or there were cloud server(s) in the last seven days.
OR
- The **VPN access to local site** option is enabled and either the Site-to-site Open VPN tunnel is established or there are data reported from the VPN appliance for the last 7 days.

All the rest of the tenants are considered as inactive tenants. For such tenants the system performs the following:

- Deletes the VPN gateway and all cloud resources related to the tenant.
- Unregisters the VPN appliance.

The inactive tenants are rolled back to their state before the connectivity was configured.

Initial connectivity configuration

This section describes connectivity configuration scenarios.

Configuring Cloud-only mode

To configure a connection in the cloud-only mode

1. In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
2. Select **Cloud-only** and click **Configure**.
As a result, the VPN gateway and cloud network with the defined address and mask are deployed on the cloud site.

To learn how to manage your networks in the cloud and set up the VPN gateway settings, refer to "[Managing cloud networks](#)".

Configuring Site-to-site Open VPN

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Requirements for the VPN appliance

System requirements

- 1 CPU
- 1 GB RAM
- 8 GB disk space

Ports

- TCP 443 (outbound) – for VPN connection
- TCP 80 (outbound) – for automatic [update of the appliance](#)

Ensure that your firewalls and other components of your network security system allow connections through these ports to any IP address.

Configuring a Site-to-site Open VPN connection

The VPN appliance extends your local network to the cloud through a secure VPN tunnel. This kind of connection is often referred to as a "Site-to-site" (S2S) connection. You can follow the procedure below or watch the [video tutorial](#).

To configure a connection through the VPN appliance

1. In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
2. Select **Site-to-site Open VPN connection**, and click **Configure**.

The system starts deploying the VPN gateway in the cloud. This will take some time. Meanwhile, you can proceed to the next step.

Note

The VPN gateway is provided without additional charge. It will be deleted if the Disaster Recovery functionality is not used, i.e. no primary or recovery server is present in the cloud for seven days.

3. In the **VPN appliance** block, click **Download and deploy**. Depending on the virtualization platform you are using, download the VPN appliance for VMware vSphere or Microsoft Hyper-V.
4. Deploy the appliance and connect it to the production networks.

In vSphere, ensure that **Promiscuous mode** and **Forged transmits** are enabled and set to **Accept** for all virtual switches that connect the VPN appliance to the production networks. To access these settings, in vSphere Client, select the host > **Summary** > **Network**, and then select the switch > **Edit settings...** > **Security**.

In Hyper-V, create a **Generation 1** virtual machine with 1024 MB of memory. Also, we recommend that you enable **Dynamic Memory** for the machine. Once the machine is created, go to **Settings** > **Hardware** > **Network Adapter** > **Advanced Features** and select the **Enable MAC address spoofing** check box.

5. Power on the appliance.
6. Open the appliance console and log in with the "admin"/"admin" user name and password.
7. [Optional] Change the password.
8. [Optional] Change the network settings if needed. Define which interface will be used as the WAN for Internet connection.
9. Register the appliance in the Cyber Protection service by using the credentials of the company administrator.

These credentials are only used once to retrieve the certificate. The data center URL is predefined.

Note

If two-factor authentication is configured for your account, you will also be prompted to enter the TOTP code. If two-factor authentication is enabled but not configured for your account, you cannot register the VPN appliance. First, you must go to the Cyber Protect console login page and complete the two-factor authentication configuration for your account. For more details on two-factor authentication, go to the Management Portal Administrator's Guide.

Once the configuration is complete, the appliance will have the **Online** status. The appliance connects to the VPN gateway and starts to report information about networks from all active interfaces to the Cyber Disaster Recovery Cloud service. The Cyber Protect console shows the interfaces, based on the information from the VPN appliance.

Configuring Multi-site IPsec VPN

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can configure a Multi-site IPsec VPN connection in the following two ways:

- from the **Disaster Recovery > Connectivity** tab.
- by applying a protection plan on one or several devices, and then manually switching from the automatically created Site-to-site Open VPN connection to a Multi-site IPsec VPN connection, configuring the Multi-site IPsec VPN settings, and reassigning IP addresses.

To configure a Multi-site IPsec VPN connection from the Connectivity tab

1. In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
2. In the **Multi-site VPN connection** section, click **Configure**.
A VPN gateway is deployed on the cloud site.
3. [Configure the Multi-site IPsec VPN settings](#).

To configure a Multi-site IPsec VPN connection from a protection plan

1. In the Cyber Protect console, go to **Devices**.
2. Apply a protection plan to one or multiple devices from the list.
The recovery server and the cloud infrastructure settings are automatically configured for Site-to-site Open VPN connectivity.
3. Go to **Disaster Recovery > Connectivity**.
4. Click **Show properties**.
5. Click **Switch to Multi-site IPsec VPN**.
6. [Configure the Multi-site IPsec VPN settings](#).
7. [Reassign the IP addresses](#) of the cloud network and cloud servers.

Configuring the Multi-site IPsec VPN settings

Note

The availability of this feature depends on the service quotas that are enabled for your account.

After you configure a Multi-site IPsec VPN, you must configure the cloud site and the local sites settings on the **Disaster Recovery > Connectivity** tab.

Prerequisites

- Multi-site IPsec VPN connectivity is configured. For more information about configuring the Multi-site IPsec VPN connectivity, see "Configuring Multi-site IPsec VPN" (p. 700).
- Each local IPsec VPN gateway has a public IP address.

- Your cloud network has enough IP addresses for the cloud servers that are copies of your protected machines (in the production network), and for the recovery servers (with one or two IP addresses, depending on your needs).
- [If you use a firewall between the local sites and the cloud site] The following IP protocols and UDP ports are allowed on the local sites: IP Protocol ID 50 (ESP), UDP Port 500 (IKE), and UDP Port 4500.
- The NAT-T configuration on the local sites is disabled.

To configure a Multi-site IPsec VPN connection

1. Add one or more networks to the cloud site.
 - a. Click **Add Network**.

Note

When you add a cloud network, a corresponding test network is added automatically with the same network address and mask for performing test failovers. The cloud servers in the test network have the same IP addresses as the ones in the cloud production network. If you need to access a cloud server from the production network during a test failover, when you create a recovery server, assign it a second test IP address.

- b. In the **Network address** field, type the IP address of the network.
 - c. In the **Network mask** field, type the mask of the network.
 - d. Click **Add**.
2. Configure the settings for each local site that you want to connect to the cloud site, following the recommendations for the local sites. For more information about these recommendations, see "General recommendations for local sites" (p. 702).
 - a. Click **Add Connection**.
 - b. Enter a name for the of the local VPN gateway.
 - c. Enter the public IP address of the local VPN gateway.
 - d. [Optional] Enter a description of the local VPN gateway.
 - e. Click **Next**.
 - f. In the **Pre-shared key** field, type the pre-shared key, or click **Generate a new pre-shared key** to use an automatically generated value.

Note

You must use the same pre-shared key for the local and the cloud VPN gateways.

- g. Click **IPsec/IKE security settings** to configure the settings. For more information about the settings that you can configure, see "IPsec/IKE security settings" (p. 702).

Note

You can use the default settings, which are populated automatically, or use custom values. Only IKEv2 protocol connections are supported. The default **Startup action** when establishing the VPN is **Add** (your local VPN gateway initiates the connection), but you can change it to **Start** (the cloud VPN gateway initiates the connection) or **Route** (suitable for firewalls that support the route options).

h. Configure the **Network policies**.

The network policies specify the networks to which the IPsec VPN connects. Type the IP address and mask of the network using the CIDR format. The local and cloud network segments should not overlap.

i. Click **Save**.

General recommendations for local sites

Note

The availability of this feature depends on the service quotas that are enabled for your account.

When you configure the local sites for your Multi-site IPsec VPN connectivity, consider the following recommendations:

- For each IKE Phase, set at least one of the values that are configured in the cloud site for the following parameters: Encryption algorithm, Hash algorithm, and Diffie-Hellman group numbers.
- Enable Perfect forward secrecy with at least one of the values for Diffie-Hellman group numbers that is configured in the cloud site for IKE Phase 2.
- Configure the same **Lifetime** value for IKE Phase 1 and IKE Phase 2 as in the cloud site.
- Configurations with NAT traversal (NAT-T) are not supported. Disable the NAT-T configuration on the local site. Otherwise, the additional UDP encapsulation cannot be negotiated.
- The **Startup action** configuration defines which side initiates the connection. The default value **Add** means that the local site initiates the connection, and cloud site is waiting for the connection initiation. Change the value to **Start** if you want the cloud site to initiate the connection, or to **Route** if you want both sides to be able to initiate the connection (suitable for firewalls that support the route option).

For more information and configuration examples for different solutions, see:

- [This series of knowledge base articles](#)
- [This video example](#)

IPsec/IKE security settings

Note

The availability of this feature depends on the service quotas that are enabled for your account.

The following table provides more information about the Psec/IKE security parameters.

Parameter	Description
Encryption algorithm	The encryption algorithm that will be used to ensure that data is not viewable while in transit. By default, all algorithms are selected. You must configure at least one of the selected algorithms on your local gateway device for each IKE phase.
Hash algorithm	The hash algorithm that will be used to verify the data integrity and authenticity. By default, all algorithms are selected. You must configure at least one of the selected algorithms on your local gateway device for each IKE phase.
Diffie-Hellman group numbers	<p>The Diffie-Hellman group numbers define the strength of the key that is used in the Internet Key Exchange (IKE) process.</p> <p>Higher group numbers are more secure but require additional time for the key to compute.</p> <p>By default, all groups are selected. You must configure at least one of the selected groups on your local gateway device for each IKE phase.</p>
Lifetime (seconds)	<p>The lifetime value determines the duration of a connection instance with a set of encryption/authentication keys for user packets, from successful negotiation to expiry.</p> <p>Range for Phase 1: 900-28800 seconds with default 28800.</p> <p>Range for Phase 2: 900-3600 seconds with default 3600.</p> <p>The lifetime for Phase 2 must be less than the lifetime for Phase 1.</p> <p>The connection is re-negotiated through the keying channel before it expires, see Rekey margin time. If the local and the remote side do not agree on the lifetime, a clutter of superseded connections will occur on the side with the longer lifetime. See also Rekey margin time and Rekey fuzz.</p>
Rekey margin time (seconds)	The margin time before connection expiration or keying-channel expiration, during which the local side of the VPN connection attempts to negotiate a replacement. The exact time of the rekey is randomly selected based on the value of Rekey fuzz . Relevant only locally, the remote side does

Parameter	Description
	not need to agree on it. Range: 900-3600 seconds. The default value is 3600.
Replay window size (packet)	<p>The IPsec replay window size for this connection.</p> <p>The default -1 uses the value configured with charon.replay_window in the strongswan.conf file.</p> <p>Values larger than 32 are supported only when using the Netlink backend.</p> <p>A value of 0 disables the IPsec replay protection.</p>
Rekey fuzz (%)	<p>The maximum percentage by which marginbytes, marginpackets and margintime are randomly increased to randomize rekeying intervals (important for hosts with many connections).</p> <p>The Rekey fuzz value can exceed 100%. The value of marginTYPE, after the random increase, must not exceed lifeTYPE, where TYPE is one of bytes, packets or time.</p> <p>The value 0% disables randomization. Relevant only locally, the remote side does not need to agree on it.</p>
DPD timeout (seconds)	Time after which a dead peer detection (DPD) timeout occurs. You can specify value 30 or higher. The default value is 30.
Dead peer detection (DPD) timeout action	<p>The action to take after a dead peer detection (DPD) timeout occurs.</p> <p>Restart - Restart the session when DPD timeout occurs.</p> <p>Clear - End the session when DPD timeout occurs.</p> <p>None - Take no action when DPD timeout occurs.</p>
Startup action	<p>Determines which side initiates the connection and establishes the tunnel for the VPN connection.</p> <p>Add - your local VPN gateway initiates the connection.</p> <p>Start - the cloud VPN gateway initiates the connection.</p> <p>Route - suitable for VPN gateways that support the route option. The tunnel is up only when there is traffic initiated from either the local VPN gateway,</p>

Parameter	Description
	or the cloud VPN gateway.

Recommendations for the Active Directory Domain Services availability

If your protected workloads need to authenticate in a domain controller, we recommend that you have an Active Directory Domain Controller (AD DC) instance at the Disaster Recovery site.

Active Directory Domain Controller for L2 Open VPN connectivity

With the L2 Open VPN connectivity, the IP addresses of the protected workloads are retained in the cloud site during a test failover or a production failover. Therefore, the AD DC during a test failover or a production failover has the same IP address as in the local site.

With custom DNS you can set your own custom DNS server for all cloud servers. For more information, see "Configuring custom DNS servers" (p. 713).

Active Directory Domain Controller for L3 IPsec VPN connectivity

With L3 IPsec VPN connectivity, the IP addresses of the protected workloads are not retained in the cloud site. Therefore, we recommend that you have an additional dedicated AD DC instance as a primary server in the cloud site before you perform a production failover.

The recommendations for a dedicated AD DC instance that is configured as a primary server in the cloud site are the following:

- Turn off Windows firewall.
- Join the primary server to the Active Directory service.
- Ensure that the primary server has Internet access.
- Add the Active Directory feature.

With custom DNS you can set your own custom DNS server for all cloud servers. For more information, see "Configuring custom DNS servers" (p. 713).

Configuring Point-to-site remote VPN access

Note

The availability of this feature depends on the service quotas that are enabled for your account.

If you need to connect to your local site remotely, you can configure the Point-to-site connection to the local site. You can follow the procedure below or watch the [video tutorial](#).

Prerequisites

- A Site-to-site Open VPN connectivity is configured.
- The VPN appliance is installed on the local site.

To configure the Point-to-site connection to the local site

1. In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**.
3. Enable the **VPN access to local site** option.
4. Ensure that your user who needs to establish the Point-to-site connection to the local site has:
 - a user account in Cyber Protect Cloud. These credentials are used for authentication in the VPN client. Otherwise, [create a user account in Cyber Protect Cloud](#).
 - a "Company Administrator" or "Cyber Protection" user role.
5. Configure the OpenVPN client:
 - a. Download the OpenVPN client version 2.4.0 or later from the following location <https://openvpn.net/community-downloads/>.
 - b. Install the OpenVPN client on the machine from which you want to connect to the local site.
 - c. Click **Download configuration for OpenVPN**. The configuration file is valid for users in your organization with the "Company Administrator" or "Cyber Protection" user role.
 - d. Import the downloaded configuration to OpenVPN.
 - e. Log in to the OpenVPN client with your Cyber Protect Cloud user credentials (see step 4 above).
 - f. [Optional] If two-factor authentication is enabled for your organization, then you should provide the [one-time generated TOTP code](#).

Important

If you enabled two-factor authentication for your account, you need to re-generate the configuration file and renew it for your existing OpenVPN clients. Users must re-log in to Cyber Protect Cloud to set up two-factor authentication for their accounts.

As a result, your user will be able to connect to machines on the local site.

Network management

This section describes network management scenarios.

Managing networks

Note

Some features might require additional licensing, depending on the applied licensing model.

Site-to-site Open VPN connection

To add a network on the local site and extend it to the cloud

1. On the VPN appliance, set up the new network interface with the local network that you want to extend in the cloud.
2. Log in to the VPN appliance console.
3. In the **Networking** section, set up network settings for the new interface.

```

Disaster Recovery VPN Appliance
Registered by:
9.0.1.234
[dagmy@mailinator.com]

[Appliance Status]
DHCP: Enabled
VPN tunnel: Connected
VPN Service: Started
WAN interface: ens160
Internet: Available
Gateway: Available

[WAN interface Settings]
IP address: 172.16.1.110
Network mask: 255.255.255.0
Default gateway: 172.16.1.1
Preferred DNS server: 172.16.1.1
Alternate DNS server:
MAC address: 00:50:56:91:90:66

Commands:
Register
Networking
Change password
Restart the VPN service
Run Linux shell command
Reboot

```

The VPN appliance starts to report information about networks from all active interfaces to Cyber Disaster Recovery Cloud. The Cyber Protect console shows the interfaces based on the information from the VPN appliance.

To delete a network extended to the cloud

1. Log in to the VPN appliance console.
2. In the **Networking** section, select the interface that you want to delete, and then click **Clear network settings**.
3. Confirm the operation.

As a result, the local network extension to the cloud via a secure VPN tunnel will be stopped. This network will operate as an independent cloud segment. If this interface is used to pass the traffic from (to) the cloud site, all of your network connections from (to) the cloud site will be disconnected.

To change the network parameters

1. Log in to the VPN appliance console.
2. In the **Networking** section, select the interface that you want to edit.
3. Click **Edit network settings**.
4. Select one of the two possible options:
 - For automatic network configuration via DHCP, click **Use DHCP**. Confirm the operation.
 - For manual network configuration, click **Set static IP address**. The following settings are available for editing:
 - **IP address**: the IP address of the interface in the local network.
 - **VPN gateway IP address**: the special IP address which is reserved for the cloud segment of network for the proper Cyber Disaster Recovery Cloud service work.
 - **Network mask**: network mask of the local network.
 - **Default gateway**: default gateway on the local site.
 - **Preferred DNS server**: primary DNS server on the local site.
 - **Alternate DNS server**: secondary DNS server on the local site.

```
Disaster Recovery VPN Appliance
Registered by:                               9.0.1.234
                                              [ldagny@mailinator.com]

Command: Networking \ configure ens160

Usage:
<Up>, <Down> - to select parameter
<Esc> - to cancel the command

IP address:
VPN gateway IP address:
Network mask:
Default gateway:
Preferred DNS server:
Alternate DNS server:
```

- Make the necessary changes and confirm them by pressing Enter.

Cloud-only mode

You can have up to 23 networks in the cloud.

To add a new cloud network

1. Go to **Disaster Recovery > Connectivity**.
2. On **Cloud site**, click **Add cloud network**.
3. Define the cloud network parameters: the network address and mask. When ready, click **Done**.

As a result, the additional cloud network with the defined address and mask will be created on the cloud site.

To delete a cloud network

Note

You cannot delete a cloud network if there is at least one cloud server in it. First, delete the cloud server, and then delete the network.

1. Go to **Disaster Recovery > Connectivity**.
2. On **Cloud site**, click the network address that you want to delete.
3. Click **Delete** and confirm the operation.

To change cloud network parameters

1. Go to **Disaster Recovery > Connectivity**.
2. On **Cloud site**, click the network address that you want to edit.
3. Click **Edit**.
4. Define the network address and mask, and click **Done**.

IP address reconfiguration

For proper disaster recovery performance, the IP addresses assigned to the local and cloud servers must be consistent. If there is any inconsistency or mismatch in IP addresses, you will see the exclamation mark next to the corresponding network in **Disaster Recovery > Connectivity**.

Some of the commonly known reasons of IP address inconsistency are listed below:

1. A recovery server was migrated from one network to another or the network mask of the cloud network was changed. As a result, cloud servers have the IP addresses from networks to which they are not connected.
2. The connectivity type was switched from one without Site-to-site connection to a Site-to-site connection. As a result, a local server is placed in the network different from the one that was created for the recovery server on the cloud site.
3. The connectivity type was switched from Site-to-site Open VPN to Multi-site IPsec VPN, or from Multi-site IPsec VPN to Site-to-site Open VPN. For more information about this scenario, see [Switching connections](#) and [Reassigning IP addresses](#).
4. Editing the following network parameters on the VPN appliance site:
 - Adding an interface via the network settings
 - Editing the network mask manually via the interface settings
 - Editing the network mask via DHCP
 - Editing the network address and mask manually via the interface settings
 - Editing the network mask and address via DHCP

As a result of the actions listed above, the network on the cloud site may become a subset or superset of the local network, or the VPN appliance interface may report the same network settings for different interfaces.

To resolve the issue with network settings

1. Click the network that requires IP address reconfiguration.
You will see a list of servers in the selected network, their status, and IP addresses. The servers whose network settings are inconsistent are marked with the exclamation mark.
2. To change network settings for a server, click **Go to server**. To change network settings for all servers at once, click **Change** in the notification block.
3. Change the IP addresses as needed by defining them in the **New IP** and **New test IP** fields.
4. When ready, click **Confirm**.

Move servers to a suitable network

When you create a disaster recovery protection plan and apply it on selected devices, the system checks devices IP addresses and automatically creates cloud networks if there are not existing cloud networks where IP address fits. By default, the cloud networks are configured with maximum ranges recommended by IANA for private use (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16). You can narrow your network by editing the network mask.

In case if the selected devices was on the multiple local networks, the network on the cloud site may become a superset of the local networks. In this case, to reconfigure cloud networks:

1. Click the cloud network that requires network size reconfiguration and then click **Edit**.
2. Reconfigure the network size with the correct settings.
3. Create other required networks.
4. Click the notification icon next to the number of devices connected to the network.

5. Click **Move to a suitable network**.
6. Select the servers that you want to move to suitable networks and then click **Move**.

Managing the VPN appliance settings

Note

The availability of this feature depends on the service quotas that are enabled for your account.

In the Cyber Protect console (**Disaster Recovery** > **Connectivity**), you can:

- Download log files.
- Unregister the appliance (if you need to reset the VPN appliance settings or switch to the cloud-only mode).

To access these settings, click the **i** icon in the **VPN appliance** block.

In the VPN appliance console, you can:

- Change the password for the appliance.
- View/change the network settings and define which interface to use as the WAN for the Internet connection.
- Register/change the registration account (by repeating the registration).
- Restart the VPN service.
- Reboot the VPN appliance.
- Run the Linux shell command (only for advanced troubleshooting cases).

Reinstalling the VPN gateway

If there is an issue with the VPN gateway which you cannot resolve, you might want to reinstall the VPN gateway. Possible issues include the following:

- The VPN gateway is in **Error** status.
- The VPN gateway is in **Pending** status for a long time.
- The VPN gateway status is undetermined for a long time.

Reinstalling the VPN gateway process includes the following automatic actions: deleting the existing VPN gateway virtual machine completely, installing a new virtual machine from the template, and applying the settings of the previous VPN gateway on the new virtual machine.

Prerequisites:

One of the connectivity types to the cloud site must be set.

To reinstall the VPN gateway

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click the gear icon of the VPN gateway, and select **Reinstall VPN gateway**.

3. In the **Reinstall VPN gateway** dialog, enter your login.
4. Click **Reinstall**.

Enabling and disabling the Site-to-site connection

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can enable the Site-to-site connection in the following cases:

- If you need the cloud servers on the cloud site to communicate with servers on the local site.
- After a failover to the cloud, the local infrastructure is recovered, and you want to fail back your servers to the local site.

To enable the site-to-site connection

1. Go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**, and then enable the **Site-to-site connection** option.

As a result, the site-to-site VPN connection is enabled between the local and cloud sites. The Cyber Disaster Recovery Cloud service gets the network settings from the VPN appliance and extends the local networks to the cloud site.

If you do not need cloud servers on the cloud site to communicate with servers on the local site, you can disable the Site-to-site connection.

To disable the site-to-site connection

1. Go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**, and then disable the **Site-to-site connection** option.

As a result, the local site is disconnected from the cloud site.

Switching the Site-to-site connection type

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can easily switch from a Site-to-site Open VPN connection to a Multi-site IPsec VPN connection, and from a Multi-site IPsec VPN connection to a Site-to-site Open VPN connection.

When you switch the connectivity type, the active VPN connections are deleted, but the cloud servers and network configurations are preserved. However, you will still need to reassign the IP addresses of the cloud networks and servers.

The following table compares the basic characteristics of the Site-to-site Open VPN connection and the Multi-site IPsec VPN connection.

	Site-to-site Open VPN	Multi-site IPsec VPN
Local site support	Single	Single, Multiple
VPN Gateway mode	L2 Open VPN	L3 IPsec VPN
Network segments	Extends the local network to the cloud network	Local networks and cloud network segments should not overlap
Supports Point-to-Site access to local site	Yes	No
Supports Point-to-Site access to cloud site	Yes	Yes
Requires a public IP offering item	No	Yes

To switch from a Site-to-site Open VPN connection to a Multi-site IPsec VPN connection

1. In the Cyber Protect console, go to **Disaster Recovery** -> **Connectivity**.
2. Click **Show properties**.
3. Click **Switch to multi-site IPsec VPN**.
4. Click **Reconfigure**.
5. [Reassign the IP addresses](#) of the cloud network and cloud servers.
6. [Configure the Multi-site IPsec connection settings](#).

To switch from a Multi-site IPsec VPN connection to a Site-to-site Open VPN connection

1. In the Cyber Protect console, go to **Disaster Recovery** -> **Connectivity**.
2. Click **Show properties**.
3. Click **Switch to site-to-site Open VPN**.
4. Click **Reconfigure**.
5. [Reassign the IP addresses](#) of the cloud network and cloud servers.
6. [Configure the Site-to-site connection settings](#).

Reassigning IP addresses

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You must reassign the IP addresses of the cloud networks and the cloud servers in order to complete the configuration in the following cases:

- After you switch from Site-to-site Open VPN to Multi-site IPsec VPN, or the opposite.
- After you apply a protection plan (if the Multi-site IPsec VPN connectivity is configured).

To reassign the IP address of a cloud network

1. In the **Connectivity** tab, click the IP address of the cloud network.
2. In the **Network** pop-up, click **Edit**.
3. Type the new the network address and network mask.
4. Click **Done**.

After you reassign the IP address of a cloud network, you must reassign the cloud servers that belong to the reassigned cloud network.

To reassign the IP address of a server

1. In the **Connectivity** tab, click the IP address of the server in the cloud network.
2. In the **Servers** pop-up, click **Change IP address**.
3. In the **Change IP address** pop-up, type the new IP address of the server, or use the automatically generated IP address which is part of the reassigned cloud network.

Note

Cyber Disaster Recovery Cloud automatically assigns IP addresses from the cloud network to all cloud servers that were part of the cloud network before the reassignment of the network IP address. You can use the suggested IP addresses to reassign the IP addresses of all the cloud servers at once.

4. Click **Confirm**.

Configuring custom DNS servers

Note

The availability of this feature depends on the service quotas that are enabled for your account.

When you configure a connectivity, Cyber Disaster Recovery Cloud creates your cloud network infrastructure. The cloud DHCP server automatically assigns default DNS servers to the recovery servers and primary servers, but you can change the default settings and configure custom DNS servers. The new DNS settings will be applied at the time of the next request to the DHCP server.

Prerequisites:

One of the connectivity types to the cloud site must be set.

To configure a custom DNS server

1. In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**.
3. Click **Default (Provided by Cloud Site)**.
4. Select **Custom servers**.
5. Type the IP address of the DNS server.

- [Optional] If you want to add another DNS server, click **Add**, and type the DNS server IP address.

Note

After you add the custom DNS servers, you can also add the default DNS servers. In that way, if the custom DNS servers are unavailable, Cyber Disaster Recovery Cloud will use the default DNS servers.

- Click **Done**.

Deleting custom DNS servers

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can delete DNS servers from the custom DNS list.

Prerequisites:

Custom DNS servers are configured.

To delete a custom DNS server

- In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
- Click **Show properties**.
- Click **Custom servers**.
- Click the delete icon next to the DNS server.

Note

The delete operation is disabled when only one custom DNS server is available. If you want to delete all custom DNS servers, select **Default (provided by Cloud Site)**.

- Click **Done**.

Configuring local routing

In addition to your local networks that are extended to the cloud through the VPN appliance, you may have other local networks that are not registered in the VPN appliance but have servers which need to communicate with cloud servers. To establish the connectivity between such local servers and cloud servers, you need to configure the local routing settings.

To configure local routing

- Go to **Disaster Recovery > Connectivity**.
- Click **Show properties**, and then click **Local routing**.

3. Specify the local networks in the CIDR notation.
4. Click **Save**.

As a result, the servers from the specified local networks can communicate with the cloud servers.

Allowing DHCP traffic over L2 VPN

If devices on your local site get their IP address from a DHCP server, you can protect the DHCP server with Disaster Recovery, fail it over to the cloud, and then allow the DHCP traffic to run over L2 VPN. Thus, your DHCP server will be running in the cloud, but will continue assigning IP addresses to your local devices.

Prerequisites:

A Site-to-site L2 VPN connectivity type to the cloud site must be set.

To allow the DHCP traffic via the L2 VPN connection

1. Go to **Disaster Recovery > Connectivity** tab.
2. Click **Show Properties**.
3. Enable the **Allow DHCP traffic via L2 VPN** switch.

Managing point-to-site connection settings

Note

The availability of this feature depends on the service quotas that are enabled for your account.

In the Cyber Protect console, go to **Disaster Recovery > Connectivity** and then click **Show properties** in the upper right corner.

The screenshot displays the Acronis Cyber Protect Cloud console interface. On the left is a dark blue navigation sidebar with the following menu items: Manage account, DASHBOARD, DEVICES, PLANS, DISASTER RECOVERY, Servers, Connectivity (highlighted in blue), Runbooks, ANTI-MALWARE PROTECTION, and SOFTWARE MANAGEMENT. The main content area is titled 'Connectivity' and shows a diagram of a VPN tunnel connecting a 'Local site' and a 'Cloud site'. The Local site contains an 'Appliance' with status 'Online' and IP address '172.16.1.110'. The Cloud site contains a 'VPN gateway' with status 'Online' and 'Internet access: Enabled'. A 'Point-to-site' connection is shown between them, with a status of '172.16.1.0/24'. A 'Properties' panel is open on the right, showing 'Site-to-site' connection settings. The 'Site-to-site' section has a toggle for 'Site-to-site connection' which is turned on. Below it are options for 'Download VPN appliance' and 'Local routing'. The 'Point-to-site' section has a toggle for 'VPN access to local site' which is also turned on. Other options in this section include 'Re-generate configuration file', 'Download configuration for OpenVPN', and 'How to connect?'. A button labeled 'Add cloud network' is visible at the bottom of the VPN gateway card.

VPN access to local site

This option is used for managing VPN access to the local site. By default it is enabled. If it is disabled, then the Point-to-site access to the local site will be not allowed.

Download configuration for OpenVPN

This will download the configuration file for the OpenVPN client. The file is required to establish a Point-to-site connection to the cloud site.

Re-generate configuration

You can re-generate the configuration file for the OpenVPN client.

This is required in the following cases:

- If you suspect that the configuration file is compromised.
- If two-factor authentication was enabled for your account.

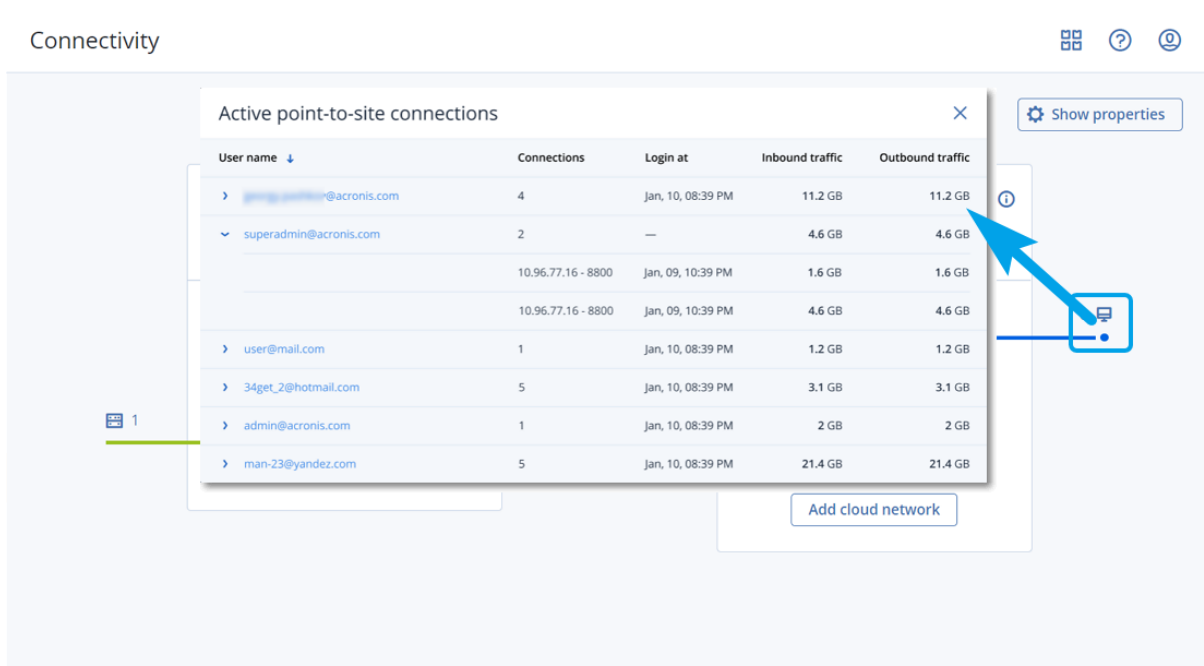
As soon as the configuration file is updated, connecting by means of the old configuration file becomes not possible. Make sure to distribute the new file among the users who are allowed to use the Point-to-site connection.

Active point-to-site connections

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can view all active point-to-site connections in **Disaster recovery > Connectivity**. Click the machine icon on the blue **Point-to-site** line and you will see the detailed information about active point-to-site connections grouped by the user name.



Working with logs

Disaster Recovery collects logs for the VPN appliance and the VPN gateway. The logs are saved as .txt files, which are compressed in a .zip archive. You can download and extract the archive, and use the information for troubleshooting or monitoring purposes.

The following list describes the log files that are part of the .zip archive, and the information that they contain.

dnsmasq.config.txt - The file contains information about the configuration of the service that provides DNS and DHCP addresses.

dnsmasq.leases.txt - The file contains information about the current DHCP address leases.

dnsmasq_log.txt - The file contains logs of the dnsmasq service.

eatables.txt - The file contains information about the firewall tables.

free.txt - The file contains information about the free memory.

ip.txt - The file contains the logs from the configuration of the network interfaces, including their names which can be used in the configuration of the **Capturing network packets** settings.

NetworkManager_log.txt - The file contains logs from the NetworkManager service.

NetworkManager_status.txt - The file contains information about the status of the NetworkManager service.

openvpn@p2s_log.txt - The file contains logs from the OpenVPN service.

openvpn@p2s_status.txt - The file contains information about the status of the VPN tunnels.

ps.txt - The file contains information about the currently running processes on the VPN gateway or VPN appliance.

resolv.conf.txt - The file contains information about the configuration of the DNS servers.

routes.txt - The file contains information about the networking routes.

uname.txt - The file contains information about the current version of the kernel of the operating system.

uptime.txt - The file contains information about the length of period for which the operating system has not been restarted.

vpnserver_log.txt - The file contains logs from the VPN service.

vpnserver_status.txt - The file contains information about the status of the VPN server.

For more information about log files that are specific to the IPsec VPN connectivity, see "Multi-site IPsec VPN log files" (p. 722).

Downloading the logs of the VPN appliance

You can download and extract the archive that contains the logs of the VPN appliance, and use the information for troubleshooting or monitoring purposes.

To download the logs of the VPN appliance

1. On the **Connectivity** page, click the gear icon next to the VPN appliance.
2. Click the **Download log**.
3. [Optional] Select **Capture network packets**, and configure the settings. For more information, see "Capturing network packets" (p. 719).
4. Click **Done**.
5. When the .zip archive is ready for download, click **Download log**, and save it locally.

Downloading the logs of the VPN gateway

You can download and extract the archive that contains the logs of the VPN gateway, and use the information for troubleshooting or monitoring purposes.

To download the logs of the VPN gateway

1. On the **Connectivity** page, click the gear icon next to the VPN gateway.
2. Click the **Download log**.
3. [Optional] Select **Capture network packets**, and then configure the settings. For more information, see "Capturing network packets" (p. 719).
4. Click **Done**.
5. When the .zip archive is ready for download, click **Download log**, and save it locally.

Capturing network packets

To troubleshoot and analyze the communication between the local production site and a primary or recovery server, you can choose to collect network packets on the VPN gateway or VPN appliance.

After collecting 32000 network packets, or reaching time limit, capturing network packets stops, and the results are written in a .libpcap file that is added to the logs .zip archive.

The following table provides more information about the **Capture network packets** settings that you can configure.

Setting	Description
Network interface name	The network interface on which to capture network packets. If you want to capture network packets on all network interfaces, select Any .
Time limit (seconds)	The time limit for capturing network packets. The maximum value you can set is 1800.
Filtering	<p>An extra filter to apply on the captured network packets.</p> <p>You can enter a string containing protocols, ports, directions, and their combinations, separated by space, such as: "and", "or", "not", "(", ")", "src", "dst", "net", "host", "port", "ip", "tcp", "udp", "icmp", "arp", "esp".</p> <p>If you want to use brackets, surround them with spaces. You can also enter IP addresses and network addresses, for example: "icmp or arp" and "port 67 or 68".</p> <p>For more information about the values that you can enter, see the Linux tcpdump help.</p>

Troubleshooting the IPsec VPN configuration

Note

The availability of this feature depends on the service quotas that are enabled for your account.

When you configure or use the IPsec VPN connection, you might experience problems.

You can learn more about the problems that you encountered in the IPsec log files, and check the Troubleshooting IPsec VPN configuration issues topic for possible solutions of some of the common problems that might occur.

Troubleshooting IPsec VPN configuration issues

Note

The availability of this feature depends on the service quotas that are enabled for your account.

The following table describes the IPsec VPN configuration problems that occur most often, and explains how to troubleshoot them.

Problem	Possible solution
<p>I see the following error message: IKE phase 1 negotiation error. Check the IPsec IKE settings on the Cloud and the Local sites.</p>	<p>Click Retry and check if a more specific error message appears. For example, a more specific error message may be an error message about an algorithm mismatch or an incorrect Pre-shared key.</p> <hr/> <p>Note For security reasons, the following restrictions apply to the IPsec VPN connectivity:</p> <ul style="list-style-type: none"> • IKEv1 is called for deprecation in RFC8247 and is not supported due to security risks. Only IKEv2 protocol connections are supported. • The following Encryption algorithms are not considered secure and are not supported: DES, and 3DES. • The following Hash algorithms are not considered secure and are not supported: SHA1, and MD5. • Diffie-Hellman group number 2 is not considered secure and is not supported.
<p>The connection between my local site and the cloud site stays in status Connecting.</p>	<p>Check:</p> <ul style="list-style-type: none"> • If the UDP port 500 is open (when you use a firewall). • The connectivity between the local site and the cloud site. • If the IP address of the local site is correct.
<p>The connection between my local site and the cloud site stays in status Waiting for a connection.</p>	<p>You see this status when the Startup action for cloud site is set to Add, which means that the cloud site is waiting for the local site to initiate the connection.</p> <p>Initiate connection from the local site.</p>
<p>The connection between my local site and the cloud site stays in status Waiting for traffic.</p>	<p>You see this status when the Startup action for cloud site is set to Route.</p> <p>If you are expecting a connection from the local site, do the following:</p> <ul style="list-style-type: none"> • From the local site, try to ping the virtual machine in the cloud site. This is a standard behavior necessary for establishing a tunnel for

Problem	Possible solution
	<p>some devices, for example Cisco ASA. (Route mode)</p> <ul style="list-style-type: none"> • Ensure that the local site established a tunnel by setting the Startup action of the local site to Start.
<p>The connection between my local site and the cloud site is established, but I can see that one or more of the network policies are down.</p>	<p>This issue may be due to the following reasons:</p> <ul style="list-style-type: none"> • Network mapping in the cloud IPsec site is different from the network mapping in the local site. Ensure that the network mappings and the sequence of the network policies in the local and cloud sites match exactly. • This state is correct when the Startup action of the local site and/or of the cloud site is set to Route (for example, on Cisco ASA devices), and currently there is no traffic. You can try to ping to make sure that the tunnel is established. If the ping is not working, check the network mapping on the local and the cloud site.
<p>I want restart a specific IPsec connection.</p>	<p>To restart a specific IPsec connection:</p> <ol style="list-style-type: none"> 1. In the Disaster recovery > Connectivity screen, click the IPsec connection. 2. Click Disable connection. 3. Click the IPsec connection again. 4. Click Enable connection.

Downloading the IPsec VPN log files

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can find additional information about the IPsec connectivity in the log files on the VPN server. The log files are compressed in a .zip archive that you can download and extract.

Prerequisites

Multi-site IPsec VPN connectivity is configured.

To download the .zip archive with the log files

1. In the Cyber Protect console, go to **Disaster Recovery > Connectivity**.
2. Click the gear icon next to the VPN gateway of the cloud site.
3. Click **Download log**.

4. Click **Done**.
5. When the .zip archive is ready for download, click **Download log**, and save it locally.

Multi-site IPsec VPN log files

Note

The availability of this feature depends on the service quotas that are enabled for your account.

The following list describes the IPsec VPN log files that are part of the zip archive, and the information that they contain.

- `ip.txt` - The file contains the logs from the configuration of the network interfaces. You must see two IP addresses - a public IP address, and a local IP address. If you do not see these IP addresses in the log, there is a problem. Contact the Support team.

Note

The mask for the public IP address must be 32.

- `swanctl-list-loaded-config.txt` - The file contains information about all IPsec sites. If you do not see a site in the file, then the IPsec configuration was not applied. Try to update the configuration and save it, or contact the Support team.
- `swanctl-list-active-sas.txt` - The file contains connections and policies that are in status active or a connecting.

Setting up recovery servers

This section describes the concepts of failover and failback, creation of a recovery server, and the disaster recovery operations.

Creating a recovery server

To create a recovery server that will be a copy of your workload, follow the procedure below. You can also watch the [video tutorial](#) that demonstrates the process.

Important

When you perform a failover, you can select only recovery points that were created after the creation of the recovery server.

Prerequisites

- A protection plan must be applied to the original machine that you want to protect. This plan must back up the entire machine, or only the disks, required for booting up and providing the necessary services, to a cloud storage.
- One of the connectivity types to the cloud site must be set.

To create a recovery server

1. On the **All devices** tab, select the machine that you want to protect.
2. Click **Disaster recovery**, and then click **Create recovery server**.
3. Select the number of virtual cores and the size of RAM.

Note

You can see the compute points for every option. The number of compute points reflects the cost of running the recovery server per hour. For more information, see "Compute points" (p. 683).

4. Specify the cloud network to which the server will be connected.
5. Select the **DHCP** option.

DHCP option	Description
Provided by cloud site	Default setting. The IP address of the server will be provided by an automatically configured DHCP server in the cloud.
Custom	The IP address of the server will be provided by your own DHCP server in the cloud.

6. [Optional] Specify the **MAC address**.

The MAC address is a unique identifier that is assigned to the network adapter of the server. If you use custom DHCP, you can configure it to always assign a specific IP addresses to a specific MAC address. In that way you will ensure that the recovery server always gets the same IP address. You can run applications that have licenses that are registered with the MAC address.

7. Specify the IP address that the server will have in the production network. By default, the IP address of the original machine is set.

Note

If you use a DHCP server, add this IP address to the server exclusion list in order to avoid IP address conflicts.

If you use a custom DHCP server, you must specify the same IP address in **IP address in production network** as the one configured in the DHCP server. Otherwise, test failover will not work properly, and the server will not be reachable via a public IP address.

8. [Optional] Select the **Test IP address** check box, and then specify the IP address.

This will give you the capability to test a failover in the isolated test network and to connect to the recovery server via RDP or SSH during a test failover. In the test failover mode, the VPN gateway will replace the test IP address with the production IP address by using the NAT protocol.

If you leave the check box cleared, the console will be the only way to access the server during a test failover.

Note

If you use a DHCP server, add this IP address to the server exclusion list, in order to avoid IP address conflicts.

You can select one of the proposed IP addresses or type in a different one.

9. [Optional] Select the **Internet access** check box.

This will enable the recovery server to access the Internet during a real or test failover. By default, the TCP port 25 is open for outbound connections to public IP addresses.

10. [Optional] Set the **RPO threshold**.

The RPO threshold defines the maximum time interval allowed between the last suitable recovery point for a failover and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

11. [Optional] Select the **Use public IP address** check box.

Having a public IP address makes the recovery server available from the Internet during a failover or test failover. If you leave the check box cleared, the server will be available only in your production network.

The **Use public IP address** option requires the **Internet access** option to be enabled.

The public IP address will be shown after you complete the configuration. By default, TCP port 443 is open for inbound connections to public IP addresses.

Note

If you clear the **Use Public IP address** check box or delete the recovery server, its public IP address will not be reserved.

12. [Optional] [If the backups for the selected machine are encrypted by using encryption as a machine property], specify the password that will be automatically used when creating a virtual machine for the recovery server from the encrypted backup.

- a. Click **Specify**, and then enter the password for the encrypted backup and define a name for the credentials.

By default, you will see the most recent backup in the list.

- b. [Optional] To view all the backups, select **Show all backups**.

- c. Click **Done**.

Note

Although the password that you specify will be stored in a secure credentials store, saving passwords might be against your compliance obligations.

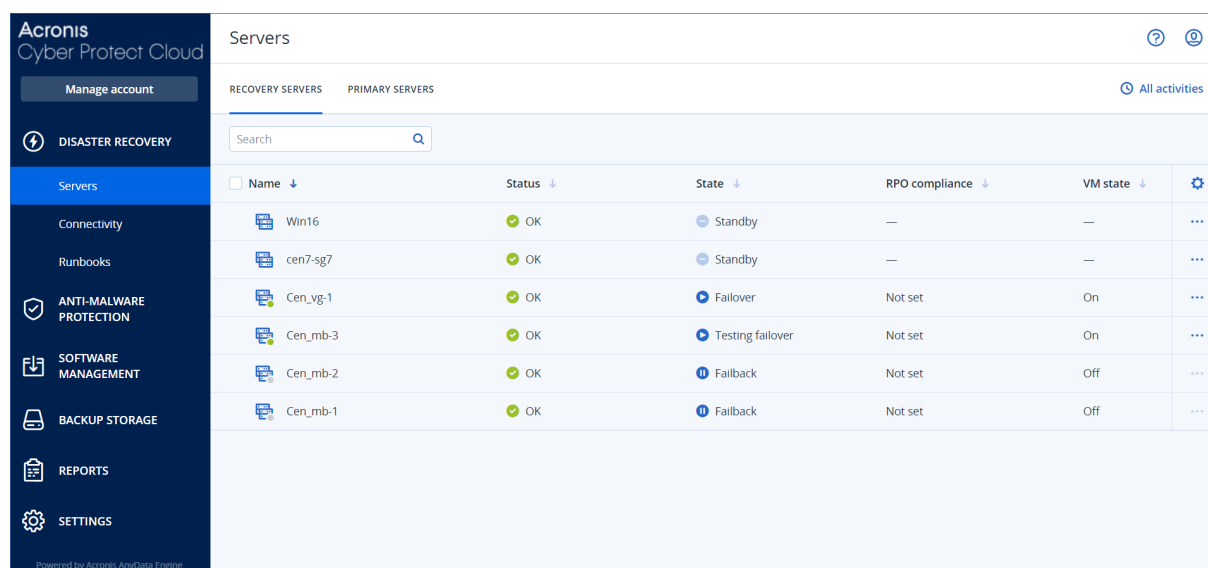
13. [Optional] Change the recovery server name.

14. [Optional] Type a description for the recovery server.

15. [Optional] Click the **Cloud firewall rules** tab to edit the default firewall rules. For more information, see "Setting firewall rules for cloud servers" (p. 749).

16. Click **Create**.

The recovery server appears in the **Disaster Recovery > Servers > Recovery servers** tab of the Cyber Protect console. You can view its settings by selecting the original machine and clicking **Disaster recovery**.



Name	Status	State	RPO compliance	VM state
Win16	OK	Standby	—	—
cen7-sg7	OK	Standby	—	—
Cen_vg-1	OK	Failover	Not set	On
Cen_mb-3	OK	Testing failover	Not set	On
Cen_mb-2	OK	Fallback	Not set	Off
Cen_mb-1	OK	Fallback	Not set	Off

How failover works

Production failover

Note

The availability of this feature depends on the service quotas that are enabled for your account.

When a recovery server is created, it stays in the **Standby** state. The corresponding virtual machine does not exist until you start a failover. Before starting a failover process, you must create at least one disk image backup (with bootable volume) of the original machine.

When starting the failover process, you select the recovery point (backup) of the original machine from which a virtual machine with the predefined parameters will be created. The failover operation uses the "run VM from a backup" functionality. The recovery server gets the transition state **Finalization**. This process implies transferring the server's virtual disks from the backup storage ('cold' storage) to the disaster recovery storage ('hot' storage).

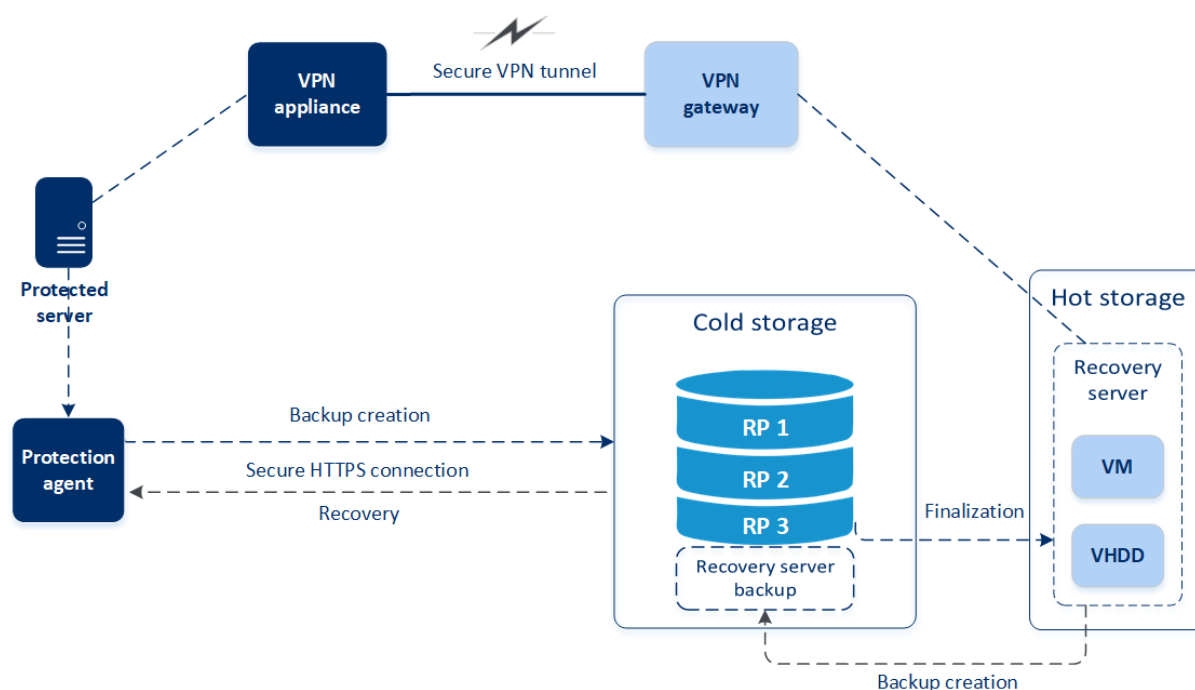
Note

During the **Finalization**, the server is accessible and operable, although the performance is lower than normal. You can open the server console by clicking the **Console is ready** link. The link is available in the **VM State** column on the **Disaster Recovery > Servers** screen, and in the server's **Details** view.

When the **Finalization** is completed, the server performance reaches its normal value. The server state changes to **Failover**. The workload is now switched from the original machine to the recovery server in the cloud site.

If the recovery server has a protection agent inside, the agent service is stopped in order to avoid interference (such as starting a backup or reporting outdated statuses to the backup component).

On the diagram below, you can see both the failover and failback processes.



Test failover

During a **test failover**, the virtual machine is not finalized. This means that the agent reads the virtual disks' content directly from the backup, performing random access to different parts of the backup, so its performance might be slower than the normal performance. For more information about the test failover process, see "Performing a test failover" (p. 726).

Automated test failover

When automated test failover is configured, it is performed once a month without any manual interaction. For more information, see "Automated test failover" (p. 728) and "Configuring automated test failover" (p. 729).

Performing a test failover

Performing a test failover means starting a recovery server in a test VLAN that is isolated from your production network. You can test several recovery servers at a time and check their interaction. In the test network, the servers communicate using their production IP addresses, but they cannot initiate TCP or UDP connections to the workloads in your local network.

During test failover, the virtual machine (recovery server) is not finalized. The agent reads the content of the virtual disks directly from the backup and randomly accesses different parts of the backup. This might make the performance of the recovery server in the test failover state slower than its normal performance.

Though performing a test failover is optional, we recommend that you make it a regular process with a frequency that you find adequate in terms of cost and safety. A good practice is creating a runbook – a set of instructions describing how to spin up the production environment in the cloud.

Important

You must [create a recovery server](#) in advance to protect your devices from a disaster.

You can perform failover only from recovery points that were created after the recovery server of the device was created.

At least one recovery point must be created before failing over to a recovery server. The maximum number of recovery points that is supported is 100.

To perform a test failover

1. Select the original machine or select the recovery server that you want to test.
2. Click **Disaster Recovery**.
The description of the recovery server opens.
3. Click **Failover**.
4. Select the failover type **Test failover**.
5. Select the recovery point (backup), and then click **Start**.
6. If the backup that you selected is encrypted by using encryption as a machine property:
 - a. Enter the encryption password for the backup set.

Note

The password will only be saved temporarily and will be used only for the current test failover operation. The password is automatically deleted from the credentials store if the test failover is stopped, or after the test failover completes.

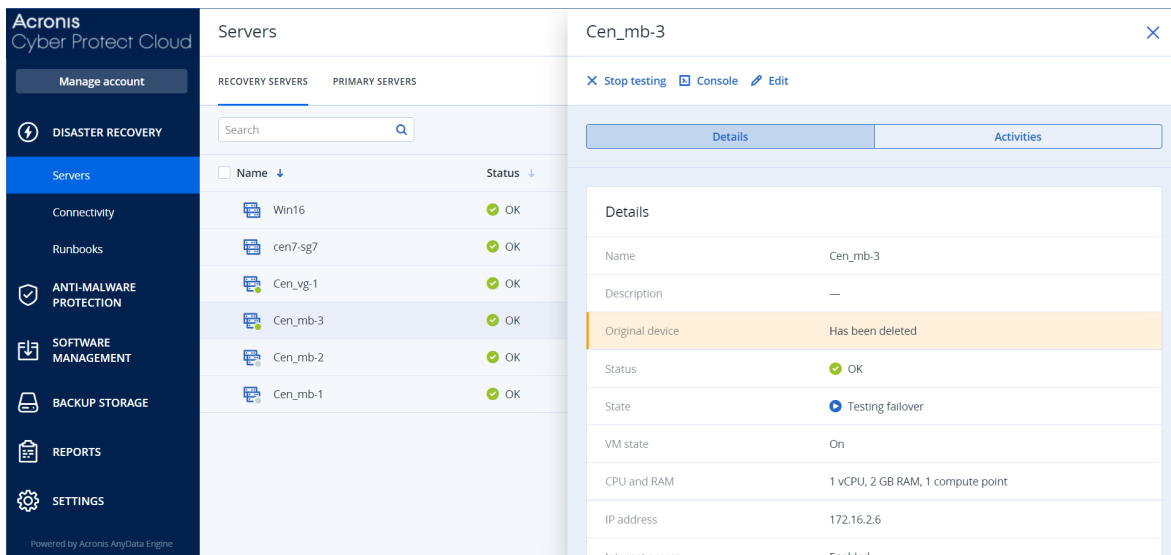
- b. [Optional] To save the password for the backup set and use it in subsequent failover operations, select the **Store the password in a secure credentials store...** check box and then, in the **Credentials name** field, enter a name for the credentials.

Important

The password will be stored in a secured credentials store and will be applied automatically in subsequent failover operations. However, saving passwords might conflict with your compliance obligations.

- c. Click **Done**.

When the recovery server starts, its state changes to **Testing failover**.



7. Test the recovery server by using any of the following methods:
 - In **Disaster Recovery > Servers**, select the recovery server, and then click **Console**.
 - Connect to the recovery server by using RDP or SSH, and the test IP address that you specified when creating the recovery server. Try the connection from both inside and outside the production network (as described in "Point-to-site connection").
 - Run a script within the recovery server.
The script may check the login screen, whether applications are started, the Internet connection, and the ability of other machines to connect to the recovery server.
 - If the recovery server has access to the Internet and a public IP address, you may want to use TeamViewer.
8. When the test is complete, click **Stop testing**.
The recovery server is stopped. All changes made to the recovery server during the test failover are not preserved.

Note

The **Start server** and **Stop server** actions are not applicable for test failover operations, both in runbooks and when starting a test failover manually. If you try executing such an action, it will fail with the following error message:

Failed: The action is not applicable to the current server state.

Automated test failover

With automated test failover, the recovery server is tested automatically once a month without any manual interaction.

The automated test failover process consists of the following parts:

1. creating a virtual machine from the last recovery point
2. taking a screenshot of the virtual machine

3. analyzing if the operating system of the virtual machine starts successfully
4. notifying you about the test failover status

Note

Automated test failover consumes compute points.

You can configure the automated test failover in the recovery server's settings. For more information, see "Configuring automated test failover" (p. 729).

Note that, in very rare cases, automated test failover might be skipped and might not be performed at the scheduled time. This is because production failover has higher priority than automated test failover, so the hardware resources (CPU and RAM) allocated for automated test failover might be temporarily limited to ensure that there are enough resources for a concurrent production failover.

If automated test failover is skipped for some reason, an alert will be raised.

Note

Automated test failover will fail if the backups of the original machine are encrypted by using encryption as a machine property, and the encryption password is not specified when creating the recovery server. For more information about specifying the encryption password, see "Creating a recovery server" (p. 722).

Configuring automated test failover

By configuring automated test failover, you can test your recovery server every month without performing any manual actions.

To configure automated test failover

1. In the console, go to **Disaster recovery > Servers > Recovery servers**, and then select the recovery server.
2. Click **Edit**.
3. In the **Automated test failover** section, in the **Schedule** field, select **Monthly**.
4. [Optional] In **Screenshot timeout**, change the default value of the maximum time period (in minutes) for the system to try performing automated test failover.
5. [Optional] If you want to save the **Screenshot timeout** value as the default and have it populated automatically when you enable automated test failover for the other recovery servers, select **Set as default timeout**.
6. Click **Save**.

Viewing the automated test failover status

You can view the details of a completed automated test failover, such as status, start time, end time, duration, and the screenshot of the virtual machine.

To view the automated test failover status of a recovery server

1. In the console, go to **Disaster recovery** > **Servers** > **Recovery servers** and then select the recovery server.
2. In the **Automated test failover** section, check the details of the last automated test failover.
3. [Optional] Click **Show screenshot** to view the screenshot of the virtual machine.

Disabling automated test failover

You can disable automated test failover if you want to save resources or you do not need automated test failover to be performed for a certain recovery server.

To disable automated test failover

1. In the console, go to **Disaster recovery** > **Servers** > **Recovery servers**, and then select the recovery server.
2. Click **Edit**.
3. In the **Automated test failover** section, in the **Schedule** field, select **Never**.
4. Click **Save**.

Performing a failover

Note

The availability of this feature depends on the service quotas that are enabled for your account.

A failover is a process of moving a workload from your premises to the cloud, and also the state when the workload remains in the cloud.

When you start a failover, the recovery server starts in the production network. To avoid interference and unwanted issues, ensure that the original workload is not online and cannot be accessed via VPN.

To avoid a backup interference into the same cloud archive, manually revoke the protection plan from the workload that is currently in **Failover** state. For more information about revoking plans, see [Revoking a protection plan](#).

Important

You must [create a recovery server](#) in advance to protect your devices from a disaster.

You can perform failover only from recovery points that were created after the recovery server of the device was created.

At least one recovery point must be created before failing over to a recovery server. The maximum number of recovery points that is supported is 100.

You can follow the instructions below or watch the [video tutorial](#).

To perform a failover

1. Ensure that the original machine is not available on the network.
2. In the Cyber Protect console, go to **Disaster recovery > Servers > Recovery servers** and select the recovery server.
3. Click **Failover**.
4. Select the type of failover **Production failover**.
5. Select the recovery point (backup), and then click **Start**.
6. [If the backup that you selected is encrypted by using encryption as a machine property]
 - a. Enter the encryption password for the backup set.

Note

The password will only be saved temporarily and will be used only for the current failover operation. The password is automatically deleted from the credentials store after the failover operation completes and the server returns to the **Standby** state.

- b. [Optional] To save the password for the backup set and use it in subsequent failover operations, select the **Store the password in a secure credentials store...** check box and then, in the **Credentials name** field, enter a name for the credentials.

Important

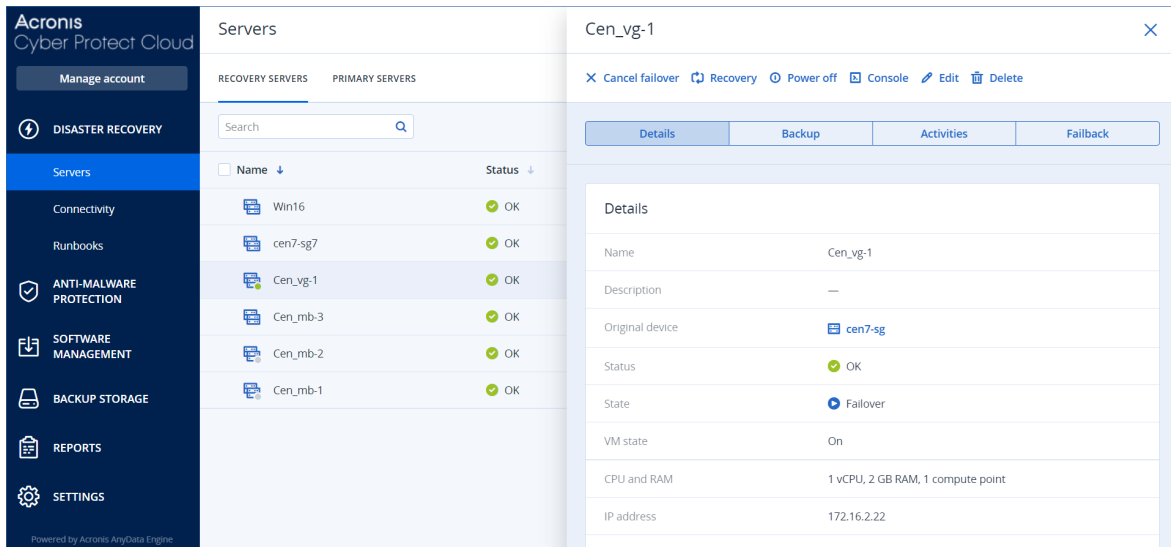
The password will be stored in a secured credentials store and will be applied automatically in subsequent failover operations. However, saving passwords might conflict with your compliance obligations.

- c. Click **Done**.

When the recovery server starts, its state changes to **Finalization**, and after some time to **Failover**.

Important

It is critical to understand that the server is available in both the **Finalization** and **Failover** states. During the **Finalization** state, you can access the server console by clicking the **Console is ready** link. The link is available in the **VM State** column on the **Disaster Recovery > Servers** screen, and in the server's **Details** view. For details, see "How failover works" (p. 725).



7. Ensure that the recovery server is started by viewing its console. Click **Disaster Recovery > Servers**, select the recovery server, and then click **Console**.
8. Ensure that the recovery server can be accessed using the production IP address that you specified when creating the recovery server.

Once the recovery server is finalized, a new protection plan is automatically created and applied to it. This protection plan is based on the protection plan that was used for creating the recovery server, with certain limitations. In this plan, you can change only the schedule and retention rules. For more information, refer to "[Backing up the cloud servers](#)".

If you want to cancel failover, select the recovery server and click **Cancel failover**. All changes starting from the failover moment - except the recovery server backups - will be lost. The recovery server will return back to the **Standby** state.

If you want to perform failback, select the recovery server and click **Failback**.

How to perform failover of servers using local DNS

If you use DNS servers on the local site for resolving machine names, then after a failover the recovery servers, corresponding to the machines relying on the DNS, will fail to communicate because the DNS servers used in the cloud are different. By default, the DNS servers of the cloud site are used for the newly created cloud servers. If you need to apply custom DNS settings, contact the support team.

How to perform failover of a DHCP server

Your local infrastructure may have the DHCP server located on a Windows or Linux host. When such a host is failed over to the cloud site, the DHCP server duplication issue occurs because the VPN gateway in the cloud also performs the DHCP role. To resolve this issue, do one of the following:

- If only the DHCP host was failed over to the cloud, while the rest local servers are still on the local site, then you must log in to the DHCP host in the cloud and turn off the DHCP server on it. Thus, there will be no conflicts and only the VPN gateway will work as the DHCP server.

- If your cloud servers already got the IP addresses from the DHCP host, then you must log in to the DHCP host in the cloud and turn off the DHCP server on it. You must also log in to the cloud servers and renew the DHCP lease to assign new IP addresses allocated from the correct DHCP server (hosted on the VPN gateway).

Note

The instructions are not valid when your cloud DHCP server is configured with the **Custom DHCP** option, and some of the recovery or primary servers get their IP address from this DHCP server.

How failback works

Note

The availability of this feature depends on the service quotas that are enabled for your account.

A failback is a process of moving the workload from the cloud back to a physical or virtual machine on your local site. You can perform a failback on a recovery server in **Failover** state, and continue using the server on your local site.

You can perform automated failover to a virtual or physical target machine on your local site. During the failback, you can transfer the backup data to your local site while the virtual machine in the cloud continues to run. This technology helps you to achieve a very short downtime period, which is estimated and displayed in the Cyber Protect console. You can view it and use this information to plan your activities and, if necessary, warn your clients about an upcoming downtime period.

The failback processes to target virtual machines and target physical machines are slightly different. For more information about the phases of the failback process, see "Failback to a target virtual machine" (p. 733) and "Failback to a target physical machine" (p. 738).

In specific cases when you cannot use the automated failback procedure, you can perform a manual failback. For more information, see "Manual failback" (p. 742).

Note

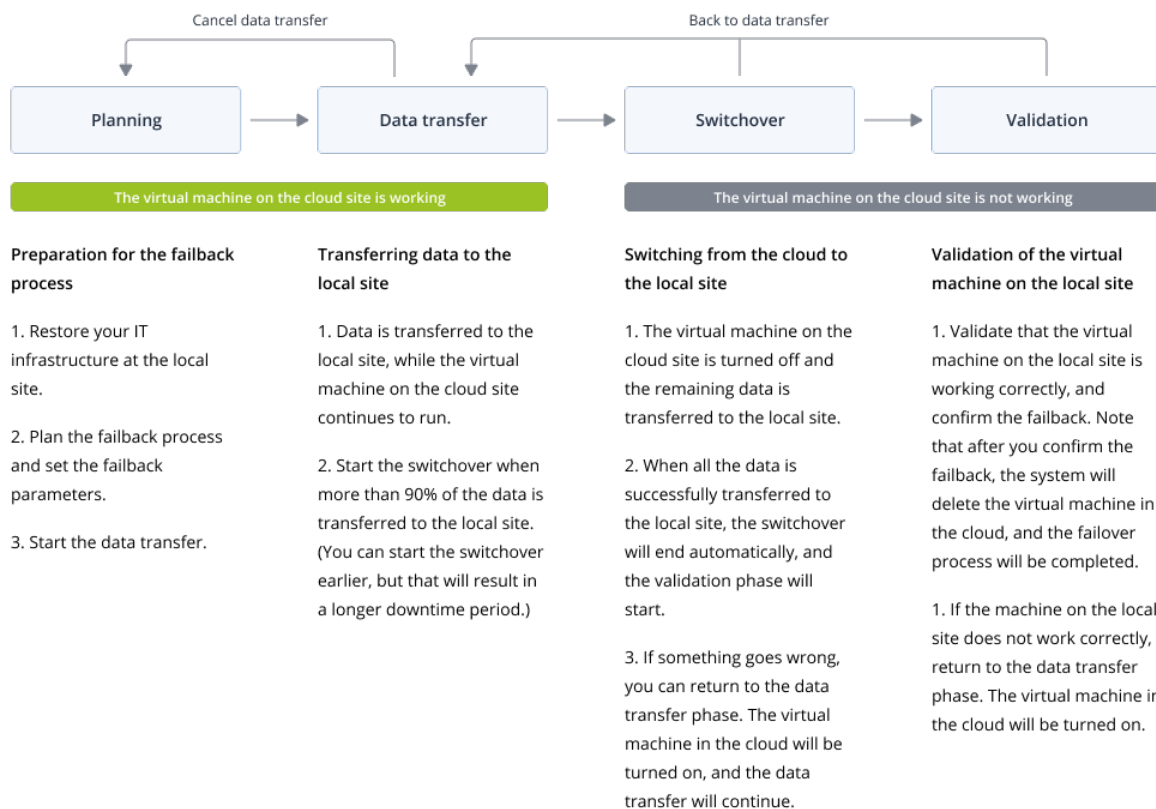
Runbook operations support the failback in manual mode only. This means that if you start the failback process by executing a runbook that includes a **Failback server** step, the procedure will require a manual interaction: you must manually recover the machine, and confirm or cancel the failback process from the **Disaster Recovery > Servers** tab.

Failback to a target virtual machine

Note

The availability of this feature depends on the service quotas that are enabled for your account.

The failback process to a target virtual machine consists of four phases.



1. **Planning.** During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.

Note

To minimize the total time for the failback process, we recommend that you start the data transfer phase immediately after you set up your local servers, and then continue with the configuration of the network and the rest of the local infrastructure during the data transfer phase.

2. **Data transfer.** During this phase, the data is transferred from the cloud site to the local site while the virtual machine in the cloud continues to run. You can start the next phase - switchover - at any time during the data transfer phase, but you should consider the following relations. The longer you remain in the data transfer phase,
 - the longer the virtual machine in the cloud continues to run.
 - the more data will be transferred to your local site.
 - the higher the cost you will pay (you spend more compute points).
 - the shorter the downtime period that you will experience during the switchover phase.
 If you want to minimize the downtime, start the switchover phase after more than 90% of the data is transferred to the local site.

If you can afford to experience a longer downtime period, and do not want to spend more compute points for running the virtual machine in the cloud, you can start the switchover phase earlier.

If you cancel the failback process during the data transfer phase, the transferred data will not be deleted from the local site. To avoid potential issues, manually delete the transferred data before you start a new failback process. The following data transfer process will start from the beginning.

3. **Switchover.** During this phase, the virtual machine in the cloud is turned off, and the remaining data - including the last backup increment - is transferred to the local site. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process.

You can view the estimated time to finish (downtime period) of this phase in the Cyber Protect console. When all the data is transferred to the local site (there is no data loss, and the virtual machine on the local site is an exact copy of the virtual machine in the cloud), the switchover phase completes. The virtual machine on the local site is recovered, and the validation phase starts automatically.

4. **Validation.** During this phase, the virtual machine on the local site is ready and you can turn it on. You can verify if the virtual machine is working correctly, and:
 - If everything is working as expected, confirm the failback. After the failback confirmation, the virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state. This is the end of the failback process.
 - If something is wrong, you can cancel the switchover and return to the data transfer phase.

Performing failback to a virtual machine

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can perform failback to a target virtual machine on your local site.

Prerequisites

- The agent that you will use to perform failback is online and is not currently used for another failback operation.
- Your Internet connection is stable.
- There is at least one full backup of the virtual machine in the cloud.

To perform a failback to a virtual machine

1. In the Cyber Protect console, go to **Disaster recovery > Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.
4. In the **Failback parameters** section, select **Virtual machine** as a **Target**, and configure the other parameters.

Note that by default, some of the **Failback parameters** are populated automatically with suggested values, but you can change them.

The following table provides more information about the **Failback parameters**.

Parameter	Description
Backup size	<p>Amount of data that will be transferred to your local site during the failback process.</p> <p>After you start the failback process to a target virtual machine, the Backup size will be increasing during the data transfer phase, because the virtual machine in the cloud will continue to run and generate new data.</p> <p>To calculate the estimated downtime period during the failback process to a target virtual machine, take 10% of the Backup size value (as we recommend that you start the switchover phase after 90% of the data is transferred to your local site), and divide it by the value of your Internet speed.</p> <hr/> <p>Note</p> <p>The value of the Internet speed will decrease when you perform several failback processes at the same time.</p> <hr/>
Target	Type of workload on your local site to which you will recover the cloud server: Virtual machine or Physical machine .
Target machine location	Failback location: a VMware ESXi host or a Microsoft Hyper-V host. You can select from all the hosts that have an agent which is registered with the Cyber Protection service.
Agent	<p>Agent which will perform the failback operation.</p> <p>You can use one agent to perform one failback operation at the same time.</p> <p>You can select an agent that is online and is not currently used for another failback process, has a version which supports the failback functionality, and has rights to access the backup.</p> <p>Note that you can install several agents on VMware ESXi hosts, and start a separate failback process using each of them. These failback processes can be performed at the same time.</p>
Target machine settings	<p>Virtual machine settings:</p> <ul style="list-style-type: none"> • Virtual processors. Select the number of virtual processors. • Memory. Select how much memory the virtual machine will have. • Units. Select the units for the memory. • [Optional] Network adapters. To add a network adapter, click Add, and select a network in the Network field. <p>When you are ready with the changes, click Done.</p>
Path	(For Microsoft Hyper-V hosts) Folder on the host where your machine will

Parameter	Description
	be stored. Ensure that there is enough free memory space on the host for the machine.
Datastore	(For VMware ESXi hosts) Datastore on the host where your machine will be stored. Ensure that there is enough free memory space on the host for the machine.
Provisioning mode	Method of allocation of the virtual disk. For Microsoft Hyper-V hosts: <ul style="list-style-type: none"> • Dynamically expanding (default value). • Fixed size. For Microsoft Hyper-V hosts: <ul style="list-style-type: none"> • Thin (default value). • Thick.
Target machine name	Name of the target machine. By default, the target machine name is the same as the recovery server name. The target machine name must be unique on the selected Target machine location .

- Click **Start data transfer**, and then in the confirmation window, click **Start**.

The **Data transfer** phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows how much data is already transferred to the local site, and the total amount of data that must be transferred. The total amount of data includes the data from the last backup before the data transfer phase was started, and the backups of the newly generated data (backup increments), as the virtual machine continues to run during the data transfer phase. For this reason, both values of the Progress parameter increase with time.
Downtime estimation	This parameter shows how much time the virtual machine in the cloud will be unavailable if you start the switchover phase now. The value is calculated based on the values of the Progress parameter, and decreases with time.

- Click **Switchover** and then, in the confirmation window, click **Switchover** again.

The switchover phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows the progress of restoring the machine on the local

Field	Description
	site.
Estimated time to finish	This parameter shows the approximate time when the switchover phase will be completed and you will be able to start the machine on the local site.

Note

If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

- After the **Switchover** phase completes, validate that the virtual machine on your local site is working as expected.
- Click **Confirm failback**, and then in the confirmation window, click **Confirm** to finalize the process.

The virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state.

Note

Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Failback to a target physical machine

Note

The availability of this feature depends on the service quotas that are enabled for your account.

The automatic failback process to a target physical machine consists of the following phases:

- Planning.** During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.
- Data transfer.** During this phase, the data is transferred from the cloud site to the local site while the virtual machine in the cloud continues to run. You can start the next phase - switchover - at any time during the data transfer phase, but you should consider the following relations.

The longer you remain in the data transfer phase,

- the longer the virtual machine in the cloud continues to run.
- the more data will be transferred to your local site.
- the higher the cost you will pay (you spend more compute points).
- the shorter the downtime period that you will experience during the switchover phase.

If you want to minimize the downtime, start the switchover phase after more than 90% of the data is transferred to the local site.

If you can afford to experience a longer downtime period, and do not want to spend more compute points for running the virtual machine in the cloud, you can start the switchover phase earlier.

Note

The data transfer process uses a flashback technology. This technology compares the data that is available on the target machine to the data of the virtual machine in the cloud. If part of the data is already available on the target machine, it will not be transferred again. This technology makes the data transfer phase faster.

For this reason, we recommend that you restore the server to the original machine on your local site.

3. **Switchover.** During this phase, the virtual machine in the cloud is turned off, and the remaining data - including the last backup increment - is transferred to the local site. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process.
4. **Validation.** During this phase, the physical machine on the local site is ready, and you can reboot it using a Linux-based bootable media. You can verify if the virtual machine is working correctly, and:
 - If everything is working as expected, confirm the failback. After the failback confirmation, the virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state. This is the end of the failback process.
 - If something is wrong, you can cancel the failover and return to the planning phase.

Note

After the bootable media has been rebooted, you will not be able to use it again. If, at the validation phase, you discover something wrong, you must register a new bootable media and start the failback process again.

However, as flashback technology will be used, the data that is already on the local site will not be transferred again, and the failback process will be much faster.

Performing failback to a physical machine

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can perform automatic failback to a target physical machine on your local site.

Note

The data transfer process uses a flashback technology. This technology compares the data that is available on the target machine to the data of the virtual machine in the cloud. If part of the data is already available on the target machine, it will not be transferred again. This technology makes the data transfer phase faster.

For this reason, we recommend that you restore the server to the original machine on your local site.

Prerequisites

- The agent that you will use to perform failback is online and is not currently used for another failback operation.
- Your Internet connection is stable.
- A registered bootable media is available. For more information, see "Creating bootable media to recover operating systems" in the Cyber Protection User Guide.
- The target physical machine is the original machine on your local site, or has the same firmware as the original machine.
- There is at least one full backup of the virtual machine in the cloud.

To perform a failback to a physical machine

1. In the Cyber Protect console, go to **Disaster recovery > Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.
4. In the **Target** field, select **Physical machine**.
5. In the **Target bootable media** field, click **Specify**, select the bootable media, and then click **Done**.

Note

We recommend that you use ready-made bootable media as it is already configured. For more information, see "Creating bootable media to recover operating systems" in the Cyber Protection User Guide.

6. [Optional] To change the default disk mapping, in the **Disk mapping** field, click **Specify**, map the disks of the backup to the disks of the target machine, and then click **Done**.
7. Click **Start data transfer** and then, in the confirmation window, click **Start**.

The data transfer phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows how much data is already transferred to the local site, and the total amount of data that must be transferred. The total amount of data includes the data from the last backup before the

Field	Description
	<p>data transfer phase was started, and the backups of the newly generated data (backup increments), as the virtual machine continues to run during the data transfer phase. For this reason, the Progress values increase with time.</p> <p>As the system uses a flashback technology during the data transfer and does not transfer the data that is already available on the target machine, the progress might be faster than what is initially calculated by the console.</p>
Downtime estimation	<p>This parameter shows how much time the virtual machine in the cloud will be unavailable if you start the switchover phase now. The value is calculated based on the values of the Progress parameter, and decreases with time.</p> <p>As the system uses a flashback technology during the data transfer and does not transfer the data that is already available on the target machine, the downtime might be much shorter than the value that is initially displayed in the console.</p>

- Click **Switchover** and then, in the confirmation window, click **Switchover** again.

The switchover phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows the progress of restoring the machine on the local site.
Estimated time to finish	This parameter shows the approximate time when the switchover phase will be completed and you will be able to start the machine on the local site.

Note

If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

- After the **Switchover** phase completes, reboot the bootable media, and then verify that the physical machine on your local site is working as expected.
For more information, see "Recovering disks using bootable media" in the Cyber Protection User Guide.
- Click **Confirm fallback** and then, in the confirmation window, click **Confirm** to finalize the process.
The virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state.

Note

Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Manual failback

Note

We recommend that you use the failback process in a manual mode only when you are advised to do so by the Support team.

You can also start a failback process in a manual mode. In this case, the data transfer from the backup in the cloud to the local site will not be done automatically. It must be done manually after the virtual machine in the cloud is powered off. This makes the failback process in a manual mode much slower, and you should expect a longer downtime period.

The failback process in a manual mode consists of the following phases:

1. **Planning.** During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.
2. **Switchover.** During this phase, the virtual machine in the cloud is turned off, and the newly generated data is backed up. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process. When the backup is complete, you recover the machine to the local site manually. You can either recover the disk by using bootable media, or recover the entire machine from the cloud backup storage.
3. **Validation.** During this phase, you verify that the physical or virtual machine at the local site is working correctly, and confirm the failback. After the confirmation, the virtual machine on the cloud site is deleted, and the recovery server returns to the **Standby** state.

Performing manual failback

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can perform a manual failback to a target physical or virtual machine on your local site.

To perform a manual failback

1. In the Cyber Protect console, go to **Disaster recovery > Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.

4. In the **Target** field, select **Physical machine**.
5. Click the gear icon, and then enable the **Use manual mode** switch.
6. [Optional] Calculate the estimated downtime period during the failback process, by dividing the **Backup size** value by the value of your Internet speed.

Note

The value of the Internet speed will decrease when you perform several failback processes at the same time.

7. Click **Switchover**, and then in the confirmation window, click **Switchover** again.
The virtual machine on the cloud site is turned off.

Note

If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

8. Recover the server from the cloud backup to the physical or virtual machine on your local site.
For more information, see "Recovering a machine" in the Cyber Protection User Guide.
9. Ensure that the recovery is completed and the recovered machine works properly, and click **Machine is restored**.
10. If everything is working as expected, click **Confirm failback**, and then in the confirmation window, click **Confirm** again.
The recovery server and recovery points become ready for the next failover. To create new recovery points, apply a protection plan to the new local server.

Note

Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Working with encrypted backups

You can create recovery servers from the encrypted backups. For your convenience, you can set up an automatic password application to an encrypted backup during the failover to a recovery server.

When creating a recovery server, you can [specify the password to be used for automatic disaster recovery operations](#). It will be saved to the Credentials store, a secure storage of credentials that can be found in **Settings > Credentials** section.

One credential can be linked to several backups.

To manage the saved passwords in the Credentials store

1. Go to **Settings > Credentials**.
2. To manage a specific credential, click the icon in the last column. You can view the items linked to this credential.
 - To unlink the backup from the selected credential, click the recycle bin icon near the backup. As a result, you will have to specify the password manually during the failover to the recovery server.
 - To edit the credential, click **Edit**, and then specify the name or password.
 - To delete the credential, click **Delete**. Note that you will have to specify the password manually during the failover to the recovery server.

Operations with Microsoft Azure virtual machines

Note

Some features might require additional licensing, depending on the applied licensing model.

You can perform failover of Microsoft Azure virtual machines to Acronis Cyber Protect Cloud. For more information, see "Performing a failover" (p. 730).

After that, you can perform failback from Acronis Cyber Protect Cloud back to Azure virtual machines. The failback process is same as the failback process to a physical machine. For more information, see "Prerequisites" (p. 740).

Note

To register a new Azure virtual machine for failing back, you can use the Acronis Backup VM extension that is available in Azure.

You can configure a Multisite IPsec VPN connectivity between Acronis Cyber Protect Cloud and the Azure VPN gateway. For more information, see "Configuring Multi-site IPsec VPN" (p. 700).

Setting up primary servers

This section describes how to create and manage your primary servers.

Creating a primary server

Prerequisites

- One of the connectivity types to the cloud site must be set.

To create a primary server

1. Go to **Disaster Recovery > Servers > Primary servers** tab.
2. Click **Create**.
3. Select a template for the new virtual machine.

- Select the flavor of the configuration (number of virtual cores and the size of RAM). The following table shows the maximum total amount of disk space (GB) for each flavor.

Type	vCPU	RAM (GB)	Maximum total amount of disk space (GB)
F1	1	2	500
F2	1	4	1000
F3	2	8	2000
F4	4	16	4000
F5	8	32	8000
F6	16	64	16000
F7	16	128	32000
F8	16	256	64000

Note

You can see the compute points for every option. The number of compute points reflects the cost of running the primary server per hour. For more information, see "Compute points" (p. 683).

- [Optional] Change the virtual disk size. If you need more than one hard disk, click **Add disk**, and then specify the new disk size. Currently, you can add no more than 10 disks for a primary server.
- Specify the cloud network in which the primary server will be included.
- Select the **DHCP** option.

DHCP option	Description
Provided by cloud site	Default setting. The IP address of the server will be provided by an automatically configured DHCP server in the cloud.
Custom	The IP address of the server will be provided by your own DHCP server in the cloud.

- [Optional] Specify the **MAC address**.
The MAC address is a unique identifier that is assigned to the network adapter of the server. If you use custom DHCP, you can configure it to always assign a specific IP addresses to a specific MAC address. This ensures that the primary server always gets the same IP address. You can run applications that have licenses that are registered with the MAC address.
- Specify the IP address that the server will have in the production network. By default, the first free IP address from your production network is set.

Note

If you use a DHCP server, add this IP address to the server exclusion list in order to avoid IP address conflicts.

If you use a custom DHCP server, you must specify the same IP address in **IP address in production network** as the one configured in the DHCP server. Otherwise, test failover will not work properly, and the server will not be reachable via a public IP address.

10. [Optional] Select the **Internet access** check box.

This will enable the primary server to access the Internet. By default, TCP port 25 is open for outbound connections to public IP addresses.

11. [Optional] Select the **Use public IP address** check box.

Having a public IP address makes the primary server available from the Internet. If you leave the check box cleared, the server will be available only in your production network.

The public IP address will be shown after you complete the configuration. By default, TCP port 443 is open for inbound connections to public IP addresses.

Note

If you clear the **Use Public IP address** check box or delete the recovery server, its public IP address will not be reserved.

12. [Optional] Select **Set RPO threshold**.

RPO threshold defines the maximum allowable time interval between the last recovery point and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

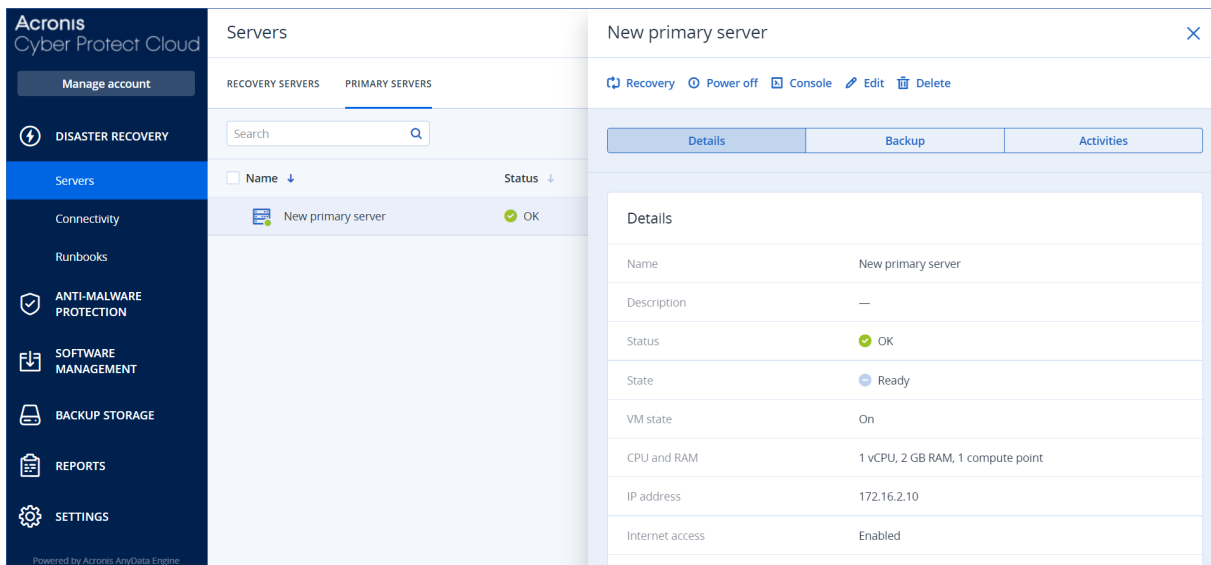
13. Define the primary server name.

14. [Optional] Specify a description for the primary server.

15. [Optional] Click the **Cloud firewall rules** tab to edit the default firewall rules. For more information, see "Setting firewall rules for cloud servers" (p. 749).

16. Click **Create**.

The primary server becomes available in the production network. You can manage the server by using its console, RDP, SSH, or TeamViewer.



Operations with a primary server

The primary server appears in the **Disaster Recovery > Servers > Primary servers** tab in the Cyber Protect console.

To start or stop the server, click **Power on** or **Power off** on the primary server panel.

To edit the primary server settings, stop the server, and then click **Edit**.

To apply a protection plan to the primary server, select it and on the **Plan** tab click **Create**. You will see a predefined protection plan where you can change only the schedule and retention rules. For more information, refer to "[Backing up the cloud servers](#)".

Managing the cloud servers

To manage the cloud servers, go to **Disaster Recovery > Servers**. There are two tabs there: **Recovery servers** and **Primary servers**. To show all optional columns in the table, click the gear icon.

You can find the following information about each cloud server by selecting it.

Column name	Description
Name	A cloud server name defined by you
Status	The status reflecting the most severe issue with a cloud server (based on the active alerts)
State	A cloud server state
VM state	The power state of a virtual machine associated with a cloud server
Active	The location where a cloud server is hosted. For example, Cloud .

location	
RPO threshold	The maximum time interval allowed between the last suitable recovery point for failover and the current time. The value can be set within 15-60 minutes, 1-24 hours, 1-14 days.
RPO compliance	<p>The RPO compliance is the ratio between the actual RPO and RPO threshold. The RPO compliance is shown if the RPO threshold is defined.</p> <p>It is calculated as follows:</p> <p>RPO compliance = Actual RPO / RPO threshold</p> <p>where</p> <p>Actual RPO = current time - last recovery point time</p> <p>RPO compliance statuses</p> <p>Depending on the value of the ratio between the actual RPO and RPO threshold, the following statuses are used:</p> <ul style="list-style-type: none"> • Compliant. The RPO compliance < 1x. A server meets the RPO threshold. • Exceeded. The RPO compliance <= 2x. A server violates the RPO threshold. • Severely exceeded. The RPO compliance <= 4x. A server violates the RPO threshold more than 2x times. • Critically exceeded. The RPO compliance > 4x. A server violates the RPO threshold more than 4x times. • Pending (no backups). The server is protected with the protection plan but the backup is being created and not completed yet.
Actual RPO	The time passed since the last recovery point creation
Last recovery point	The date and time when the last recovery point was created

Firewall rules for cloud servers

You can configure firewall rules to control the inbound and outbound traffic of the primary and recovery servers on your cloud site.

You can configure inbound rules after you provision a public IP address for the cloud server. By default, TCP port 443 is allowed, and all other inbound connections are denied. You can change the default firewall rules, and add or remove Inbound exceptions. If a public IP is not provisioned, you can only view the inbound rules, but cannot configure them.

You can configure outbound rules after when you provision Internet access for the cloud server. By default, TCP port 25 is denied, and all other outbound connections are allowed. You can change the default firewall rules, and add or remove outbound exceptions. If Internet access is not provisioned, you can only view the outbound rules, but cannot configure them.

Note

For security reasons, there are predefined firewall rules that you cannot change.

For inbound and outbound connections:

- Permit ping: ICMP echo-request (type 8, code 0) and ICMP echo-reply (type 0, code 0)
- Permit ICMP need-to-frag (type 3, code 4)
- Permit TTL exceeded (type 11, code 0)

For inbound connections only:

- Non-configurable part: Deny all

For outbound connections only:

- Non-configurable part: Reject all
-

Setting firewall rules for cloud servers

You can edit the default firewall rules for the primary and recovery servers in the cloud.

To edit the firewall rules of a server on your cloud site

1. In the Cyber Protect console, go to **Disaster Recovery > Servers**.
2. If you want to edit the firewall rules of a recovery server, click the **Recovery servers** tab.
Alternatively, if you want to edit the firewall rules of a primary server, click the **Primary servers** tab.
3. Click the server, and then click **Edit**.
4. Click the **Cloud firewall rules** tab.
5. If you want to change the default action for the inbound connections:
 - a. In the **Inbound** drop-down field, select the default action.

Action	Description
Deny all	Denies any inbound traffic. You can add exceptions and allow traffic from specific IP addresses, protocols, and ports.
Allow all	Allows all inbound TCP and UDP traffic. You can add exceptions and deny traffic from specific IP addresses, protocols, and ports.

Note

Changing the default action invalidates and removes the configuration of existing inbound rules.

- b. [Optional] If you want to save the existing exceptions, in the confirmation window, select

Save filled-in exceptions.

- c. Click **Confirm**.
- 6. If you want to add an exception:
 - a. Click **Add exception**.
 - b. Specify the firewall parameters.

Firewall parameter	Description
Protocol	Select the protocol for the connection. The following options are supported: <ul style="list-style-type: none">• TCP• UDP• TCP+UDP
Server port	Select the ports to which the rule applies. You can specify the following: <ul style="list-style-type: none">• a specific port number (for example, 2298)• a range of port numbers (for example, 6000-6700)• any port number. Use * if you want the rule to apply to any port number.
Client IP address	Select the IP addresses to which the rule applies. You can specify the following: <ul style="list-style-type: none">• a specific IP address (for example, 192.168.0.0)• a range of IP addresses using the CIDR notation (for example, 192.168.0.0/24)• any IP address. Use * if you want the rule to apply to any IP address.

- 7. If you want to remove an existing inbound exception, click the bin icon next to it.
- 8. If you want to change the default action for the outbound connections:
 - a. In the **Outbound** drop-down field, select the default action.

Action	Description
Deny all	Denies any outbound traffic. You can add exceptions and allow traffic to specific IP addresses, protocols, and ports.
Allow all	Allows all outbound traffic. You can add exceptions and deny traffic from specific IP addresses, protocols, and ports.

Note

Changing the default action invalidates and removes the configuration of existing outbound rules.

- b. [Optional] If you want to save the existing exceptions, in the confirmation window, select **Save filled-in exceptions**.
 - c. Click **Confirm**.
9. If you want to add an exception:
- a. Click **Add exception**.
 - b. Specify the firewall parameters.

Firewall parameter	Description
Protocol	Select the protocol for the connection. The following options are supported: <ul style="list-style-type: none"> • TCP • UDP • TCP+UDP
Server port	Select the ports to which the rule applies. You can specify the following: <ul style="list-style-type: none"> • a specific port number (for example, 2298) • a range of port numbers (for example, 6000-6700) • any port number. Use * if you want the rule to apply to any port number.
Client IP address	Select the IP addresses to which the rule applies. You can specify the following: <ul style="list-style-type: none"> • a specific IP address (for example, 192.168.0.0) • a range of IP addresses using the CIDR notation (for example, 192.168.0.0/24) • any IP address. Use * if you want the rule to apply to any IP address.

10. If you want to remove an existing outbound exception, click the bin icon next to it.
11. Click **Save**.

Checking the cloud firewall activities

After an update of the configuration of the firewall rules of a cloud server, a log of the update activity becomes available in the Cyber Protect console. You can view the log and check the following information:

- user name of the user who updated the configuration
- date and time of the update
- firewall settings for inbound and outbound connections

- the default actions for inbound and outbound connections
- the protocols, ports and IP addresses of the exceptions for inbound and outbound connections

To view the details about a cloud firewall rules configuration change

1. In the Cyber Protect console, click **Monitoring > Activities**.
2. Click the corresponding activity, and click **All Properties**.
The description of the activity should be **Updating cloud server configuration**.
3. In the **context** field, inspect the information that you are interested in.

Backing up the cloud servers

Primary and recovery servers are backed up agentless on the cloud site. These backups have the following restrictions.

- The only possible backup location is the cloud storage. Primary servers are backed up to the **Primary servers backup** storage.

Note

Microsoft Azure backup locations are not supported.

- A backup plan cannot be applied to multiple servers. Each server must have its own backup plan, even if all of the backup plans have the same settings.
- Only one backup plan can be applied to a server.
- Application-aware backup is not supported.
- Encryption is not available.
- Backup options are not available.

When you delete a primary server, its backups are also deleted.

A recovery server is backed up only in the failover state. Its backups continue the backup sequence of the original server. When a failback is performed, the original server can continue this backup sequence. So, the backups of the recovery server can only be deleted manually or as a result of applying the retention rules. When a recovery server is deleted, its backups are always kept.

Note

The backup plans for cloud servers are performed according to UTC time.

Orchestration (runbooks)

Note

Some features might require additional licensing, depending on the applied licensing model.

A runbook is a set of instructions describing how to spin up the production environment in the cloud. You can create runbooks in the Cyber Protect console. To access the **Runbooks** screen, select **Disaster recovery > Runbooks**.

Why use runbooks?

With runbooks, you can:

- Automate a failover of one or multiple servers
- Automatically check the failover result by pinging the server IP address and checking the connection to the port you specify
- Set the sequence of operations for servers running distributed applications
- Include manual operations in the workflow
- Verify the integrity of your disaster recovery solution, by executing runbooks in the test mode.

Creating a runbook

A runbook consists of steps that are executed consecutively. A step consists of actions that start simultaneously.

You can follow the instruction below or watch the [video tutorial](#).

To create a runbook

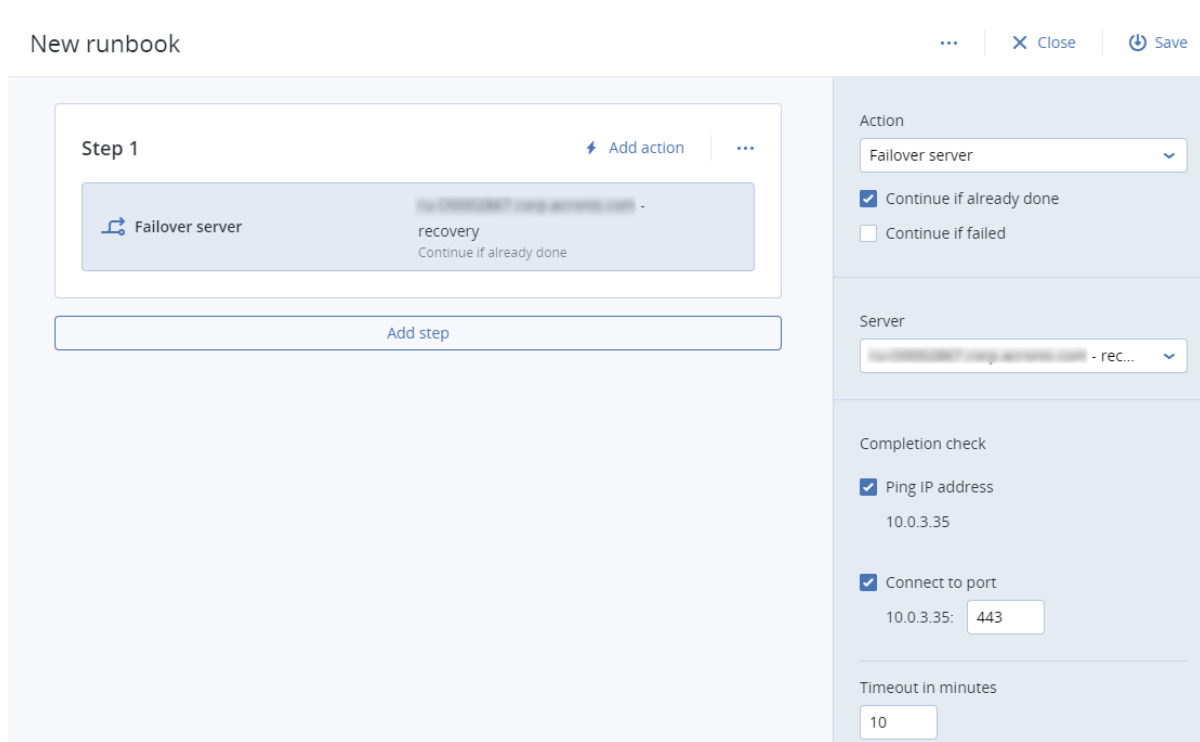
1. In the Cyber Protection console, go to **Disaster recovery > Runbooks**.
2. Click **Create runbook**.
3. Click **Add step**.
4. Click **Add action**, and then select the action that you want to add to the step.

Action	Description
Failover server	<p>Performs a failover of a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these parameters, see "Runbook parameters" (p. 755).</p> <hr/> <p>Note If the backup of the server that you select is encrypted by using encryption as a machine property, the Failover server action will be paused and will be changed automatically to Interaction required. To proceed with the execution of the runbook, you will have to provide the password for the encrypted backup.</p> <hr/>
Failback server	<p>Performs a failback of a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see "Runbook parameters" (p. 755).</p>

Action	Description
	<p>Note Runbook operations support the failback in manual mode only. This means that if you start the failback process by executing a runbook that includes a Failback server step, the procedure will require a manual interaction: you must manually recover the machine, and confirm or cancel the failback process from the Disaster Recovery > Servers tab.</p>
<p>Start server</p>	<p>Starts a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see "Runbook parameters" (p. 755).</p> <p>Note The Start server action is not applicable for test failover operations in runbooks. If you try executing such an action, it will fail with the following error message: Failed: The action is not applicable to the current server state.</p>
<p>Stop server</p>	<p>Stops a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see "Runbook parameters" (p. 755).</p> <p>Note The Stop server action is not applicable for test failover operations in runbooks. If you try executing such an action, it will fail with the following error message: Failed: The action is not applicable to the current server state.</p>
<p>Manual operation</p>	<p>A manual operation requires an interaction from a user. To define this action, you must enter a description.</p> <p>When a runbook sequence reaches a manual operation, the runbook will be paused and will not proceed until a user performs the required manual operation, such as clicking the confirmation button.</p>
<p>Execute runbook</p>	<p>Executes another runbook. To define this action, you must choose a runbook.</p> <p>A runbook can include only one execution of a given runbook. For example, if you added the action "execute Runbook A", you can add the action "execute Runbook B", but cannot add another action "execute Runbook A".</p>

5. Define the runbook parameters for the action. For more information about these parameters, see "Runbook parameters" (p. 755).
6. [Optional] To add a description of the step:
 - a. Click the ellipsis icon, and then click **Description**.
 - b. Enter a description of the step.
 - c. Click **Done**.
7. Repeat steps 3-6 until you create the desired sequence of steps and actions.
8. [Optional] To change the default name of the runbook:

- a. Click the ellipsis icon.
 - b. Enter the name of the runbook.
 - c. Enter a description of the runbook.
 - d. Click **Done**.
9. Click **Save**.
 10. Click **Close**.



Runbook parameters

Runbook parameters are specific settings that you must configure to define a runbook action. There are two categories of runbook parameters - action parameters and completion check parameters.

Action parameters define the runbook behavior depending on the action initial state or result.

Completion check parameters ensure that the server is available and provides the necessary services. If a completion check fails, the action is considered failed.

The following table describes the configurable runbook parameters for each action.

Runbook parameter	Category	Available for action	Description
Continue if already done	Action parameter	<ul style="list-style-type: none"> • Failover server • Start server • Stop server • Failback server 	This parameter defines the runbook behavior when the required action is already done (for example, a failover has already been performed or a server is already running). When enabled, the

Runbook parameter	Category	Available for action	Description
			runbook issues a warning and proceeds. When disabled, the action fails, and then the runbook fails too. By default, this parameter is enabled.
Continue if failed	Action parameter	<ul style="list-style-type: none"> • Failover server • Start server • Stop server • Failback server 	This parameter defines the runbook behavior when the required action fails. When enabled, the runbook issues a warning and proceeds. When disabled, the action fails, and then the runbook fails too. By default, this parameter is disabled.
Ping IP address	Completion check	<ul style="list-style-type: none"> • Start server 	The software will ping the production IP address of the cloud server until the server replies or the timeout expires, whichever comes first.
Connect to port (443 by default)	Completion check	<ul style="list-style-type: none"> • Failover server • Start server 	The software will try to connect to the cloud server by using its production IP address and the port you specify, until the connection is established or the timeout expires, whichever comes first. This way, you can check if the application that listens on the specified port is running.
Timeout in minutes	Completion check	<ul style="list-style-type: none"> • Failover server • Start server 	The default timeout is 10 minutes.

Operations with runbooks

Note

The availability of this feature depends on the service quotas that are enabled for your account.

To access the list of operations, hover on a runbook and click the ellipsis icon. When a runbook is not running, the following operations are available:

- **Execute**
- **Edit**
- **Clone**
- **Delete**

Executing a runbook

Every time you click **Execute**, you are prompted for the execution parameters. These parameters apply to all failover and failback operations included in the runbook. The runbooks specified in the **Execute runbook** operations inherit these parameters from the main runbook.

- **Failover and failback mode**

Choose whether you want to run a test failover (by default) or a real (production) failover. The failback mode will correspond to the chosen failover mode.

- **Failover recovery point**

Choose the most recent recovery point (by default) or select a point in time in the past. If the latter is the case, the recovery points closest before the specified date and time will be selected for each server.

Stopping a runbook execution

During a runbook execution, you can select **Stop** in the list of operations. The software will complete all of the already started actions except for those that require user interaction.

Viewing the execution history

When you select a runbook on the **Runbooks** tab, the software displays the runbook details and execution history. Click the line corresponding to a specific execution to view the execution log.

The screenshot shows a web interface for managing runbooks. On the left is a list of runbooks, with 'Rb0 000' selected. The main panel displays the details and execution history for this runbook.

Runbooks

Search

Execute Edit Clone Delete

Rb0 000

Details

Name Rb0 000

Description -

Execution history

Start and end time	Result	Mode
Aug 14, 5:30 PM - Aug 14, 10:27 PM	Failed	Production
Aug 14, 5:23 PM - Aug 14, 5:25 PM	Failed	Production
Aug 4, 2:45 AM - Aug 4, 2:46 AM	Completed	Test
Jul 30, 4:18 PM - Jul 30, 4:18 PM	Completed	Test
Jul 30, 4:16 PM - Jul 30, 4:16 PM	Completed	Test

Configuring your antivirus and antimalware protection

Note

On Windows machines, the antimalware protection and URL filtering features require the installation of Agent for Antimalware protection and URL filtering. It will be installed automatically for protected workloads if the **Antivirus & Antimalware protection** or the **URL filtering** module is enabled in their protection plans.

Antimalware protection in Cyber Protection provides you with the following benefits:

- Top protection on all the stages: proactive, active, and reactive.
- Four different antimalware technologies inside to provide the best of the breed multi-layered protection.
- Management of Microsoft Security Essentials and Microsoft Defender Antivirus.

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Important

EICAR test file is detected only when the **Advanced Antimalware** option is enabled in the protection plan. However, not detecting the EICAR file does not affect the antimalware capabilities of Cyber Protection.

Supported platforms

Active protection, antivirus and antimalware features are supported on the following platforms.

Operating system	Version/Distribution
Windows	Windows 7 Service Pack 1 and later Windows Server 2008 R2 Service Pack 1 and later Note For Windows 7, you must install the following updates from Microsoft before installing the protection agent. <ul style="list-style-type: none">• Windows 7 Extended Security Updates (ESU)• KB4474419• KB4490628 For more information on the required updates, refer to this knowledge base article .
Linux	Red Hat Linux 7.x, 8.x, 9.x

Operating system	Version/Distribution
	CloudLinux 6.10, 7.x, 8.x CentOS 6.5 and later 6.x versions, 7.x, 8.x Ubuntu 16.04, 18.04, 20.04, 22.04, 22.10 Debian 8.x, 9.x, 10.x, 11.x Oracle Linux 7.x, 8.x, 9.x SUSE Enterprise Linux 15.x openSUSE Leap 15.x
macOS	macOS 10.13.x and later

Supported features per platform

Note

Antimalware protection for Linux and macOS is available with the Advanced Antimalware pack.

Feature set	Windows	Linux	macOS
Antivirus and Antimalware protection			
Fully-integrated Active Protection functionality	Yes	No	No
Real-time antimalware protection	Yes	Yes, with the Advanced Antimalware pack	Yes, with the Advanced Antimalware pack
Advanced real-time antimalware protection with local signature-based detection	Yes	Yes	Yes
Static analysis for portable executable files	Yes	No	Yes*
On-demand antimalware scanning	Yes	Yes**	Yes
Network folder protection	Yes	Yes	No
Server-side protection	Yes	No	No
Scan of archive files	Yes	No	Yes
Scan of removable drives	Yes	No	Yes
Scan of new and changed files only	Yes	No	Yes

Feature set	Windows	Linux	macOS
Antivirus and Antimalware protection			
File/folder exclusions	Yes	Yes	Yes***
Processes exclusions	Yes	No	Yes
Behavioral analysis engine	Yes	No	Yes
Exploit prevention	Yes	No	No
Quarantine	Yes	Yes	Yes
Quarantine auto clean-up	Yes	No	Yes
URL filtering (http/https)	Yes	No	No
Corporate-wide whitelist	Yes	No	Yes
Firewall management****	Yes	No	No
Microsoft Defender Antivirus management*****	Yes	No	No
Microsoft Security Essentials management	Yes	No	No
Registering and managing Antivirus and Antimalware protection via Windows Security Center	Yes	No	No
For more information about the supported operating systems and their versions, see "Supported platforms" (p. 758).			

* Static analysis for portable executable files is supported only for scheduled scans on macOS.

** Start conditions are not supported for on-demand scanning on Linux.

*** File/folder exclusions are only supported for the case when you specify files and folders that will not be scanned by real-time protection or scheduled scans on macOS.

**** Firewall management is supported on Windows 8 and later. Windows Server is not supported.

***** Microsoft Defender Antivirus management is supported on Windows 8.1 and later.

Feature set	Windows	Linux	macOS
Active Protection			
Process Injects detection	Yes	No	No
Automatic recovery of affected files from the local cache	Yes	Yes	Yes

Feature set	Windows	Linux	macOS
Active Protection			
Self-defense for Acronis backup files	Yes	No	No
Self-defense for Acronis software	Yes	No	Yes (Only Active Protection and antimalware components)
Trusted/blocked process management	Yes	No	Yes
Processes/folders exclusions	Yes	Yes	Yes
Ransomware detection based on a process behavior (AI-based)	Yes	Yes	Yes
Cryptomining process detection based on process behavior	Yes	No	No
External drives protection (HDD, flash drives, SD cards)	Yes	No	Yes
Network folder protection	Yes	Yes	Yes
Server-side protection	Yes	No	No
Zoom, Cisco Webex, Citrix Workspace, and Microsoft Teams protection	Yes	No	No
For more information about the supported operating systems and their versions, see "Supported platforms" (p. 758).			

Antivirus and antimalware protection

Note

Some features might require additional licensing, depending on the applied licensing model.

The **Antivirus & Antimalware** module protects your Windows, Linux, and macOS machines from all recent malware threats. See the full list of supported antimalware features in "Supported platforms" (p. 758).

Antivirus & Antimalware protection is supported and registered in Windows Security Center.

Antimalware features

- Detection of malware in files in the real-time protection and on-demand modes
- Detection of malicious behavior in processes (for Windows)
- Blocking access to malicious URLs (for Windows)
- Placing dangerous files to the quarantine
- Adding trusted corporate applications to the allowlist

Scanning types

You can configure antivirus and antimalware protection to run constantly in the background or on demand.

Real-time protection

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Real-time protection checks all files that are being executed or opened on a machine to prevent malware threats.

To prevent potential compatibility and performance issues, real-time protection cannot work in parallel with other antivirus solutions that also use real-time protection features. The statuses of other installed antivirus solutions are determined through Windows Security Center. If the Windows machine is already protected by another antivirus solution, real-time protection is automatically turned off.

To enable real-time protection, disable or uninstall the other antivirus solution. Real-time protection can replace Microsoft Defender real-time protection automatically.

Note

On machines running Windows Server operating systems, Microsoft Defender will not be turned off automatically when real-time protection is enabled. An administrator must turn off the Microsoft Defender manually to avoid potential compatibility issues.

You can choose one of the following scan modes:

- **Smart on-access** detection means that the antimalware program runs in the background and actively and constantly scans your machine system for viruses and other malicious threats for the entire duration that your system is powered on. Malware will be detected in both cases when a file is being executed and during various operations with the file such as opening it for reading or editing.
- **On-execution** detection means that only executable files will be scanned at the moment they are run to ensure they are clean and will not cause any damage to your machine or data. Copying of an infected file will remain unnoticed.

Scheduled scan

Antimalware scanning is performed according to a schedule.

You can choose one of the following scan modes.

- **Quick scan**—Checks only workload system files.
- **Full scan**—Checks all files on your workload.
- **Custom scan**—Checks files/folders that were added by the administrator to the Protection plan.

After antimalware scanning completes, you can see details about the workloads that were affected by threats in the **Monitoring > Overview > Recently affected** widget.

Antivirus and antimalware protection settings

This section describes the features that you can configure in the **Antivirus & Antimalware protection** module in a protection plan. To learn how to create a protection plan, see "Creating a protection plan" (p. 209).

The following features can be configured in the Antivirus & Antimalware protection module for a protection plan:

- [Active Protection](#)
- [Advanced Antimalware](#)
- [Network folder protection](#)
- [Server-side protection](#)
- [Self protection](#)
- [Cryptomining process detection](#)
- [Quarantine](#)
- [Behavior engine](#)
- [Exploit prevention](#)
- [Real-time protection](#)
- [Schedule scan](#)
- [Exclusions](#)

Note

Not all the operating systems support the Antivirus & Antimalware protection features. For more information about the supported operating systems and features, see "Supported platforms" (p. 758). Some features require a certain license to be available in your protection plan. For more information, see

Active Protection

Active Protection protects your system from malicious software known as ransomware that encrypts files and demands a ransom for the encryption key.

Default setting: **Enabled**.

Note

A protection agent must be installed on the protected machine. For more information about the supported operating systems and features, see "Supported platforms" (p. 758).

To configure Active Protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Active Protection**.
3. In the **Action on detection** section, select one of the available options:
Default setting: **Revert using cache**
 - **Notify only**—The software generates an alert about the process suspected of ransomware activity.
 - **Stop the process**—The software generates an alert and stops the process suspected of ransomware activity.
 - **Revert using cache**—The software generates an alert, stops the process, and reverts the file changes by using the service cache.
4. Click **Done** to apply the selected options to your protection plan.

Advanced Antimalware

This engine uses an enhanced database of virus signatures to improve the efficiency of antimalware detection in both quick and full scans.

Important

This feature is available only if you have the Advanced Security protection pack enabled. For more information, see <https://www.acronis.com/en-us/products/cloud/cyber-protect/security/>

Note

The availability of this feature depends on the service quotas that are enabled for your account.

To configure Advanced Antimalware

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. In the **Advanced Antimalware** section, use the toggle to enable the local signature-based engine.

Note

Antivirus and Antimalware protection for macOS and Linux also requires the local signature-based engine. For Windows, Antivirus and Antimalware protection is available with or without this engine.

Network folder protection

The **Network folder protection** feature defines whether Antivirus & Antimalware protection protects network folders that are mapped as local drives. The protection applies to folders shared via SMB or NFS protocols.

To configure Network folder protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Network folder protection**.
3. Add the files where you want to backup the network folders:
 - For example, If your workload is Windows, in the **Windows** field, enter the path for the Windows file where you want to backup the network folders. Default value: C:\ProgramData\Acronis\Restored Network Files.
 - For example, If your workload is macOS, in the **macOS** field, enter the path for the macOS files where you want to backup the network folders. Default value: /Library/Application Support/Acronis/Restored Network Files/.

Note

Enter the path of a local folder. Network folders, including folders on mapped drives, are not supported as backup destinations for the network folders.

4. Click **Done** to apply the selected options to your protection plan.

Server-side protection

This feature defines whether Active protection protects network folders that are shared by you from the external incoming connections from other servers in the network that may potentially bring threats.

Default setting: **Off**.

Note

Server-side protection is not supported for Linux.

To set trusted connections

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Server-side protection**.
3. Use the **Server-side protection** toggle to enable it.
4. Select the **Trusted** tab.
5. In the **Trusted connections** field, click **Add** to define connections that will be allowed to modify data.

6. In the **ComputerName/Account** field, type the name of the computer and the account of the machine where the protection agent is installed. For example, MyComputer\TestUser.
7. In the **Host name** field, type the host name of the machine that is allowed to connect to the machine using the protection agent.
8. Click the check mark to the right to save the connection definition.
9. Click **Done**.

To set blocked connections

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Server-side protection**.
3. Use the **Server-side protection** toggle to enable it.
4. Select the **Blocked** tab.
5. In the **Blocked connections** field, click **Add** to define connections that will not be allowed to modify data.
6. In the **ComputerName/Account** field, type the name of the computer and the account of the machine where the protection agent is installed. For example, MyComputer\TestUser.
7. In the **Host name** field, type the host name of the machine that is allowed to connect to the machine using the protection agent.
8. Select the check box to the right to save the connection definition.
9. Click **Done**.

Self-protection

Self-protection prevents unauthorized changes to the software's own processes, registry records, executable and configuration files, and backups located in your local folders.

Administrators can enable **Self-protection**, without enabling **Active Protection**.

Default setting: **Off**.

Note

Self-protection is not supported for Linux.

To enable Self-protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Self-protection**.
3. Use the **Self-protection** toggle to enable it.

To enable Password protection

1. Once the **Self-protection** feature is enabled, you can enable the **Password protection** feature by using the toggle.
2. Click **Generate new password** to generate a password to modify or delete local agents.
3. Click **Copy**, and then paste it in a safe place because this will be requested when you want to modify the components list locally.

Important

The password will not be available after you close the window. To get this password applied to devices, the protection plan settings must be saved.

4. Click **Close**.

Password protection prevents unauthorized users or software from uninstalling Agent for Windows or modifying its components. These actions are only possible with a password that an administrator can provide.

A password is never required for the following actions:

- Updating the installation by running the setup program locally
- Updating the installation by using the Cyber Protect console
- Repairing the installation

Default setting: **Disabled**

For more information about how to enable **Password protection**, refer to [Preventing unauthorized uninstallation or modification of agents](#).

Cryptomining process detection

Cryptomining malware degrades the performance of useful applications, increases electricity bills, and may cause system crashes and even hardware damage due to abuse. The **Cryptomining process detection** feature protects your devices against cryptomining malware to prevent unsanctioned using of computer resources.

Administrators can enable **Cryptomining process detection**, without enabling **Active Protection**.

Default setting: **Enabled**.

Note

Cryptomining process detection is not supported for Linux.

To configure network folder protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Cryptomining process detection**.
3. Use the **Detect cryptomining processes** toggle to enable or disable the feature.
4. Select what to do with processes suspected of cryptomining activities:

Default setting: **Stop the process**

- **Notify only**—The software generates an alert.
 - **Stop the process** — The software generates an alert and stops the process.
5. Click **Done** to apply the selected options to your protection plan.

Quarantine

Quarantine is a folder used to isolate suspicious (probably infected) or potentially dangerous files.

To configure Quarantine

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Quarantine**.
3. In the **Remove quarantined files after** field, you can define the period in days after which the quarantined files will be removed.

Default setting: **30 days**

4. Click **Done**.

For more information about this feature, refer to [Quarantine](#).

Behavior-engine

The **Behavior engine** feature protects a system from malware by using behavioral heuristic to identify malicious processes.

Default setting: **Enabled**.

Note

Behavior engine is not supported for Linux.

To configure Network folder protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Behavior engine**.
3. Use the **Behavior engine** toggle to enable or disable the feature.
4. In the **Action on detection** section select the action that the software will perform when a malware activity is detected:

Default setting: **Quarantine**

- **Notify only**—The software generates an alert about the process suspected of malware activity.
- **Stop the process**—The software generates an alert and stops the process suspected of malware activity.

- **Quarantine**—The software generates an alert, stops the process, and moves the executable files to the quarantine folder.
5. Click **Done** to apply the selected options to your protection plan.

Exploit prevention

Important

This feature is available only if you have the Advanced Security protection pack enabled. For more information, see <https://www.acronis.com/en-us/products/cloud/cyber-protect/security/>

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Exploit prevention detects and prevents infected processes from spreading and exploiting the software vulnerabilities on a systems. When an exploit is detected, the software can generate an alert and stop the process suspected of exploit activities.

Exploit prevention is available only with agent versions 12.5.23130 (21.08, released in August 2020) or later.

Default setting: **Enabled** for newly created protection plans, and **Disabled** for existing protection plans, created with previous agent versions.

Note

Exploit prevention is not supported for Linux.

You can select what should the program do when an exploit is detected, and which exploit prevention methods are applied by the program.

To configure Exploit prevention

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Exploit prevention**.
3. In the **Action on detection** section, select one of the available options:
Default setting: **Stop the process**
 - **Notify only**
The software will generate an alert about the process suspected of exploit activities.
 - **Stop the process**
The software will generate an alert and stop the process suspected of exploit activities.
4. In the **Enabled exploit prevention techniques** section, select from the available options that you want to be applied:
Default setting: **All methods are enabled**

- **Memory protection**

Detects and prevents suspicious modifications of the execution rights on memory pages. Malicious processes apply such modifications to page properties, to enable the execution of shell codes from non-executable memory areas like stack and heaps.

- **Return-oriented programming (ROP) protection**

Detects and prevents attempts for use of the ROP exploit technique.

- **Privilege escalation protection**

Detects and prevents attempts for elevation of privileges made by an unauthorized code or application. Privilege escalation is used by malicious code to gain full access of the attacked machine, and then perform critical and sensitive tasks. Unauthorized code is not allowed to access critical system resources or modify system settings.

- **Code injection protection**

Detects and prevents malicious code injection into remote processes. Code injection is used to hide malicious intent of an application behind clean or benign processes, to evade detection by antimalware products.

5. Click **Done** to apply the selected options to your protection plan.

Note

Processes that are listed as trusted processes in the Exclusions list will not be scanned for exploits.

Allowing processes to modify backups

The **Allow specific processes to modify backups** setting is only available when the **Self-protection** setting is enabled.

It applies to files that have extensions .tibx, .tib, .tia, and are located in local folders.

This setting lets you specify the processes that are allowed to modify the backup files, even though these files are protected by self-protection. This is useful, for example, if you remove backup files or move them to a different location by using a script.

If this setting is disabled, the backup files can be modified only by processes signed by the backup software vendor. This allows the software to apply retention rules and to remove backups when a user requests this from the web interface. Other processes, no matter suspicious or not, cannot modify the backups.

If this setting is enabled, you can allow other processes to modify the backups. Specify the full path to the process executable, starting with the drive letter.

Default setting: **Disabled**.

Real-time protection

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Real-time protection constantly checks your computer system for viruses and other malicious threats for the entire time that your system is powered on unless paused by the computer user.

Default setting: **Enabled**.

Important

This feature is available only if you have the Advanced Security protection pack enabled. For more information, see <https://www.acronis.com/en-us/products/cloud/cyber-protect/security/>

To configure Real-time protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.

2. Click **Real-time protection**.

3. In the **Action on detection** drop-down, select one of the available options:

Default setting: **Quarantine**

- **Notify only**

The software generates an alert about the process suspected of ransomware activity.

- **Block and notify**

The software blocks the process and generates an alert about the process suspected of malware activities.

- **Quarantine**

4. The software generates an alert, stops the process, and moves the executable file to the quarantine folder.

5. In the **Scan mode** section, select the action that the software will perform when a virus or other malicious threat is detected:

Default setting: **Smart on-access**

- **Smart on-access**—Monitors all system activities and automatically scans files when they are accessed for reading or writing, or whenever a program is launched.

- **On-execution**—Automatically scans only executable files when they are launched to ensure that they are clean and will not cause any damage to your computer or data.

6. Click **Done**.

Schedule scan

On-demand scanning checks your computer system for viruses according to the specified schedule. A full scan checks all the files on your machine, while a quick scan checks only the machine system files.

To configure Schedule scan

Default settings:

- **Custom scan** is disabled.
 - **Quick** and **Full** scan are scheduled.
1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
 2. Click **Schedule scan**.
 3. Use the toggle to enable the type of scan that you want to apply for your machine.

Available types of scan:

- **Full** — takes much longer to finish in comparison to the quick scan because every file will be checked.
- **Quick** — only scans the common areas where malware normally resides on the machine.
- **Custom** — Checks the files/folders that were selected by the administrator of the Protection plan.

Note

You can schedule all three scans - **Quick**, **Full**, and **Custom** - in one protection plan.

To configure custom scan

- Use the **Custom scan toggle** to enable or disable this type of scan.
- In the **Action on detection** drop-down list, select one of the available options:

Default setting: **Quarantine**

Quarantine

The software generates an alert and moves the executable file to the quarantine folder.

Notify only

The software generates an alert about the process that is suspected to be malware.

Field	Description
Schedule the task run using the following events	<p>This setting defines when the task will run.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Schedule by time – This is the default setting. The task will run according to the specified time. • When user logs in to the system – By default, a login of any user will trigger the task. You can modify this setting so that only a specific user account can trigger the task. • When user logs off the system – By default, a logoff of any user will trigger the task. You can modify this setting so that only a specific user account can trigger the task.

Field	Description
	<p>Note The task will not run at system shutdown. Shutting down and logging off are different events in the scheduling configuration.</p> <ul style="list-style-type: none"> • On the system startup – The task will run when the operating system starts. • On the system shutdown – The task will run when the operating system shuts down.
Schedule type	<p>The field appears if in Schedule the task run using the following events you have selected Schedule by time.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Monthly – Select the months and the weeks or days of the month when the task will run. • Daily – This is the default setting. Select the days of the week when the task will run. • Hourly – Select the days of the week, repetition number, and the time interval in which the task will run.
Start at	<p>The field appears if in Schedule the task run using the following events you have selected Schedule by time</p> <p>Select the exact time when the task will run.</p>
Run within a date range	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time.</p> <p>Set a range in which the configured schedule will be effective.</p>
Specify a user account whose login to the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs in to the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the login of any user to trigger the task. • The following user - Use this option if you want only the login of a specific user account to trigger the task.
Specify a user account whose logout from the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs off the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the logout of any user to trigger the task. • The following user - Use this option if you want only the logout of a specific user account to trigger the task.

Field	Description
Start conditions	<p>Defines all conditions that must be met simultaneously for the task to run.</p> <p>Start conditions for antimalware scans are similar to the start conditions for the Backup module that are described in "Start conditions".</p> <p>You can define the following additional start conditions:</p> <ul style="list-style-type: none"> • Distribute task start time within a time window – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM. • If the machine is turned off, run missed tasks at the machine startup • Prevent the sleep or hibernate mode during task running – This option is effective only for machines running Windows. • If start conditions are not met, run the task anyway after – Specify the period after which the task will run, regardless of the other start conditions. <hr/> <p>Note Start conditions are not supported for Linux.</p>

- Select the **Scan only new and changed files** check box if you want to scan only newly created and modified files.

Default setting: **Enabled**

- Two additional options displayed for **Custom scan** for **Full scan** only:

1. **Scan archive files**

Default setting: **Enabled**.

Max recursion depth

Default setting: **16**

How many levels of embedded archives can be scanned. For example, MIME document > ZIP archive > Office archive > document content.

Max size

Default setting: **100**

Maximum size of an archive file to be scanned.

2. **Scan removable drives**

Default setting: **Disabled**

- **Mapped (remote) network drives**
- **USB storage devices** (such as pens and external hard drives)
- **CDs/DVDs**

Note

Scan removable drives is not supported for Linux.

Protection exclusions

Protection exclusions enable you to eliminate false positives when a trusted program is considered ransomware or malware. You can define trusted and blocked items by adding them to the protection exclusions list.

In the trusted items list, you can add files, processes and folders to consider them as safe in the system, and to prevent any future detections for these.

In the blocked items list, you can add processes and hashes. This option guarantees that those processes will be blocked, and your workload will be safe.

Protection exclusion item	Blocked	Trusted
Hash	<p>When a hash is added to the blocked list, the system will stop the process, based on the provided hash.</p> <p>For example, when you add this MD5 hash, 938c2cc0dcc05f2b68c4287040cfcf71, the process associated with this hash will be blocked.</p>	<p>When a hash is added to the trusted list, the system will know what processes have to be ignored by monitoring, based on the provided hash.</p> <p>For example, when you add this MD5 hash, 938c2cc0dcc05f2b68c4287040cfcf71, the process associated with this hash will be trusted and excluded from monitoring.</p>
Process	<p>When a process is added to the blocked list, the system will know that those processes must to be monitored, and the processes will always be blocked.</p> <p>For example, if you add this path C:\Users\user1\application\npp\Installer.exe to the blocked list, this specific process will be blocked, and when you will try to open it, it will not be allowed to start.</p>	<p>When a process is added to the trusted list, the system will know that those processes have to be excluded from monitoring.</p> <hr/> <p>Note Processes signed by Microsoft are always trusted.</p> <hr/> <p>For example, if you add this path C:\Users\user1\application\npp\Installer.exe, this specific process will be excluded from monitoring, and</p>

Protection exclusion item	Blocked	Trusted
		antivirus will not interfere with this process.
File/folder		When a file or a folder is added to the trusted list, the system will know that those files or folders should always be considered safe, and there is no need for those to be scanned/monitored.

To specify the items that will always be trusted

1. Open the protection plan.
2. Expand the **Antivirus and Antimalware protection** module.
3. Select the **Exclusions** option.
The **Protection exclusions** window opens.
4. In the **Trusted items** section, click **Add** to select from the available options:
 - To trust files, folders, or processes, select the **File/folder/process** option. The **Add file/folder/process** window opens.
 - In the **File/process/folder** field, enter the path for each process, folder, or file on a new line. In the **Description** section, enter a short description so that you can recognize your change in the list of trusted items.
 - Select the **Add as file/folder** checkbox to trust the file/folder.
Examples of folder description: D:\folder\, /home/Folder/folder2, F:\
 - Select the **Add as process** checkbox to trust a process. The selected processes will be excluded from monitoring.

Note

Specify the full path to the process executable, starting with the drive letter. For example, C:\Windows\Temp\er76s7sdkh.exe.

Note

Local network paths are supported. e.g: \\localhost\folderpath\file.exe

- Select the **Hash** option to add MD5 hashes to the list of trusted items. The **Add hash** window opens.
 - Here you can insert MD5 hashes on separate lines to be included as trusted in the Protection exclusions list. Based on these hashes, Cyber Protection will exclude the processes described by the MD5 hashes from being monitored.

Default setting: No exclusions are defined by default.

To specify the items that will always be blocked

1. Open the protection plan.
2. Expand the **Antivirus and Antimalware protection** module.
3. Select the **Protection exclusions** option. The **Protection exclusions** window opens.
In the **Blocked items** section, click **Add** to select from the available options:
 - To block processes, select the **Process** option. The **Add process** window opens.
 - In the **Process** field, enter the path for each process on a new line. In the **Description** field, enter a short description so that you can recognize your change in the list of blocked items.

Note

These processes will not be able to start as long as Active Protection is enabled on the machine.

- To block hashes, select the **Hash** option. The **Add hash** window is displayed.
 - In the **Hash** field, enter the hash for each process on a new line. In the **Description** field, enter a short description so that you can recognize your change in the list of blocked items.

Default setting: No exclusions are defined by default.

Wildcards

For specifying folders, you can use the wildcard characters * and ?. The asterisk (*) substitutes for zero or more characters. The question mark (?) substitutes for exactly one character. Environment variables, such as %AppData%, cannot be used.

You can use a wildcard (*) to add items to the exclusion lists.

- Wildcards can be used in the middle or the end of a description.

Examples of accepted wildcards in descriptions:

C:*.pdf

D:\folders\file.*

C:\Users*\AppData\Roaming

- Wildcards cannot be used at the beginning of a description.

Examples of unaccepted wildcards in descriptions:

*.docx

*:\folder\

Variables

You can also use variables to add items to the Protection exclusions list, with the following limitations:

- For Windows, only SYSTEM variables are supported. User specific variables, for example, %USERNAME%, %APPDATA% are not supported. Variables with {username} are not supported.

For more information, see <https://ss64.com/nt/syntax-variables.html>.

- For macOS, environment variables are not supported.
- For Linux, environment variables are not supported.

Examples of supported formats:

- %WINDIR%\Media
- %public%
- %CommonProgramFiles%\Acronis\

Description

You can use the **Description** field to make notes on the exclusions that you added in the protection exclusions list. Some suggestions on the notes you may make:

- Reasons and purposes for the exclusion.
- Actual file name of a hash exclusion.
- Time stamps.

If there are multiple items added in a single entry, there can only be 1 comment captured for the multiple items.

Active Protection in the Cyber Backup Standard edition

In Cyber Backup Standard edition, Active Protection is a separate module in the protection plan. Thus, it can be configured separately and applied to different devices or group of devices.

In all other editions of the Cyber Protection service, Active Protection is part of the **Antivirus & Antimalware** module of the protection plan.

Default setting: **Enabled**.

Note

A protection agent must be installed on the protected machine. For more information about the supported operating systems and features, see "Supported platforms" (p. 758).

How it works

Active Protection monitors processes running on the protected machine. When a third-party process tries to encrypt files or mine cryptocurrency, Active Protection generates an alert and performs additional actions, as specified in the protection plan.

In addition, Active Protection prevents unauthorized changes to the backup software's own processes, registry records, executable and configuration files, and backups located in local folders.

To identify malicious processes, Active Protection uses behavioral heuristics. Active Protection compares the chain of actions performed by a process with the chains of events recorded in the database of malicious behavior patterns. This approach enables Active Protection to detect new malware by its typical behavior.

Active Protection settings in Cyber Backup Standard

In the Cyber Backup Standard edition, you can configure the following Active Protection features:

- [Action on detection](#)
- [Self-protection](#)
- [Network folder protection](#)
- [Server-side protection](#)
- [Cryptomining process detection](#)
- [Exclusions](#)

Note

Active Protection for Linux supports the following settings: Action on detection, Network folder protection, and Exclusions. Network folder protection is always on and not configurable.

Action on detection

In the **Action on detection** section, select one of the available options:

- **Notify only**
The software will generate an alert about the process suspected of ransomware activity.
- **Stop the process**
The software will generate an alert and stop the process suspected of ransomware activity.
- **Revert using cache**
The software will generate an alert, stop the process, and revert the file changes by using the service cache.

Default setting: **Revert using cache**.

Self-protection prevents unauthorized changes to the software's own processes, registry records, executable and configuration files, and backups located in your local folders.

Administrators can enable **Self-protection**, without enabling **Active Protection**.

Default setting: **Off**.

Note

Self-protection is not supported for Linux.

To enable Self-protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Self-protection**.
3. Use the **Self-protection** toggle to enable it.

To enable Password protection

1. Once the **Self-protection** feature is enabled, you can enable the **Password protection** feature by using the toggle.
2. Click **Generate new password** to generate a password to modify or delete local agents.
3. Click **Copy**, and then paste it in a safe place because this will be requested when you want to modify the components list locally.

Important

The password will not be available after you close the window. To get this password applied to devices, the protection plan settings must be saved.

4. Click **Close**.

Password protection prevents unauthorized users or software from uninstalling Agent for Windows or modifying its components. These actions are only possible with a password that an administrator can provide.

A password is never required for the following actions:

- Updating the installation by running the setup program locally
- Updating the installation by using the Cyber Protect console
- Repairing the installation

Default setting: **Disabled**

For more information about how to enable **Password protection**, refer to [Preventing unauthorized uninstallation or modification of agents](#).

Network folder protection

The **Protect network folders mapped as local drives** setting defines whether Active protection protects from local malicious processes network folders that are mapped as local drives.

This setting applies to folders shared via SMB or NFS protocols.

If a file was originally located on a mapped drive, it cannot be saved to the original location when extracted from the cache by the **Revert using cache** action. Instead, it will be saved to the folder specified in this setting. The default folder is C:\ProgramData\Acronis\Restored Network Files for Windows, and Library/Application Support/Acronis/Restored Network Files/ for macOS. If this folder does not exist, it will be created. If you want to change this path, specify a local folder.

Network folders, including folders on mapped drives, are not supported.

Default setting: **On**.

This feature defines whether Active protection protects network folders that are shared by you from the external incoming connections from other servers in the network that may potentially bring threats.

Default setting: **Off**.

Note

Server-side protection is not supported for Linux.

To set trusted connections

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Server-side protection**.
3. Use the **Server-side protection** toggle to enable it.
4. Select the **Trusted** tab.
5. In the **Trusted connections** field, click **Add** to define connections that will be allowed to modify data.
6. In the **ComputerName/Account** field, type the name of the computer and the account of the machine where the protection agent is installed. For example, MyComputer\TestUser.
7. In the **Host name** field, type the host name of the machine that is allowed to connect to the machine using the protection agent.
8. Click the check mark to the right to save the connection definition.
9. Click **Done**.

To set blocked connections

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Server-side protection**.
3. Use the **Server-side protection** toggle to enable it.
4. Select the **Blocked** tab.
5. In the **Blocked connections** field, click **Add** to define connections that will not be allowed to modify data.
6. In the **ComputerName/Account** field, type the name of the computer and the account of the machine where the protection agent is installed. For example, MyComputer\TestUser.
7. In the **Host name** field, type the host name of the machine that is allowed to connect to the machine using the protection agent.
8. Select the check box to the right to save the connection definition.
9. Click **Done**.

Cryptomining malware degrades the performance of useful applications, increases electricity bills, and may cause system crashes and even hardware damage due to abuse. The **Cryptomining process detection** feature protects your devices against cryptomining malware to prevent unsanctioned using of computer resources.

Administrators can enable **Cryptomining process detection**, without enabling **Active Protection**.
Default setting: **Enabled**.

Note

Cryptomining process detection is not supported for Linux.

To configure network folder protection

1. In the **Create protection plan** window, expand the **Antivirus & Antimalware protection** module.
2. Click **Cryptomining process detection**.
3. Use the **Detect cryptomining processes** toggle to enable or disable the feature.
4. Select what to do with processes suspected of cryptomining activities:
Default setting: **Stop the process**
 - **Notify only**—The software generates an alert.
 - **Stop the process** — The software generates an alert and stops the process.
5. Click **Done** to apply the selected options to your protection plan.

Protection exclusions enable you to eliminate false positives when a trusted program is considered ransomware or malware. You can define trusted and blocked items by adding them to the protection exclusions list.

In the trusted items list, you can add files, processes and folders to consider them as safe in the system, and to prevent any future detections for these.

In the blocked items list, you can add processes and hashes. This option guarantees that those processes will be blocked, and your workload will be safe.

Protection exclusion item	Blocked	Trusted
Hash	<p>When a hash is added to the blocked list, the system will stop the process, based on the provided hash.</p> <p>For example, when you add this MD5 hash, 938c2cc0dcc05f2b68c4287040cfcf71, the process associated with this hash will be blocked.</p>	<p>When a hash is added to the trusted list, the system will know what processes have to be ignored by monitoring, based on the provided hash.</p> <p>For example, when you add this MD5 hash, 938c2cc0dcc05f2b68c4287040cfcf71, the process associated with this hash will be trusted and excluded from monitoring.</p>
Process	<p>When a process is added to the blocked list, the system will know that</p>	<p>When a process is added to the trusted list, the system will know that those</p>

Protection exclusion item	Blocked	Trusted
	<p>those processes must to be monitored, and the processes will always be blocked.</p> <p>For example, if you add this path C:\Users\user1\application\nppInstaller.exe to the blocked list, this specific process will be blocked, and when you will try to open it, it will not be allowed to start.</p>	<p>processes have to be excluded from monitoring.</p> <hr/> <p>Note Processes signed by Microsoft are always trusted.</p> <hr/> <p>For example, if you add this path C:\Users\user1\application\nppInstaller.exe, this specific process will be excluded from monitoring, and antivirus will not interfere with this process.</p>
File/folder		<p>When a file or a folder is added to the trusted list, the system will know that those files or folders should always be considered safe, and there is no need for those to be scanned/monitored.</p>

To specify the items that will always be trusted

1. Open the protection plan.
2. Expand the **Antivirus and Antimalware protection** module.
3. Select the **Exclusions** option.
The **Protection exclusions** window opens.
4. In the **Trusted items** section, click **Add** to select from the available options:
 - To trust files, folders, or processes, select the **File/folder/process** option. The **Add file/folder/process** window opens.
 - In the **File/process/folder** field, enter the path for each process, folder, or file on a new line. In the **Description** section, enter a short description so that you can recognize your change in the list of trusted items.
 - Select the **Add as file/folder** checkbox to trust the file/folder.
Examples of folder description: D:\folder\, /home/Folder/folder2, F:\
 - Select the **Add as process** checkbox to trust a process. The selected processes will be excluded from monitoring.

Note

Specify the full path to the process executable, starting with the drive letter. For example, C:\Windows\Temp\er76s7sdkh.exe.

Note

Local network paths are supported. e.g: \\localhost\folderpath\file.exe

- Select the **Hash** option to add MD5 hashes to the list of trusted items. The **Add hash** window opens.
 - Here you can insert MD5 hashes on separate lines to be included as trusted in the Protection exclusions list. Based on these hashes, Cyber Protection will exclude the processes described by the MD5 hashes from being monitored.

Default setting: No exclusions are defined by default.

To specify the items that will always be blocked

1. Open the protection plan.
2. Expand the **Antivirus and Antimalware protection** module.
3. Select the **Protection exclusions** option. The **Protection exclusions** window opens.

In the **Blocked items** section, click **Add** to select from the available options:

 - To block processes, select the **Process** option. The **Add process** window opens.
 - In the **Process** field, enter the path for each process on a new line. In the **Description** field, enter a short description so that you can recognize your change in the list of blocked items.

Note

These processes will not be able to start as long as Active Protection is enabled on the machine.

- To block hashes, select the **Hash** option. The **Add hash** window is displayed.
 - In the **Hash** field, enter the hash for each process on a new line. In the **Description** field, enter a short description so that you can recognize your change in the list of blocked items.

Default setting: No exclusions are defined by default.

Wildcards

For specifying folders, you can use the wildcard characters * and ?. The asterisk (*) substitutes for zero or more characters. The question mark (?) substitutes for exactly one character. Environment variables, such as %AppData%, cannot be used.

You can use a wildcard (*) to add items to the exclusion lists.

- Wildcards can be used in the middle or the end of a description.

Examples of accepted wildcards in descriptions:

C:*.pdf

D:\folders\file.*

C:\Users*\AppData\Roaming

- Wildcards cannot be used at the beginning of a description.

Examples of unaccepted wildcards in descriptions:

*.docx

*:\folder\

Variables

You can also use variables to add items to the Protection exclusions list, with the following limitations:

- For Windows, only SYSTEM variables are supported. User specific variables, for example, %USERNAME%, %APPDATA% are not supported. Variables with {username} are not supported. For more information, see <https://ss64.com/nt/syntax-variables.html>.
- For macOS, environment variables are not supported.
- For Linux, environment variables are not supported.

Examples of supported formats:

- %WINDIR%\Media
- %public%
- %CommonProgramFiles%\Acronis\

Description

You can use the **Description** field to make notes on the exclusions that you added in the protection exclusions list. Some suggestions on the notes you may make:

- Reasons and purposes for the exclusion.
- Actual file name of a hash exclusion.
- Time stamps.

If there are multiple items added in a single entry, there can only be 1 comment captured for the multiple items.

URL filtering

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Malware is often distributed by malicious or infected sites and uses the so called [Drive-by download](#) method of infection.

The URL filtering functionality allows you to protect machines from threats like malware and phishing coming from the Internet. You can protect your organization by blocking user access to the websites that may have malicious content.

The URL filtering also allows you to control web usage to comply with the external regulations and internal company policies. You can configure access to the websites depending on the category they relate to. The URL filtering supports currently 44 website categories and allows to manage access to them.

Currently, the HTTP/HTTPS connections on Windows machines will be checked by the protection agent.

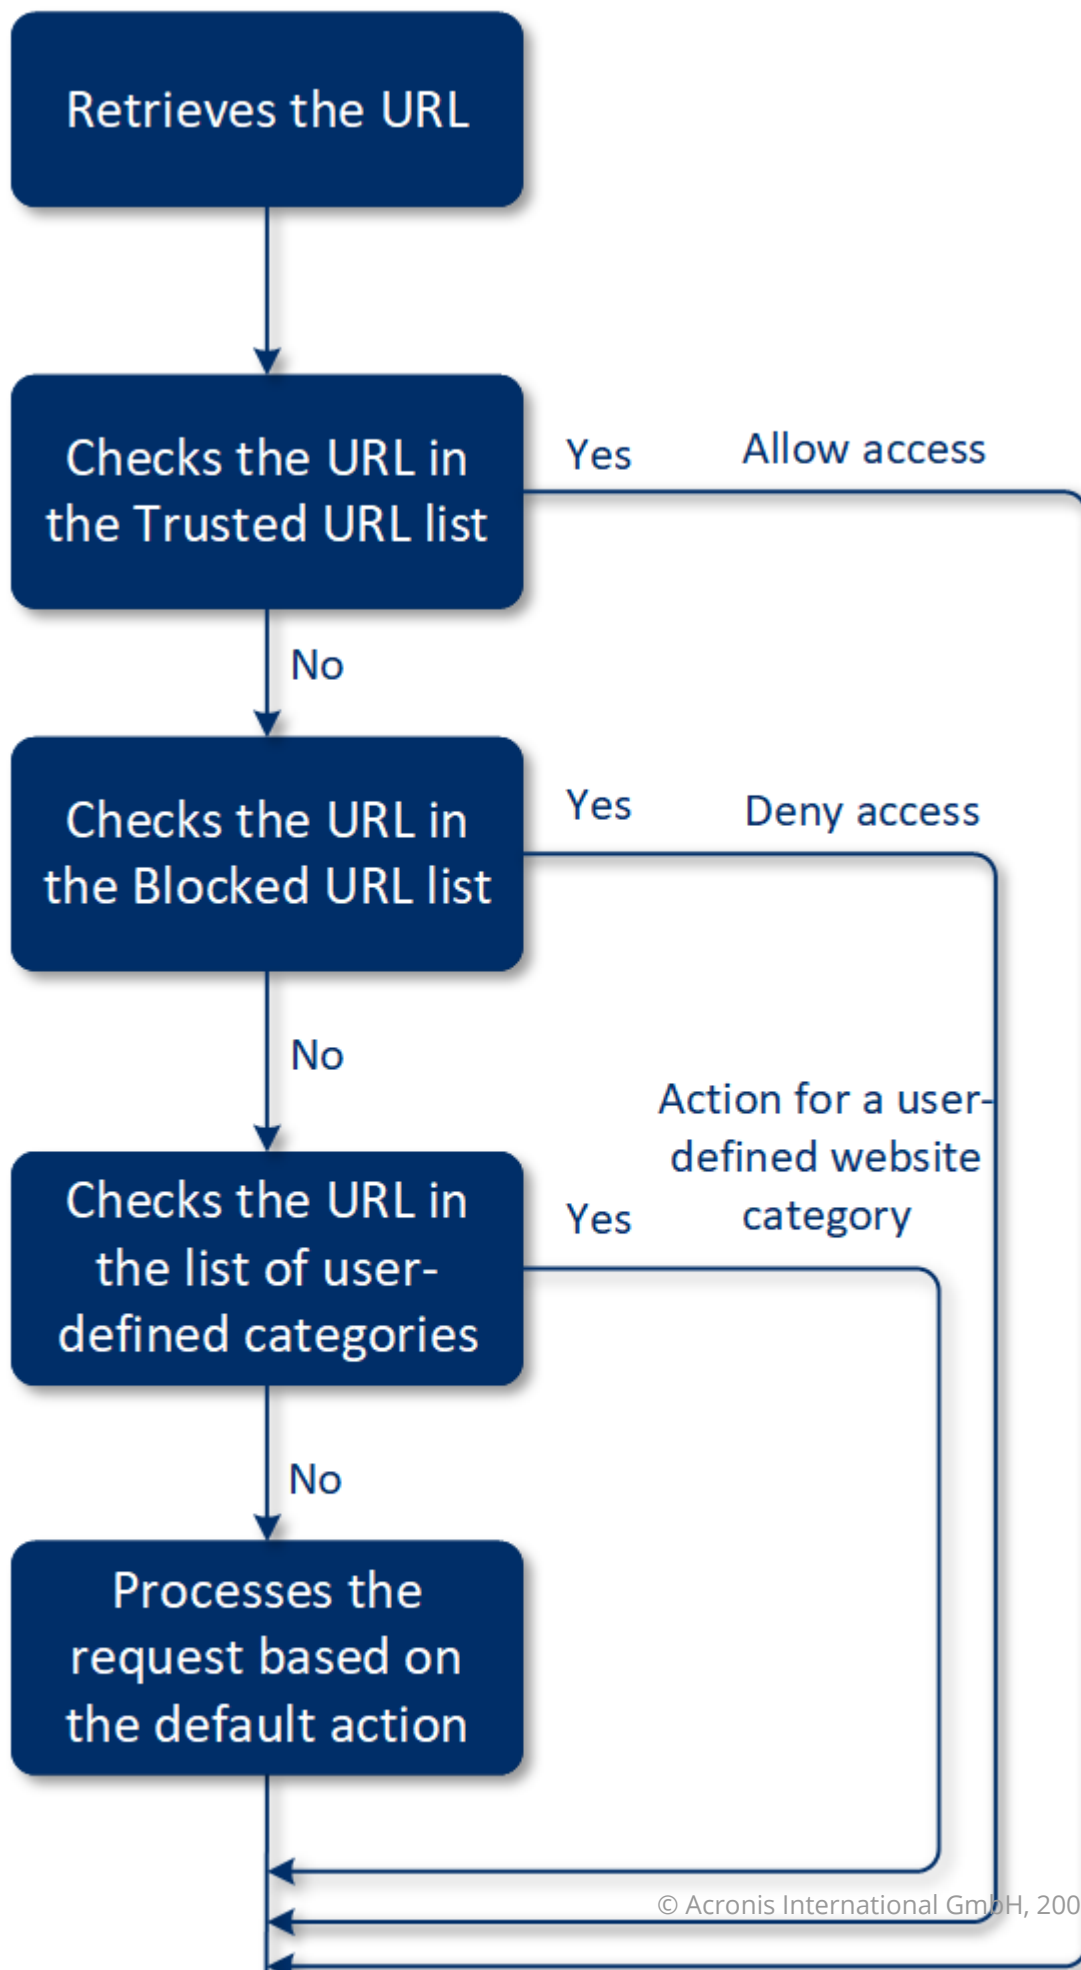
The URL filtering feature requires an internet connection to function.

Note

To prevent possible compatibility issues with protection agent builds 15.0.26692 (release C21.03 HF1) and earlier, the URL filtering functionality will be automatically disabled if another antivirus solution is detected, or if the Windows Security Center service is not present on the system. In later protection agents, the compatibility issues are resolved so URL filtering is always enabled according to the policy.

How it works

A user enters a URL link in a browser. The Interceptor gets the link and sends it to the protection agent. The agent gets the URL, parses it, and then checks the verdict. The Interceptor redirects a user to the page with the message with available actions to manually proceed to the requested page.

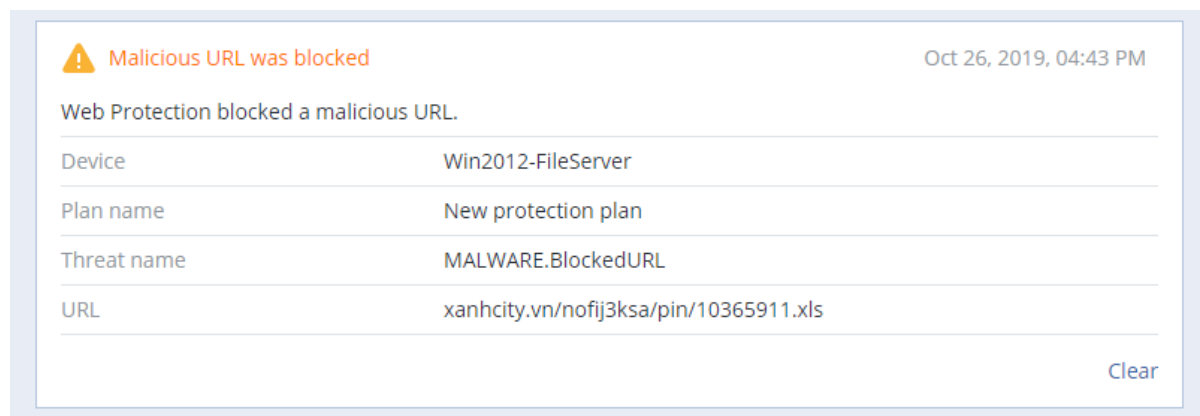


URL filtering configuration workflow

Generally, the URL filtering configuration consists of the following steps:

1. You [create a protection plan](#) with the enabled **URL filtering** module.
2. Specify the URL filtering settings (see below).
3. Assign the protection plan to the machines.

To check which URLs have been blocked, go to **Monitoring > Alerts**.



URL filtering settings

The following settings can be specified for the URL filtering module.

Malicious website access

Specify which action will be performed when a user opens a malicious website:

- **Notify only**—the software generates an alert about the process suspected of ransomware activity.
- **Block** —blocks the access to the malicious website. A user will not be able to access the website and a warning alert will be generated.
- **Always ask user**—asks the user whether to proceed to the website anyway or return back.

Categories to filter

There are 44 website categories for which you can configure access:

- **Allow** – allow access to websites related to the selected category.
- **Deny** – deny access to websites related to the selected category.

By default all categories are allowed.

Show all notifications for blocked URLs by categories – if enabled, you will get all notifications shown in the tray for blocked URLs by categories. If a website has several sub-domains, then the system also generates notifications for them, therefore the number of notifications may be big.

In the table below, you can find category descriptions:

	Website category	Description
1	Advertising	This category covers domains whose main purpose is to serve advertisements.
2	Message boards	This category covers forums, discussion boards, and question-answer type websites. This category does not cover the specific sections on company websites where customers ask questions.
3	Personal websites	This category covers personal websites, as well as all types of blogs: individual, group, and even company ones. A blog is a journal published on the World Wide Web. It consists of entries ("posts"), typically displayed in reverse chronological order so that the most recent post appears first.
4	Corporate/business websites	This is a broad category that covers corporate websites that typically do not belong to any other category.
5	Computer software	This category covers websites offering computer software, typically either open-source, freeware, or shareware. It may also cover some online software stores.
6	Medical drugs	This category covers websites related to medicine/alcohol/cigars that have discussions on the use or selling of (legal) medical drugs or paraphernalia, alcohol, or tobacco products. Note that illegal drugs are covered in the Narcotics category.
7	Education	This category covers websites belonging to official educational institutions, including those that are outside of the .edu domain. It also includes educational websites, such as an encyclopedia.
8	Entertainment	This category covers websites that provide information related to artistic activities and museums, as well as websites that review or rate content such as movies, music, or art.
9	File sharing	This category covers file-sharing websites where a user can upload files and share them with others. It also covers torrent-sharing websites and torrent trackers.
10	Finance	This category covers websites belonging to all banks around the world that provide online access. Some credit unions and other financial institutions are covered as well. However, some local banks may be left uncovered.
11	Gambling	This category covers gambling websites. These are the "online casino" or "online lottery" type website, which typically requires payment before a user can gamble for money in online roulette, poker, blackjack, or similar games. Some of them are legitimate, meaning there is a chance to win; and some are fraudulent, meaning that there is no chance to win. It also detects "beating tips and cheats" websites that describe the ways to

		make money on gambling and online lottery websites.
12	Games	<p>This category covers websites that provide online games, typically based on Adobe Flash or Java applets. It does not matter for detection whether the game is free or requires a subscription, however, casino-style websites are detected in the Gambling category.</p> <p>This category does not cover:</p> <ul style="list-style-type: none"> • Official websites of companies that develop video games (unless they produce online games) • Discussion websites where games are discussed • Websites where non-online games can be downloaded (some of them are covered in the Illegal category) • Games that require a user to download and run an executable, like World of Warcraft; those can be prevented by different means like a firewall
13	Government	This category covers government websites, including government institutions, embassies, and office websites.
14	Hacking	This category covers websites that provide the hacking tools, articles, and discussion platforms for hackers. It also covers websites offering exploits for common platforms that facilitate Facebook or Gmail account hacking.
15	Illegal activities	<p>This category is a broad category related to hate, violence and racism, and it is intended to block the following categories of websites:</p> <ul style="list-style-type: none"> • Websites belonging to terrorist organizations • Websites with racist or xenophobic content • Websites discussing aggressive sports, and/or promoting violence
16	Health and fitness	This category covers websites associated with medical institutions, websites related to disease prevention and treatment, websites that offer information or products about weight loss, diets, steroids, anabolic or HGH products, as well as websites providing information on plastic surgery.
17	Hobbies	This category covers websites that present resources related to activities typically performed during an individual's free time, such as collecting, arts and crafts, and cycling.
18	Web hosting	This category covers free and commercial website hosting services that allow private users and organizations to create and publish web pages.
19	Illegal downloads	<p>This category covers websites related to software piracy, including:</p> <ul style="list-style-type: none"> • Peer-to-peer (BitTorrent, emule, DC++) tracker websites that are known in helping to distribute copyrighted content without the copyright holder's consent • Warez (pirated commercial software) websites and discussion boards

		<ul style="list-style-type: none"> • Websites providing users with cracks, key generators, and serial numbers to facilitate the use of software illegally <p>Some of these websites may also be detected as pornography or alcohol/cigars, since they often use porn or alcohol advertisements to earn money.</p>
20	Instant messaging	This category covers instant messaging and chat websites that allow users to chat in real-time. It will also detect yahoo.com and gmail.com since they both contain an embedded instant messenger service.
21	Jobs/employment	This category covers websites presenting job boards, job-related classified advertisements, and career opportunities, as well as aggregators of such services. It does not cover recruiting agencies or the "jobs" pages on regular company websites.
22	Mature content	This category covers the content that was labeled by a website creator as requiring a mature audience. It covers a wide range of websites from the Kama Sutra book and sex education websites, to hardcore pornography.
23	Narcotics	This category covers websites sharing information about recreational and illegal drugs. This category also covers websites covering development or growing drugs.
24	News	This category covers news websites that provide text and video news. It strives to cover both global and local news websites; however, some small local news websites may not be covered.
25	Online dating	<p>This category covers online dating websites – paid and free - where users can search for other people by using some criteria. They may also post their profiles to let others search them. This category includes both free and paid online dating websites.</p> <p>Because most of the popular social networks can be used as online dating websites, some popular websites like Facebook are also detected in this category. We recommend that you use this category with the Social networks category.</p>
26	Online payments	This category covers websites offering online payments or money transfers. It detects popular payment websites like PayPal or Moneybookers. It also heuristically detects the webpages on the regular websites that ask for the credit card information, allowing detection of hidden, unknown, or illegal online stores.
27	Photo sharing	This category covers photo-sharing websites whose primary purpose is to let users upload and share photos.
28	Online stores	This category covers known online stores. A website is considered an online store if it sells goods or services online.

29	Pornography	This category covers websites containing erotic content and pornography. It includes both paid and free websites. It covers websites that provide pictures, stories, and videos, and it will also detect pornographic content on mixed-content websites.
30	Portals	This category covers websites that aggregate information from multiple sources and various domains, and that usually offer features such as search engines, e-mail, news, and entertainment information.
31	Radio	This category covers websites that offer Internet music streaming services, from online radio stations to websites that provide on-demand (free or paid) audio content.
32	Religion	This category covers websites promoting religion or a sect. It also covers the discussion forums related to one or multiple religions.
33	Search engines	This category covers search engine websites, such as Google, Yahoo, and Bing.
34	Social networks	This category covers social network websites. This includes MySpace.com, Facebook.com, Bebo.com, etc. However, specialized social networks, like YouTube.com, will be listed in the Video/Photo category.
35	Sport	This category covers websites that offer sports information, news, and tutorials.
36	Suicide	This category covers websites promoting, offering, or advocating suicide. It does not cover suicide prevention clinics.
37	Tabloids	This category is mainly designed for soft pornography and celebrity gossip websites. A lot of the tabloid-style news websites may have subcategories listed here. Detection for this category is also based on heuristics.
38	Waste of time	This category covers websites where individuals tend to spend a lot of time. This can include websites from other categories such as social networks or entertainment.
39	Traveling	This category covers websites that present travel offers and travel equipment, as well as travel destination reviews and ratings.
40	Videos	This category covers websites that host various videos or photos, either uploaded by users or provided by various content providers. This includes websites like YouTube, Metacafe, Google Video, and photo websites like Picasa or Flickr. It will also detect videos embedded in other websites or blogs.
41	Violent cartoons	This category covers websites discussing, sharing, and offering violent cartoons or manga that may be inappropriate for minors due to violence, explicit language, or sexual content. This category doesn't cover the websites that offer mainstream cartoons

		such as “Tom and Jerry”.
42	Weapons	This category covers websites offering weapons for sale or exchange, manufacture, or usage. It also covers the hunting resources and the usage of air and BB guns, as well as melee weapons.
43	Email	This category covers websites that provide email functionality as a web application.
44	Web proxy	<p>This category covers websites that provide web proxy services. This is a “browser inside a browser” type website when a user opens a web page, enters the requested URL into a form, and clicks “Submit”. The web proxy site downloads the actual page and shows it inside the user browser.</p> <p>These are the following reasons this type is detected (and might need to be blocked):</p> <ul style="list-style-type: none"> • For anonymous browsing. Since requests to the destination web server are made from the proxy web server, only its IP address is visible and if the server administrators trace the user, the trace will end on web proxy – which may or may not keep logs necessary to locate the original user. • For location spoofing. User IP addresses are often used for profiling the service by the source location (some national government websites may only be available from local IP addresses), and using those services might help the user to spoof their true location. • For accessing prohibited content. If a simple URL filter is used, it will only see the web proxy URLs and not the actual servers that the user visits. • For avoiding company monitoring. A business policy might require monitoring employee Internet usage. By accessing everything through a web proxy, a user might escape monitoring that will not provide correct information. <p>Since the SDK analyzes the HTML page (if provided), and not just URLs, for some categories the SDK will still be able to detect the content. Other reasons, however, cannot be avoided just by using the SDK.</p>

URL exclusions

URLs that are known as safe can be added to the list of the trusted domain. URLs that represent a threat can be added to the list of the blocked domain.

To specify the URLs that will always be trusted or blocked

1. In the URL filtering module of a protection plan, click **URL exclusions**.

The **URL exclusions** window opens.

The following options are displayed:

Trusted items—Click **Add** to select from the available options:

- **Domain**—When you select this option, the **Add domain** window opens.
 - In the **Domain** field, enter each domain on a new line. In the **Description** field, enter a short description so that you can recognize your change in the list of trusted items.
- **Process**—When you select this option, the **Add process** window is displayed.
 - In the **Process** field, enter the path for each process on a new line. In the **Description** section, enter a short description so that you can recognize your change in the list of trusted items.

Blocked items—Click **Add**. The **Add domain** window is displayed.

In the **Domain** field, enter each domain on a new line. In the **Description** field, enter a short description so that you can recognize your change in the list of blocked items.

Note

Local network paths are supported. For example, \\localhost\folderpath\file.exe.

Description

You can use the **Description** field to make notes on the exclusions that you added in the URL exclusions list. Some suggestions on the notes you may make:

- Reasons and purposes for the exclusion.
- Time stamps.

If there are multiple items added in a single entry, there can only be 1 comment captured for the multiple items.

Microsoft Defender Antivirus and Microsoft Security Essentials

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Microsoft Defender Antivirus

Microsoft Defender Antivirus is a built-in antimalware component of Microsoft Windows that is delivered starting from Windows 8.

The Microsoft Defender Antivirus (WDA) module allows you to configure Microsoft Defender Antivirus security policy and track its status via the Cyber Protect console.

This module is applicable for the workloads on which Microsoft Defender Antivirus is installed.

Microsoft Security Essentials

Microsoft Security Essentials is a built-in antimalware component of Microsoft Windows that is delivered with Windows versions earlier than 8.

The Microsoft Security Essentials module allows you to configure Microsoft Security Essentials security policy and track its status via the Cyber Protect console.

This module is applicable for the workloads on which Microsoft Security Essentials is installed.

The settings for Microsoft Security Essentials are similar to the settings for Microsoft Defender Antivirus, but you cannot configure real-time protection, and cannot define exclusions via the Cyber Protect console.

Schedule scan

Specify the schedule for scheduled scanning.

Scan mode:

- **Full** – a full check of all files and folders additionally to the items scanned in the quick scan. It required more machine resources for execution compared to the quick scan.
- **Quick** – a quick check of the in-memory processes and folders where malware is typically found. It required less machine resources for execution.

Define the time and day of week when the scan will be performed.

Daily quick scan – define the time for the daily quick scan.

You can set the following options depending on your needs:

Start the scheduled scan when the machine is on but not in use

Check for the latest virus and spyware definitions before running a scheduled scan

Limit CPU usage during the scan to

For more details about the setting for Microsoft Defender Antivirus, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#scheduled-scans-settings>

Default actions

Define the default actions to be performed for the detected threats of different severity levels:

- **Clean** – clean up the detected malware on a workload.
- **Quarantine** – put the detected malware in the quarantine folder but do not remove it.
- **Remove** – remove the detected malware from a workload.
- **Allow** – do not remove or quarantine the detected malware.
- **User defined** – a user will be prompted to specify the action to be performed with the detected malware.
- **No action** – no actions will be taken.
- **Block** – block the detected malware.

For more details about the default actions settings for Microsoft Defender Antivirus, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#default-actions-settings>

Real-time protection

Enable **Real-time protection** to detect and stop malware from installing or running on workloads.

Scan all downloads – if selected, scanning is performed for all downloaded files and attachments.

Enable behavior monitoring – if selected, behavior monitoring will be enabled.

Scan network files – if selected, network files will be scanned.

Allow full scan on mapped network drives – if selected, mapped network drives will be fully scanned.

Allow email scanning – if enabled, the engine will parse the mailbox and mail files, according to their specific format, in order to analyze the mail bodies and attachments.

For more details about the real-time protection settings for Microsoft Defender Antivirus, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#real-time-protection-settings>

Advanced

Specify the advanced scan settings:

- **Scan archive files** – include archived files such as .zip or .rar files into scanning.
- **Scan removable drives** – scan removable drives during full scans.
- **Create a system restore point** – in some cases an important file or registry entry could be removed as "false positive", then you will be able to recover from a restore point.
- **Remove quarantined files after** – define the period after which the quarantined files will be removed.
- **Send file samples automatically when a further analysis is required:**
 - **Always prompt** – you will be asked for confirmation before file sending.
 - **Send safe samples automatically** – most samples will be sent automatically except files that may contain personal information. Such files will require additional confirmation.
 - **Send all samples automatically** – all samples will be sent automatically.
- **Disable Windows Defender Antivirus GUI** – if selected, the WDA user interface will not be available to a user. You can manage the WDA policies via Cyber Protect console.
- **MAPS (Microsoft Active Protection Service)** – online community that helps you choose how to respond to potential threats.
 - **I don't want to join MAPS** – no information will be sent to Microsoft about the software that was detected.
 - **Basic membership** – basic information will be sent to Microsoft about the software that was detected.

- **Advanced membership** – more detailed information will be sent to Microsoft about the software that was detected.

For more details, refer to <https://www.microsoft.com/security/blog/2015/01/14/maps-in-the-cloud-how-can-it-help-your-enterprise/>

For more details about the advanced settings for Microsoft Defender Antivirus, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#advanced-settings>

Exclusions

You can define the following files and folders to be excluded from scanning:

- **Processes** – any file that the defined process reads from or writes to will be excluded from scanning. You need to define a full path to the executable file of the process.
- **Files and folders** – the specified files and folders will be excluded from scanning. You need to define a full path to a folder or file, or define the file extension.

For more details about the exclusion settings for Microsoft Defender Antivirus, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#exclusion-settings>

Firewall management

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Firewall management allows you to easily configure firewall settings on protected workloads.

This functionality in Cyber Protect is provided through a built-in Microsoft Defender Firewall component of Microsoft Windows. Microsoft Defender Firewall blocks unauthorized network traffic flowing into or out of your workloads.

Firewall management is applicable for the workloads on which Microsoft Defender Firewall is installed.

Supported Windows operating systems

The following Windows operating systems are supported for the firewall management:

Windows

- Windows 8
- Windows 8.1
- Windows 10
- Windows 11

Windows Server is not supported.

Enabling and disabling firewall management

You can enable firewall management when [creating a protection plan](#). You can change an existing protection plan to enable or disable firewall management.

To enable or disable firewall management

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Do one of the following to open the protection plan panel:
 - If you are going to create a new protection plan, select a machine to protect, click **Protect**, and then click **Create plan**.
 - If you are going to change an existing protection plan, select a protected machine, click **Protect**, click the ellipsis (...) next to the name of the protection plan, and then click **Edit**.
3. In the protection plan panel, navigate to the **Firewall management** area, and enable or disable **Firewall management**.
4. Do one of the following to apply your changes:
 - If creating a protection plan, click **Create**.
 - If editing a protection plan, click **Save**.

Microsoft Defender Firewall status in the **Firewall management** area of the protection plan panel is displayed as **On** or **Off**, depending on whether you enabled or disabled the firewall management.

You might also access the protection plan panel from the [Management tab](#). However, this capability is not available in all editions of the Cyber Protection service.

Quarantine

Quarantine is a special isolated folder on a machine's hard disk where the suspicious files detected by Antivirus and Antimalware protection are placed to prevent further spread of threats.

Quarantine allows you to review suspicious and potentially dangerous files from all machines and decide whether they should be removed or restored. The quarantined files are automatically removed if the machine is removed from the system.

How do files get into the quarantine folder?

1. You configure the protection plan and define the default action for infected files – to place in Quarantine.
2. The system during the scheduled or on-access scanning detects malicious files, places them in the secure folder - Quarantine.
3. The system updates the quarantine list on machines.
4. Files are automatically cleaned up from the quarantine folder after the time period defined in the **Remove quarantined files after** setting in the protection plan.

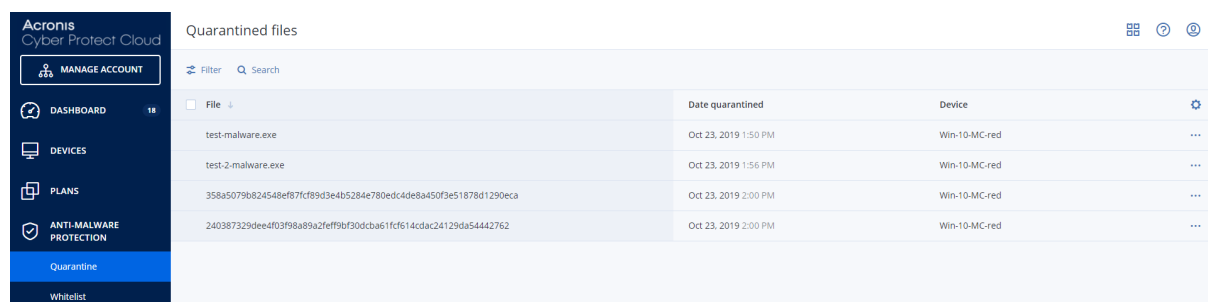
Managing quarantined files

To manage the quarantined files, go to **Antimalware protection > Quarantine**. You will see a list with quarantined files from all machines.

Name	Description
File	The file name.
Date quarantined	The date and time when the file was placed in Quarantine.
Device	The device on which the infected file was found.
Threat name	The threat name.
Protection plan	The protection plan according to which the suspicious file was placed in Quarantine.

You have two possible actions with quarantined files:

- **Delete** – permanently remove a quarantined file from all machines. You can delete all files with the same file hash. You can restore all files with the same file hash. Group the files by hash, select needed files and then delete them.
- **Restore** – restore a quarantined file to the original location without any modifications. If currently there is a file with the same name in the original location, then it will be overwritten with the restored file. Note that the restored file will be added to the allowlist and skipped during further antimalware scans.



File	Date quarantined	Device
test-malware.exe	Oct 23, 2019 1:50 PM	Win-10-MC-red
test-2-malware.exe	Oct 23, 2019 1:56 PM	Win-10-MC-red
358a50796824548e871cf89d3e4b5284e780edc4de8a450f3e51878d1290eca	Oct 23, 2019 2:00 PM	Win-10-MC-red
240387329dee4f03f98a89a2ff9bf90dcb611cf614cdac24129da54442762	Oct 23, 2019 2:00 PM	Win-10-MC-red

Quarantine location on machines

The default location for quarantined files is:

- For a Windows machine: %programdata%\Acronis\NGMP\quarantine
- For a Mac machine: /Library/Application Support/Acronis/NGMP/quarantine
- For a Linux machine: /var/lib/Acronis/NGMP/quarantine

The quarantine storage is under the service provider's self-defense protection.

Self-service custom folder on-demand

You can select custom folders on the workload and scan them directly from the context menu.

To access the Scan with Cyber Protect option in the context menu

For workloads with Antivirus and Antimalware enabled in the protection plan, right-click the files/folders on which you want to scan.

Note

This option is available only to administrators of the workload.

Corporate whitelist

An antivirus solution might identify legitimate corporate-specific applications as suspicious. To prevent these false positives detections, the trusted applications are manually added to a whitelist, which is time consuming.

Note

Corporate whitelist does not affect antimalware scans of backups.

Cyber Protection can automate this process: backups are scanned by the Antivirus and Antimalware protection module and the scanned data are analyzed, so that such applications are moved to the whitelist, and false positive detections are prevented. Also, the company-wide whitelist improves the further antimalware scanning performance.

The whitelist is created for each customer, and is based only on this customer's data.

The whitelist can be enabled and disabled. When it is disabled, the files added to it are temporarily hidden.

Note

Only accounts with the administrator role (for example, Cyber Protection administrator; company administrator; partner administrator who acts on behalf of a company administrator; unit administrator) can configure and manage the whitelist. This functionality is not available for a read-only administrator account or a user account.

Automatic adding to the whitelist

1. Run a cloud scanning of backups on at least two machines. You can do this by using the [backup scanning plans](#).
2. In the whitelist settings, enable the **Automatic generation of whitelist** switch.

Manual adding to the whitelist

Even when the **Automatic generation of whitelist** switch is disabled, you can add files to the whitelist manually.

1. In the Cyber Protect console, go to **Antimalware protection > Whitelist**.
2. Click **Add file**.
3. Specify the path to the file, and then click **Add**.

Adding quarantined files to the whitelist

You can add files that are quarantined to the whitelist.

1. In the Cyber Protect console, go to **Antimalware protection > Quarantine**.
2. Select a quarantined file, and then click **Add to whitelist**.

Whitelist settings

When you enable the **Automatic generation of whitelist** switch, you must specify one of the following levels of heuristic protection:

- **Low**
Corporate applications will be added to the whitelist only after a significant amount of time and checks. Such applications are more trusted. However, this approach increases the possibility of false positive detections. The criteria to consider a file as clean and trusted are high.
- **Default**
Corporate applications will be added to the whitelist according to the recommended protection level, to reduce possible false positive detections. The criteria to consider a file as clean and trusted are medium.
- **High**
Corporate applications will be added to the whitelist faster, to reduce possible false positive detections. However, this does not guarantee that the software is clean, and it might later be recognized as suspicious or malware. The criteria to consider a file as clean and trusted are low.

Viewing details about items in the whitelist

You can click an item in the whitelist to view more information about it and to analyze it online.

If you are unsure about an item that you added, you can check it in the VirusTotal analyzer. When you click **Check on VirusTotal**, the site analyzes suspicious files and URLs to detect types of malware by using the file hash of the item that you added. You can view the hash in the **File hash (MD5)** string.

The **Machines** value represents the number of machines where such hash was found during backup scanning. This value is populated only if an item came from Backup scanning or Quarantine. This field remains empty if the file has been added manually to the whitelist.

Antimalware scan of backups


With an antimalware scan of backups, you can prevent recovery of infected files by checking whether your backups are free of malware. Antimalware scans are performed by a cloud agent that resides in the Cyber Protection data center and no local computing resources are used.

To perform an antimalware scan, you need to configure a backup scanning plan. For more information about how to do this, refer to "Backup scanning plans" (p. 192).

Every backup scanning plan creates a scanning task for the cloud agent and adds this task to a queue, which is one per data center. Scanning tasks are processed according to their order in the queue. Also, the scanning time depends on the backup size. That is why there is a delay between creating a backup scanning plan and completing the scan.

The backups that you selected for scanning can be in one of the following states:

- Not scanned
- No malware
- Malware detected



You can check the results of a backup scan in the **Backup scanning details (threats)** widget. You can find it in the Cyber Protect console, on the **Monitoring > Overview** tab.

Limitations


- Antimalware scan is supported for **Entire machine** or **Disks/volumes** backups of the following workloads:
 - Windows machines on which a protection agent is installed.
 - Windows virtual machines that are backed up at the hypervisor level (agentless backup) by Agent for Hyper-V and Agent for VMware (Windows).
- Antimalware scan is not supported for backups created by virtual appliances, such as Agent for VMware (Virtual appliance), Agent for Virtuozzo, Agent for Scale Computing HC3.
- Only volumes with the NTFS file system, and GPT or MBR partitioning are scanned.
 - Only the default cloud storage is supported as backup location. Local storages and partner-owned cloud storages are not supported.
 - When you select backups to scan, you can select backup sets that include a Continuous data protection (CDP) backup. However, only non-CDP backups in these backup sets will be scanned.

For more information about the CDP backups, refer to "Continuous data protection (CDP)" (p. 386).

- When you perform safe recovery of an entire machine, you can select a backup set that includes a CDP backup. However, this recovery operation will not use the data in the CDP backup. To recover the CDP data, run an additional **Files/folders** recovery operation.

Working with Advanced protection features

By default, Cyber Protect includes features that cover most of the cyber security threats. You can use these features without an additional fee. In addition, you can enable advanced features to boost the protection of your workloads.

- If an Advanced protection feature is available for you to use, it appears in the protection plan marked with the Advanced feature icon .
- If an Advanced protection feature is not available for you, contact your administrator to enable the required Advanced protection pack.
- If the administrator enabled you to buy extra security packs, you can select to enable the Advanced features. A message will prompt you to a screen to inform you that extra billing applies.

Note

If at least one feature is enabled, you will have to purchase the corresponding Advanced protection pack.

Note

If all the Advanced features are disabled on your protection plan, the corresponding Advanced protection pack will be disabled.

Advanced protection pack	Advanced protection features
Advanced Backup	Protects your workloads continuously and ensures that even last-minute changes of your work will not be lost. Features include: <ul style="list-style-type: none">• One-click recovery• Continuous data protection• Backup support for Microsoft SQL Server clusters and Microsoft Exchange clusters – Always On Availability Groups (AAG) and Database Availability Groups (DAG)• Backup support for MariaDB, MySQL, Oracle DB, and SAP HANA• Data protection map and compliance reporting• Off-host data processing• Backup frequency for Microsoft 365 and Google Workspace workloads• Remote operations with bootable media• Direct backup to Microsoft Azure public cloud storage
Advanced Security + EDR	Protects your workloads continuously from all malware threats. Features include: <ul style="list-style-type: none">• Manage incidents in a centralized Incident page• Visualize the scope and impact of incidents• Recommendations and remediation steps• Check for publicly disclosed attacks on your workloads using Threat feeds

	<ul style="list-style-type: none"> • Store security events for 180 days • Antivirus and antimalware protection with local signature-based detection (with real-time protection) • Exploit prevention • URL filtering • Endpoint firewall management • Forensic backup, scan backups for malware, safe recovery, corporate allowlist • Smart protection plans (integration with CPOC alerts) • Centralized backup scanning for malware • Remote wipe • Microsoft Defender Antivirus • Microsoft Security Essentials
Advanced Management	<p>Allows you to patch vulnerabilities on the protected workloads. Features include:</p> <ul style="list-style-type: none"> • Patch management • Disk health • Software inventory • Fail-safe patching • Cyber Scripting • Remote assistance • File transfer and sharing • Selecting a session to connect • Observing workloads in multi-view • Connection modes: control, view-only, and curtain • Connection via the Quick Assist application • Remote connection protocols: NEAR and Apple Screen Sharing • Session recording for NEAR connections • Screenshot transmission • Session history report • 24 monitors • Threshold-based monitoring • Anomaly-based monitoring
Advanced Data Loss Prevention	<p>Prevents leakage of sensitive information from the protected workloads. Features include:</p> <ul style="list-style-type: none"> • Content-aware prevention of data loss from workloads via peripheral devices and network communication • Pre-built automatic detection of personally identifiable information (PII), protected health information (PHI), and Payment Card Industry Data Security Standard (PCI DSS) data, as well as documents in the “Marked as Confidential” category • Automatic data loss prevention policy creation with optional end user assistance • Adaptive data loss prevention enforcement with automatic learning-based policy adjustment • Cloud-based centralized audit logging, alerting, and end user notifications

Advanced Data Loss Prevention

The Advanced Data Loss Prevention module analyzes the content and context of data transfers on protected workloads and prevents leakage of sensitive data through peripheral devices or network transfers within and outside the company network based on data flow policy.

Advanced Data Loss Prevention features can be included in any protection plan for a customer tenant if the Protection service and the Advanced Data Loss Prevention pack are enabled for this customer.

Before you start using the Advanced Data Loss Prevention module, verify that you read and understand the basic concepts and logic of Advanced DLP management that are described in the [Fundamentals guide](#).

You might also want to review the [Technical Specifications](#) document.

Creating the data flow policy and policy rules

The key principle of data loss prevention demands that users of a corporate IT system should be allowed to handle sensitive data only to the extent necessary to perform their job duties. Any other sensitive data transfers - irrelevant to the business processes - should be blocked. Therefore it is crucial to distinguish between business-related and rogue data transfers, or flows.

The data flow policy contains rules that specify which data flows are allowed and which are prohibited, thus preventing unauthorized transfers of sensitive information when the Data Loss Prevention module is enabled in a protection plan and running in Enforcement mode.

Each sensitivity category in the policy contains one default rule, marked with an asterisk (*) and one or more explicit (non-default) rules that define the data flows for specific users or groups. Read more about the types of policy rules in the [Fundamentals guide](#).

The data flow policy is usually created automatically while Advanced Data Loss Prevention is running in observation mode. The time required for building a representative data flow policy is approximately one month, but it could differ, depending on the business processes in your organization. The data flow policy can also be created, configured, or edited manually by a company or unit administrator.

To start the automatic creation of data flow policy

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Management > Protection plans**.
3. Click **Create plan**.
4. Expand the **Data Loss Prevention** section and click the **Mode** row.

- In the Mode dialog, select **Observation mode**, and select how to process data transfers:

Option	Description
Allow all	All transfers of sensitive data from user workloads are treated as necessary for the business process and safe. A new rule is created for every detected data flow that does not match an already defined rule in the policy.
Justify all	All transfers of sensitive data from user workloads are treated as necessary for the business process, but risky. Therefore, for every intercepted transfer of sensitive data to any recipient or destination both inside and outside the organization that does not match a previously created data flow rule, the user must provide a one-time business justification. When the justification is submitted, a new data flow rule is created in the data flow policy.
Mixed	The Allow all logic is applied for all internal sensitive data flows, and the Justify all logic is applied for all external data flows. Note For more information about internal and external data see Automated detection of destination

- Save the protection plan and apply it to the workloads from which you want to collect data to build the policy.

Note

Data leakage is not prevented during observation mode.

To configure the data flow policy manually

- In the Cyber Protect console, navigate to **Protection > Data flow policy**.
- Click **New data flow rule**.
The New data flow rule pane expands on the right.
- Select a sensitivity category, add a sender and a recipient, and define the permission for data transfers for the selected category, sender, and recipient.

Option	Description
Allow	Allow this sender to transfer data of this sensitivity category to this recipient.
Exception	Do not allow this sender to transfer data of this sensitivity category to this recipient, but allow the sender to submit an exception to the rule for a specific transfer. When this sender tries to transfer data of this sensitivity category to this recipient, block the transfer and ask the sender to submit an exception to allow this transfer. When the exception is submitted, the data transfer is allowed to proceed. Important All subsequent data transfers between this sender and recipient for this sensitivity category will be allowed for five minutes after the exception is submitted.

Option	Description
Deny	Do not allow this sender to transfer data of this sensitivity category to this recipient, and do not allow the sender to request an exception to the rule.

4. (Optional) Select an action that should be executed when the rule is triggered.

Action	Description
Write in log	Store an event record in the audit log when the rule is triggered. We recommend to select this action for rules with Exception permission.
Generate an alert	Generate an alert in the Cyber Protect Alerts tab when the rule is triggered. If notifications are enabled for the administrator, an email notification will be sent as well.
Notify the end user when a data transfer is denied	Notify the user in real time with an on-screen warning when they trigger the rule.

5. Click **Save**.
6. Repeat steps 2 to 5 to create multiple rules of different sensitivity categories and options, and verify that the resulting rules correspond to the options that you selected.

Data flow policy structure

In the **Data flow policy** view, policy rules are grouped according to the category of sensitive data that they control. The sensitivity category identifier is displayed right above the group of policy rules.

- Sensitive
 - Protected Health Information (PHI)
 - Personally Identifiable Information (PII)
 - Payment Card Industry Data Security Standard (PCI DSS),
 - Marked as Confidential
- Non-sensitive

For more information on the data flow policy concept and features, see the [Fundamentals guide](#).

Rule structure

Each policy rule consists of the following elements.

- **Sensitivity Category**
 - **Protected Health Information (PHI)**
 - **Personally Identifiable Information (PII)**
 - **Payment Card Industry Data Security Standard (PCI DSS)**
 - **Marked as Confidential**

See "Sensitive data definitions" (p. 818)

- **Sender** - specifies the initiator of a data transfer controlled by this rule. It may be a single user, a list of users, or user group.
 - **Any internal** - a user group that includes all internal users of the organization.
 - **Contact / From organization** - a Windows account in the organization, recognized by Advanced Data Loss Prevention, as well as all other accounts (including those used by third-party communication applications) that a given Windows account has used earlier.
 - **Contact / Custom identity** - identifier of an internal user specified in one of the following formats: email, Skype ID, ICQ identifier, IRC identifier, Jabber e-mail, Mail.ru Agent e-mail, Viber phone number, Zoom e-mail.
The following wild cards can be used for specifying a group of contacts:
 - * - any number of symbols
 - ? - any single symbol
- **Recipient** - specifies the destination of a data transfer controlled by this rule. It may be a single user, a list of users, or user group, as well as other types of destinations specified below.
 - **Any** - any of the recipient types supported by Advanced DLP.
 - **Contact / Any contact** - any internal or external contact.
 - **Contact / Any internal contact** - any contact of an internal user (see "Automated detection of destination" (p. 818)).
 - **Contact / Any external contact** - any contact of an external person or entity.
 - **Contact / From organization** - the same principle as described in the Sender field.
 - **Contact / Custom identity** - the same principle as described in the Sender field.
 - **File sharing services** - the identifier of a controlled file sharing service.
 - **Social network** - the identifier of a controlled social network.
 - **Host / Any host** - any computer recognized by Advanced DLP as internal or external.
 - **Host / Any internal host** - any computer recognized by Advanced DLP as internal.
 - **Host / Any external host** - any computer recognized by Advanced DLP as external.
 - **Host / Specific host** - a computer identifier specified as a host name (e.g. FQDN) or IP address (IPv4 or IPv6).
 - **Device / Any device** - any peripheral device connected to the workload.
 - **Device / External storage** - a removable storage or redirected mapped drive connected to the workload.
 - **Device / Encrypted removable** - a removable storage device encrypted with BitLocker To Go.
 - **Device / Redirected clipboard** - a redirected clipboard connected to the workload.
 - **Printers** - any local or network printer connected to the workload.
- **Permission** - a preventive control enforced over a data transfer controlled by this rule. Described in more detail in topic [Permissions in data flow policy rules](#).
- **Action** - a non-preventive action performed when this rule is triggered. By default this field is set to "No action". The options are:
 - **Write in log** - store an event record in the audit log when the rule is triggered.
 - **Notify the end user when a data transfer is denied** - notify user with a real-time onscreen

warning when they trigger the rule.

- **Generate an alert** - alert the administrator when the rule is triggered.

Warning!

When **No action** is selected and the rule is triggered:

- no event record is added to the audit log;
 - no alert is sent to the administrator;
 - no onscreen notification is displayed to the end user.
-

What triggers a policy rule?

A data transfer matches a data flow policy rule if all of the following conditions are true:

- All senders of this data transfer are listed or belong to a user group specified in the **Sender** field of the rule.
- All recipients of this data transfer are listed or belong to a user group specified in the **Recipient** field of the rule.
- The data being transferred matches the **Sensitivity category** of the rule.

Adjusting the permissions in data flow policy rules

Advanced Data Loss Prevention supports three types of permissions in data flow policy rules. The permissions are configured individually in each rule of the policy.

Allow (permissive) Data transfers that match the combination of sensitivity category, sender, and recipient defined in the rule are allowed.

Exception (prohibitive) Data transfers that match the combination of sensitivity category, sender, and recipient defined in the rule are not allowed, but the sender can submit an exception to the rule to allow a specific transfer.

Important

All subsequent data transfers between this sender and recipient for this sensitivity category will be allowed for five minutes after the exception is submitted.

Deny (prohibitive) Data transfers that match the combination of sensitivity category, sender, and recipient defined in the rule are not allowed, and the sender does not have the option to submit an exception.

In addition, a priority flag can be assigned to the **Allow** and **Exception** permissions to increase the policy management flexibility. With this setting, you can override the permissions set for specific groups in other data flow rules in the policy. You can use it to apply a group data flow rule only to some of its members. To achieve this, you must create a data flow rule for specific users that you want to exclude from the group rules, and then prioritize their permissions over the data flow restrictions configured in the rules for the group to which these users belong. For information on permission priorities when combining rules, see "Combining data flow policy rules" (p. 811).

Important

Before switching a company or unit policy from Observation to Enforcement mode, it is crucial to adjust the default rules for each sensitive data category from the permissive to a prohibitive state. Default rules are marked with an asterisk (*) in the **Data flow policy** view. Read more about the types of policy rules in the [Fundamentals guide](#).

To edit permissions in policy rules

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > Data flow policy**.
3. Select the policy rule that you wish to edit and click **Edit** above the rules list.
The **Edit data flow rule** window opens.
4. In the **Permission** section, select **Allow**, **Exception**, or **Deny**.
5. (Optional) To prioritize the **Allow** or **Exception** permission of this rule over the permissions in other rules, select the **Prioritize** check box.
You do not need to use this check box to prioritize a data flow rule over the default Any > Other rule, because it has the lowest priority in the policy by default.
For information on permission priorities when combining rules, see "Combining data flow policy rules" (p. 811).
6. (Optional) Select an action to be executed when the rule is triggered.
7. Save the changes to the policy rule.

Combining data flow policy rules

When a data transfer matches more than one rule, the permissions and actions configured for all rules are combined and applied as follows.

Permissions

If a data transfer matches more than one rule and these rules have different permissions for the same data category, the overriding rule is the one with higher priority permission, according to the following permission priority list (in descending order):

1. Exception with the **Prioritized** flag
2. Allow with the **Prioritized** flag
3. Deny
4. Exception
5. Allow

If a data transfer matches more than one rule and these rules have different permissions for different data categories, the following logic is applied for the override:

1. The most restrictive rule permission is defined for each of the sensitivity categories that the data transfer matches.
2. The most restrictive of the rule permissions defined in point 1 is enforced.

Example

A file transfer matches three rules in different sensitivity categories as follows:

Sensitivity category	Permission
PII	Allow - Prioritized
PHI	Exception - Prioritized
PCI	Deny

The permission that will be applied is Deny.

Actions

If a data transfer matches more than one rule and these rules have different options configured in the **Action** field, all configured actions in all triggered rules are performed.

Policy review and management

Before the automatically created baseline data flow policy is enforced, it has to be reviewed, validated, and approved by the client, because it is the client who inherently knows all the specifics of their business processes and can assess whether they are consistently interpreted in the baseline policy. Also, the client can identify inaccuracies, which are then fixed by the partner administrator.

During the policy review, the partner administrator presents the baseline data flow policy to the client, who reviews each data flow in the policy and validates its consistency with their business processes. The validation does not require any technical skills, because the representation of policy rules in the Cyber Protect console is intuitively clear: each rule describes who are the sender and the recipient of a sensitive data flow.

Based on client's instructions, the partner administrator manually adjusts the baseline policy by editing, deleting, and creating data flow policy rules. After client's approval, the reviewed policy is enforced on protected workloads by switching the protection plan applied to these workloads to the Enforcement mode.

Before enforcing a reviewed policy, it is important to change the **Allow** permission in all automatically created default policy rules for sensitive data categories to **Deny** or **Exception**. The **Deny** permission cannot be overridden by users, while the **Exception** permission blocks a transfer matching the rule but allows users to override the block in an emergency situation by submitting a business-related exception.

Data flow policy renewal

When the business process of the company or its unit is considerably changed, their DLP policies have to be renewed in order to make them consistent with the changes in sensitive data flows of the updated business process. A policy renewal is also required if an employee's job role is changed – in this case, the part of the unit policy used to protect employee workload has also to be renewed.

The Advanced DLP policy management workflow allows administrators to automate policy renewals for the entire company, a unit, a user, or a part of users in a unit.

Renewing the policy for a company or unit

All options of the Observation mode can be used to renew the company or unit-wide policy, as well as a part of a unit policy for one or more users in the unit.

To renew the policy for a company or unit

The renewal process consists of the following steps that must be performed by a Company administrator or a Partner who manages the company workloads.

1. Delete all non-default rules in the enforced policy.
2. To start the renewal, switch the protection plan with Advanced DLP applied to the company or unit to one of the observation mode options, depending on which one is the optimal for this particular company or unit, and then apply the plan to all workloads in the company or the unit.
3. When the renewal period ends, review the new company or unit policy with the client, adjust if necessary, and get an approval by the client.
4. Switch the protection plan applied to the company or unit workloads to an appropriate enforcement mode option, which the client considers as optimal for preventing data leakage from the unit's workloads.

Renewing the policy for one or more users in the company or unit

User-level policies can be renewed by using any option of the Observation mode, as well as the adaptive enforcement mode.

Using the Observation mode for renewing a user policy

Using the observation mode for renewing a policy for a user or a part of users in the company (or unit) has the following specifics: the data flow policy enforced for the entire company (or unit) is not enforced over user's data transfers during the renewal period. As a result, new individual rules for the user can be created during the renewal that could contradict with or match existing group rules in the enforced policy for the company (or unit). After the renewal is completed and the policy is re-enforced over the user's data transfers, whether these new individual rules created for the user will be actually applied or not to the user's data transfers depends on their priorities in comparison with other rules in the policy that these data transfers match.

To renew the policy for a user through Observation mode

The renewal process consists of the following steps that must be performed by a Company administrator or a Partner who manages the company workloads.

1. Delete all non-default rules in the policy enforced for the company (or unit) that have the user as their single sender.
2. Remove the user from the sender lists of all non-default data flow rules in the enforced policy.

3. Create a new protection plan with Advanced DLP in observation mode and apply it to the user's workload to start the renewal (observation) period.
The duration of the renewal period depends on how long it could take for the user to have performed all or 90-95% of their regular business activities that involve transferring sensitive data from their workloads.
4. When the renewal period ends, review the new rules related to this user that have been added to the enforced policy, adjust them if necessary, and get them approved by the client.
5. Switch the protection plan applied to the user's workload to the **Strict enforcement** mode or the **Adaptive enforcement** mode - depending on which option the client considers as optimal for preventing data leakage from the user's workload.
Alternatively, you can re-apply to the user's workload the protection plan applied to the company (or unit).

Using the Adaptive enforcement mode for renewing a user policy

Policy renewal for a single user or a part of all users in the company (or unit) can be performed by using the Adaptive enforcement mode of a protection plan with Advanced DLP applied to the user's workload.

Note

This policy renewal method has the following specifics: the enforced company (unit) policy rules for sender groups with the user's membership (i.e. Any internal) are also enforced over data transfers from this user during the renewal. As a result, the renewal will not create new individual rules for the user that would contradict with or match these already existing policy rules for sender groups. Which of these two methods is more effective for user policy renewals for a particular client depends on its specific IT security requirements

To renew the policy for a user through Adaptive enforcement mode

The renewal process consists of the following steps that must be performed by a Company administrator or a Partner who manages the company workloads.

1. Delete all non-default rules in the policy enforced for the company (unit) that have the user as their single sender.
2. Remove the user from the sender lists of all non-default data flow rules in the enforced policy.
3. For all default rules in the policy enforced for the company (or unit), set their permission to **Exception**, and select the **Write in log** action in the **Action** field.
4. If the protection plan currently applied to the user's workload is set to the **Strict enforcement** mode, create a new protection plan with Advanced DLP and apply it to the user's workload in the **Adaptive enforcement** mode to start the renewal period.
The duration of the renewal period depends on how long it could take for the user to have performed all or 90-95% of their regular business activities that involve transferring sensitive data from their workloads.
5. When the renewal period ends, review the new rules related to this user that have been added to the enforced policy, adjust them if necessary, and get them approved by the client.

6. Switch the protection plan applied to the user's workload to the **Strict enforcement** mode or leave it in the **Adaptive enforcement** mode - depending on which option the client considers as optimal for preventing data leakage from the user's workload.
Alternatively, you can re-apply to the user's workload the protection plan applied to the company (or unit).

Enabling Advanced Data Loss Prevention in protection plans

Advanced Data Loss Prevention features can be included in any protection plan for a customer tenant if the Protection service and the Advanced Data Loss Prevention pack are enabled for this customer.

Advanced DLP is the advanced module of the Data loss prevention feature group. The Advanced DLP features and Device control can be used independently or together (in a single protection plan, or in two plans protecting the same workload). If used together, their functional capabilities are coordinated as follows.

- Device control stops controlling user access to those local channels in which Advanced DLP inspects the content of transferred data. However, Device control retains the control over the following device types if they are configured to Read-only or Denied access:
 - Removable
 - Encrypted removable
 - Mapped drive

For example, if you have both Device control and Advanced DLP enabled in a single protection plan or in two plans protecting the same workload, and you have the Read-only access configured for USB devices in Device control, the Read-only access will be applied to all USB devices, except for the ones in the allowlist, regardless of the access settings in the Advanced DLP module. If the default, Enable access is configured in Device control, the access setting in Advanced DLP will be applied.

- User access to the following local channels and peripherals in the allowlist is enforced by Device Control:
 - Optical drives
 - Floppy drives
 - MTP-connected mobile devices
 - Bluetooth adapters
 - Windows clipboard
 - Screenshot captures
 - USB devices and device types (except for Removable storage and Encrypted)

To create a protection plan with Advanced DLP

1. Navigate to **Management > Protection plans**.
2. Click **Create plan**.
3. Expand the **Data Loss Prevention** section and click the **Mode** row.
The **Mode** dialog opens.

- To start the creation or renewal of the data flow policy, select **Observation mode** and then select how to process data transfers:

Option	Description
Allow all	All transfers of sensitive data from user workloads are treated as necessary for the business process and safe. A new rule is created for every detected data flow that does not match an already defined rule in the policy.
Justify all	All transfers of sensitive data from user workloads are treated as necessary for the business process, but risky. Therefore, for every intercepted transfer of sensitive data to any recipient or destination both inside and outside the organization that does not match a previously created data flow rule, the user must provide a one-time business justification. When the justification is submitted, a new data flow rule is created in the data flow policy.
Mixed	The Allow all logic is applied for all internal transfers of sensitive data, and the Justify all logic is applied for all external transfers of sensitive data. For definition of internal destinations, see "Automated detection of destination" (p. 818)

Warning!

- Select **Observation mode** only if you do not have a data flow policy created before or if you are renewing the policy. Before you start the policy renewal, see "Data flow policy renewal" (p. 812).
 - Data leakage is not prevented in the Observation mode. See [Observation mode](#) in the Fundamentals guide.
-

- To enforce the existing data flow policy, select **Enforcement mode**, and then select how strictly to enforce the data flow policy rules:

Option	Description
Strict enforcement	The data flow policy is enforced as is and will not be extended with new permissive policy rules when previously unobserved sensitive data flows are detected. See Strict enforcement in the Fundamentals guide.
Adaptive enforcement (Enforcement with learning)	The enforced policy continues its automatic adaptation to those business operations that were not performed during the observation period or to changes in business processes. This mode allows the enforced data flow policy to expand based on newly learned data flows detected on the workloads. See Adaptive enforcement in the Fundamentals guide.

Important

Before switching a company or unit policy from Observation to Enforcement mode, it is crucial to adjust the default rules for each sensitive data category from the permissive to a prohibitive state. Default rules are marked with an asterisk (*) in the **Data flow policy** view. Read more about the types of policy rules in the [Fundamentals guide](#).

- Click **Done** to close the Mode dialog.

5. (Optional) To configure optical character recognition, allowlists, and more protection options, click **Advanced Settings**.
For information on available options, see "Advanced settings" (p. 817).
6. Save the protection plan and apply it to the workloads that you want to protect.

Advanced settings

You can use the advanced settings in protection plans with Advanced Data Loss Prevention to increase the quality of data content inspection in channels controlled by Advanced Data Loss Prevention, as well as exclude from any preventive controls data transfers to peripheral device types in the allowlist, categories of network communications, destination hosts, as well as data transfers initiated by applications in the allowlist. You can configure the following advanced settings:

- **Optical character recognition**

This setting turns on or off optical character recognition (OCR) in order to extract pieces of text in 31 language for further content inspection from graphical files and images in documents, messages, scans, screenshots, and other objects.

- **Transfer of password-protected data**

The content of password-protected archives and documents cannot be inspected. With this setting, Advanced DLP allows the administrator to select whether outgoing transfers of password-protected data are to be allowed or blocked.

- **Prevent data transfer on errors**

Sometimes, the analysis of content that is being sent might fail or another control error might occur in DLP agent operations. If this option is enabled, the transfer will be blocked. If the option is disabled, the transfer will be allowed despite the error.

- **Allowlist for device types and network communications**

Data transfers to the types of peripheral devices and in network communications checked in this list are allowed regardless of their data sensitivity and the enforced data flow policy.

Warning!

This option is used if issues with a specific Device type or Protocol occur. Do not enable it unless advised by a Support representative.

- **Allowlist for remote hosts**

Data transfers to destination hosts specified in this list are allowed regardless of their data sensitivity and the enforced data flow policy.

- **Allowlist for applications**

Data transfers performed by applications specified in this list are allowed regardless of their data sensitivity and the enforced data flow policy.

The **Security level** indicator of Advanced settings displayed in the **Create protection plan** view and in the "Details" view of a protection plan has the following logic of level indication:

- **Basic** indicates that none of the advanced settings is turned on.
- **Moderate** indicates that one or more settings are turned on, but the combination of **OCR**, **Transfer of password-protected data**, and **Prevent data transfer on errors** is not activated.
- **Strict** indicates that at least the combination of **OCR**, **Transfer of password-protected data**, and **Prevent data transfer on errors settings** is activated.

Automated detection of destination

In Mixed Observation mode, Advanced Data Loss Prevention applies different rules depending on the destination of the detected data transfer - internal or external. The logic for determining a destination as internal is described below. All other destinations are considered external.

For each intercepted data transfer, Advanced Data Loss Prevention detects automatically if the destination HTTP, FTP, or SMB server is internal by performing a DNS request and comparing the FQDN names of the machine where the Data Loss Prevention agent runs and the remote server. If the DNS request fails, it also checks if the protected workload and the remote server are in the same network. Servers that have the same domain name (or are in the same subnetwork) as the machine where the Data Loss Prevention agent runs are considered internal.

For email communication, Advanced Data Loss Prevention treats as internal transfers all emails sent from a corporate email address by using the corporate mail server if the recipient email is on the same domain as the sender email, and the recipient mail server name is the same.

Non-corporate emails are treated as external communication unless the recipient account is known. Known email addresses are updated as Data Loss Prevention monitors the user activity on the network and updates the database at the back end with data for email addresses associated with the user.

Communications via messengers are treated as external communications unless the recipient account is known. Known accounts are updated as Data Loss Prevention monitors the user activity on the network and updates the database at the back end with data for accounts associated with the user.

Sensitive data definitions

This topic describes the logic of identifying sensitive data during content analysis.

To reduce the number of false positives, identical matches are counted as one match for all groups of the described logical expressions.

Important

The logical expressions used for content identification are provided for information only and do not describe the solution in full detail.

Protected Health Information (PHI)

Supported languages

- US, UK, English-International
- Finnish
- Italian
- French
- Polish
- Russian
- Hungarian
- Norwegian
- Spanish

Data considered Protected Health Information

The following data is considered protected health information.

- First names and last names
- Address (street, city, county, precinct, zip code, and their equivalent geocodes)
- Phone numbers
- Email addresses
- Social security numbers
- Health plan beneficiary numbers
- Bank account numbers
- URLs
- IP address numbers
- ICD-10-CM codes
- ICD-10-PCS-and-GEMs
- HIPAA
- Other health-care related
- Credit card numbers

Logical expression used for content detection

The logical expression consists of the following strings that are joined by the logical operator OR. The OR operator is used to join different data groups in the list above if the AND logical operator is not specified explicitly. The numbers in brackets represent the number of detected instances that would return a positive detection result.

- **Social Security Numbers** (5)
- (First names and Last names (3) OR Address (3) OR Phone Numbers (3) OR Email Address (3) OR Bank Account Numbers (3) OR Credit Card Numbers (3)) AND (Social security numbers

(3) OR Health plan beneficiary numbers (3) * OR ICD-10-CM codes (3) OR ICD-10-PCS-and-GEMs
(3) OR HIPAA (3) OR * Other Health-care related (3))

Personally Identifiable Information (PII)

Supported languages

- US, UK, English-International
- Bulgarian
- Chinese
- Czech
- Danish
- Dutch
- Finnish
- French
- German
- Hungarian
- Indonesian
- Italian
- Korean
- Malay
- Norwegian
- Polish
- Portuguese (Brazil)
- Portuguese (Portugal)
- Romanian
- Russian
- Serbian
- Singapore
- Spanish
- Swedish
- Taiwan
- Turkish
- Thai
- Japanese

Data considered Personally Identifiable Information (PII)

- First names and last names
- Address (street, city, county, zip code)
- Bank account numbers

- Personal and fiscal ID numbers
- Passport numbers
- Social security numbers
- Phone numbers
- Car plate numbers
- Driving license numbers
- Identifiers and serial numbers
- IP addresses
- Email addresses
- Credit card numbers

Logical expression used for content detection

Logical expression for all supported languages except Japanese

The logical expression consists of the following strings joined by the logical operator OR or AND. The numbers in brackets represent the number of detected instances that would return a positive detection result.

- Personal and fiscal ID numbers (5)
- First names and Last names (3) AND (Credit Card Number (3) OR Social Security Number (3) OR Bank Account Number (3) OR Personal and fiscal ID numbers (3) OR Driving license numbers (3) OR Passport Numbers (3) OR Social security numbers (3) OR IP Addresses (3) OR Car plate numbers (3) OR Identifiers and serial numbers)
- Phone Numbers (3) AND (Credit Card Number (3) OR Social Security Number (3) OR Bank Account Number (3) OR Address (3) OR Personal and fiscal ID numbers (3) OR Driving license numbers (3) OR Passport Numbers (3) OR Social security numbers (3) OR Car plate numbers (3) OR Identifiers and serial numbers (3))
- (First names and Last names (30) OR Address (30)) AND (Email Addresses (30) OR Phone Numbers (30) OR IP Addresses (30))
- Email Addresses (3) AND (Credit Card Number (3) OR Social Security Number (3) OR Bank Account Number (3) OR Personal and fiscal ID numbers (3) OR Driving license numbers (3) OR Passport Numbers (3) OR Social security numbers (3) OR Car plate numbers (3) OR Identifiers and serial numbers (3))
- Email Address (30) AND (Address (30) OR Phone Numbers (30))
- First names and Last names (30) AND Address (30)
- Phone Numbers (30) AND Address (30)
- First names and Last names (3) AND Bank Account Numbers (3)
- Phone Numbers (3) AND (Credit Card Number (3) OR Bank Account Number (3) OR Social security numbers (3) OR Personal and fiscal ID numbers (3) OR Driving license numbers (3) OR Passport Numbers (3))

Logical expression for Japanese

Note

Only unique matches are counted by content detection.

The logical expression consists of the following strings joined by the logical operator OR. The operator OR is used to join different groups if logical operator AND is not explicitly specified.

- Social security numbers (5)
- First names and Last names (3) AND (Credit Card Number (3) OR Bank Account Number (3) OR Driving license numbers (3) OR Passport Numbers (3) OR Social security numbers (3))
- First names and Last names (30) AND (Email Addresses (30) OR Phone Numbers (30) OR IP Addresses (30) OR Address (30))
- Address (3) AND (Credit Card Number (3) OR Bank Account Number (3) OR Driving license numbers (3) OR Passport Numbers (3) OR Social security numbers (3))
- Email Address (3) AND (Credit Card Number (3) OR Bank Account Number (3) OR Social security numbers (3) OR Driving license numbers (3))
- Address (5) AND (Email Address (5) OR First names and Last names (5) OR Phone Numbers (5) OR IP Addresses (5))
- First names and Last names (3) AND Bank Account Numbers (3)
- Phone Numbers (3) AND (Credit Card Number (3) OR Bank Account Number (3) OR Address (3) OR Social security numbers (3) OR Driving license numbers (3))

Payment Card Industry Data Security Standard (PCI DSS)

Supported languages

This sensitivity group is language - independent. The PCI DSS data is in English in all countries.

Data considered PCI DSS

- Cardholder data
 - Primary Account Number (PAN)
 - Cardholder Name
 - Expiration date
 - Service code
- Sensitive Authentication Data
 - Full track data (magnetic-stripe data or equivalent on a chip)
 - CAV2/CVC2/CVV2/CID
 - PINs/PIN blocks

Logical expression used for content detection

The logical expression consists of the following strings joined by the logical operator OR. The numbers in brackets represent the number of detected instances that would return a positive detection result.

- Credit Card Number (5)
- Credit Card Number (3) AND (American Name (Ex) (3) OR American Name (3) OR PCI DSS Keywords (3) OR Date (month/year) (3))
- Credit Card Dump (5)

Marked as Confidential

Data marked as confidential is detected through keywords group.

The Match condition is weight-based, and every word has weight == 1. The content detection is considered positive when Match if weight > 3.

Supported languages

- English
- Bulgarian
- Chinese Simplified
- Chinese Traditional
- Czech
- Danish
- Dutch
- Finnish
- French
- German
- Hungarian
- Indonesian
- Italian
- Japanese
- Korean
- Malay
- Norwegian
- Polish
- Portuguese - Brazil
- Portuguese - Portugal
- Russian
- Serbian
- Spanish

- Swedish
- Turkish

Keyword groups

The keyword group for each language contains the country-specific equivalents of the following keywords that are used for the English language (case-insensitive).

- confidential
- internal distribution
- not for distribution
- do not distribute
- not for public
- not for external distribution
- for internal use only
- highly qualified documentation
- private
- privileged information
- for internal use only
- for official use only

Data Loss Prevention events

Advanced Data Loss prevention generates events in the DLP events view as follows.

- During observation mode, events are generated for all justified data transfers.
- During enforcement mode, events are generated based on the **Write in log** action configured for each policy rule that is triggered.

To view the events for a rule in the data flow policy

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > Data flow policy**.
3. Locate the rule for which you want to view the events and click the ellipsis at the end of the rule line.
4. Select **View events**.

To view details about an event in the DLP events view

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > DLP events**.
3. Click an event in the list to view more details about it.
The Event details pane expands to the right.

4. Scroll down and up in the Event details pane to view the available information.
The details that are displayed in the pane depend on the type of rule and rule settings that triggered the event.

To filter events in the DLP events list

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > DLP events**.
3. In the upper left, click **Filter**.
4. Select sensitivity category, workload, action type, user, and channel from the drop-down menus.
You can select more than one item in the drop-down menus. Filtering applies the logical operator OR between items in the same menu, but the logical operator AND is used between items from different menus.
For example, if you select **PHI** and **PII** sensitivity category, the result will return all events that contain PHI or PII, or both. If you select sensitivity category **PHI** and action **Write access**, only events that match both categories will appear in the filtered result.
5. Click **Apply**.
6. To view all events again, click **Filter**, then **Reset to default**, and finally click **Apply**.

To search for events in the DLP events list

1. Repeat steps 1-2 from the procedure above.
2. From the drop-down list to the right of Filter, select a category in which you want to search: **Sender, Destination, Process, Message subject, or Reason**.
3. In the text box, enter the phrase you are interested in and confirm by pressing Enter on the keyboard.
Only events matching the phrase you entered appear in the list.
4. To reset the list of events, click the **X** sign in the search text box and press Enter.

To view the list of events related to specific rules in the data flow policy

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > Data flow policy**.
3. Select the check box in front of the name of the policy rule you are interested in.
You can select multiple policy rules if needed.
4. Click **View events**.
The view switches to **Protection > DLP events** and the events that are related to the policy rules that you selected appear in the list.

Advanced Data Loss Prevention widgets on the Overview dashboard

The **Overview** dashboard provides a number of customizable widgets that give an overview of operations related to the Cyber Protection service, including Advanced Data loss Prevention. You can find the following Advanced Data Loss Prevention widgets on the **Overview** dashboard under **Monitoring**.

- **Sensitive data transfers** - shows a total number of sensitive data transfer operations to internal and external recipients. The chart is divided by the type of permission: allowed, justified or blocked. You can customize this widget by selecting the desired time range (1 day, 7 days, 30 days, or this month).
- **Outbound sensitive data categories** - shows a total number of sensitive data transfers to external recipients. The chart is divided by sensitive categories: Protected Health Information (PHI), Personally Identifiable Information (PII), PCI DSS and Marked as Confidential (Confidential).
- **Top senders of outbound sensitive data** - shows a total number of sensitive data transfers from the organization to external recipients and a list of the top five users with the largest number of transfers (along with these numbers). This statistic includes both allowed and justified transfers. You can customize this widget by selecting the desired time range (1 day, 7 days, 30 days, or this month).
- **Top senders of blocked sensitive data transfers** - shows a total number of blocked sensitive data transfers and a list of the top five users with the largest number of attempted transfers (along with these numbers). You can customize this widget by selecting the desired time range (1 day, 7 days, 30 days, or this month).
- **Recent DLP events** - shows details of recent Data loss prevention events for the selected time range. You can customize this widget using the following options:
 - **Range (date posted)** (1 day, 7 days, 30 days, or this month).
 - Name of the **workload**
 - **Operation status** (allowed, justified, or blocked)
 - **Sensitivity** (PHI, PII, Confidential, PCI DSS)
 - **Destination type** (external, internal)
 - **Grouping** (workload, user, channel, destination type)

The widgets are updated every five minutes. The widgets have clickable elements that enable you to investigate and troubleshoot issues. You can download the current state of the dashboard or send it via email in the .pdf or/and .xlsx format.

Custom sensitivity categories

Custom sensitive data categories may help an organization to protect intellectual property and confidential data specific for that organization by expanding Advanced DLP built-in catalog of compliance regulatory-related content definitions.

To create custom sensitive category

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > Data Loss Prevention > Data classifiers**.
3. Select **Sensitivity category**.
4. You will see a list of sensitivities, both built-in (such as Protected health information or Personally Identifiable Information) and custom ones.
5. Click **Create Sensitivity** in the top right corner of the window.
6. Enter its name in the next window.

7. New custom sensitivities are always disabled by default. You can enable them once you configure all their parameters.
8. After creating new sensitivity, you will need to set up its content detectors. Click an arrow to expand the contents of your new sensitivity and select **Add content detector**.
9. In the next window you can either use any of the existing content detectors (by clicking the checkmark next to their name and then clicking **Add** in the lower right corner) or define a new one.
10. Instead of creating a new sensitivity from scratch you can also reuse existing one (either built-in or existing custom sensitivity) by cloning it and adjusting its parameters.
 - To clone an existing sensitivity, click a checkmark next to its name and then select **Clone** from the Action drop down menu (indicated by an ellipsis) in the top left corner. You can select multiple items at a time to clone more than one sensitivity.
 - In the next window you can select which parameters of the existing sensitivity you wish to retain by clicking the checkmarks next to each parameter.

Note

Copying of built-in sensitivities inside one tenant will create a new sensitivity that consists of same detectors (they become Custom once copied)

To create new content detector

1. Log in to the Cyber Protect console as an administrator.
2. Navigate to **Protection > Data Loss Prevention > Data classifiers**.
3. Select **Content detectors**.
4. You will see a list of content detectors, both built-in and custom ones.
5. Click **Create content detector** in the top right corner of the window.
6. A drop-down menu will open, where you can select the type of detector you want to create – at this point only **File type** content detector is available, more to come in the future updates.
7. In the following window you can configure the content detector.

Type of content detector	Description
File type content detector	a. There are two lists: Supported file types and Selected file types . By clicking a “plus” icon to the right of the supported file type you will move it to the Selected file types list. You can also select multiple supported file types by clicking on the checkmarks next to their names and then using Add selected button in the top right corner. b. To remove a file type from the Selected file types list, click on a trashcan icon to the right of its name. You can also remove multiple file types at once using checkmarks and Remove selected button.
Keywords content	a. When creating new keywords content detector, you will need to import keywords from a file. After successful importing you can either merge new keywords with the

Type of content detector	Description
detector	list of existing ones or replace the existing ones with imported keywords. b. You also need to determine if you want the content detector to match all keywords from the list, any keyword from the list or a custom number of keywords.

8. Instead of creating a new content detector from scratch you can also reuse an existing one (either built-in or existing custom sensitivity) by cloning it and adjusting its parameters.
 - To clone an existing content detector, click a checkmark next to its name and then select **Clone** from the Action drop down menu (indicated by an ellipsis) in the top left corner. You can select multiple items at a time to clone more than one content detector.

Note

Copying of built-in content detector causes the detector to become custom.

Organization map

Note

This functionality is accessible only to Company Administrator users.

The organization map is a database that contains data for users and all their accounts used to transfer data through instant messaging, email, or any other means, that have been intercepted by Advanced DLP.

The organization map provides means to create and manage user groups in Advanced DLP, and to manage users and accounts associated to users in Advanced DLP. User groups can then be used for group-based DLP policy management.

To locate the Organization map

- In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.

How does it work?

Note

The organization map is populated while the Advanced DLP module operates in Observation mode.

For every data transfer intercepted by the DLP Agent, the following attributes are collected in the back end.

Attribute	Description	Label in the UI
Organization Unit	A manually created group. The Organization unit can have one or more nested Organization Units.	Group name, as defined
Security ID	A unique security identifier.	On the user details page > SID
	A user-friendly display name derived from the account names for the user. This name is not always available in Organization map.	Name
PCUserName	The name of the user of the endpoint (workload). A user name can be assigned to only one Organization Unit.	User name
Device (Workload)	The name of the endpoint (workload).	Workload
Account	Accounts that were used by a user for communication via instant messaging and email, and have been intercepted by the DLP Agent. For example, if the agent detects that username 'PC\John' uses john@gmail.com to send an email - this account is linked to PC\John user name.	Accounts

In the organization map, you can view and search accounts, users, and groups, and create, edit, and delete groups.

To search for specific accounts

As part of incident investigation, Administrator users might need to find the owner of a specific account that was involved in a potential data breach.

1. In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.
2. In the **Search** text box above the users list, start typing or paste the account. The list is filtered as you type.

To search for a specific user name

1. In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.
2. To search in a specific group, click the group name in the list.
3. In the **Search** text box above the users list, start typing or paste a user name. The list is filtered as you type.

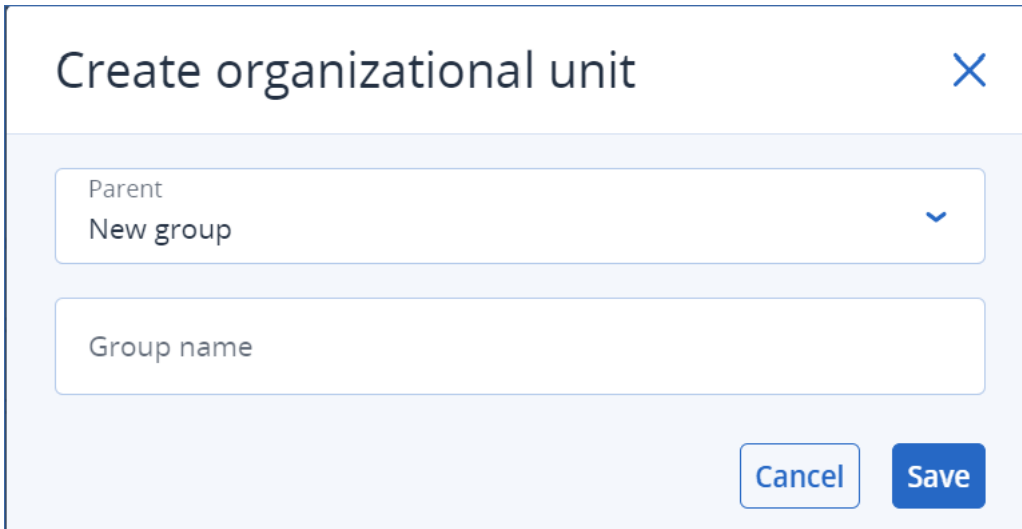
To view the accounts used by a particular user name

1. Locate the user in the users list.
2. Click the three dots at the end of the user row and select **View**.

3. In the user details dialog, locate the **Associated accounts** section.
4. You can add comments in the Description text box.

To create a user group

1. In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.
2. In the lower left section of the groups list, click **Create group**.
The Create organizational unit dialog opens.



The screenshot shows a dialog box titled "Create organizational unit" with a close button (X) in the top right corner. The dialog contains a "Parent" dropdown menu with "New group" selected, and a "Group name" text input field. At the bottom right, there are "Cancel" and "Save" buttons.

3. From the Parent drop-down menu, select the context for the new group.

Note

You cannot change the parent later. The group will remain nested in this context.

4. Enter a group name and click **Save**.

To add a user to a group

1. In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.
2. In the users list, locate the user that you want to add and select the check box in the beginning of the user row.
The **Move selected** and **Delete selected** buttons appear above the users list.
3. Click **Move selected**.
The Move user dialog opens.
4. Select a new parent for the selectet user and click **Save**.

Note

A user can belong to only one group.

To delete an account associated to a user

1. Locate the user in the users list.
2. Click the three dots at the end of the user row and select **View**.
3. In the user details dialog, locate the **Associated accounts** section.
4. Locate the account that you want to delete and click the three dots next to it.
5. From the drop-down list, select **Delete**.

To rename a user group

1. In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.
2. Click the three dots next to the name of the group and click **Rename**.

To delete a user group

1. In the Cyber Protect Cloud console, navigate to **Protection > Data loss prevention > Organization map**.
2. Click the three dots next to the name of the group and click **Delete**.
All users from the group are moved to the parent entity.

Known issues and limitations

- [DEVLOCK-1558] Acronis Cyber Protect Agents setup hangs at upgrading due to deadlock in DeviceLock agent (A timeout (600000 milliseconds) was reached...)
- [DEVLOCK-4028] There is no control for the group chats in the Zoom desktop agent.
- [DEVLOCK-3717] Automatic internal/external destination detection for HTTP within a domain with hostname access does not work correctly (detected as External).
- [DEVLOCK-2728] BSOD in process of copying many files to Removable and SMB in parallel (DeviceLockDMK.sys/DeviceLockDriver.sys).
- [DEVLOCK-4016] Friendly name and sender ID not captured for GMX Web Mail and Web.de Mail in case of draft creation.
- [DEVLOCK-4447] There is no Justification dialog for naver.com WebMail in case of draft creation.
- [DEVLOCK-1033] DeviceLockDriver: potential bugcheck DRIVER_POWER_STATE_FAILURE caused by a deadlock during IRP_MN_QUERY_DEVICE_RELATIONS processing.

Endpoint Detection and Response (EDR)

Note

This functionality is part of the Advanced Security + EDR protection pack, which in turn is part of the Cyber Protection service. Note that when you add EDR functionality to a protection plan, you may be subject to additional charges.

EDR detects suspicious activity on the workload, including attacks that have gone unnoticed. EDR then generates incidents, which provide a step-by-step overview of each attack, helping you understand how an attack happened and how to prevent it from happening again. With easy-to-

understand interpretations of each stage in the attack, the time spent on investigating attacks can be reduced to a matter of minutes.

Why you need Endpoint Detection and Response (EDR)

In today's ever-expanding world of cyber threats and malicious attacks, prevention no longer guarantees 100% protection. Some attacks will always make it through prevention layers and successfully penetrate the network. Conventional solutions can't see when this happens, leaving attackers free to dwell in your environment for days, weeks, or months.

Existing EDR solutions do help prevent these "silent failures" by finding and removing attackers quickly. However, they typically require a high level of security expertise or expensive Security Operation Center (SOC) analysts, and analysis of incidents can be extremely time-consuming.

The Acronis Advanced Security + EDR functionality overcomes these limitations by detecting attacks that have gone unnoticed, and helping you understand how an attack happened and how to prevent it from happening again. In turn, this reduces the time spent on investigating attacks.

Here's why you need EDR:

- **Full visibility:** Understand what happened and how it happened, even for attacks that have gone unnoticed. The evolution of each attack is also visually mapped out, step-by-step (from the initial point of entry to viewing the data that was targeted and/or exfiltrated), enabling you to quickly understand the scope and impact of an incident. For more information, see "How to investigate incidents in the cyber kill chain" (p. 844).
- **Minimize investigation time:** Reduce incident investigation time from hours to just a matter of minutes. EDR details each step of the attack in clear, easy-to-understand human language, in turn helping reduce the need for expensive experts or additional headcount. For more information, see "Investigating incidents" (p. 843)
- **Check for known threats on your workloads:** You can automatically search your workloads for threats from malware, vulnerabilities, and other types of global events that may affect your data protection. These threats are referred to as Incidents of Compromise (IOCs), and are based on threat data received from the Cyber Protection Operations Center (CPOC). For more information, see "Check for indicators of compromise (IOCs) from publicly known attacks on your workloads" (p. 854).
- **Respond faster to incidents:** With access to all post-breach activities and a breakdown of each step of the kill chain, you can perform a number of actions to remediate each attack point. Among other things, you can investigate using remote control and forensic backup (this feature is not available in the Early Access version), quarantine workloads, and kill malware processes. You can also recover business operations using Cyber Disaster Recovery Cloud. For more information, see "Remediating incidents" (p. 857).
- **Report on your security posture with confidence:** With EDR enabled, you can eliminate much of the insecurity and fear of the impact cyber attacks can have on your business. In addition, incident-related information is stored for 180 days, which can be used for auditing purposes.

Features

Endpoint Detection and Response (EDR) includes the following features:

- [Receive alert notifications when a breach happens](#)
- [Manage your incidents in the Incident page](#)
- [Easy to understand visualization of the attack storyline](#)
- [Recommendations and remediation steps](#)
- [Check for publicly disclosed attacks on your workloads using threat feeds](#)
- [Quick glance overview in the dashboard](#)
- [Store security events for 180 days](#)

Receive alert notifications when a breach happens

EDR provides alert notifications whenever an incident occurs. These alerts are highlighted in the main menu of the Cyber Protect console. You can then investigate an alert by clicking the **Investigate incident** button, which redirects you to the incident investigation screen (otherwise known as the cyber kill chain).

For more information, see "Reviewing incidents" (p. 837).

Manage your incidents in the Incident page

EDR enables you to manage all your incidents in the Incidents page (accessed from the Protection menu in the Cyber Protect console). The Incidents page, which can be filtered according to your requirements, ensures you can quickly and easily understand the current status of your incidents, including their severity, workload affected, and positivity level. You can also navigate directly to the cyber kill chain to view the attack storyline, node-by-node.

For more information about the Incidents page, see "Reviewing incidents" (p. 837).

Easy to understand visualization of the attack storyline

EDR provides a visual representation of an attack in an easy readable format. This ensures that even non-security personnel can digest the objectives and severity of any attack. There's really no need for a Security Operation Center (SOC) service or to hire security experts; EDR details how exactly an attack happened, including:

- How the attacker got in
- How the attacker hid their tracks
- What harm was caused
- How the attack spread

For more information, see "How to investigate incidents in the cyber kill chain" (p. 844).

Recommendations and remediation steps

EDR provides clear and easy to implement recommendations for resolving attacks on a workload. To resolve an attack quickly, click the **Remediate entire incident** button to view and follow recommended steps for mitigating the incident. These recommended steps enable you to rapidly resume operations affected by an attack. However, if you want to take more granule remediation steps, you can navigate to each node and remediate it with the relevant action.

For more information, see "Remediating incidents" (p. 857).

Check for publicly disclosed attacks on your workloads using threat feeds

EDR includes the ability to review existing, known attacks in threat feeds against your workloads. These threat feeds are automatically generated based on threat data received from the Cyber Protection Operations Center (CPOC); EDR enables you to verify whether or not a threat is impacting your workload, and then take the necessary steps to nullify the threat.

For more information, see "Check for indicators of compromise (IOCs) from publicly known attacks on your workloads" (p. 854).

Quick glance overview in the dashboard

EDR provides a range of statistics within the Cyber Protect console dashboard. You can view:

- The current threat status, including the number of incidents that need to be investigated.
- The evolution of attacks by severity, indicating possible attack campaigns.
- The efficiency rate of closing down incidents.
- The most targeted tactics used to attack your customers.
- The network status of the workload, meaning whether it is isolated or connected.

Store security events for 180 days

EDR collects workload and application events and stores them for 180 days. Events that pre-date the 180-day period are deleted (event deletion is based on age and not according to storage space).

Note that even when EDR is switched off, all previously collected events for a workload are retained, and will be available for incident investigation.

Software requirements

Endpoint Detection and Response (EDR) supports the following operating systems:

- Microsoft Windows 7 Service Pack 1 and later
- Microsoft Windows Server 2008 R2 and later

Enabling Endpoint Detection and Response (EDR) functionality

You can enable EDR in any protection plan.

To enable EDR

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Select the relevant protection plan from the displayed list, and in the right sidebar, click **Edit**. Alternatively, you can create a new protection plan and continue to the next step. For further information about working with protection plans, see "Protection plans and modules" (p. 208).
3. In the protection plan sidebar, enable the **Endpoint Detection and Response (EDR)** module by clicking the switch next to the module name.

Protection plan [↗](#) Cancel Save

Backup 🟢 >
Entire machine to Cloud storage, Monday to Friday at 11:00 PM

Endpoint Detection and Response (EDR) 🟡 🔴
Disabled

Antivirus & Antimalware protection 🟢 >
Notify only, Self-protection on

4. In the displayed dialog, click **Enable**. Note that when EDR is enabled, other protection modules are also enabled, as shown in the displayed dialog.

Endpoint Detection and Response ✕

Endpoint Detection and Response (EDR) detects suspicious or malicious activity on the workload, generating incidents upon detection. When you enable this feature, you also automatically enable the following modules:

- Antivirus & Antimalware protection
 - Real-time protection
 - Behavior engine
 - Exploit prevention
 - Active protection
 - Network folder protection
 - Cryptomining process detection
- URL filtering

Cancel Enable

Note

If any one of **Active protection**, **Behavior engine**, **Exploit prevention**, or **URL filtering** are switched **Off**, **Endpoint Detection and Response (EDR)** is also switched **Off**.

5. The **Advanced Security + EDR** pack icon, as shown below, is added to the list of protection packs required for the implementation of the protection plan, depending on additional packs you select.



How to use Endpoint Detection and Response (EDR)

EDR enables you to detect attacks that have gone unnoticed, while helping you understand how an attack happened and how to prevent it from happening again. With easy-to-understand interpretations of each stage in the attack, the time spent on investigating attacks can be reduced to a matter of minutes.

The table below describes the general workflow when working with EDR. Initially, you will review and prioritize any new incidents, investigate them further in the cyber kill chain, and then take the relevant remediation actions.

Step	How to use EDR
STEP 1: Review incidents	<p>In the EDR incident list:</p> <ul style="list-style-type: none">• Understand the security posture of an organization: how many incidents need to be investigated?• Understand which are the most critical incidents, and prioritize their investigation according to their severity.• Understand which incidents are new or ongoing.
STEP 2: Investigate incidents	<p>In the EDR cyber kill chain:</p> <ul style="list-style-type: none">• Understand the objectives of the attacker and view the attack techniques used.• Verify how likely any incident is a true malicious attack.• Verify whether or not a threat feed is impacting your workload.• See what response actions have already been applied to an incident.
STEP 3: Remediate incidents	<p>In the relevant EDR remediation sections:</p> <ul style="list-style-type: none">• Quickly and easily remediate an entire incident by applying global response actions.• Remediate individual attack points within an incident.• Apply actions to prevent the attack (or future attacks) from spreading or affecting workloads that have not yet been targeted by the attacker.

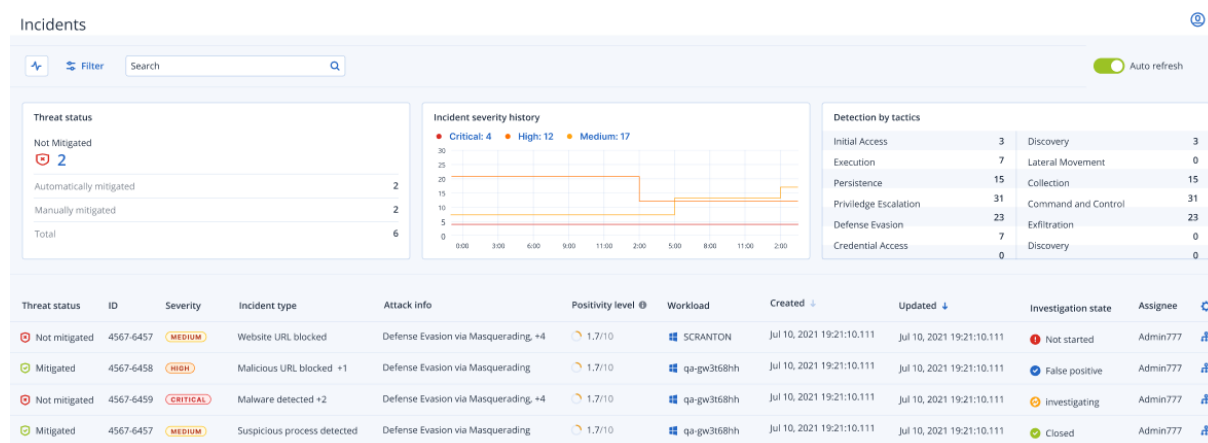
Reviewing incidents

Endpoint Detection and Response (EDR) provides an incident list that includes both prevention (or malware) and suspicious detections on a workload. The incident list gives you a quick-glance overview of any attacks or threats that are affecting your workloads, including threats that are yet to be mitigated.

From the incident list, you can quickly determine:

- The security posture of an organization: how many incidents need to be investigated?
- Which are the most critical incidents, and prioritize their investigation according to their severity.
- Which incidents are new or ongoing.

The incident list, as shown below, is accessed from the **Protection** menu in the Cyber Protect console. For further information about reviewing the incidents in the incident list, see "Viewing which incidents are currently not mitigated" (p. 839) To learn more about when an incident is created, see [What exactly are incidents?](#).



Note

The Cyber Protect console must be open in order for you to receive incident notifications.

What exactly are incidents?

Incidents, or security incidents, can be thought of as *containers* of at least one prevention or suspicious detection point (or a mix), and include all the related events and detections of a single attack. These security incidents can also include additional benign events that give further context into what happened.

This enables you to view attack events together in one single incident, and understand the logical steps that the attacker performed. In addition, it helps speed up the investigation time for an attack.

When EDR is [enabled in the protection plan](#), security incidents are created when:

- **A prevention layer stops something:** These incidents are automatically closed by the system, according to the protection plan settings. However, you can investigate what exactly the malware

did before it was stopped. For example, ransomware is stopped when it starts to encrypt files, but prior to that it could have stolen credentials or installed a service.

- **Suspicious activity is detected by EDR:** These are detections that should be investigated and remediated. By reviewing the visually enhanced cyber kill chain (for more information, see "How to investigate incidents in the cyber kill chain" (p. 844)), you can easily apply the relevant remediation actions.

Prioritize which incidents need immediate attention

The Cyber Protect console incident list can be accessed at any time from the **Protection** menu in the Cyber Protect console. The incident list gives you a quick-glance overview of any attacks or threats, enabling you to prioritize incidents that require attention.

Important

To ensure your workloads remain secure, *always* analyze and prioritize the incidents that are ongoing or not mitigated.

How to analyze which security incidents need immediate attention

The incident list enables you to analyze and prioritize the listed incidents that require attention. You can:

- **View which incidents are currently not mitigated:** Quickly understand from the incident list if any attacks are currently in progress. Any incidents that are not mitigated, as indicated in the **Threat status** column, should be looked at immediately (by default, the incident list is filtered to display these incidents).
- **Understand the scope and impact of incidents:** Based on your filtering of newly opened or ongoing attacks, understand the severity for the filtered incidents as well as the impact on your business.

Once you have a refined list of the most important incidents, you can then analyze incident details to get a better understanding of a specific incident, as well as the techniques used by the attacker to achieve their objective. For more information, see "Analyze incident details" (p. 841).

Threat status	ID	Severity	Incident type	Attack info	Positivity level	Workload	Created	Updated	Investigation state	Assignee
Not mitigated	4567-6457	MEDIUM	Website URL blocked	Defense Evasion via Masquerading, +4	1,7/10	SCRANTON	Jul 10, 2021 19:21:10.111	Jul 10, 2021 19:21:10.111	Not started	Admin777
Mitigated	4567-6458	HIGH	Malicious URL blocked +1	Defense Evasion via Masquerading	1,7/10	qa-gw3t68hh	Jul 10, 2021 19:21:10.111	Jul 10, 2021 19:21:10.111	False positive	Admin777
Not mitigated	4567-6459	CRITICAL	Malware detected +2	Defense Evasion via Masquerading, +4	1,7/10	qa-gw3t68hh	Jul 10, 2021 19:21:10.111	Jul 10, 2021 19:21:10.111	Investigating	Admin777
Mitigated	4567-6457	MEDIUM	Suspicious process detected	Defense Evasion via Masquerading	1,7/10	qa-gw3t68hh	Jul 10, 2021 19:21:10.111	Jul 10, 2021 19:21:10.111	Closed	Admin777

Note

By default, the incident list is sorted according to the **Updated** column, which details the date and time the incident was last updated with new detections recorded inside the incident. Note that any existing incident can be updated at any time, even if the incident was previously closed. You can also filter the list to show newly opened or ongoing attacks according to your requirements, as described in the procedure below.

To filter the incident list

1. At the top of the Incident list, click **Filter** to filter the displayed list of incidents. For example, if you select a start and end date in the **Created** field, the incident list and widgets display the relevant incidents created during the defined time period.

Threat status
Not Mitigated

Incident type
All

Investigation state
All

Updated
Last month

Severity
All

Attack info
All

Positivity level

– 1 + – 10 +

Clear Apply

2. When done, click **Apply**.


Viewing which incidents are currently not mitigated

You can view the current threat status for incidents in the **Threat status** column, which shows if the incident is **Mitigated** or **Not mitigated**. The threat status is automatically defined by EDR; any incident that is not mitigated should be investigated as soon as possible.

You can then refine the displayed incident list further by applying filters. For example, if you want to filter the list according to threat status *and* a specific level of severity, select the relevant filter

options. Once you have filtered the incidents that are of interest to you, you can then investigate them, as described in "Investigating incidents" (p. 843).

You can also use the **Threat status** widget, as shown below, for a quick glance overview of the current threat status. Note that the data displayed in this widget reflects the filters you applied; see "To filter the incident list" (p. 839).

Threat status	
Not Mitigated	
 2	
Automatically mitigated	2
Manually mitigated	2
Total	6

Understanding the scope and impact of incidents

You can quickly understand the scope and impact of incidents by reviewing the **Severity**, **Attack info**, and **Positivity level** columns. As mentioned above, after you have determined which incidents are currently in progress you can then filter these additional columns to do the following:

- Review which incidents are more critical in the **Severity** column. The severity of an incident can be one of **Critical**, **High**, or **Medium**.
 - **Critical**: There is a severe risk of malicious cyber activity with the risk of compromising critical hosts in your environment.
 - **High**: There is a high risk of malicious cyber activity with the risk of severe damage in your environment.
 - **Medium**: There is an increased risk of malicious cyber activity.

Note


When determining the severity, the EDR algorithm takes into consideration the workload type as well as the scope of each step of the attack. For example, an incident which includes steps related to credential theft is set to **Critical**.

- Understand why an incident was created in the **Incident type** column. The incident type can include any one or more of the following:
 - **Ransomware detected**
 - **Malware detected**
 - **Suspicious process detected**
 - **Malicious process detected**
 - **Suspicious URL blocked**
 - **Malicious URL blocked**

- Determine which attack techniques are in use in the **Attack info** column, and understand if there is a common theme or pattern to the attacks.
- Confirm how likely an incident is a true malicious attack; the **Positivity level** column includes a score of between 1-10 (the higher the score, the more likely the attack is a true malicious attack).

After you have found the incidents that need immediate attention, you can then investigate them, as described in "Investigating incidents" (p. 843)

You can also use the **Severity history** and **Detection by tactics** widgets for a quick glance overview of the severity and attack techniques.



The **Detection by tactics** widget displays the various attack techniques used, with values in green or red indicating the increase or decrease over the previous specified time range. This widget provides an aggregated view of all the objectives in the filtered incidents, giving you a quick overview of the impact on your customers.

Detection by tactics			
Initial Access	3	Discovery	3
Execution	7	Lateral Movement	0
Persistence	15	Collection	15
Privilege Escalation	31	Command and Control	31
Defense Evasion	23	Exfiltration	23
Credential Access	7	Discovery	0
Impact	0	Resouce Development	0

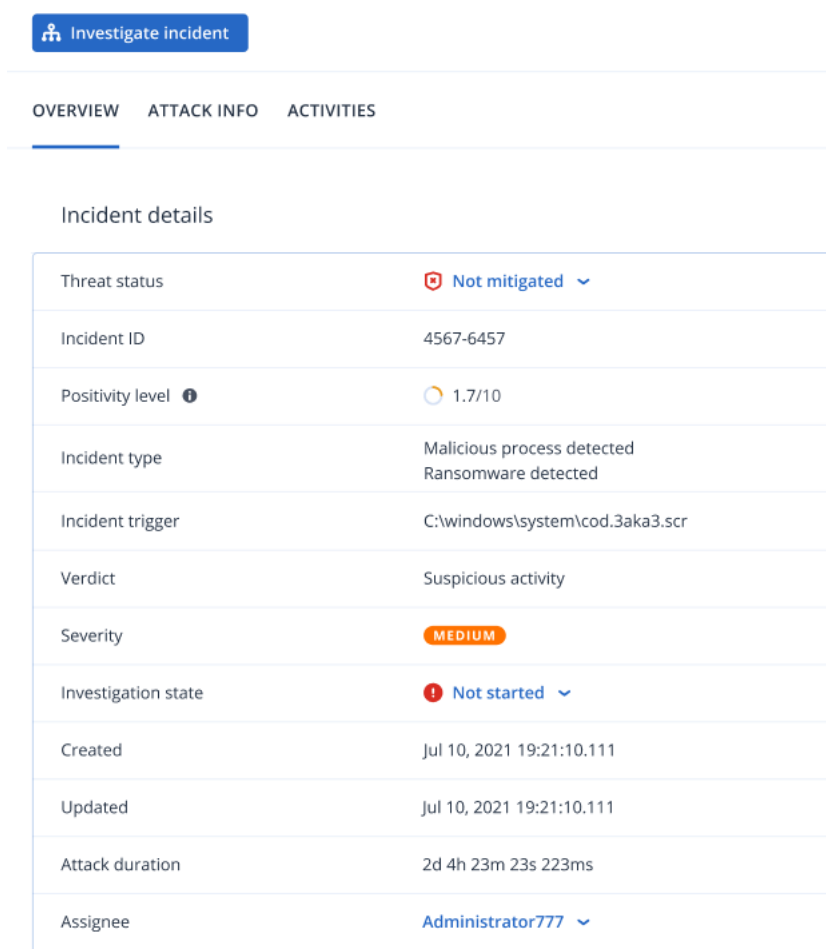
Analyze incident details

During the [incident review stage](#), you can also analyze the details of each incident from the Endpoint Detection and Response (EDR) incident list. These details enable you to drill-down into the entire









incident and understand how and when it occurred. In addition, you can assign an incident to specific users for investigation, and set the investigation status.

To analyze incident details

1. In the Cyber Protect console, go to **Protection > Incidents**. The Incident list is displayed.
2. Click on the incident you want to review. The details for the selected incident are displayed.
3. In the displayed **Overview** tab, you can review the incident and workload details, including the current threat status and severity. You can also define the **Investigation state** (select from one of **Investigating**, **Not started** (the default state), **False positive**, or **Closed**), and select a user to assign the incident to (in the **Assignee** drop-down list, select the relevant user).



The screenshot shows the 'Investigate incident' interface. At the top, there is a blue button with a magnifying glass icon and the text 'Investigate incident'. Below this, there are three tabs: 'OVERVIEW' (selected), 'ATTACK INFO', and 'ACTIVITIES'. The main content area is titled 'Incident details' and contains a table with the following information:

Threat status	 Not mitigated 
Incident ID	4567-6457
Positivity level 	 1.7/10
Incident type	Malicious process detected Ransomware detected
Incident trigger	C:\windows\system\cod.3aka3.scr
Verdict	Suspicious activity
Severity	 MEDIUM
Investigation state	 Not started 
Created	Jul 10, 2021 19:21:10.111
Updated	Jul 10, 2021 19:21:10.111
Attack duration	2d 4h 23m 23s 223ms
Assignee	Administrator777 

4. Click the **Attack Info** tab to review details of the attack and the techniques used in the attack. Click the link next to each listed attack technique to review further information about the technique on [MITRE.org](https://www.mitre.org).
5. Click the **Activities** tab to review any action taken in the cyber kill chain to mitigate an incident. For more information, see "How to investigate incidents in the cyber kill chain" (p. 844). For example, if a patch was run on the workload, you can see who initiated the patch, how long it took, and any errors that occurred during the implementation of the patch.

- Click **Investigate incident** to access the cyber kill chain where you can investigate the incident node-by-node. For more information, see "How to investigate incidents in the cyber kill chain" (p. 844).

Investigating incidents

Endpoint Detection and Response (EDR) enables you to investigate an entire incident, including all of the attack stages and objects (processes, registries, scheduled tasks, and domains) impacted by an attack. These objects are represented by nodes in the easy-to-understand cyber kill chain, as shown below. Use the cyber kill chain to quickly understand what exactly happened, and when it happened.

The screenshot displays the EDR interface for investigating an incident. At the top, there are filters for Threat status (Not mitigated), Severity (CRITICAL), Investigation state (Not started), Positivity level (10 / 10), Incident type (Malicious process detected), Created (Jan 10, 2022 12:21:10:111 AM), and Updated (Jan 10, 2022 12:21:10:111 AM). Below this is a navigation bar with 'CYBER KILL CHAIN' and 'ACTIVITIES' tabs. The main area is divided into three sections: a legend on the left, a central kill chain diagram, and a detailed view of the selected process on the right.

Legend:

- Workload: 1
- Process: 4
- File: 91
- Registry: 49
- Involved: 166
- Malicious threat: 1
- Incident trigger: 1

Attack stages:

- Persistence:**
 - Jan 10, 2022 07:11:29:530 AM: Process (Bundler.exe) is adding the file to be executed when the user logs in.

Kill Chain Diagram:

- Process: Bundler.exe (selected)
- Parent process: Postinstall.exe (Create process)
- Parent process: Conhost.exe (Create process)
- Files read by Bundler.exe:
 - Imm32.dll
 - Bundler.exe
 - SortDefault.nls
 - kernel.appcore...
 - cryptbase.dll
 - msi.dll
 - msi.dll
 - version.dll

Process Details (Bundler.exe):

- Type: Process
- Name: Bundler.exe
- PID: 9248
- State: Stopped
- Path: C:\users\autotest\appdata\localtemp\{a2ee5cde-9a05-43ee-825c-aa8485bccb88}\Bundler.exe
- Command Line: -q -burn,elevated BurnPipe,{C1191EE9-D128-4175-9E4D-EA3E88D7EF04} {A238B294-7BEE-...

Each and every step of an attack is viewed in the cyber kill chain, which provides you with a detailed interpretation of how and why the incident happened. The cyber kill chain uses easy to understand sentences and graphs that help explain each step of the attack, in turn helping to minimize investigation time.

You can quickly understand the scope and impact of an incident, with the attack evolution mapped to the [MITRE framework](#). This enables you to analyze what happened in each step of an attack, including:

- The initial point of entry
- How the attack was executed
- Any escalations of privileges
- Avoidance detection techniques
- Lateral movements to other workloads
- Credential theft
- Exfiltration attempts


Note

Each object impacted in the attack, whether it is a process, registry, scheduled task or domain, is represented by a node in the cyber kill chain.

How to investigate incidents in the cyber kill chain

You can investigate each and every step of an attack in the cyber kill chain. Follow the cyber kill chain's easy to comprehend sentences and graphs to understand each step of the attack, which in turn help you to minimize investigation time.

To begin an investigation in the cyber kill chain

1. In the Cyber Protect console, go to **Protection > Incidents**.
2. In the displayed list of incidents, click  in the far right column of the incident you want to investigate. The cyber kill chain for the selected incident is displayed.

3. View a summary of the incident in the threat status bar at the top of the page. The threat status bar includes the following information:
 - Current threat status: The threat status is automatically defined by the system. Any incident that is **Not mitigated** should be investigated as soon as possible.

Important

An incident is set to **Mitigated** when a restore from backup has been successfully completed or when all detections have been successfully remediated by a stop process, quarantine, or rollback action.

An incident is set to **Not mitigated** when a restore from backup has not been successfully completed or when at least one detection has not been successfully remediated by a stop process, quarantine, or rollback action.

You can also manually set the threat status to **Mitigated** or **Not mitigated**. When selecting either status, you are prompted to enter a comment. This comment is saved as part of the investigation activities, and can be viewed in the **Activities** tab. Note that EDR can still revert the threat status to **Mitigated** or **Not mitigated** if new detections were discovered for the incident or response actions were run and were completed successfully.

- Incident severity: **Critical, High, or Medium**. For more information, see "Reviewing incidents" (p. 837).
- Current investigation state: One of **Investigating, Not started** (the default state), **False positive**, or **Closed**. You should change the state when you start investigating the incident so that other colleagues are aware of any changes to the incident.
- Positivity level: Indicates how likely an incident is a true malicious attack, between a range of 1-10. For more information, see "Reviewing incidents" (p. 837).
- Incident type: One or more of **Ransomware detected, Malware detected, Suspicious process detected, Malicious process detected, Suspicious URL blocked, and Malicious URL blocked**.
- When the incident was created and updated: Date and time the incident was detected, or when the incident was last updated with new detections recorded inside the incident.

Threat status Not mitigated	Severity CRITICAL	Investigation state Not started	Positivity level 10 / 10	Incident type Malicious process detected	Created Jan 10, 2022 12:21:10:111 AM	Updated Jan 10, 2022 12:21:10:111 AM
--------------------------------	----------------------	------------------------------------	-----------------------------	---	---	---

4. Click the **Legend** tab to view the various nodes that make up the kill chain graph, and define which nodes to view. For further information, see "Understanding and customizing the cyber kill chain view" (p. 846).
5. Investigate and remediate the incident by performing the following steps. Note that this is the typical workflow for investigating and remediating an incident, but may vary according to each incident and your own requirements.
 - a. Investigate each stage of the attack in the **Attack stages** tab. For further information, see "How to navigate attack stages" (p. 848).
 - b. Click **Remediate entire incident** to apply remediation actions. For further information, see "Remediate an entire incident" (p. 857).

You can also remediate individual nodes in the cyber kill chain, as described in "Response actions for individual cyber kill chain nodes" (p. 862).




- c. Review actions taken to mitigate the incident in the **Activities** tab. For further information, see "Understand the actions taken to mitigate an incident" (p. 852).

Understanding and customizing the cyber kill chain view





To understand the nodes impacted in the cyber kill chain, access the legend. The legend displays all of the nodes involved in an incident, enabling you to understand how the various nodes have been impacted by the attacker. You can also define the nodes you want to hide or display in the cyber kill chain.

To access the legend






1. Click the arrow icon to the right of the Legend section.
The Legend section expands, as shown below.

CYBER KILL CHAIN	ACTIVITIES
Legend 	
 Workload	1
 Process	3
 File	51 
 Network	11 
 Registry	21 
 Involved	92 
 Malicious threat	3 
 Incident trigger	1

2. There are four main colors used in the legend, which enable you to quickly understand what happened to each node in the cyber kill chain, as shown below. These color-coded nodes are also included in the attack stages, as described in "How to navigate attack stages" (p. 848).

-  Involved
-  Suspicious activity
-  Malicious threat
-  Incident trigger

To hide or display nodes in the cyber kill chain

1. In the expanded Legend section, ensure  is displayed next to the nodes you want to display in the cyber kill chain. If the displayed icon is , click the icon to change it to .
2. To hide a node in the cyber kill chain, click . The icon changes to  and the node is not displayed in the cyber kill chain.

Investigate the attack stages of an incident

The attack stages of an incident provide easy to understand interpretations of every incident.

Each attack stage summarizes what exactly happened, and what were the objects (referred to as *nodes* in the cyber kill chain) targeted. For example, if a downloaded file was masquerading as something else, the attack stage will indicate this, and include links to the relevant node in the cyber kill chain which you can investigate, and to the relevant MITRE ATT&CK technique.

Each stage of the attack provides you with the information you need to resolve three crucial questions:

- What was the attacker's objective?
- How did the attacker achieve this objective?
- Which nodes were targeted?

More importantly, the interpretation provided ensures the time spent on investigating an incident is greatly reduced, as you no longer need to go through each security event from a timeline or graph node and then try to create an interpretation of the attack.

The attack stages also include information about compromised files that contain sensitive information, such as credit card numbers and social security numbers, as shown in the **Collection** stage in the example below.

For more information, see "What information is included in an attack stage?" (p. 848).

Attack stages

Execution ⓘ

- Jun 15, 2021, 09:38:11:374395 AM +03:00
User pbeesly, with standard privileges, on workload SCRANTON, executes a suspicious file `[?][cod.3aka3.scr]`

Defense Evasion ⓘ

- Jun 15, 2021, 09:38:11:374395 AM +03:00
To trick user pbeesly, the file was masquerading as a benign doc file, by the name `rcs.3aka.doc`

Command And Control ⓘ

- Jun 15, 2021, 09:38:11:374395 AM +03:00
To control workload SCRANTON, once `[?][cod.3aka3.scr]` is executed, a TCP connection is established on an unusual port 1234 to a unknown domain 192.168.0.5

Collection ⓘ

- Jun 15, 2021, 09:38:52:669601 AM +03:00
The adversary collects
`*.doc,*.xps,*.xls,*.ppt,*.pps,*.wps,*.wpd,*.ods,*.odt,*.lwp,*.jtd,*.p...`
files containing sensitive information credit card numbers, social security numbers and more from `$env:USERPROFILE` and compresses them into an archive `draft.zip` via a powershell script

Exfiltration ⓘ

- Jun 15, 2021, 09:39:23:725078 AM +03:00
The adversary is trying to steal data - previously created archive file `draft.zip` is exfiltrated via an existing TCP connection 192.168.0.5 established on an unusual port port:1234

How to navigate attack stages

Attack stages are listed in chronological order. Scroll down to see the complete list of attack stages for the incident.

To investigate a specific attack stage further, click anywhere in the attack stage to navigate to the relevant node in the cyber kill chain graph. For more information about navigating the cyber kill chain graph and specific nodes, see "Investigate individual nodes in the cyber kill chain" (p. 850).

What information is included in an attack stage?

Each attack stage provides an easy to understand interpretation of the attack, in easily readable human language. This interpretation is composed of a number of elements, as shown below and described in the following table.

Credential Access ⓘ

• Jun 15, 2021, 10:16:44:191934 AM +03:00

The adversary accessed credentials stored in Chrome web browser by executing a known malicious tool `chromepass.exe` masqueraded as legitimate Microsoft sysinternals tool

`accesschk.exe`

• Jun 15, 2021, 10:17:05:500810 AM +03:00

The adversary searched for private key certificate files `*.pfx` under Downloads folder by invoking malicious powershell script `C:\Program Files\SysinternalsSuite\readme.ps1` loaded previously

Attack stage element	Description
Header	<p>Describes what the attacker tried to do, and their objective (in the example above, Credential Access), with a link to a known MITRE ATT&CK technique. Click the link to learn more on the MITRE ATT&CK website.</p> <hr/> <p>Note If an attack stage is not a known MITRE ATT&CK technique, the header text won't be linked. This is relevant for generic techniques such as files detected in a random folder.</p> <hr/>
Timestamp	The time the attack stage occurred.
Technique	<p>How the attacker technically achieved their objective, and what objects (registry entries, files, or scheduled tasks) were affected.</p> <p>Included in the text description of the attack technique are color-coded links to each affected node in the cyber kill chain, as shown in the example above. These color-coded links enable you to navigate quickly to the affected node and to investigate what exactly happened. The colors used in an attack stage indicate the following:</p> <ul style="list-style-type: none">● Involved● Suspicious activity● Malicious threat★ Incident trigger <p>Looking at the above legend, we can see that</p>

Attack stage element	Description
	<p>the Credential Access example attack stage has a link to a malware node <code>accesschk.exe</code> and a suspicious file node <code>*.pfx</code> (click on the links to jump to the corresponding node in the cyber kill chain). For more information about navigating these nodes and the actions available, see "Investigate individual nodes in the cyber kill chain" (p. 850).</p> <p>Note that attack stages also include links to file nodes that have information about compromised files which contain sensitive information, such as protected health information (PHI), credit card numbers and social security numbers.</p>

Note


Each attack stage is a single detection event. The content listed in each stage (header, timestamp, technique) is generated according to specific parameters in the detection event, and which are based on attack stage templates stored by Endpoint Detection and Response (EDR).

Investigate individual nodes in the cyber kill chain

In addition to [reviewing the attack stages](#), you can also navigate through each of the attack nodes in the cyber kill chain. This enables you to drill-down to specific nodes in the cyber kill chain and to investigate and remediate each node as required.

For example, you can determine how likely an incident is a true malicious attack. Based on your investigation, you can also apply a number of response actions to the node, including isolate a workload or quarantine a suspicious file.

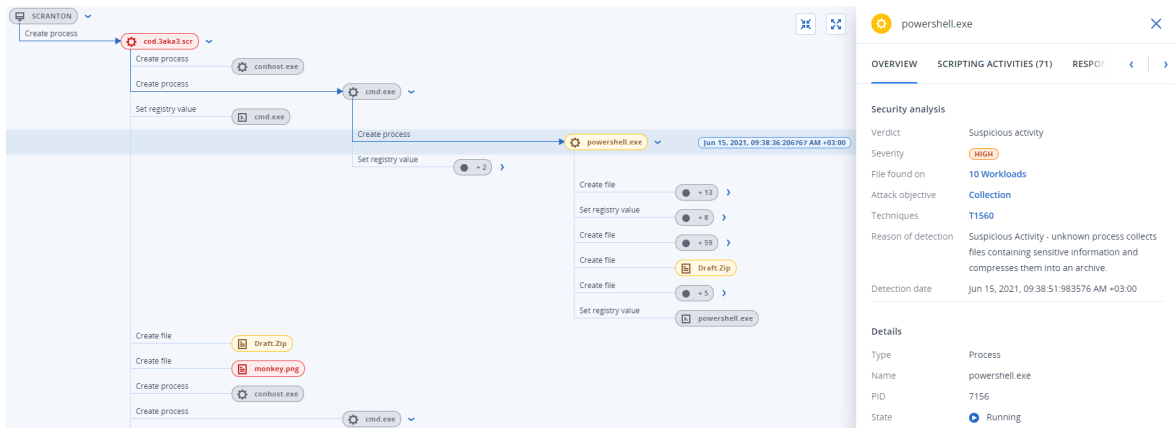
To investigate individual nodes in the cyber kill chain

1. In the Cyber Protect console, go to **Protection > Incidents**.
2. In the displayed list of incidents, click  in the far right column of the incident you want to investigate. The cyber kill chain for the selected incident is displayed.
3. Navigate to the relevant node, and click it to display the sidebar for the node.

Note

Click the node to expand it and display associated nodes.

For example, clicking the **powershell.exe** node in the example below opens the sidebar for the node. You can also click the arrow icon next to the node to view the associated nodes, including files and registry values, that may be affected by the **powershell.exe** node. In turn, you can click on these associated nodes to investigate further.




4. Investigate the information included in the sidebar tabs:

- **Overview:** Includes two main sections that provide a security summary of the attacked node.
 - **Security analysis:** Provides an analysis of the attacked node, including the EDR verdict on the threat (such as suspicious activity), the objective of the attack according to MITRE attack techniques (click on the link to go to the [MITRE website](#)), the reason for detection, and the number of workloads that may be affected by the attack (click the **n Workloads** link to view the affected workloads).

Note

The **n Workloads** link means that the specific malicious or suspicious object has been *found* on other workloads. It does not mean that the attack is happening on these other workloads, but that there is an indicator of compromise on these other workloads. The attack may have already happened (and created another incident), or the attacker is preparing to hit these other workloads using their attack "toolkit".

- **Details:** Includes details about the node, including its type, name and current state, path to the node, and any file hashes and digital signatures (such as MD5 and certificate serial numbers).
- **Scripting Activities:** Includes details of any scripts invoked or loaded in the attack. Click  to copy the script to your clipboard for further investigation.

Note

The **Scripting Activities** tab is only displayed for process nodes that run commands or scripts (such as cmd or PowerShell commands).

- **Response Actions:** Includes a number of sections that provide additional investigation, remediation and prevention actions, depending on the node type. For example, for workload nodes, you can define a number of responses that include a forensic backup and a restore from backup. Alternatively, for malicious or suspicious nodes, you can stop or quarantine the node, rollback changes made by the attack, and add it to a protection plan allowlist or blocklist. For more information about applying response actions to specific nodes, see "Response actions for individual cyber kill chain nodes" (p. 862).

- **Activities:** Displays the actions applied to the incident in chronological order. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Understand the actions taken to mitigate an incident


After you have [reviewed an incident](#) and [investigated how the attack occurred](#), you will typically [apply response actions](#). Once you have applied response actions, these actions can be viewed in a number of places to get a better understanding of what steps have been taken to mitigate the incident.

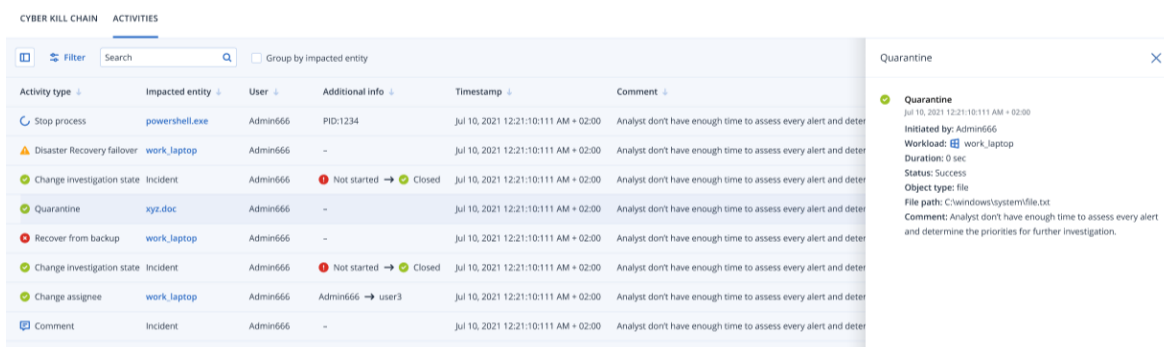
Note

Incidents created by prevention layers automatically apply the actions configured in the protection plan. For detection points, you need to define the relevant response actions to mitigate each attack scenario.

To understand the response actions taken, you can view all the response actions applied to an entire incident, or view the actions applied to a specific node in the incident cyber kill chain.




To view all response actions applied to an incident

1. In the Cyber Protect console, go to **Protection > Incidents**.
2. In the displayed list of incidents, click  in the far right column of the incident you want to investigate. The cyber kill chain for the selected incident is displayed.
3. Click the **Activities** tab.
The list of [response actions](#) already applied to the incident is displayed.



The screenshot shows the 'ACTIVITIES' tab in the Cyber Protect console. It displays a table of activities with columns for Activity type, Impacted entity, User, Additional info, Timestamp, and Comment. A sidebar on the right provides details for the selected 'Quarantine' activity, including its timestamp, initiator (Admin666), workload (work_laptop), duration, status (Success), and file path (C:\windows\system\file.txt).

Activity type	Impacted entity	User	Additional info	Timestamp	Comment
Stop process	powershell.exe	Admin666	PID:1234	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Disaster Recovery failover	work_laptop	Admin666	-	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Change investigation state	Incident	Admin666	Not started → Closed	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Quarantine	xyz.doc	Admin666	-	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Recover from backup	work_laptop	Admin666	-	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Change investigation state	Incident	Admin666	Not started → Closed	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Change assignee	work_laptop	Admin666	Admin666 → user3	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter
Comment	Incident	Admin666	-	Jul 10, 2021 12:21:10:111 AM + 02:00	Analyst don't have enough time to assess every alert and deter

4. You can perform a number of actions on the displayed list:
 - Click on an activity row to display more information about the selected activity. The information is displayed in a sidebar, as shown in Step 3, and includes details on who initiated the action, its status, file path, and any comments added by the initiator.
 - Use the **Search** box to search for a specific action.
 - Click **Filter** to apply filters to the list.
 - Select the **Group by impacted entity** check box to group relevant actions according to entity.
 - Click  to show / hide the list of completed actions.
Ensure  is displayed next to the actions you want to display. If you want to hide an action from the displayed list, click again to change it to .

Completed actions

Remediated

Isolated workloads ⓘ	1/1	🔍
Connected to network	2/3	🔍
Patched	2/3	🔍
Restarted workload	2/3	🔍
Stopped process	2/3	🔍
Quarantined	2/3	🔍
Rollback changes ⓘ	2/3	🔍
Deleted	2/3	🔍

Recovered

Recovered from backup	2/3	🔍
Disaster recovery failover	2/3	🔍

Prevent

Added to allowlist	2/3	🔍
Added to blocklist	2/3	🔍

Investigation

Forensic backup	2/3	🔍
Remote desktop connection	2/3	🔍

Other

Comments	2/3	🔍
Change investigation state	2/3	🔍
Change threat status	2/3	🔍
Change assignee	2/3	🔍

To view response actions applied to a specific node

1. In the cyber kill chain, click on a node to view the sidebar for that node.
2. Click the **Activities** tab.

ACTIVITIES (71) RESPONSE ACTIONS **ACTIVITIES** < | >

✓ **Patch**
Jun 22, 2021, 06:45:23:111 AM +02:00
Initiated by: Admin
Workload: SCRANTON
Duration: 1h 43 min
Status: Success
Patches: -

- 2021-01 Update for Windows 10 Version 2004 for x64-based Systems (KB4589212)
- 2021-06 Cumulative Update for .NET Framework 3.5 and 4.8 for Windows 10 Version 2004 for x64 (KB5003254)
- Microsoft Silverlight (KB4481252)

Comment: Analyst don't have enough time to assess every alert and determine the priorities for further investigation.

✓ **Remote desktop connection**
Jun 22, 2021, 06:45:23:111 AM +02:00
Initiated by: Admin

3. To get a complete understanding of what actions were applied and why, you may need to scroll through the applied response actions for the node. For example, for remote desktop connection actions, you can view who started the action and when, the duration of the action, and its overall status (if it succeeded, failed, or succeeded with errors).

Check for indicators of compromise (IOCs) from publicly known attacks on your workloads

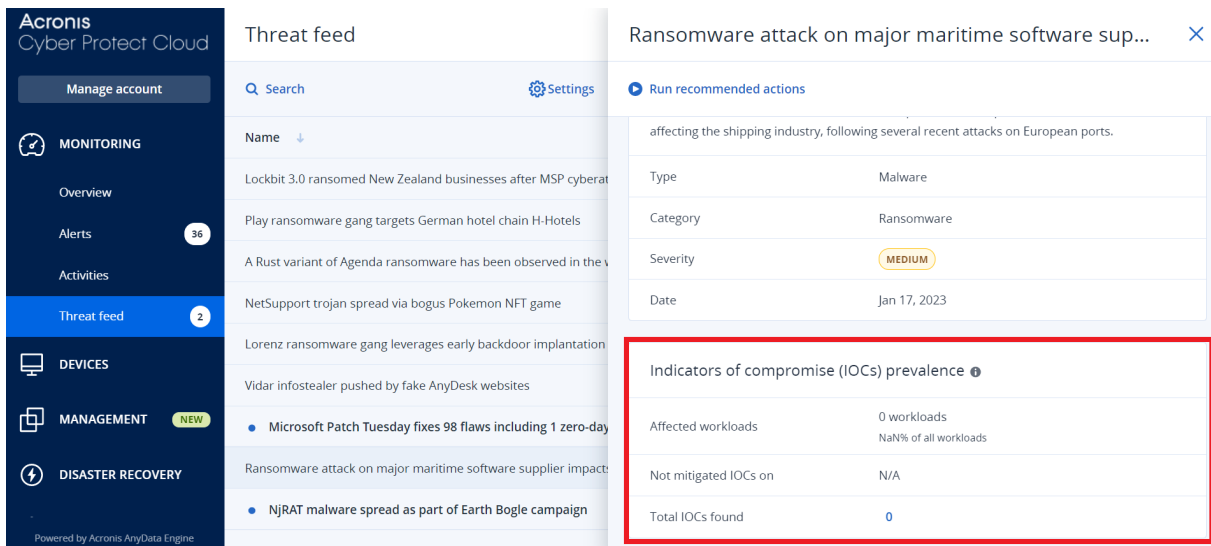
Endpoint Detection and Response (EDR) includes the ability to review existing, known attacks in threat feeds against your workloads. These [threat feeds](#) are automatically generated based on threat data received from the Cyber Protection Operations Center (CPOC); EDR enables you to verify whether or not a threat is impacting your workload, and then take the necessary steps to nullify the threat.

You can access threat feeds from the **Monitoring** menu in the Cyber Protect console. For more information, see "Threat feed" (p. 288).

To review specific threat details and confirm if they impact your workloads, click on a threat feed. You can view the number of IOCs detected and workloads affected, and drilldown to workloads that contain unmitigated IOCs.

Note

If the protection plan does not have EDR enabled, this additional threat feed functionality, as shown below, is not displayed.



Define threat feed settings

You can define a number of threat feed settings to automatically locate and mitigate any known threats.

To define threat feed settings

1. In the Cyber Protect console, go to **Monitoring > Threat feed**.
2. On the displayed Threat feed page, click **Settings**.
3. In the displayed dialog, select any of the following options:

Option	Description
Search for indicators of compromise (IOCs)	Click the switch to enable the automatic search for IOCs on your workloads. When this option is enabled, the Action on detection and Generate alert options are also displayed.
Action on detection	From the dropdown list, select the action to be taken on the relevant files when a threat is discovered on a workload: <ul style="list-style-type: none"> • No action • Quarantine • Delete • Isolate workloads
Generate alert	Select the checkbox to generate an alert if an IOC is found on a workload. The alert will be displayed in the Alerts page.

4. Click **Apply**.

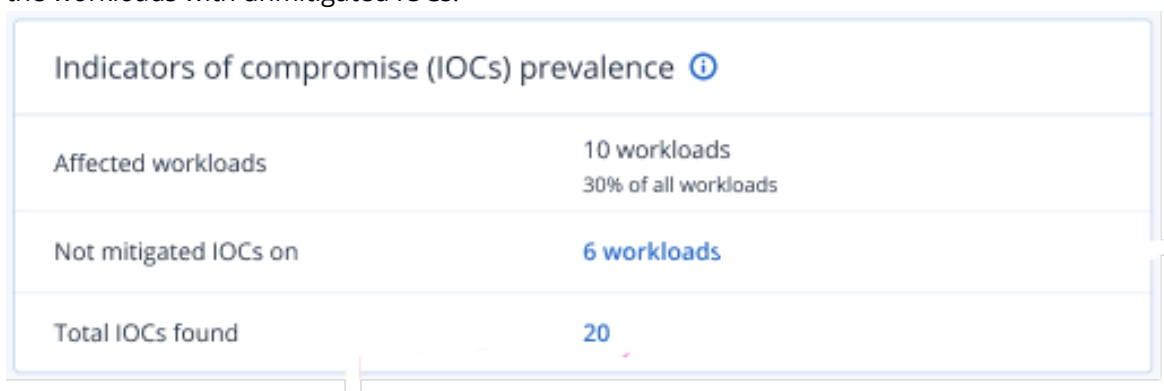
Review and mitigate IOCs on affected workloads

When Endpoint Detection and Response (EDR) is enabled in a protection plan, you can view any known threats that are affecting workloads in the protection plan. You can also mitigate any

remaining indicators of compromise (IOCs) that were not automatically mitigated. For information on how to automatically mitigate IOCs, see "Define threat feed settings" (p. 855).

To review and mitigate affected workloads

1. In the Cyber Protect console, go to **Monitoring > Threat feed**.
2. Click on a thread to display the details for that threat.
3. In the **Indicators of compromise (IOCs) prevalence** section, click the **n workloads** link to view the workloads with unmitigated IOCs.



Indicators of compromise (IOCs) prevalence ⓘ	
Affected workloads	10 workloads 30% of all workloads
Not mitigated IOCs on	6 workloads
Total IOCs found	20

4. In the displayed Workloads page, click on the relevant workload and review its details. You can run specific functionality on the workload, including defining additional URLs to filter (see "URL filtering" (p. 785)), and blocking malicious processes (refer to the Exclusions section in "Antivirus and antimalware protection settings" (p. 763)).
For example, if a threat feed indicates that a workload has been affected by an IOC, first locate and analyze the IOC, as described in "Review and analyze discovered IOCs" (p. 856). Then go to the protection plan for the workload and define additional protection, such as blocking malicious file hashes or processes.

Review and analyze discovered IOCs

In addition to [reviewing any workloads affected by known threats](#), you can also review and analyze specific indicators of compromise (IOCs). This enables you to view the individual workloads that are affected by an IOC, and mitigate the IOC.


To review and analyze IOCs

1. In the Cyber Protect console, go to **Monitoring > Threat feed**.
2. Click on a thread to display the details for that threat.
3. In the **Indicators of compromise (IOCs) prevalence** section, click the **Total IOCs found** link. The Found indicators page is displayed.

Found indicators



File name	File hash	Threat status	Workload	File path
randomware.exe	Show	Quarantined	qa-gw3t68hh	C:\Users\nikolatesla\Documents\terr
randomware.exe	Show	Quarantined	MF_2012_R2	C:\Users\mariecurie\Documents\terr
paint.exe	Show	Not mitigated	vm-Win-2012-ABA12	C:\Users\davinci\Pictures\Download:
hellorworld.exe	Show	Not mitigated	qa-gw3t68hh	C:\Users\nikolatesla\Documents\terr
hellorworld.exe	Show	Not mitigated	vm-Win-2012-ABA12	C:\Users\mariecurie\Documents\terr
services.exe	Show	Not mitigated	qa-gw3t68hh	C:\Users\nikolatesla\Documents\terr

- (Optional) Use the **Filter** option to filter the list of IOCs according to their status. You can also use the **Search** option to search for specific IOCs.
- To view the workload affected by an IOC, click the link in the **Workload** column. You can then perform a variety of actions on the workload, such as run patch management, or modify a protection plan.
- (Optional) In the **File hash** column, click **Show** to display the file hashes found for a specific IOC. In the displayed dialog, click  to copy the file hash of the IOC to a text editor.

Remediating incidents

Endpoint Detection and Response (EDR) enables you to remediate entire incidents, or the individual attack points of an incident.

By [remediating an entire incident](#), you can choose the remediation(s) that you want to execute globally on the incident. If you need to manage the incident in more granular detail, you can [remediate individual attack points](#) as required. For example, you may want to isolate the network of a workload to stop lateral movement or command and control (C&C) activities; this ensures that even though the workload is isolated, all Acronis Cyber Protect technologies are still functional and an investigation can be launched.

EDR ensures effective remediation by:


- Mitigating - to ensure the threat is stopped.
- Recovering - to ensure services are back online immediately.
- Preventing - to ensure techniques used in an attack are prevented in future attacks.

Remediate an entire incident

By remediating an entire incident, you can quickly and easily choose the remediation(s) that you want to execute globally on the incident. Endpoint Detection and Response (EDR) guides you through the remediation process, step by step.

If you need to manage your network and the incident in more granular detail, see "Response actions for individual cyber kill chain nodes" (p. 862).

To remediate an entire incident

1. In the Cyber Protect console, go to **Protection > Incidents**.
2. In the displayed list of incidents, click  in the far right column of the incident you want to investigate. The cyber kill chain for the selected incident is displayed.
3. Click **Remediate entire incident**. The Remediate entire incident dialog is displayed.

Remediate entire incident ✕

Analyst verdict

True positive False positive

Remediation actions

Step 1 – Stop threats
Stops all processes related to the threat.


Step 2 – Quarantine threats
After being stopped, all malicious or suspicious processes and files are quarantined.

Step 3 – Rollback changes
Rollback first deletes any new registry entries, scheduled tasks or files created by the threat (and any of its children threats). Next, rollback reverts any modifications made by the threat (or its children) to the registry, scheduled tasks and/or files existing on the workload prior to the attack.
To optimize speed, rollback tries to recover items from the local cache. Items that fail to be recovered will be recovered by the system from backup images.

Allow this response action to access encrypted backups using your stored credentials

Affected items: [Show \(40\)](#)

Recover workload
If any of the above selected remediation steps fail completely or partially.

Recovery point: [20 Jan, 2021, 6:45:23 AM](#) 

Items to be recovered: **Entire workload**

Prevention actions

Add to blocklist
Adds all threats from the incident to the blocklist in the selected protection plans. This action will prevent these threats from future executions.

Patch workload
Prevents further attacks by patching software that contains vulnerabilities used by attackers in order to get a foothold on the workload.

Change investigation state of the incident to: Closed

Comment
Analyst don't have enough time to assess every alert and determine the priorities for further investigation. Automatic alerts triage presents a clear story that analysts can easily read and understand. It reduces the time spent for triaging alerts and enables faster incident response.

4. In the **Analyst verdict** section, based on your [investigation of the incident](#), select one of the following:

- **True positive:** Select if you are sure the attack is a legitimate attack. Once selected, you then add remediation and prevention actions, as described in the following steps.
- **False positive:** Select if you are sure the attack is not a genuine attack. In this mode, you can define how to prevent this from happening again, such as by adding the incident to a protection plan allowlist.

Note

After selecting **False positive**, you can only define prevention actions. For more information, see "Remediate a false positive incident" (p. 861).

5. In the **Remediation actions** section, perform the following remediation steps. Note that they must be performed in sequential order; for example, you cannot select Step 2 before Step 1 is completed.
 - a. **Step 1 - Stop threats:** Select the check box to stop all processes related to the threat.
 - b. **Step 2 - Quarantine threats:** Once the threat is stopped, select the check box to quarantine all malicious and suspicious processes and files.
 - c. **Step 3 - Rollback changes:** After threats have been quarantined, select the check box to delete any new registry entries, scheduled tasks or files created by the threat (and any of its children threats). The rollback process then reverts any modifications made by the threat (or its children) to the registry, scheduled tasks and/or files existing on the workload prior to the attack. To optimize speed, the rollback process tries to recover items from the local cache. Items that fail to be recovered will be recovered by the system from backup images.

Note

The rollback process recovers from items in the local cache only. Rollback from backup archives will be available in future releases.

Select the **Allow this response action to access encrypted backups using your stored credentials** check box if access to the relevant backups is encrypted. EDR accesses the stored user credentials to decrypt the encrypted archives and search for the relevant files. You can also click **Affected items** to view all items (files, registry, or scheduled tasks) affected by the rollback, the actions applied (**Delete**, **Recover**, or **None**), and if the items are being restored from the local cache or backup images.

Affected items ✕

Name ↓	Type ↓	Path ↓	Action ↓	Recover from
xyz.doc	File	C:\windows\system\vchost.xyz.doc	Recover	local cache
xyz.doc	Registry	C:\windows\system\vchost.xyz.doc	Delete	-
xyz.doc	File	C:\windows\system\vchost.xyz.doc	Recover	Recovery points (12 Jan, 2021, 6:45:23 AM)
xyz.doc	Windows Scheduled Task	C:\windows\system\vchost.xyz.doc	None	-
xyz.doc	File	C:\windows\system\vchost.xyz.doc	Recover	Recovery points (12 Jan, 2021, 6:45:23 AM)
xyz.doc	Registry	C:\windows\system\vchost.xyz.doc	Recover	Recovery points (12 Jan, 2021, 6:45:23 AM)

- d. **Recover workload:** Select the check box to recover a workload if any of the above remediation steps fail, whether completely or partially.

Recover workload
 If any of the above selected remediation steps fail completely or partially.

Recover workload from backup
 Disaster recovery failover

Recovery point: 20 Jan, 2021, 6:45:23 AM [✎](#)

Select one of the following recovery options:

- **Recover workload from backup:** Enables you to recover a workload from a specific recovery point. Click the recovery point edit icon to select from a list of recovery backups.
- **Disaster recovery failover:** Enables you to run disaster recovery, if you have this functionality enabled in your protection plan. We recommend that you use this option for critical workloads, such as AD servers, or database servers. For more information, see "Implementing disaster recovery" (p. 679).

6. In the **Prevention actions** section, select the relevant remediation steps:

- **Add to blocklist:** Select the check box and from the displayed protection plan list, select the relevant protection plans. This prevention action ensures all detections of the incident will be blocked from being executed for the selected protection plans.
- **Patch workload:** Select the check box to patch any vulnerable software and prevent attackers from gaining access to the workload. You can then select the relevant action to perform once the patch is complete (**Do not restart**, **Restart**, or **Restart only if required**), depending if the user is logged in or not.

You can also select the **Do not restart while backup is in progress** check box to ensure the workload is not restarted during backup.

Patch workload
Prevents further attacks by patching software that contains vulnerabilities used by attackers in order to get a foothold on the workload.

If user is logged out

Do not restart Restart Restart only if required

If user is logged in

Do not restart Restart Restart only if required

Do not restart while backup is in progress

7. Select the **Change investigation state of the incident to: Closed** check box. If not selected, the investigation state remains in its previous state.
8. Click **Remediate**. The remediation actions you selected are executed, step by step, with the progress of each remediation step shown in the Remediate entire incident dialog. Once clicked, the button displays **Go to activities**. Click **Go to activities** to review all response actions applied to the incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Remediate a false positive incident

If you are sure an attack is not a genuine attack, in other words a false positive, you can define how to prevent the incident from happening again. For example, you can add the incident to a protection plan allowlist.

To remediate a false positive incident

1. In the cyber kill chain for the selected incident, click **Remediate entire incident**. The Remediate entire incident dialog is displayed.

2. In the **Analyst verdict** section, select **False positive**.

Remediate entire incident ✕

Analyst verdict

True positive False positive

Prevention actions

Add to allowlist
Adds all detections from the incident to the allowlist in the selected protection plans. This action will consider those processes and URLs safe and will prevent them from being detected.

Protection plan
My protection plan ▼

Change investigation state of the incident to: False positive

Comment
Analyst don't have enough time to assess every alert and determine the priorities for further investigation. Automatic alerts triage presents a clear story that analysts can easily read and understand. It reduces the time spent for triaging alerts and enables faster incident response.

3. In the **Prevention actions** section, select the **Add to allowlist** check box. From the displayed protection plan list, select the relevant protection plans. This prevention action ensures all detections of the incident will be prevented from being detected for the selected protection plans.
4. Select the **Change investigation state of the incident to: False positive** check box.
5. Click **Remediate**.
Once clicked, the button displays **Go to activities**. Click **Go to activities** to review the response actions applied to the incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Response actions for individual cyber kill chain nodes

If you need to manage the incident in more granular detail, you can apply various response actions to individual cyber kill chain nodes. These response actions enable you to quickly and easily remediate any node.

Note

To apply global response actions to an entire incident, see "Remediate an entire incident" (p. 857).

Response actions are divided into the following categories, although not all nodes include all of the following categories:

- **Remediate:** Actions in this category enable you to apply an immediate response to the attack, and include managing network isolation for a workload, and the deletion and quarantining of files, processes, and registry values.
- **Investigate:** Actions in this category (applicable to workloads only) enable you to run a Forensic backup, or remote desktop connection for a more in-depth investigation.
- **Investigate:** Actions in this category (applicable to workloads only) enable you to run a remote desktop connection for a more in-depth investigation.
- **Recovery:** Actions in this category (applicable to workloads only) enable you to respond to intensive attacks by running a recovery from backup, or Disaster Recovery failover.
- **Prevent:** Actions in this category enable you to prevent future threats or false positives by adding them to a protection plan allowlist or blocklist.

Note

If an incident is closed, you cannot apply a response action to a node. However, you can reopen a closed incident by [changing its investigation state](#) to **Investigating**. When reopened, you can then apply response actions.

The following table describes each of the node types in the cyber kill chain, the applicable categories for each node, and the response actions available.

Node	Category	Response Actions
Workload	Remediate	<ul style="list-style-type: none"> • Manage network isolation • Restart workload
	Investigate	<ul style="list-style-type: none"> • Forensic backup • Remote desktop connection
	Investigate	<ul style="list-style-type: none"> • Remote desktop connection
	Recovery	<ul style="list-style-type: none"> • Recovery from backup • Disaster Recovery failover
	Prevent	<ul style="list-style-type: none"> • Patch

Node	Category	Response Actions
Process	Remediate	<ul style="list-style-type: none"> • Stop process • Quarantine
	Prevent	<ul style="list-style-type: none"> • Add to allowlist • Add to blocklist
File	Remediate	<ul style="list-style-type: none"> • Delete • Quarantine
	Prevent	<ul style="list-style-type: none"> • Add to allowlist • Add to blocklist
Registry	Remediate	<ul style="list-style-type: none"> • Delete
Network	Prevent	<ul style="list-style-type: none"> • Add to allowlist • Add to blocklist

Define response actions for an affected workload

As part of your response to an attack, you can apply the following actions to affected workloads:

- **Manage network isolation:** Enables you to manage the network isolation of a workload to stop lateral movement or Command and Control (C&C) activities. For more information, see "Manage the network isolation of a workload" (p. 865).
- **Patch:** Enables you to patch a workload to prevent future vulnerability exploitations in future potential attacks. For more information, see "Patch a workload" (p. 868).
- **Restart workload:** Enables you to immediately restart a workload, or restart the workload according to a predefined timeout period. For more information, see "Restart a workload" (p. 869).
- **Forensic backup:** Enables you to do an on-demand forensic backup for audit or further investigation purposes. For more information, see "Run an on-demand forensic backup on a workload" (p. 870).
- **Remote desktop connection:** Enables you to remotely access the workload under investigation. For more information, see "Remote connection to a workload" (p. 871).
- **Recovery from backup:** Enables you to recover your entire machine from backup or specific files or folders. For more information, see "Recovery from backup" (p. 872).

- **Disaster Recovery failover:** Enables you to run "Implementing disaster recovery" (p. 679). Note that your workload must have a subscription for Advanced Disaster Recovery. For more information, see "Disaster Recovery failover" (p. 873).

Manage the network isolation of a workload

EDR enables you to manage the network isolation of a workload to stop lateral movement or Command and Control (C&C) activities. There are a number of isolation options to choose from, according to your requirements. Note that all Acronis Cyber Protect technologies are functional even if a workload is isolated, ensuring that an investigation can be fully carried out.

To isolate a workload from the network

1. In the cyber kill chain, click the workload node you want to remediate.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Remediate** section, click **Manage network isolation**.

REMEDiate

▼ **Manage network isolation**

Network status Connected

Do you want to isolate the network of workload `work_laptop`?

Immediate action after isolation

Isolate only ▼

Message to display

Comment (optional)

Isolate
Manage network exclusions

Note

The **Network Status** value indicates if the workload is currently connected or not. If the value displays **Isolated**, you can reconnect the isolated workload to the network, as described in the procedure below. If the workload is offline you can still isolate the workload; when the workload goes back online it is automatically put into the **Isolated** state.

4. In the **Immediate action after isolation** drop-down list, select from one of:
 - **Isolate only**
 - **Isolate and backup workload**
 - **Isolate and backup workload with forensic data**
 - **Isolate and power off workload**

For more information about defining where to backup the workload and encryption options, see "Managing the backup and recovery of workloads and files" (p. 375).

5. [Optional] In the **Message to display** field, add a message to display to end users when they access the isolated workload. For example, you can inform users that the workload is now isolated and that network access in and out of the workload is currently not available. Note that this message is also displayed as a tray monitor notification, and remains displayed until the user dismisses the message.
6. [Optional] In the **Comment** field, add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
7. Click **Manage network exclusions** to add ports, URLs, host names, and IP addresses that will have access to the workload during the isolation. For more information, see [how to manage network exclusions](#).
8. Click **Isolate**.
The workload is isolated. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Note

The workload is also shown as **Isolated** under the **Workloads** menu in the Cyber Protect console. You can also isolate single or multiple workloads from the **Workloads > Workloads with agents** menu; select the relevant workload(s) and in the right sidebar select **Manage network isolation**. In the displayed dialog, you can manage network exclusions and click **Isolate** or **Isolate all** to isolate the selected workload(s).

To connect an isolated workload back to the network

1. In the cyber kill chain, click the workload node you want to reconnect.

Note

If the isolated workload is currently offline you can still reconnect it back to the network; when the workload goes back online it is automatically put into the **Connected** state.

2. In the displayed sidebar, click the **Response Actions** tab.
 3. In the **Remediate** section, click **Manage network isolation**.
 4. Select from one of the following:
 - **Connect to network immediately**: The workload is reconnected to the network.
 - **Recover workload from backup before connecting to network**: Select a recovery point from which to recover the workload.
 - a. In the **Recovery point** field, click **Select**.
 - b. In the displayed sidebar, select the relevant recovery point.
 - c. Click **Recover > Entire workload** to recover all the files and folders on the workload.
- Or

Click **Recover > Files/folders** to recover specific files and folders on the workload. You are then prompted to select the relevant files or folders. Once selected, you can view the list of items by clicking the relevant value in the **Items to be recovered** field.

Manage network isolation

Workload status **Isolated**

Do you want to connect work_laptop to the network? All network access to the machine will no longer be restricted.

Connection method
Recover workload from backup before connecting to netwo...

Recovery point **20 Jan, 2021, 6:45:23 AM**

Items to be recovered **32**

Recover to C:\Program Files\Applications\Backup

Message to display

Comment (optional)

Recover and connect Manage network exclusions

Note

If the recovery point you select is encrypted, you will be prompted for the password.

5. [Optional] Select the **Automatically restart the workload if required** check box. This option is relevant only if you selected **Recover > Entire workload** in Step 4.
6. [Optional] In the **Message to display** field, add a message to display to end users when they access the connected workload. For example, you can inform users that a backup was restored to the workload and that network access in and out of the workload is resumed.
7. [Optional] In the **Comment** field, add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
8. Click **Connect** if you selected **Connect to network immediately** in Step 4.

Or

Click **Recover and connect** if you selected **Recover workload from backup before connecting to network** in Step 4.

The workload is reconnected to the network and all network access to the workload is no longer restricted.

Note

You can also connect single or multiple isolated workloads from the **Workloads > Workloads with agents** menu in the Cyber Protect console; select the relevant workload(s) and in the right sidebar select **Manage network isolation**. In the displayed dialog, click **Connect** or **Connect all** to reconnect the selected workload(s) to the network.

To manage network exclusions

Note

Even if all Acronis Cyber Protect technologies are working when the workload is in isolation, there may be scenarios in which you need additional network connections to be established (for example, you may need to upload a file from the workload to a shared directory). In these scenarios, you can add a network exclusion, but make sure any threats are removed before you add the exclusion.

1. In the **Remediate** section of the **Response actions** tab, click **Manage network exclusions**.
2. In the Network exclusions sidebar, add the relevant exclusions. For each of the options available (Ports, URL address, and Hostname / IP address), do the following:
 - a. Click **Add** and then enter the relevant port(s), URL addresses, or Hostname / IP addresses.
 - b. In the **Traffic direction** drop-down list, select one of **Incoming and outgoing connections**, **Incoming connections only**, or **Outgoing connections only**.
 - c. Click **Add**.
3. Click **Save**.

Patch a workload

EDR automatically detects if a workload requires a patch, and enables you to patch the workload to prevent vulnerability exploitations in future potential attacks. Note that this feature is available only if the partner's workload has a subscription for Advanced Management.

To patch a workload

1. In the cyber kill chain, click the workload node you want to patch.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Remediate** section, click **Patch**.
4. In the **Patches to install** field, click **Select**. In the displayed dialog, select the relevant patches and then click **Select**.
5. In the **Post-installation options** field, click the displayed link. The Post-installation options dialog is displayed.

Post-installation options ×

Choose what to do after patch installation

If user is logged out

Do not restart
 Restart
 Restart only if required

If user is logged in

Do not restart
 Restart
 Restart only if required

Schedule restart
 Right after patch installation ▼

Allow snoozing
 Allow unlimited snoozing ▼

Reminder interval ▼ Time unit
 15 ▼ Minute(s) ▼

Do not restart while backup is in progress

6. Select the action to perform after the patch is installed:
 - **If user is logged out:** Select one of **Do not restart**, **Restart**, or **Restart only if required**.
 - **If user is logged in:** Select one of **Do not restart**, **Restart**, or **Restart only if required**.
 When you select **Restart**, you can also define the following:
 - Schedule the restart.
 - Allow snoozing, including the defined intervals between snoozes.
7. [Optional] Select the **Do not restart while backup is in progress** check box to ensure the workload is not restarted if a backup is currently in progress.
8. Click **Save**.
9. In the **Response Actions** tab, click **Patch**.
 The selected patch is run. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Restart a workload

As part of your remediation response to an attack, EDR enables you to immediately restart a workload, or restart the workload according to a predefined timeout period.

To restart a workload

1. In the cyber kill chain, click the workload node you want to set a restart schedule for.
2. In the displayed sidebar, click the **Response Actions** tab.

3. In the **Remediate** section, click **Restart workload**.

REMEDiate

- > Manage network isolation
- > Patch

▼ Restart workload

Do you want to restart the workload **work_laptop**? Note that any unsaved changes will be lost.

Restart timeout **3 minutes** ▼

Fail if end-user is logged in

Message to display to users when they access the isolated workload **work_laptop**: **Restart immediately** minutes. Any unsaved work will be lost.

Comment (optional)

Restart

4. In the **Restart timeout** field, click the displayed link, and then select one of the following:
 - **Set timeout:** In the Restart timeout dialog, set the restart period for the workload, and then click **Save**.
 - **Restart immediately:** Select to restart the workload immediately.
5. [Optional] Select the **Fail if end-user is logged in** check box to ensure the workload is not restarted if the user is logged in.
6. In the **Message to display** field, add a message to display to users when they access the isolated workload.
7. [Optional] In the **Comment** field, add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
8. Click **Restart**.

The workload is set to restart according to the schedule defined. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Run an on-demand forensic backup on a workload

As part of your investigation into an attack, EDR enables you to run an on-demand forensic backup for audit or further investigation purposes. Note that this feature is available only if the partner's workload has a subscription for Advanced Backup.

To run a forensic backup

1. In the cyber kill chain, click the workload node you want to run a forensic backup on.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Investigate** section, click **Forensic backup**.

INVESTIGATE

> Remote desktop connection

▼ Forensic backup

Backup name	New forensic backup	✎
Forensic options	Raw memory dump, Snapshot on	
Where to back up	Cloud storage	
Encryption	<input checked="" type="checkbox"/>	

Comment (optional)

4. [Optional] In the **Backup name** field, click the edit icon to edit the backup name.
5. In the **Forensic options** field, click the displayed link. In the displayed Forensic options dialog, select one of the following:

- **Collect raw memory dump**
- **Collect kernel memory dump**

You can also select the **Snapshot of running processes** check box to add information about the processes running at the moment the backup starts. This information is stored in a backup image.

Click **Save** to close the Forensic options dialog.

6. In the **Where to back up** field, click the displayed link to define a location for the backup.
7. [Optional] Click the **Encryption** option to enable encryption. In the displayed dialog, enter the password for the encrypted backup and select the relevant encryption algorithm.
8. [Optional] In the **Comment** field, add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
9. Click **Run**.

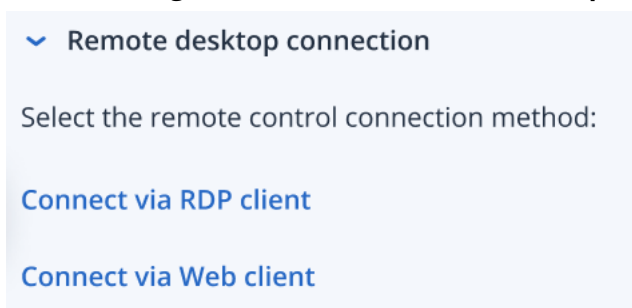
The forensic backup is started. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Remote connection to a workload

As part of your investigation into an attack, EDR enables you to remotely access the workload under investigation.

To remotely connect to a workload

1. In the cyber kill chain, click the workload node you want to remotely connect to.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Investigate** section, click **Remote desktop connection**.



4. Select one of the following remote connection methods:
 - **Connect via RDP client:** This method will prompt you to download and install the Remote Desktop Connection Client. You can then [remotely connect to a workload](#) from the Cyber Protect console.
 - **Connect via Web client:** This method does not require the installation of an RDP client on your workload. You are redirected to the login screen where your credentials to the remote machine have to be entered.


When the remote connection is started, this action can be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Recovery from backup

As part of your recovery response to an attack, EDR enables you to recover your entire machine from backup or specific files or folders.

To recover your workload from backup

1. In the cyber kill chain, click the workload node you want to recover.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Recovery** section, click **Recover from backup**.



4. In the **Recovery point** field, click **Select** and then perform the following steps:
 - a. In the displayed sidebar, select the relevant recovery point.
 - b. Click **Recover > Entire workload** to recover all the files and folders on the workload.

Or

Click **Recover > Files/folders** to recover specific files and folders on the workload. You are then prompted to select the relevant files or folders. Once selected, you can view the items selected for recovery by clicking the relevant value in the **Items to be recovered** field.

Note

If the recovery point you select is encrypted, you will be prompted for the password.

5. [Optional] Select the **Automatically restart the workload** check box. This option is relevant only if you selected **Recover > Entire workload** in Step 4.
6. [Optional] In the **Comment** field, add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
7. Click **Start recovery**.
The process to recover the workload starts. The progress for this action can be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Disaster Recovery failover

As part of your recovery response to an attack, EDR enables you to run "Implementing disaster recovery" (p. 679), which allows you to switch the workload to the recovery server. Note that your workload must have a subscription for Advanced Disaster Recovery.

To run Disaster Recovery failover

1. In the cyber kill chain, click the workload node you want to recover.
2. In the displayed sidebar, click the **Response Actions** tab.

3. In the **Recovery** section, click **Disaster Recovery failover**.

RECOVERY

› Recovery from backup

▼ Disaster Recovery failover ↑

Are you sure you want to switch the workload from the original workload to the recovery server?

Recovery server name	Cloud storage
IP address	192.168.1.2
Internet access	Enabled
Public IP address	-
Recovery point	06 Jan, 2021, 6:45:23 AM

Comment (optional)

Failover

4. In the **Recovery point** field, perform the following steps:
 - a. Click the current recovery point date to select a recovery point.
 - b. In the displayed sidebar, select the relevant recovery point.

Note

If you have an Advanced Disaster Recovery subscription, you can select the relevant recovery server (the offline VM) created in [Disaster Recovery](#). If you do not have a subscription, you will be prompted to configure Disaster Recovery.

5. [Optional] In the **Comment** field, add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
6. Click **Failover**.

The workload is switched to the recovery server. This action can be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Define response actions for a suspicious process

As part of your remediation response to an attack, you can apply the following actions to suspicious processes:

- Stop a process (see below)
- Quarantine a process (see below)
- Roll back changes made by a process (see below)

- Add the process to a protection plan allowlist or blocklist (see "Add or remove a process, file or network in the protection plan blocklist or allowlist" (p. 880))

To stop a suspicious process

1. In the cyber kill chain, click the process node you want to remediate.

Note

Windows critical processes or non-running processes cannot be stopped and are disabled in the cyber kill chain.

2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Remediate** section, click **Stop process**.

REMEDiate

▼ Stop process

Do you want to end the process **powershell.exe** running on **work_laptop**? Ending this process will close the related application and you will lose any unsaved data.

Stop process

Stop process tree

Comment (optional)

Stop

4. Select one of the following:
 - **Stop process** (stops the specific process)
 - **Stop process tree** (stops the specific process and all child processes)
5. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
6. Click **Stop**. The process is stopped.

Note

The related application is closed and any unsaved data will be lost.

This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

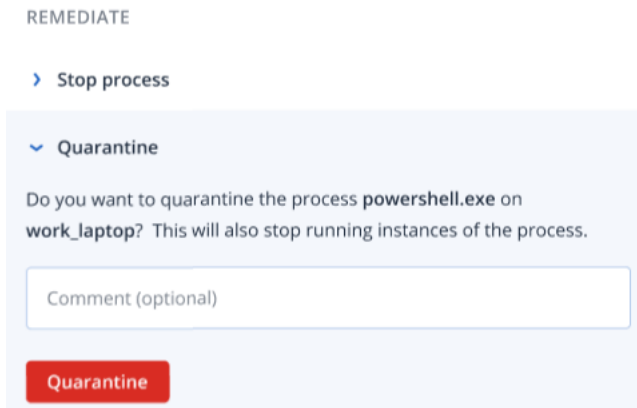
To quarantine a suspicious process

1. In the cyber kill chain, click the process node you want to quarantine.

Note

Windows critical processes cannot be quarantined and are disabled in the cyber kill chain.

2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Remediate** section, click **Quarantine**.



4. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
5. Click **Quarantine**. The process is stopped and then quarantined.

Note

The process is added to and managed in the quarantine section available under [antimalware protection](#).

This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

To rollback changes

1. In the cyber kill chain, click the process node you want to rollback changes for.

Note

This action is available for detection nodes (shown as red or yellow nodes) only.

2. In the displayed sidebar, click the **Response Actions** tab.

- In the **Remediate** section, click **Rollback changes**.

REMEDIALTE

- › Stop process
- › Quarantine
- ▼ Rollback changes

Do you want to rollback any changes made by the process **powershell.exe**?

Rollback first deletes any new registry, scheduled tasks or files created by the threat (and any of its children threats). Next, rollback reverts any modifications made by the threat (or its children) to the registry, scheduled tasks and/or files existing on the workload prior to the attack.

To optimize speed, rollback tries to restore items from the local cache. Items that fail to be restored will be restored by the system from backup images.

Affected items **6**

Comment (optional)

Rollback

Note

The rollback process recovers from items in the local cache only. Rollback from backup archives will be available in future releases.

- To view the items affected by the rollback changes, click the **Affected items** link. The displayed dialog shows all items (files, registry, scheduled tasks) that the rollback will revert and with what action (**Delete**, **Recover**, or **None**). In addition, you can see whether the restored items will be recovered from the local cache or backup recovery points.

Affected items ✕

Name ↓	Type ↓	Path ↓	Action ↓	Recover from
xyz.doc	File	C:\windows\system\wchost.xyz.doc	Recover	local cache
xyz.doc	Registry	C:\windows\system\wchost.xyz.doc	Delete	-
xyz.doc	File	C:\windows\system\wchost.xyz.doc	Recover	Recovery points (12 Jan, 2021, 6:45:23 AM)
xyz.doc	Windows Scheduled Task	C:\windows\system\wchost.xyz.doc	None	-
xyz.doc	File	C:\windows\system\wchost.xyz.doc	Recover	Recovery points (12 Jan, 2021, 6:45:23 AM)
xyz.doc	Registry	C:\windows\system\wchost.xyz.doc	Recover	Recovery points (12 Jan, 2021, 6:45:23 AM)

5. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
6. Click **Rollback**. The rollback functionality reverts any registry, file or scheduled task changes made by the process in the following steps:
 - a. Any new entries (registry, scheduled tasks, files) created by the threat (and its child threats) are deleted.
 - b. Any modifications that the threat (and its child threats) made to the registry, scheduled tasks and/or files existing on the workload prior to the attack are reverted.
 - c. Rollback tries to recover items from the local cache. For items that cannot be recovered, EDR will automatically recover them from clean backup images.

The rollback action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

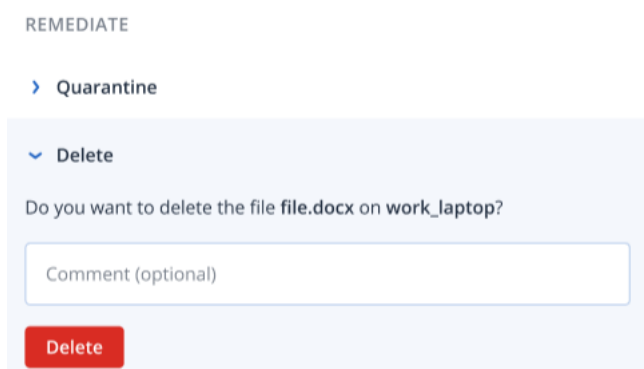
Define response actions for a suspicious file

As part of your remediation response to an attack, you can apply the following actions to suspicious files:

- Delete a file (see below)
- Quarantine a file (see below)
- Add the file to a protection plan allowlist or blocklist (see "Add or remove a process, file or network in the protection plan blocklist or allowlist" (p. 880))

To delete a suspicious file

1. In the cyber kill chain, click the file node you want to remediate.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Remediate** section, click **Delete**.



4. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.

5. Click **Delete**.

The file is deleted. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

To quarantine a suspicious file

1. In the cyber kill chain, click the file node you want to remediate.
2. In the displayed sidebar, go to **Response Actions**.
3. In the **Remediate** section, click **Quarantine**.

REMEDiate

▼ Quarantine

Do you want to quarantine the file `file.docx` on `work_laptop`?

Comment (optional)

Quarantine

4. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.

5. Click **Quarantine**.

The file is quarantined. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Define response actions for a suspicious registry entry

As part of your remediation response to an attack, you can delete suspicious registry entries.

This option is available for registry cyber kill chain nodes.

To delete a suspicious registry entry

1. In the cyber kill chain, click the node you want to remediate.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Remediate** section, click **Delete**.

REMEDiate

▼ Delete

Do you want to delete the registry `MainWindowHandle` on `work_laptop`?

Comment (optional)

Delete

4. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.
5. Click **Delete**.
The registry entry is deleted. This action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

Add or remove a process, file or network in the protection plan blocklist or allowlist

As part of your prevention response to an attack, you can add a node to your protection plan allowlist or blocklist.

You can add a node to an allowlist if you consider the node safe and want to prevent any future detections for it. Add a node to a blocklist to stop the node from running in the future.

You can also remove a node from the allowlist or blocklist to allow or prevent any future access to the node.

This option is available for the following cyber kill chain nodes:

- Process
- File
- Network

To add or remove a process, file or network in the protection plan blocklist

1. In the cyber kill chain, click the process, file, or network node you want to remediate.
2. In the displayed sidebar, click the **Response Actions** tab.
3. In the **Prevent** section, click the arrow icon next to **Blocklist**.

Blocklist

To prevent access to the file "file.docx", add it to the protection plan blocklist. If "file.docx" was previously added, you can click on Remove to remove it from the blocklist and restore access to it.

Protection plan
My protection plan

Comment (optional)

Add Remove

4. Select the relevant protection plan(s) you want to apply this action to.

5. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.

6. Click **Add**.

The action is implemented, and the process, file, or network will be prevented from launching in the future.

Alternatively, if the process, file, or network was previously added to the blocklist and you now want to remove it from the blocklist, click **Remove**. This will allow future access to the node.

The add or remove action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

To add or remove a process, file or network in the protection plan allowlist

1. In the cyber kill chain, click the process, file, or network node you want to remediate.

2. In the displayed sidebar, click the **Response Actions** tab.

3. In the **Prevent** section, click the arrow icon next to **Allowlist**.

▼ Allowlist

To allow access to the file "file.docx", add it to the protection plan allowlist. If "file.docx" was previously added, you can click on Remove to remove it from the allowlist and prevent access to it.

Protection plan
My protection plan ▼

Comment (optional)

Add Remove

4. Select the relevant protection plan(s) you want to apply this action to.

5. [Optional] Add a comment. This comment is visible in the **Activities** tab (for a single node or the entire incident), and can help you (or your colleagues) recall why you took the action when you revisit the incident.

6. Click **Add**.

The action is implemented and the process, file, or network will be prevented from detection in the future.

Alternatively, if the process, file, or network was previously added to the allowlist and you now want to remove it from the allowlist, click **Remove**. This will prevent any future access to the node.

The add or remove action can also be viewed in the **Activities** tabs of both the individual node and the entire incident. For more information, see "Understand the actions taken to mitigate an incident" (p. 852).

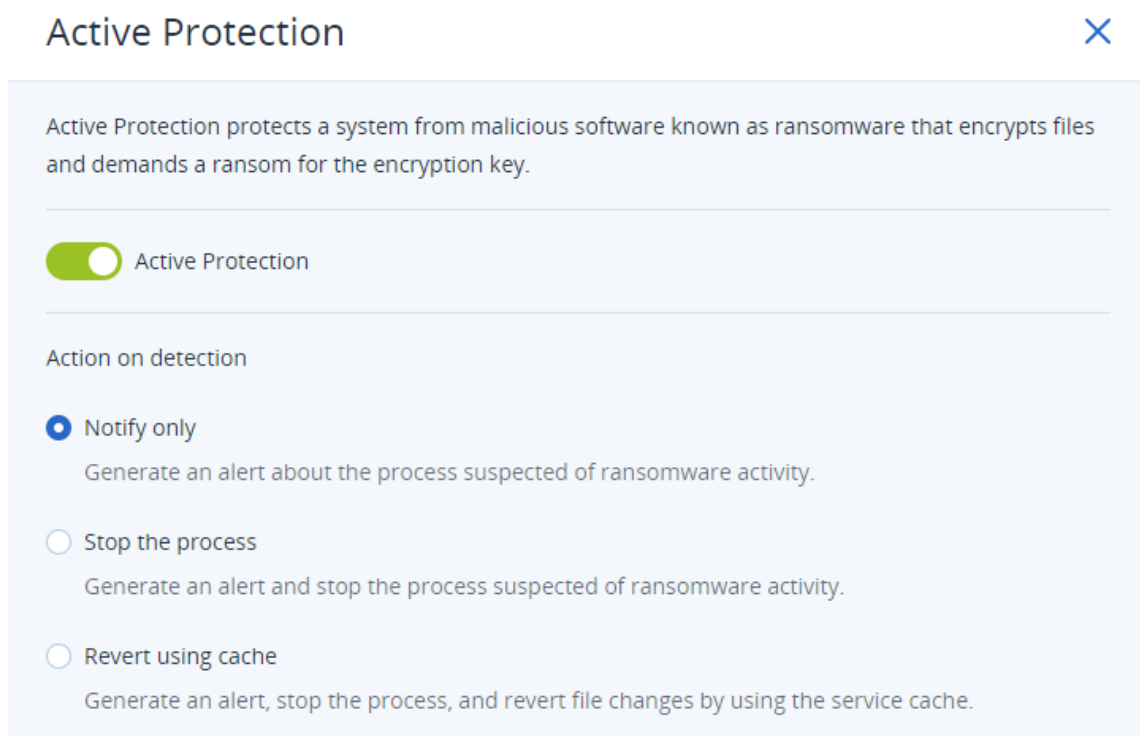
Enabling monitoring mode for Endpoint Detection and Response (EDR)

The monitoring mode in Cyber Protection enables you to use EDR in a production environment. In turn, this enables you to check for any false positives, and make necessary exclusions before fully deploying EDR.

In monitoring mode, nothing is blocked or stopped, and incidents are created, but no responses are initiated.

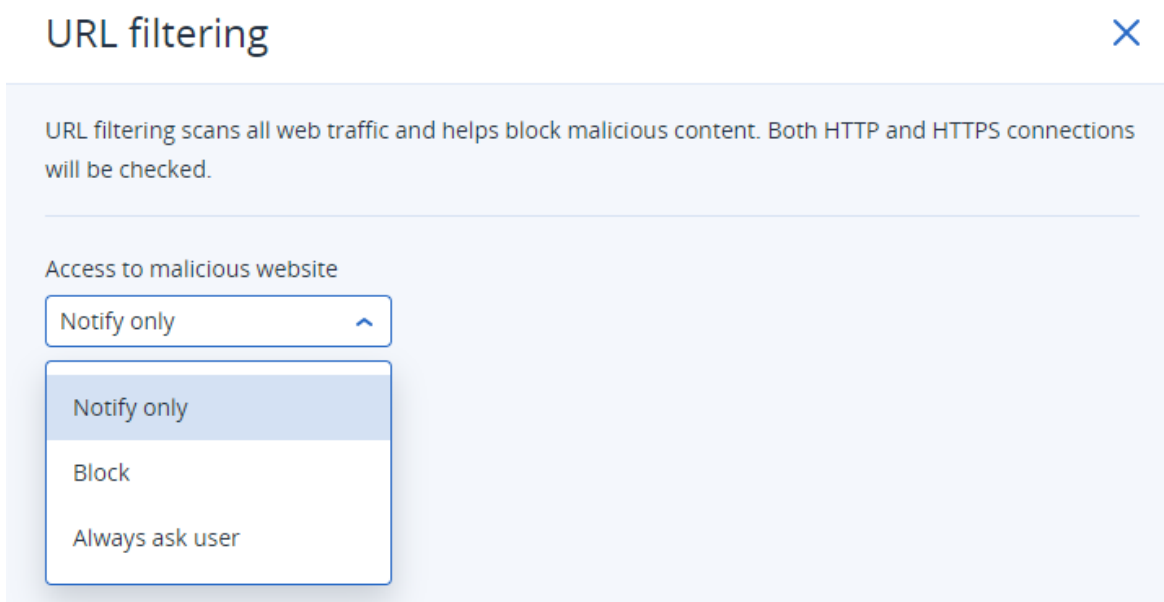
To enable the monitoring mode for EDR

1. In the relevant protection plan, ensure that EDR is enabled. For more information, see "Enabling Endpoint Detection and Response (EDR) functionality" (p. 834).
2. Expand the **Antivirus & Antimalware protection** module, and then define the following:
 - Click **Active protection**, and in the **Action on detection** section, select **Notify only**. Then click **Done**. For more information, see "Active Protection" (p. 763).



- Click **Behavior engine**, and in the **Action on detection** section, select **Notify only**. Then click **Done**. For more information, see "Behavior-engine" (p. 768).
- Click **Exploit prevention**, and in the **Action on detection** section, select **Notify only**. Then click **Done**. For more information, see "Exploit prevention" (p. 769).

- Click **Real-time protection**, and in the **Action on detection** section, select **Notify only**. Then click **Done**. For more information, see "Real-time protection" (p. 770).
 - Click **Schedule scan**, and in the **Action on detection** section, select **Notify only**. Then click **Done**. For more information, see "Schedule scan" (p. 771).
3. Expand the **URL filtering** module, and in the **Access to malicious website** drop-down list, select **Notify only**. Then click **Done**. For more information, see "URL filtering" (p. 785).



How to test if Endpoint Detection and Response (EDR) is working correctly


To ensure EDR is deployed and working, you can run a number of commands that trigger EDR detections.

Note

When EDR is deployed, you should see incidents immediately if any suspicious activity occurs. The steps below enable you to check if EDR is working if no new incidents were triggered for several days.

To test if EDR is deployed and working correctly

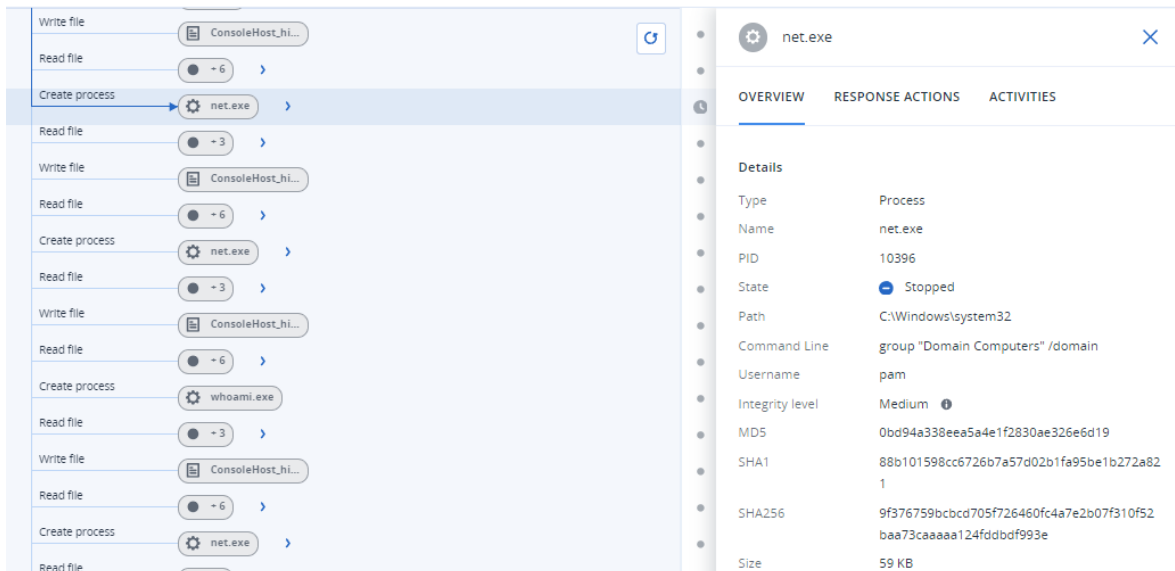
1. Login to the relevant domain-joined Active Directory user account.
2. Run the following two commands in Windows PowerShell:
 - `net group "Domain Computers" /domain`
 - `net user administrator /domain`
3. In the Cyber Protect console, go to **Protection > Incidents** to view the generated incident. You can also click on the triggered **Medium** severity type incident to display it in the EDR cyber kill chain and confirm the PowerShell commands you executed in the previous step, as shown in the example below.



4. Run the following commands in Windows PowerShell:

- `c:\>whoami`
- `c:\>net localgroup`
- `c:\>net localgroup administrators`
- `c:\>powershell -command start-process cmd -verb runas`
- `c:\WINDOWS\system32>net user administrator /active:yes`
- `c:\>powershell -command Get-Hotfix`

5. In the EDR cyber kill chain, click on the executable nodes (for example, **net.exe** or **whoami.exe**) to display the exact PowerShell commands executed on the command line. These commands are shown in the **Details** section of the **Overview** tab in the example below.



6. After you have confirmed that an EDR incident was generated, manually set the **Threat status** for the incident to **Mitigated** and the **Investigation state** to **Closed**. For more information, see "How to investigate incidents in the cyber kill chain" (p. 844). You can also enter a comment for the incident to indicate that this was a test incident.

Assessing vulnerabilities and managing patches

Vulnerability assessment (VA) is a process of identifying, quantifying, and prioritizing vulnerabilities found in the system. In the vulnerability assessment module, you can scan your machines for vulnerabilities, and check if the operating systems and installed applications are up to date and working properly.

Vulnerability assessment scanning is supported for machines with the following operating systems:

- Windows. For more information, see "Supported Microsoft and third-party products" (p. 885).
- macOS. For more information, see "Supported Apple and third-party products" (p. 887).
- Linux (CentOS 7/Virtuozzo/Acronis Cyber Infrastructure) machines. For more information, see "Supported Linux products" (p. 888).

Use the **Patch management** (PM) functionality to manage patches (updates) for applications and operating systems installed on your machines, and keep your systems up to date. In the patch management module, you can automatically or manually approve update installations on your machines.

Patch management is supported for machines with the Windows operating systems. For more information, see "Supported Microsoft and third-party products" (p. 885).

Vulnerability assessment

The vulnerability assessment process consists of the following steps:

1. You [create a protection plan](#) with the enabled vulnerability assessment module, specify the [Vulnerability assessment settings](#), and [assign the plan to machines](#).
2. The system, by schedule or on demand, sends a command to run the vulnerability assessment scanning to the protection agents installed on machines.
3. The agents get the command, start scanning machines for vulnerabilities, and generate the scanning activity.
4. After the vulnerability assessment scanning is completed, the agents generate the results and send them to the monitoring service.
5. The monitoring service processes the data from the agents and shows the results in the [vulnerability assessment widgets](#) and list of found vulnerabilities.
6. When you get a [list of found vulnerabilities](#), you can process it and decide which of the found vulnerabilities must be fixed.

You can monitor the results of the vulnerability assessment scanning in **Monitoring > Overview > Vulnerabilities / Existing vulnerabilities** widgets.

Supported Microsoft and third-party products

The following Microsoft products and third-party products for Windows operating systems are supported for vulnerability assessment and patch management:

Supported Microsoft products

Windows OS

- Windows 7 (Enterprise, Professional, Ultimate)
- Windows 8
- Windows 8.1
- Windows 10
- Windows 11

Windows Server OS

- Windows Server 2022
- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2

Microsoft Office and related components

- Microsoft Office 2019 (x64, x86)
- Microsoft Office 2016 (x64, x86)
- Microsoft Office 2013 (x64, x86)
- Microsoft Office 2010 (x64, x86)

Windows OS related components

- Internet Explorer
- Microsoft EDGE
- Windows Media Player
- .NET Framework
- Visual Studio and Applications
- Components of operating system

Server applications

- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2012
- Microsoft SQL Server 2014
- Microsoft SQL Server 2016
- Microsoft SQL Server 2017
- Microsoft SQL Server 2019
- Microsoft Exchange Server 2013
- Microsoft Exchange Server 2016

- Microsoft Exchange Server 2019
- Microsoft SharePoint Server 2016
- Microsoft SharePoint Server 2019

Supported third-party products for Windows OS

Remote work becomes more and more wide-spread across the world, therefore collaboration and communication tools, VPN clients are now important to be always up-to-date and checked on possible vulnerabilities. The Cyber Protection service supports the vulnerability assessment and patch management for such applications.

Collaboration and communication tools, VPN clients

- Microsoft Teams
- Zoom
- Skype
- Slack
- Webex
- NordVPN
- TeamViewer

For more information about the supported third-party products for Windows OS, refer to [List of third-party products supported by Patch Management \(62853\)](#).

Supported Apple and third-party products

The following Apple products and third-party products for macOS are supported for vulnerability assessment:

Supported Apple products

macOS

- macOS 10.13.x and later

macOS built-in applications

- Safari, iTunes, and others.

Supported third-party products for macOS

- Microsoft Office (Word, Excel, PowerPoint, Outlook, OneNote)
- Adobe Acrobat Reader
- Google Chrome
- Firefox
- Opera
- Zoom

- Skype
- Thunderbird
- VLC media player

Supported Linux products

The following Linux distributions and versions are supported for VA:

- Virtuozzo 7.x
- CentOS 7.x
- CentOS 8.x

Vulnerability assessment settings

To learn how to create a protection plan with the Vulnerability assessment module, refer to "[Creating a protection plan](#)". You can perform VA scanning by schedule or on demand (by using the **Run now** action in a protection plan).

You can specify the following settings in the Vulnerability assessment module.

What to scan

Define which software products you want to scan for vulnerabilities:

- Windows machines:
 - **Microsoft products**
 - **Windows third-party products** (for more information about the supported third-party products for Windows OS, refer to [List of third-party products supported by Patch Management \(62853\)](#))
- macOS machines:
 - **Apple products**
 - **macOS third-party products**
- Linux machines:
 - **Scan Linux packages**

Schedule

Define the schedule according to which to perform the vulnerability assessment scan on the selected machines:

Field	Description
Schedule the task run using the following events	<p>This setting defines when the task will run.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Schedule by time – This is the default setting. The task will run according to the specified time.

Field	Description
	<ul style="list-style-type: none"> • When user logs in to the system – By default, a login of any user will trigger the task. You can modify this setting so that only a specific user account can trigger the task. • When user logs off the system – By default, a logoff of any user will trigger the task. You can modify this setting so that only a specific user account can trigger the task. <hr/> <p>Note The task will not run at system shutdown. Shutting down and logging off are different events in the scheduling configuration.</p> <hr/> <ul style="list-style-type: none"> • On the system startup – The task will run when the operating system starts. • On the system shutdown – The task will run when the operating system shuts down.
Schedule type	<p>The field appears if in Schedule the task run using the following events you have selected Schedule by time.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Monthly – Select the months and the weeks or days of the month when the task will run. • Daily – This is the default setting. Select the days of the week when the task will run. • Hourly – Select the days of the week, repetition number, and the time interval in which the task will run.
Start at	<p>The field appears if in Schedule the task run using the following events you have selected Schedule by time</p> <p>Select the exact time when the task will run.</p>
Run within a date range	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time.</p> <p>Set a range in which the configured schedule will be effective.</p>
Specify a user account whose login to the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs in to the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the login of any user to trigger the task. • The following user - Use this option if you want only the login of a specific user account to trigger the task.
Specify a user account whose	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs off the system.</p>

Field	Description
logout from the operating system will initiate a task	<p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the logout of any user to trigger the task. • The following user - Use this option if you want only the logout of a specific user account to trigger the task.
Start conditions	<p>Defines all conditions that must be met simultaneously for the task to run.</p> <p>Start conditions for antimalware scans are similar to the start conditions for the Backup module that are described in "Start conditions".</p> <p>You can define the following additional start conditions:</p> <ul style="list-style-type: none"> • Distribute task start time within a time window – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM. • If the machine is turned off, run missed tasks at the machine startup • Prevent the sleep or hibernate mode during task running – This option is effective only for machines running Windows. • If start conditions are not met, run the task anyway after – Specify the period after which the task will run, regardless of the other start conditions. <hr/> <p>Note Start conditions are not supported for Linux.</p> <hr/>

Vulnerability assessment for Windows machines

You can scan Windows machines and third-party products for Windows for vulnerabilities.

To configure the vulnerability assessment for Windows machines

1. In the Cyber Protect console, [create a protection plan](#) and enable the **Vulnerability assessment** module.
2. Specify the vulnerability assessment settings:
 - **What to scan** – select **Microsoft products**, **Windows third-party products**, or both.
 - **Schedule** – define the schedule for performing the vulnerability assessment.

For more information about the **Schedule** options, see "Vulnerability assessment settings" (p. 888).
3. [Assign the plan to the Windows machines.](#)

After a vulnerability assessment scan, you can see a [list of found vulnerabilities](#). You can process the information and decide which of the found vulnerabilities must be fixed.

To monitor the results of the vulnerability assessment, see the **Monitoring > Overview > Vulnerabilities / Existing vulnerabilities** widgets.

Vulnerability assessment for Linux machines

You can scan Linux machines for application-level and kernel-level vulnerabilities.

To configure the vulnerability assessment for Linux machines

1. In the Cyber Protect console, [create a protection plan](#) and enable the **Vulnerability assessment** module.
2. Specify the vulnerability assessment settings:
 - **What to scan** – select **Scan Linux packages**.
 - **Schedule** – define the schedule for performing the vulnerability assessment.
For more information about the **Schedule** options, see "Vulnerability assessment settings" (p. 888).
3. [Assign the plan to the Linux machines](#).

After a vulnerability assessment scan, you can see a [list of found vulnerabilities](#). You can process the information and decide which of the found vulnerabilities must be fixed.

To monitor the results of the vulnerability assessment, see the **Monitoring > Overview > Vulnerabilities / Existing vulnerabilities** widgets.

Vulnerability assessment for macOS devices

You can scan macOS devices for operating system-level and application-level vulnerabilities.

To configure the vulnerability assessment for macOS devices

1. In the Cyber Protect console, [create a protection plan](#) and enable the **Vulnerability assessment** module.
2. Specify the vulnerability assessment settings:
 - **What to scan** – select **Apple products, macOS third-party products**, or both.
 - **Schedule** – define the schedule for performing the vulnerability assessment.
For more information about the **Schedule** options, see "Vulnerability assessment settings" (p. 888).
3. [Assign the plan to the macOS devices](#).

After a vulnerability assessment scan, you can see a [list of found vulnerabilities](#). You can process the information and decide which of the found vulnerabilities must be fixed.

To monitor the results of the vulnerability assessment, see the **Monitoring > Overview > Vulnerabilities / Existing vulnerabilities** widgets.

Managing found vulnerabilities

If the vulnerability assessment was performed at least once and some vulnerabilities were found, you can see them in **Software management > Vulnerabilities**. The list of vulnerabilities shows both vulnerabilities that have patches to be installed, and those that do not have suggested patches. You can use the filter to show only vulnerabilities with patches.

Name	Description
Name	The name of vulnerability.
Affected products	Software products for which the vulnerabilities were found.
Machines	The number of affected machines.
Severity	The severity of found vulnerability. The following levels can be assigned according to the Common Vulnerability Scoring System (CVSS): <ul style="list-style-type: none"> • Critical: 9 - 10 CVSS • High: 7 - 9 CVSS • Medium: 3 - 7 CVSS • Low: 0 - 3 CVSS • None
Patches	The number of appropriate patches.
Published	The date and time when the vulnerability was published in Common Vulnerabilities and Exposures (CVE).
Detected	The first date when an existing vulnerability was detected on machines.

You can find the description of found vulnerability by clicking its name in the list.

The screenshot shows the 'Vulnerabilities' section of the Acronis Cyber Protect Cloud interface. The table lists the following vulnerabilities:

Name	Affected products	Machines	Severity	Patches
CVE-2015-16723	Microsoft Windows 8.1	1	CRITICAL	2
CVE-2015-0016	Microsoft Windows 8.1	1	CRITICAL	1
CVE-2014-4073	Microsoft Windows 8.1	1	CRITICAL	1
CVE-2010-3190	Microsoft Visual Studio 2008	1	CRITICAL	1
CVE-2015-1756	Microsoft Windows 8.1	1	CRITICAL	1
CVE-2014-4121	Microsoft Windows 8.1	1	CRITICAL	1
CVE-2016-3236	Microsoft Windows 8.1	1	CRITICAL	1
CVE-2014-6324	Microsoft Windows 8.1	1	CRITICAL	1

To start the vulnerability remediation process

1. In the Cyber Protect console, go to **Software management > Vulnerabilities**.
2. Select the vulnerability in the list, and then click **Install patches**. The vulnerability remediation wizard will open.
3. Select the patches to be installed on the selected machines, and then click **Next**.
4. Select the machines on which you want to install the patches.
5. Select the reboot options.
 - a. Select if you want the machine to be rebooted after the patches are installed.

Option	Description
No	The machines will not be rebooted automatically after the patches are installed.
If required	The machines will be rebooted only if it is required for applying the patches.
Yes	The machines will be rebooted automatically after the patches are installed. You can also specify a reboot delay.

- b. [Optional] If you want to delay the machine reboot while a backup of the machine is in progress, select **Do not reboot until backup is finished**.
6. Click **Install patches**.

As a result, the selected patches are installed on the selected machines.

Patch management

Note

The availability of this feature depends on the service quotas that are enabled for your account.

For more information about the supported third-party products for Windows OS, refer to [List of third-party products supported by Patch Management \(62853\)](#).

Use the patch management functionality to:

- install OS-level and application-level updates
- approve patches manually or automatically
- install patches on-demand or according to a schedule
- precisely define which patches to install by different criteria: severity, category, and approval status
- perform pre-update backup to prevent possible unsuccessful updates
- define the reboot action after patch installation

Note

To work with Windows updates, the patch management feature requires that Windows updates are enabled on the workload.

Cyber Protection introduces peer-to-peer technology to minimize network bandwidth traffic. You can choose one or more dedicated agents that will download updates from the Internet and distribute them among other agents in the network. All agents will also share updates with each other as peer-to-peer agents.

The patch management workflow

The patch management workflow includes steps for configuring and applying a protection plan, running a vulnerability assessment scan, configuring patch settings, approving patches and finally, installing patches that are approved. The exact steps of the workflow are as follows.

1. Configure a protection plan that has the **Vulnerability assessment** and **Patch management** modules enabled.
2. Configure the vulnerability assessment settings. For more information about these settings, see "Vulnerability assessment settings" (p. 888).
3. Configure the patch management settings. For more information about these settings, see "Patch management settings in the protection plan" (p. 894)
4. Apply the protection plan to one or several machines.
5. Wait for a vulnerability assessment scan to be completed. The scan will start automatically, according to the schedule that is configured in the protection plan. Alternatively, you can manually start the scan on demand by clicking the **Run now** icon in the **Vulnerability assessment** module in the protection plan.
6. Approve the patches. You can define settings for automatic patch approval, which include an automatic installation of the patches on test machines. For more information, see "Automatic patch approval" (p. 901). Alternatively, you can manually approve patches by setting their approval status to **Approved**. For more information, see "Approving patches manually" (p. 906).
7. Install the patches. The approved patches can be installed automatically, according to the schedule that is configured in the protection plan. Alternatively, you can manually install patches on demand. For more information, see "Installing patches on demand" (p. 906).

You can monitor the results of the patch installation in **Monitoring > Overview > Patch installation history** widget.

Patch management settings in the protection plan

In the **Patch management** module of the protection plan, you can configure the following patch management settings:

- What updates to install for Microsoft and third-party products for Windows OS.
- When to run the automatic patch installation.
- Whether to run a pre-update backup.

For more information about creating a protection plan and enabling the **Patch management** module, see "Creating a protection plan" (p. 209).

Note

The availability of this feature depends on the service quotas that are enabled for your account.

Microsoft products

To install the Microsoft updates on the selected machines, enable the **Update Microsoft products** option.

Select the installation option:

Option	Description
All updates	Use this option if you want to install all approved updates.
Only Security and Critical updates	Use this option if you want to install all approved security and critical updates.
Updates of specific products (Automatic patch approval and testing)	Use this option if you want to define custom settings for different products. If you want to update specific products, for each product you can define which updates to install by category , severity , or approval status . If you want to configure automatic test approval and testing of the patches, select this option.

Updates of specific products (Automatic patch approval and testing)



	Products	Category	Severity	Approval status
<input type="checkbox"/>	Products	Custom	Custom	Approved
<input checked="" type="checkbox"/>	Windows 10, version 1903 and lat...	All	All	Approved
<input type="checkbox"/>	Windows Server 2016 for RS4	—	—	—
<input checked="" type="checkbox"/>	Windows Server 2016	CriticalUpdates, Securit...	All	Approved
<input checked="" type="checkbox"/>	Windows Server 2019	Updates	Critical	Approved
<input checked="" type="checkbox"/>	Windows Server, version 1903 an...	All	Critical, Unspecified	Approved

[Reset to default](#)

For Microsoft products, patch distribution uses the Windows API service. Patches and updates are not downloaded or stored internally or on distribution agents. Instead, they are downloaded from Microsoft CDN. Thus, even with the Updater role assigned, the agent cannot download and distribute patches.

Windows third-party products

To install the third-party updates for Windows OS on the selected machines, enable the **Windows third-party products** option.

Select the installation options:

Option	Description
All updates	Use this option if you want to install all approved updates. *
Only major updates	Use this option if you want to install all approved major updates.
Only minor updates	Use this option if you want to install approved minor updates.
Updates of specific products (Automatic patch approval and testing)	Use this option if you want to define custom settings for different products. If you want to update specific products then, for each product, you can define which updates to install by category, severity, or approval status . If you want to configure automatic test approval and testing of the patches, select this option.
Install the latest versions only for applications with detected vulnerabilities	Select this check box if you want to install the latest updates only for applications that have detected vulnerabilities. *

* This option requires Cyber Protect agent version 23.11.36772 or later.

Updates of specific products (Automatic patch approval and testing) ✕

	Products	Version	Severity	Approval status
<input type="checkbox"/>	Adobe AdobeReaderMUI	—	—	—
<input checked="" type="checkbox"/>	Adobe AIR	All updates	All	Approved
<input checked="" type="checkbox"/>	Adobe Flash Player for Chrome a...	Major updates	Critical, High, Unspecifi...	Approved
<input checked="" type="checkbox"/>	Adobe Flash Player for FireFox an...	Minor updates	High, Critical	Approved
<input checked="" type="checkbox"/>	Adobe Reader	All updates	All	Approved
<input type="checkbox"/>	Adobe Shockwave Player	—	—	—
<input checked="" type="checkbox"/>	Adobe Systems Incorporated Ext...	All updates	All	Approved
<input type="checkbox"/>	AdoptOpenJDK AdoptOpenJDK	—	—	—
<input type="checkbox"/>	AIMP DevTeam AIMP	—	—	—

Reset to default Cancel Save

For Windows third-party products, patches are distributed directly to the managed workloads from an internal Acronis database. In case the Updater role is assigned to an agent, this agent will be used to download and distribute patches.

Schedule

Define the schedule and conditions according to which the updates will be installed on the selected machines.

Field	Description
Schedule the task run using the following events	<p>This setting defines when the task will be run.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Schedule by time – This is the default setting. The task will run according to the specified time. • When user logs in to the system – By default, a login of any user will start the task. You can modify this setting so that only a specific user account can trigger the task. • When user logs off the system – By default, a logoff of any user will start the task. You can modify this setting so that only a specific user account can trigger the task. <hr/> <p>Note</p> <p>The task will not run at system shutdown. Shutting down and logging off are different events in the scheduling configuration.</p> <hr/> <ul style="list-style-type: none"> • On the system startup – The task will run when the operating system starts. • On the system shutdown – The task will run when the operating system shuts down.
Schedule type	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Monthly – Select the months and the weeks or days of the month when the task will run. • Daily – This is the default setting. Select the days of the week when the task will run. • Hourly – Select the days of the week, repetition number, and the time interval in which the task will run.
Start at	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time</p> <p>Select the exact time when the task will run.</p>
Configure maintenance window for patches	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time.</p> <p>Select this setting if you want the patch installation to run only during the time interval that you will specify. If the patch installation process has not</p>

Field	Description
	completed by the end time defined by the maintenance window for patches, it will be stopped automatically.
Run within a date range	<p>The field appears if, in Schedule the task run using the following events, you have selected Schedule by time.</p> <p>Set a range in which the configured schedule will be effective.</p>
Specify a user account whose login to the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs in to the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the login of any user to trigger the task. • The following user - Use this option if you want only the login of a specific user account to trigger the task.
Specify a user account whose logout from the operating system will initiate a task	<p>The field appears if, in Schedule the task run using the following events, you have selected When user logs off the system.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Any user - Use this option if you want the logout of any user to trigger the task. • The following user - Use this option if you want only the logout of a specific user account to trigger the task.
Start conditions	<p>Defines all conditions that must be met simultaneously for the task to run.</p> <p>Start conditions for antimalware scans are similar to the start conditions for the Backup module that are described in "Start conditions".</p> <p>You can define the following additional start conditions:</p> <ul style="list-style-type: none"> • Distribute task start time within a time window – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM. • If the machine is turned off, run missed tasks at the machine startup • Prevent the sleep or hibernate mode during task running – This option is effective only for machines running Windows. • If start conditions are not met, run the task anyway after – Specify the period after which the task will run, regardless of the other start conditions.

Field	Description
	<p>Note</p> <p>Start conditions are not supported for Linux.</p>
Reboot after update	<p>Define whether to reboot the machine automatically after the installation of the updates completes.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> • Never – A reboot will never be initiated after the updates. • If required – A reboot will be initiated only if it is required for applying the updates. • Always – A reboot will be always initiated after the updates. You can specify a reboot delay.
Do not reboot until backup is finished	<p>If you select this option, if a backup process is running, the reboot of the machine will be delayed until the backup is completed.</p>

Pre-update backup

Run backup before installing software updates – the system will create an incremental backup of machine before installing any updates on it. If there were no backups created earlier, then a full backup of machine will be created. It allows you to prevent such cases when the installation of updates was unsuccessful and you need to get back to the previous state. For the **Pre-update backup** option to work, the corresponding machines must have both the patch management and the backup module enabled in a protection plan and the items to back up – entire machine or boot+system volumes. If you select inappropriate items to back up, then the system will not allow you to enable the **Pre-update backup** option.

Viewing the list of available patches

After a vulnerability assessment scan completes, you can view information about the available patches in **Software management > Patches**.

To view details about a specific patch, in the list of patches, click the corresponding patch.

The following table describes the information for the patch that you can view on the screen.

Field	Description
Approval status	<p>The approval status is mainly needed for automatic approval scenarios.</p> <p>You can define one of the following statuses for a patch:</p> <ul style="list-style-type: none"> • Approved – the patch was installed on at least one machine and validated as ok • Declined – the patch is not safe and may corrupt a machine system • Pending approval – the patch status is unclear and should be validated

License agreement	<ul style="list-style-type: none"> • Agreed • Disagreed. If you disagree with the license agreement, then the patch status becomes Declined and it will not be installed
Severity	<p>The severity of the patch:</p> <ul style="list-style-type: none"> • Critical • High • Medium • Low • None
Vendor	The vendor of the patch
Affected product	Product for which the patch is applicable
Installed versions	Product versions that are already installed
Version	Version of the patch
Category	<p>The category to which the patch belongs:</p> <ul style="list-style-type: none"> • Critical update – broadly released fixes for specific problems addressing critical, non-security related bugs. • Security update – broadly released fixes for specific products addressing security issues. • Definition update – updates to virus or other definition files. • Update rollup – cumulative set of hotfixes, security updates, critical updates, and updates packaged together for easy deployment. A rollup generally targets a specific area, such as security, or a specific component, such as Internet Information Services (IIS). • Service pack – cumulative sets of all hotfixes, security updates, critical updates, and updates created since the release of the product. Service packs might also contain a limited number of customer-requested design changes or features. • Tool – utilities or features that aid in accomplishing a task or set of tasks. • Feature pack – new feature releases, usually rolled into products at the next release. • Update – broadly released fixes for specific problems addressing non-critical, non-security related bugs. • Application – patches for an application.
Release date	The date when the patch was released
Last reported	The date of the last time when the patch was reported
First installed	The date of the first successful installation of the patch on a machine
Microsoft KB	If the patch is for a Microsoft product, the field shows the KB article ID

Machines	Number of affected machines
Vulnerabilities	The number of vulnerabilities. If you click on it, you will be redirected to the list of vulnerabilities.
Size	The average size of the patch
Language	The language which is supported by the patch
Vendor site	The official site of the vendor

Configuring the patch lifetime in the list

You can keep the list of patches up to date by configuring the patch lifetime in the list on the **Patches** screen. This setting defines how long the detected available patch will be visible in the list of patches. The patch will be removed from the list after it is successfully installed on all the machines on which it was indicated as missing, or after the lifetime in the list passes.

To configure the patch lifetime in the list

1. In the Cyber Protect console, go to **Software management > Patches**.
2. Click **Settings**.
3. In **Lifetime in list**, select the appropriate option.

Option	Description
Forever	The patch will always stay in the list.
7 days	The patch will be removed from the list seven days after its first installation. For example, let us assume that you have two machines on which patches must be installed. One of them is online, and the other one is offline. The patch was installed on the first machine. After 7 days, the patch will be removed from the list of patches, even if it is not installed on the second machine (because it was offline).
30 days	The patch is removed from the list 30 days after its first installation.

Automatic patch approval

Automatic patch approval makes the process of installing updates on machines easier. With automatic patch approval, the installation of patches is not delayed by the manual patch approval process. Important updates and fixes are installed faster, which increases the reliability of your system.

You can use automatic patch approval in test scenarios for automatic installation of patches. If the patches are installed successfully on the test machines, the patches will be automatically installed on the production machines, too. For more information about this scenario, see "Use case for automatic patch approval and testing" (p. 902).

You can also use automatic patch approval in scenarios for automatic installation of patches in your production environment, skipping the testing phase. For more information about this scenario, see "Use case for automatic patch approval without testing" (p. 905).

Configuring automatic patch approval

You can configure automatic patch approval and ensure that the installation of patches is not delayed by the manual patch approval process.

To configure automatic patch approval

1. In the Cyber Protect console, go to **Software management > Patches**.
2. Click **Settings**.
3. Enable **Automatic patch approval**.
4. Configure the settings for automatic patch approval.
 - a. Select the automatic patch approval option.

Option	Description
Automatic patch approval and testing	The approval status of the patch will change to Approved when the selected number of days passes after a successful installation of the patch. We recommend that you use this setting if you want to test the patches by installing them on a test machine first, ensure that everything is working as expected, and then install the patches in your production environment.
Automatic patch approval without testing	The approval status of the patch will change to Approved when the selected number of days passes after the patch was found.

- b. Select the number of days that must pass after the condition from the automatic patch approval option is met. After this period, the approval status of the patches will automatically change from **Pending approval** to **Approved**.
5. Select **Automatically accept the license agreements**.
6. Click **Apply**.

Use case for automatic patch approval and testing

If you want to test the new patches on a test machine before installing them on your production machines, you can configure two protection plans - a plan for installation of patches for test purposes, and a plan for installation of tested patches on production machines. Thus, you will ensure that the patches that you install in your production environment are safe and your production machines work correctly after the patch installation.

The use case consists of the following stages:

1. Configure the settings for Automatic patch approval. Select the **Automatic patch approval and testing** option. For more information, see "Configuring automatic patch approval" (p. 902).
2. Configure a protection plan for test purposes (for example, 'Test patching') with the enabled **Patch management** module and apply it to the machines in the test environment. Specify the following condition of patch installation: the patch approval status must be **Pending approval**. This step is needed to validate the patches and check if the machines work properly after patch installation. For more information, see "Configuring the Test patching protection plan" (p. 903).
3. Configure a protection plan for the production environment (for example, 'Production patching') with the enabled **Patch management** module and apply it to the machines in the production environment. Specify the following condition of patch installation: the patch status must be **Approved**. For more information, see "Configuring the Production patching protection plan" (p. 904).
4. Run the Test patching plan and check the results. Leave the approval status of the machines that have no issues as **Pending approval**, but change the approval status of the machines working incorrectly to **Declined**. According to the number of days set in the **Automatic patch approval** setting, the status of the patches will automatically change from **Pending approval** to **Approved**. When you run the Production patching plan, only the **Approved** patches will be installed on the production machines. For more information, see "Running the Test patching protection plan and decline unsafe patches" (p. 905).
5. Run the Production patching plan.

Configuring the Test patching protection plan

You can configure a protection plan with patch installation settings for your machines in the test environment.

To configure the Test patching protection plan

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Click **Create plan**.
3. Enable the **Patch management** module.
4. Define which updates to install for Microsoft and third-party products, schedule, and pre-update backup. For more details about these settings, see "Patch management settings in the protection plan" (p. 894).

Important

For all the products to be updated, select the **Pending approval** approval status. Thus, the agent will install only **Pending approval** patches on the selected machines in the test environment.

Updates of specific products (Automatic patch approval and testing)



	Products	Version	Severity	Approval status
<input type="checkbox"/>	Products	Custom	Custom	Custom
<input checked="" type="checkbox"/>	Adobe Flash Player for FireFox an...	Major updates	High, Critical, Unspecifi...	Pending approval
<input checked="" type="checkbox"/>	Adobe Flash Player for Chrome a...	Major updates	Critical	Pending approval
<input checked="" type="checkbox"/>	Adobe AIR	Major updates	All	Pending approval
<input checked="" type="checkbox"/>	Adobe Reader	Minor updates	All	Pending approval
<input checked="" type="checkbox"/>	Adobe Shockwave Player	Minor updates	All	Pending approval
<input type="checkbox"/>	Oracle Java Development Kit	—	—	—
<input checked="" type="checkbox"/>	Oracle Java Runtime Environment	Minor updates	All	Pending approval
<input checked="" type="checkbox"/>	Mozilla Firefox	Major updates	All	Pending approval
<input checked="" type="checkbox"/>	Mozilla Thunderbird	Major updates	All	Pending approval

[Reset to default](#) [Cancel](#) [Save](#)

Configuring the Production patching protection plan

You can configure a protection plan with patch installation settings for your machines in the production environment.

To configure the Production patching protection plan

1. In the Cyber Protect console, go to **Management > Protection plans**.
2. Click **Create plan**.
3. Enable the **Patch management** module.
4. Define which updates to install for Microsoft and third-party products, schedule, and pre-update backup. For more details about these settings, see "Patch management settings in the protection plan" (p. 894).

Important

For all the products to be updated, set the **Approval status** to **Approved**. Thus, the agent will install only **Approved** patches on the selected machines in the production environment.

Updates of specific products (Automatic patch approval and testing)



Products		Version	Severity	Approval status
<input checked="" type="checkbox"/>	Adobe Flash Player for FireFox an...	Major updates	High, Critical, Unspecifi...	Approved
<input checked="" type="checkbox"/>	Adobe Flash Player for Chrome a...	Major updates	Critical	Approved
<input checked="" type="checkbox"/>	Adobe Air	Major updates	All	Approved
<input checked="" type="checkbox"/>	Adobe Reader	Minor updates	All	Approved
<input checked="" type="checkbox"/>	Adobe Shockwave Player	Minor updates	All	Approved
<input type="checkbox"/>	Oracle Java Development Kit	—	—	—
<input checked="" type="checkbox"/>	Oracle Java Runtime Environment	Minor updates	All	Approved
<input checked="" type="checkbox"/>	Mozilla Firefox	Major updates	All	Approved
<input checked="" type="checkbox"/>	Mozilla Thunderbird	Major updates	All	Approved

Reset to default Cancel Save

Running the Test patching protection plan and decline unsafe patches

After patches are installed on the machines in your test environment, you can check if everything is working as expected. You can leave the approval status of the machines that have no issues as **Pending approval**, but change the approval status of the machines working incorrectly to **Declined**.

To run the Test patching protection plan and decline the patches that are not safe

1. Run the Test patching protection plan (by schedule or manually).
2. Depending on the result, see which of the installed patches are safe.
3. Go to **Software management > Patches** and set the **Approval status** to **Declined** for the patches that are not safe.

Use case for automatic patch approval without testing

If you want to automatically install new patches on your production machines as soon as possible, without installing them on test machines first, you can configure only one protection plan.

The use case consists of the following stages:

1. Configure the settings for Automatic patch approval. Select the **Automatic patch approval without testing** option. For more information, see "Configuring automatic patch approval" (p. 902).
2. Configure a protection plan for the production environment (for example, 'Production patching') with the enabled **Patch management** module and apply it to the machines in the production environment. Specify the following condition of patch installation: the patch status must be **Approved**. For more information, see "Configuring the Production patching protection plan" (p.

904).

3. Run the Production patching plan.

Approving patches manually

You can manually approve a patch and speed up its installation by skipping the testing phase.

Prerequisites

- A protection plan that has the **Patch management** module enabled is applied to at least one Windows machine.
- There are patches that are still not installed on the machine or machines on which the protection plan is applied.

To manually approve patches

1. In the Cyber Protect console, go to **Software management > Patches**.
2. Select the patches that you want to install, and then accept their license agreements.
3. Set the **Approval status** of the patches to **Approved**.

The approval status of the patches is set to **Approved**. The patches will be automatically installed on the machines based on the schedule defined in the protection plan. If you want to install the patches immediately, follow the procedure that is described in "Installing patches on demand" (p. 906).

Installing patches on demand

You can manually install patches on demand when you do not want to wait for the scheduled installation time.

You can start the manual patch installation from three screens: **Patches**, **Vulnerabilities**, and **All devices**.

To manually install a patch

From Patches

1. In the Cyber Protect console, go to **Software management > Patches**.
2. Accept the license agreements for the patches that you want to install.
3. In the **Install patches** wizard, select the patches that you want to install, and then click **Install**.
4. Select the machines on which you want to install the patches.
5. Select the reboot options.

- a. Select if you want the machine to be rebooted after the patches are installed.

Option	Description
No	The machines will not be rebooted automatically after the patches are installed.
If required	The machines will be rebooted only if it is required for applying the patches.
Yes	The machines will be rebooted automatically after the patches are installed. You can also specify a reboot delay.

- b. [Optional] If you want to delay the machine reboot while a backup of the machine is in progress, select **Do not reboot until backup is finished**.

6. Click **Install patches**.

From Vulnerabilities

1. In the Cyber Protect console, go to **Software management > Vulnerabilities**.
2. Perform the remediation process, as described in "Managing found vulnerabilities" (p. 892).

From All devices

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the machine on which you want to install the patches.
3. Click **Patch**.
4. Select the patches that you want to install, and then click **Next**.
5. Select the reboot options.

- a. Select if you want the machine to be rebooted after the patches are installed.

Option	Description
No	The machines will not be rebooted automatically after the patches are installed.
If required	The machines will be rebooted only if it is required for applying the patches.
Yes	The machines will be rebooted automatically after the patches are installed. You can also specify a reboot delay.

- b. [Optional] If you want to delay the machine reboot while a backup of the machine is in progress, select **Do not reboot until backup is finished**.

6. Click **Install patches**.

Managing your software and hardware inventory

Software inventory

The software inventory feature is available for devices on which the Advanced pack is enabled, or which have the (Legacy) Cyber Protect license. The feature enables you to view all the software applications that are installed on all Windows and macOS devices.

To obtain the software inventory data, you can run automatic or manual scans on the devices.

You can use the software inventory data to:

- browse and compare the information about all applications that are installed on the company devices
- determine if an application needs to be updated
- determine if an unused application needs to be removed
- ensure that the software version on multiple company devices is the same
- monitor changes in the software status between consecutive scans.

Enabling the software inventory scanning

When software inventory scanning is enabled on the devices, the system automatically collects the software data every 12 hours.

The Software inventory scanning feature is enabled by default for all devices that have the required license, but you can change the setting when necessary.

Note

Customer tenants can enable or disable the software inventory scanning. Unit tenants can only view the software inventory scanning settings, but cannot change them.

To enable the software inventory scanning

1. In the Cyber Protect console, go to **Settings**.
2. Click **Protection**.
3. Click **Inventory scanning**.
4. Enable the **Software inventory scanning** module by clicking the switch next to the module name.

To disable the software inventory scanning

1. In the Cyber Protect console, go to **Settings**.
2. Click **Protection**.
3. Click **Inventory scanning**.

4. Disable the **Software inventory scanning** module by clicking the switch next to the module name.

Running a software inventory scan manually

You can manually run a software inventory scan from the **Software inventory** screen, or from the **Software** tab in the **Inventory** screen.

Prerequisites

- The device uses Windows or macOS operating system.
- The device has the required (Legacy) Cyber Protect license or has the Advanced Management pack activated.

To run a software inventory scan from the Software inventory screen

1. In the Cyber Protect console, go to **Software management**.
2. Click **Software inventory**.
3. In the **Group by:** drop-down field, select **Devices**.
4. Find the device which you want to scan, and click **Scan now**.

To run a software inventory scan from the Software tab in the Inventory screen

1. In the Cyber Protect console, go to **Devices**.
2. Click the device which you want to scan, and click **Inventory**.
3. In the **Software** tab, click **Scan now**.

Browsing the software inventory

You can view and browse the data for all software applications that are available on all company devices.

Prerequisites

- The devices use Windows or macOS operating system.
- The devices have the required (Legacy) Cyber Protect license or have the Advanced Management pack activated.
- Software inventory scan on the devices has finished successfully.

To view all software applications that are available on all Windows and macOS company devices

1. In the Cyber Protect console, go to **Software Management**.
2. Click **Software inventory**.

By default, the data is grouped by device. The following table describes the data that is visible in the **Software inventory** screen.

Column	Description
Name	Name of the application.
Version	Version of the application.
Status	Status of the application. <ul style="list-style-type: none"> • New. • Updated. • Removed. • No Change.
Vendor	Vendor of the application.
Date installed	Date and time when the application was installed.
Last run	For macOS devices only. Date and time when the application was last active.
Location	Directory where the application is installed.
User	User who installed the application.
System type	For Windows devices only. Bit type of the application. <ul style="list-style-type: none"> • X86 for 32-bit applications. • X64 for 64-bit applications.

3. To group the data by application, in the **Group by:** drop-down field, select **Applications**.
4. To narrow the information displayed on the screen, use one or a combination of the filters.
 - a. Click **Filter**.
 - b. Select one or a combination of several filters.

The following table describes the filters in the **Software inventory** screen.

Filter	Description
Device Name	Device name. Multiple selection is possible. Use this filter if you want to compare the software on specific devices.
Application	Application name. Multiple selection is possible. Use this filter if you want to compare the data for a specific application on specific devices or on all devices.
Vendor	Vendor of the application. Multiple selection is possible. Use this filter if you want to view all applications from a specific vendor on specific devices or on all devices.
Status	Application status. Multiple selection is possible. Use this filter if you want to view all applications in the selected

Filter	Description
	status on specific devices or on all devices.
Date installed	Date when the application is installed. Use this filter if you want to view all applications that are installed on a specific date on specific devices or on all devices.
Scan date	Date of the software inventory scan. Use this filter if you want to view the information about the software on specific devices or on all devices that are scanned on that date.

- c. Click **Apply**.
5. To browse through the whole software inventory list, use the pagination in the lower left part of the screen.
 - Click the number of the page you want to open.
 - In the drop-down field, select the page number of the page you want to open.

Viewing the software inventory of a single device

You can view a list of all the software applications that are installed on a single device, as well as detailed information about the applications, such as status, version, vendor, installation date, last run, and location.

Prerequisites

- The device uses Windows or macOS operating system.
- The device has the required (Legacy) Cyber Protect license or has the Advanced Management pack activated.
- Software inventory scan on the device has finished successfully.

To view the software inventory of a single device from the Software Inventory screen

1. In the Cyber Protect console, go to **Software management**.
2. Click **Software inventory**.
3. In the **Group by:** drop-down field, select **Devices**.
4. Find the device you want to inspect using one of the following options.
 - Find the device using the **Filter**:
 - a. Click **Filter**.
 - b. In the **Device name** field, select the name of the device you want to view.
 - c. Click **Apply**.
 - Find the device using the dynamic **Search**:
 - a. Click **Search**.
 - b. Type the full device name or part of the device name.

To view the software inventory of a single device from Devices screen

1. In the Cyber Protect console, go to **Devices**.
2. Click the device which you want to view, and click **Inventory**.
3. Click the **Software** tab.

Hardware inventory

The hardware inventory feature enables you to view all the hardware components that are available on:

- physical Windows and macOS devices with a license that supports the Hardware inventory feature.
- virtual Windows and macOS machines running on the following virtualization platforms: VMware, Hyper-V, Citrix, Parallels, Oracle, Nutanix, Virtuozzo, and Virtuozzo Hybrid Infrastructure. For more information about the supported versions of the virtualization platforms, see "Supported virtualization platforms" (p. 35).

Note

The Hardware inventory feature for virtual machines is not supported in the Cyber Protect legacy editions.

The hardware inventory feature is supported only for devices on which a protection agent is installed.

To obtain the hardware inventory data, you can run automatic or manual scans on the devices.

You can use the hardware inventory data to:

- discover all hardware assets of the organization
- browse through the hardware inventory of all devices in your organization
- compare the hardware components on multiple company devices
- view detailed information about a hardware component.

Enabling the hardware inventory scanning

When hardware inventory scanning is enabled on physical devices and virtual machines, the system automatically collects the hardware data every 12 hours.

The hardware inventory scanning feature is enabled by default, but you can change the setting when necessary.

Note

Customer tenants can enable or disable the hardware inventory scanning. Unit tenants can only view the hardware inventory scanning settings, but cannot change them.

To enable the hardware inventory scanning

1. In the Cyber Protect console, go to **Settings**.
2. Click **Protection**.
3. Click **Inventory scanning**.
4. Enable the **Hardware inventory scanning** module by clicking the switch next to the module name.

To disable the hardware inventory scanning

1. In the Cyber Protect console, go to **Settings**.
2. Click **Protection**.
3. Click **Inventory scanning**.
4. Disable the **Hardware inventory scanning** module by clicking the switch next to the module name.

Running a hardware inventory scan manually

You can manually run a hardware inventory scan for a single device, and view the current data for the hardware components of the device.

Note

Hardware inventory scanning of virtual machines is supported only when the current date and time of the virtual machine corresponds to the current date and time in UTC. To ensure that the virtual machine uses the correct time settings, disable the **Time synchronization** option of the virtual machine, set the current date, time, and time zone, and then restart **Acronis Agent Core Service** and **Acronis Managed Machine Service**.

Prerequisites

- (For all devices) The device uses a Windows or macOS operating system.
- (For all devices) The devices have a license that supports the Hardware inventory feature. Note that the Hardware inventory feature for virtual machines is not supported in the (Legacy) Cyber Protect editions.
- (For all devices) A protection agent is installed on the device.
- (For virtual machines) The machine runs on one of the supported virtualization platforms. For more information, see "Hardware inventory" (p. 912).

To run a hardware inventory scan on a single device

1. In the Cyber Protect console, go to **Devices**.
2. Click the device which you want to scan, and click **Inventory**.
3. In the **Hardware** tab, click **Scan now**.

Browsing the hardware inventory

You can view and browse the data for all hardware components that are available on all company devices.

Prerequisites

- (For all devices) The devices use Windows or macOS operating system.
- (For all devices) The devices have a license that supports the Hardware inventory feature. Note that the Hardware inventory feature for virtual machines is not supported in the Cyber Protect legacy editions.
- (For all devices) A protection agent is installed on the device.
- (For all devices) Hardware inventory scan on the devices has finished successfully.
- (For virtual machines) The machine runs on one of the supported virtualization platforms. For more information, see "Hardware inventory" (p. 912).

To view all hardware components that are available on the Windows and macOS company devices

1. In the Cyber Protect console, go to **Devices**.
2. In the **View:** drop-down field, select **Hardware**.

Note

The view is a set of columns which determines what data is visible in the screen. The predefined views are **Standard** and **Hardware**. You can create and save custom views which include different sets of columns, and are more convenient for your needs.

The following table describes the data that is visible in the **Hardware** view.

Column	Description
Name	Device name.
Hardware scan status	Status of the hardware scan. <ul style="list-style-type: none">• Completed.• Not started.• Not supported. status is shown for workloads for which hardware inventory functionality is not supported, i.e. virtual machines, mobile devices, Linux devices.• Update agent. shown in case the outdated version of agent is installed on the device. Clicking on this action will redirect to Settings > Agents page, where admin can perform the agent update.• Upgrade quota. Clicking on it will open a dialog where admin can switch the current license to one of other available for tenant licenses
Processor	Models of all processors of the device.

Column	Description
Processor cores	Number of cores of all processors of the device.
Disk storage	Used storage, and total storage of all the disks of the device.
Memory	Total RAM capacity of the device.
Scan date	Date and time of the last hardware inventory scan.
Motherboard	Motherboard of the device.
Motherboard serial number	Serial number of the motherboard.
BIOS version	Version of the BIOS of the system.
Organization	Organization to which the device belongs.
Owner	Owner of the device.
Domain	Domain of the device.
Operating system	Operating system of the device.
Operating system build	Build of the operating system of the device.

3. To add columns in the table, click the column options icon, and select the columns that you want to be visible in the table.
4. To narrow the information displayed on the screen, use one or more filters.
 - a. Click **Search**.
 - b. Click the arrow, and then click **Hardware**.
 - c. Select one or a combination of several filters.

The following table describes the **Hardware** filters.

Filter	Description
Processor model	Multiple selection is possible. Use this filter if you want to view the hardware data of the devices which have the specified processor model.
Processor cores	Use this filter if you want to view the hardware data of the devices which have the specified number of processor cores.
Disk total size	Use this filter if you want to view the hardware data of the devices which have the specified total storage size.
Memory capacity	Use this filter if you want to view the hardware data of the devices which have the specified RAM capacity.

- d. Click **Apply**.
5. To sort the data in an ascending order, click a column name.

Viewing the hardware of a single device

You can view detailed information about the motherboard, processors, memory, graphics, storage drives, network, and system of a specific device.

Prerequisites

- (For all devices) The device uses Windows or macOS operating system.
- (For all devices) The devices have a license that supports the Hardware inventory feature. Note that the Hardware inventory feature for virtual machines is not supported in the Cyber Protect legacy editions.
- (For all devices) A protection agent is installed on the device.
- (For all devices) Hardware inventory scan on the device has finished successfully.
- (For virtual machines) The machine runs on one of the supported virtualization platforms. For more information, see "Hardware inventory" (p. 912).

To view the detailed information about the hardware of a specific device

1. In the Cyber Protect console, go to **Devices->All Devices**.
2. In the **View:** drop-down field, select **Hardware**.
3. Find the device you want to inspect using one of the methods described below.
 - Find the device using the **Filter:**
 - a. Click **Filter**.
 - b. Select one or a combination of several filter parameters to find the device.
 - c. Click **Apply**.
 - Find the device using the **Search:**
 - a. Click **Search**.
 - b. Type the full device name or part of the device name, and click **Enter**.
4. Click the row listing the device, and click **Inventory**.
5. Click the **Hardware** tab.

The following hardware data is available.

Hardware component	Information displayed
Motherboard	Name, manufacturer, model, and serial number of the motherboard of the device.
Processors	Manufacturer, model, max clock speed, and number of cores of each processor of the device.
Memory	Capacity, manufacturer, and serial number of the memory of the device.
Graphics	Manufacturer and model of the GPUs of the

Hardware component	Information displayed
	device.
Storage drives	Model, media type, available space and size of the storage drives of the device.
Network	Mac address, IP address, and type of the network adapters of the device.
System	Product ID, original install date, system boot time, system manufacturer, system model, BIOS version, boot device, system locale, and time zone of the system.

Connecting to workloads for remote desktop or remote assistance

The remote desktop and assistance functionality is a convenient way to connect to the workloads in your organization for remote control or remote assistance. Starting from December 2022, the functionality supports the NEAR, RDP, and Apple Screen Sharing protocols. For more information, see "Remote connection protocols" (p. 923).

You can use the remote desktop functionality to perform the following tasks.

- Connect to remote Windows, macOS, and Linux workloads by using NEAR in view-only mode.
- Connect to remote Windows workloads by using RDP.
- Connect to remote macOS workloads by using Apple Screen Sharing in view-only or curtain mode.
- Connect to managed workloads and remotely control them by using cloud remote connections.
- Connect to unmanaged workloads and remotely control them by using direct remote connections.
- Connect to unmanaged remote workloads by using Acronis Quick Assist.
- Connect to remote workloads by using different authentication methods: with remote workload credentials, by asking for permission to observe or control, or with an access code (for Quick Assist).
- Observe multiple monitors at the same time in multi-view.
- Record remote sessions (when connected via NEAR).
- View the session history report.

For more information about the features that are part of the Standard and Advanced Management packs, see "Supported remote desktop and assistance features" (p. 919).

You can use the remote assistance functionality to perform the following tasks.

- Connect to remote Windows, macOS, and Linux workloads by using NEAR in control mode.
- Connect to remote macOS workloads by using Apple Screen Sharing in control mode.
- Provide remote assistance to workloads by using cloud remote connections.
- Transfer files between the local and remote workloads.
- Perform basic management actions on the remote workload: restart, shut down, sleep, empty recycle bin, and log out the remote user.
- Monitor the remote workload by periodically taking screenshots of its desktop.

For more information about the features that are part of the Standard protection and Advanced Management, see "Supported remote desktop and assistance features" (p. 919).

Important

To activate the complete remote desktop and assistance functionality for a managed workload, you must configure and apply a remote management plan to the workload. Although you can apply only one remote management plan on a workload, depending on your needs, you can configure different remote management plans and apply them to different workloads.

For example, you can create a remote management plan that has only the RDP protocol enabled and apply it to some workloads. In that way, you will be able to remotely connect to these workloads without activating the Advanced Management license per workload, and without paying any additional fees.

On the other hand, you can create another remote management plan that has the NEAR and Apple Screen Sharing protocols enabled. In this case, the Advanced Management license per workload will be activated, and you will be charged for each workload to which this remote management plan is applied.

For more information about remote management plans and working with them, see "Remote management plans" (p. 926).

Note

The remote desktop and assistance functionality requires:

- a one-time installation of Connect Client on the managing (host) workload. The system will suggest you to download the client when you attempt performing a remote action (remote control or remote assistance) on a target workload for the first time. Alternatively, you can download Connect Client from the **Downloads** window in the Protection console. For more information about the settings that you can configure, see "Configuring the Connect Client settings" (p. 955).
- installation of Connect Agent on the managed workloads. The Connect Agent is a module that is part of the Protection agent, starting from version 15.0.31266.
- for macOS remote workloads, the required system permissions should be granted to the Connect Agent. For more information, see "Installing protection agents in macOS" (p. 83).
- running the Acronis Quick Assist application on the unmanaged workloads. You can download Acronis Quick Assist from [the website](#).

For more information about the supported platforms by each remote desktop and assistance component, see "Supported platforms" (p. 922).

Supported remote desktop and assistance features

The following table provides more information about the changes of the supported features of the remote desktop and assistance functionality that were introduced in December 2022.

Feature	Standard protection before Dec 2022	Advanced Management before Dec 2022	Standard protection after Dec 2022	Advanced Management after Dec 2022
Remote assistance via RDP for Windows	Yes	No	No	No
Share a remote connection with users	No	Yes	No	No
Remote connections				
Remote actions	No	No	Yes	Yes
Selecting a session for Windows/macOS/Linux to connect	No	No	No	Yes
Direct connect via RDP and Apple Screen Sharing	No	No	No	Yes
Multi-window control	No	No	No	Yes
Connection modes: Control/View-only/Curtain	No	No	No	Yes
Common credentials support for remote connections	No	No	Yes	Yes
Concurrent connections per technician				
via RDP	Yes	Yes	Yes	Yes
via NEAR	No	No	No	Yes
Files transfer and sharing				
from Windows to Windows/macOS/Linux	No	No	No	Yes
from macOS to Windows/macOS/Linux	No	No	No	Yes
from Linux to Windows/macOS/Linux	No	No	No	Yes
Connecting via Quick Assist application				
from Windows to	No	No	No	Yes

Feature	Standard protection before Dec 2022	Advanced Management before Dec 2022	Standard protection after Dec 2022	Advanced Management after Dec 2022
Windows/macOS/Linux				
from macOS to Windows/macOS/Linux	No	No	No	Yes
from Linux to Windows/macOS/Linux	No	No	No	Yes
Remote connections via protocols				
Remote connection via NEAR				
from Windows to Windows/macOS/Linux	No	No	No	Yes
from macOS to Windows/macOS/Linux	No	No	No	Yes
from Linux to Windows/macOS/Linux	No	No	No	Yes
Remote connection via RDP (desktop client)				
from Windows to Windows	Yes	Yes	Yes	Yes
from macOS to Windows	Yes	Yes	Yes	Yes
from Linux to Windows	No	No	Yes	Yes
Remote connection via RDP (web client)				
from Windows to Windows	Yes	Yes	Yes	Yes
from macOS to Windows	Yes	Yes	Yes	Yes
from Linux to Windows	No	No	Yes	Yes
Remote connection via Apple Screen Sharing				
from Windows/macOS/Linux to macOS	No	No	No	Yes
Session management				
Session recording	No	No	No	Yes

Feature	Standard protection before Dec 2022	Advanced Management before Dec 2022	Standard protection after Dec 2022	Advanced Management after Dec 2022
Reporting and monitoring				
Session history and search	No	No	No	Yes
Screenshot transmission	No	No	No	Yes

Supported platforms

The following table lists the supported operating systems by each of component of the remote desktop and assistance functionality.

Remote desktop component	Supported platforms
Connect Client	<ul style="list-style-type: none"> • Windows 7 or later • macOS 10.13 or later • Linux: <ul style="list-style-type: none"> openSUSE 8 Debian 9, 10 Ubuntu 18.0-20.10 Red Hat Enterprise Linux 8 CentOS 8 Fedora 31-33 SUSE Linux Enterprise Server 15 SP2 Linux Mint 20 Manjaro 20
Connect Agent	<ul style="list-style-type: none"> • Windows 7 or later • Windows Server 2008 R2 or later • macOS 10.13 or later • Linux: <ul style="list-style-type: none"> Red Hat Enterprise Linux 8, 8.1 Fedora 30 Ubuntu 18.4 LTS (Bionic Beaver) -19.04 (Disco Dingo) Debian 9, 10 CentOS 8 openSUSE 15.1
Acronis Quick Assist	<ul style="list-style-type: none"> • Windows 7 or later • Windows Server 2008 R2 or later • macOS 10.13 or later

Remote desktop component	Supported platforms
	<ul style="list-style-type: none"> • Linux: Red Hat Enterprise Linux 8, 8.1 Fedora 30 Ubuntu 18.4 LTS (Bionic Beaver) -19.04 (Disco Dingo) Debian 9, 10 CentOS 8 openSUSE 15.1

Remote connection protocols

The remote desktop functionality uses the following protocols for remote connections.

NEAR

NEAR is a highly secure protocol developed by Acronis that has the following characteristics.

- **H.264**

NEAR implements three quality modes: **Smooth**, **Balanced** and **Sharp**. In **Smooth** mode, NEAR uses hardware H.264 encoding on macOS and Windows to encode the desktop picture, and falling back to software encoder if hardware encoder is not available. The picture size is currently limited to Full HD resolution (1920x1080).

- **Adaptive codec**

In **Balanced** and **Sharp** quality modes, NEAR uses Adaptive codec, which provides full picture quality in 32 bit, compared to the 'video' mode used by H.264.

In **Balanced** mode, the picture quality is automatically adjusted according to your current network conditions and retains the current framerate.

In **Sharp** mode, the picture is full quality, but it might be with a reduced framerate, if your network, processor, or video card are overloaded.

Adaptive codec is using OpenCL on Windows and macOS when it is available in their graphics drivers.

- **Sound transfer**

NEAR is capable of capturing the remote computer sound and transfer that to host. For more information about enabling remote sound redirection on Windows, macOS, and Linux, see "Remote sound redirection" (p. 924).

- **Different login options**

You can use the following methods to log in to the remote workload.

Access code: the user who is logged in to the remote workload runs Quick Assist and tells you the access code. With this method, you always connect to the session of the currently logged in user.

Workload credentials: log in to the remote workload using administrator credentials that are registered in the workload.

Ask for permission to observe or control: the user who is logged in to the remote workload will be asked to allow or deny the connection.

- **Security**

Your data is always two-way encrypted with AES encryption in NEAR.

RDP

Remote Desktop Protocol (RDP) is a proprietary protocol developed by Microsoft that enables connecting to the remote Windows computer over a network connection.

Apple Screen Sharing

Apple Screen Sharing is a VNC client by Apple included as part of macOS version 10.5 and later.

Remote sound redirection

Connect Client supports audio streaming via the NEAR connection protocol. For more information about NEAR, see "Remote connection protocols" (p. 923).

Redirecting sound from a remote Windows workload

For Windows workloads, the remote sound should be transmitted automatically. Ensure that there are sound output devices (speakers or headphones) connected to the remote workload.

Redirecting sound from a remote macOS workload

To enable sound redirection from a macOS workload, ensure that:

- The workload has the Protection agent installed.
- The workload has a sound capture driver installed.
- The workload uses the NEAR protocol for remote connections.

Note

For macOS 10.15 Catalina, the Microphone permission must be granted to the Connect Agent. For more information about granting the Microphone permission to the Connect Agent, see "Granting the required system permissions to the Connect Agent" (p. 84).

The agent works with the following sound capture drivers: Soundflower or Blackhole.

The installation process on the newest versions is described on the Blackhole wiki page:
<https://github.com/ExistentialAudio/BlackHole/wiki/Installation>.

Note

Connect Client currently supports only the 2-channel version of Blackhole.

Alternatively, if Homebrew is installed on the workload, you can install Blackhole by running the following command:

```
brew install --cask blackhole-2ch
```

Note

While the sound of a remote macOS workload is redirected, the user who is logged in to the remote workload will not hear the sound.

Redirecting sound from a remote Linux workload

The remote sound redirection should work automatically with most Linux distributions. If the remote sound redirection is not working by default, install PulseAudio driver by running the following command:

```
sudo apt-get install pulseaudio
```

Connections to remote workloads for remote desktop or remote assistance

The remote desktop and assistance functionality provides several ways to establish remote direct or cloud connections to your workloads.

Direct connections are established via TCP/IP in the local area network (LAN) between Connect Client and the remote workload that does not have an agent installed. They do not require Internet access.

Cloud connections are established between Connect Client and the agent or Quick Assist on the workload via Acronis Cloud.

The following table provides more info about the cloud connection options.

Cloud connection	Cloud connection option	View mode	Supported remote action	Available for
via NEAR	from Connect Client to Connect Agent from Connect Client to Quick Assist	Control View-only	Remote desktop Remote assistance	managed workloads
via RDP	from Connect Client to	Control	Remote desktop	managed

Cloud connection	Cloud connection option	View mode	Supported remote action	Available for
	Connect Agent from web client to Connect Agent			workloads
via Apple Screen Sharing	from Connect Client to Connect Agent	Control View-only Curtain	Remote desktop Remote assistance	managed workloads

The following table provides more info about the direct connection options.

Direct connection	Direct connection option	Supported remote action	Available for
via RDP	from Connect Client to RDP server	Remote desktop	unmanaged workloads
via Apple Screen Sharing	from Connect Client to Apple Screen Sharing server	Remote desktop Remote assistance	unmanaged workloads

Remote management plans

Remote management plans are plans that you apply on the Protection agent to enable and configure the remote desktop and assistance functionality on your managed workloads.

If no remote management plan is applied on a workload, the remote desktop and assistance functionality will be limited to remote actions (restart, shut down, sleep, empty recycle bin, and log out remote user).

Note

The availability of the settings that you can configure in the remote management plan depends on the service pack that is applied on the tenant. To access all settings, activate the Advanced Management pack. For more information about the features that are part of the Standard and Advanced Management packs, see "Supported remote desktop and assistance features" (p. 919).

Creating a remote management plan

You can create a remote management plan, and then assign it to a workload to configure the remote desktop and assistance functionality on the managed workload.

Note

The availability of the remote management plan's settings depends on the service quota that is assigned to the tenant. If you are using the standard functionality, you can only configure connections via RDP.

Prerequisites

2FA is enabled for your user account.

To create a remote management plan

From Remote management plans

1. In the Cyber Protect console, go to **Management > Remote management plans**.
2. Create a remote management plan by using one of the two options.
 - If there are no remote management plans in the list, click **Create**.
 - If there are remote management plans in the list, click **Create plan**.
3. [Optional] To change the default name of the plan, click the pencil icon, enter the name of the plan, and then click **Proceed**.
4. Click **Connection protocols**, and enable the protocols that you want to be available in this remote management plan for remote connections - NEAR, RDP, or Apple Screen Sharing.
5. [Optional] For the NEAR protocol, in the **Security settings** section, select or clear the check boxes to enable or disable the corresponding setting, and then click **Done**.

Setting	Description	Available for
Lock the workload when the user disconnects from the console session	If you select this setting, the remote workload will be locked when you disconnect from the console session.	Windows, macOS
Allow only one user at a time to connect using NEAR or to transfer files	If you select this setting, connections using NEAR and file transfers will not be possible while there is an active remote connection to the workload.	Windows, macOS, Linux
Allow the workload's administrator to connect to any non-admin user session	If you select this setting, the administrator will be allowed to connect to any standard user session on the workload. If both Allow the workload's administrator to connect to any non-	Windows, macOS

Setting	Description	Available for
	admin user session and Allow system session creation are clear, you will only be able to connect to active administrator sessions on the remote macOS workloads.	
Allow system session creation	If you select this setting, when establishing remote connections, the administrator will connect in a new session, instead of one of the existing active sessions.	macOS
Allow clipboard synchronization	If you select this setting, you will be able to transfer data between your clipboard and the clipboard of the remote workload. For example, you will be able to copy some text from a file on the remote workload and paste it in a file on your workload, and the opposite.	Windows, macOS, Linux

- Click **Security settings**, select or clear the check boxes to enable or disable the corresponding setting, and then click **Done**.

Setting	Description
Show if the workload is controlled remotely	If you select this setting, a notification will be displayed on the desktop of the remote workload when there is an active remote desktop connection to the workload.
Ask for the user's permission to take screenshots of the workload	If you select this setting, the user of the remote workload will be notified when the administrator requests screenshot transmission from the workload.

- Click **Workload management**, select the features that you want to be available on the remote workloads, and then click **Done**.

Setting	Description	Available on
File transfer	Enables the file transfers between local and remote workloads.	Windows, macOS, Linux
Screenshot transmission	Enables the transmission of screenshots of the desktop of the remote workload to the Cyber Protect console.	Windows, macOS, Linux

8. Click **Display settings**, select or clear the check boxes to enable or disable the corresponding setting, and then click **Done**.

Note

The **Display settings** are only available for connections via NEAR.

Setting	Description	Available on
Use Desktop Deduplication for desktop capturing	Desktop duplication is one of the screen capture methods on Windows. In some environments, it might be unstable. If you do not use Desktop deduplication, you will use the basic method (BitBlt) instead- it is much slower, but more stable.	Windows
Use OpenCL acceleration	OpenCL acceleration can speed up the Adaptive codec, which is responsible for the Balanced quality mode, by running some computations on the graphics processing unit (GPU). This requires an installation of an OpenCL driver on the remote Linux. Adaptive Codec is using OpenCL on macOS and Windows, if it is available in your graphics drivers.	Linux
Use hardware H.264 encoding	NEAR supports three quality modes: Smooth , Balanced , and Sharp . Smooth mode uses	Windows, macOS

Setting	Description	Available on
	<p>hardware H.264 encoding to encode the desktop picture.</p> <p>Balanced mode uses Adaptive codec, which provides full picture quality in 32 bit, compared to the 'video' mode used by H.264. The picture quality is automatically adjusted according to your current network conditions and retains the current framerate.</p> <p>Sharp mode uses Adaptive codec, which provides full picture quality in 32 bit, compared to the 'video' mode used by H.264. The picture quality is always full, but it might be with reduced frames per seconds, if your network or processor/video card are overloaded.</p>	

9. If you want the information about the users who last logged in to the workloads to be visible in the workload's details, click **Toolbox**, select **Show last logged-in users**, and then click **Done**.
For more information about the last logged-in users, see "Find the last logged in user" (p. 374).
10. [Optional] To add workloads to the plan:
 - a. Click **Add workloads**.
 - b. Select the workloads, and then click **Add**.
 - c. If there are compatibility issues that you want to resolve, follow the procedure as described in "Resolving compatibility issues with remote management plans" (p. 938).
11. Click **Create**.

From All devices

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click the workload to which you want to apply a remote management plan.
3. Click **Protect**, and then click **Add plan**.
4. Click **Create plan**, and select **Remote management**.
5. [Optional] To change the default name of the plan, click the pencil icon, enter the name of the plan, and then click **Proceed**.

6. Click **Connection protocols**, and enable the protocols that you want to be available in this remote management plan for remote connections - NEAR, RDP, or Apple Screen Sharing.
7. [Optional] For the NEAR protocol, in the **Security settings** section, select or clear the check boxes to enable or disable the corresponding setting, and then click **Done**.

Setting	Description	Available for
Lock the workload when the user disconnects from the console session	If you select this setting, the remote workload will be locked when you disconnect from the console session.	Windows, macOS
Allow only one user at a time to connect using NEAR or to transfer files	If you select this setting, connections using NEAR and file transfers will not be possible while there is an active remote connection to the workload.	Windows, macOS, Linux
Allow the workload's administrator to connect to any non-admin user session	If you select this setting, the administrator will be allowed to connect to any standard user session on the workload. If both Allow the workload's administrator to connect to any non-admin user session and Allow system session creation are clear, you will only be able to connect to active administrator sessions on the remote macOS workloads.	Windows, macOS
Allow system session creation	If you select this setting, when establishing remote connections, the administrator will connect in a new session, instead of one of the existing active sessions.	macOS
Allow clipboard synchronization	If you select this setting, you will be able to transfer data between your clipboard and the clipboard of the remote workload. For example, you	Windows, macOS, Linux

Setting	Description	Available for
	will be able to copy some text from a file on the remote workload and paste it in a file on your workload, and the opposite.	

8. Click **Security settings**, select or clear the check boxes to enable or disable the corresponding setting, and then click **Done**.

Setting	Description
Show if the workload is controlled remotely	If you select this setting, a notification will be displayed on the desktop of the remote workload when there is an active remote desktop connection to the workload.
Ask for the user's permission to take screenshots of the workload	If you select this setting, the user of the remote workload will be notified when the administrator requests screenshot transmission from the workload.

9. Click **Workload management**, select the features that you want to be available on the remote workloads, and then click **Done**.

Setting	Description	Available on
File transfer	Enables the file transfers between local and remote workloads.	Windows, macOS, Linux
Screenshot transmission	Enables the transmission of screenshots of the desktop of the remote workload to the Cyber Protect console.	Windows, macOS, Linux

10. Click **Display settings**, select or clear the check boxes to enable or disable the corresponding setting, and then click **Done**.

Note

The **Display settings** are only available for connections via NEAR.

Setting	Description	Available on
Use Desktop Deduplication for desktop capturing	Desktop duplication is one of the screen capture methods on Windows. In some environments, it might be	Windows

Setting	Description	Available on
	unstable. If you do not use Desktop deduplication, you will use the basic method (BitBlt) instead- it is much slower, but more stable.	
Use OpenCL acceleration	OpenCL acceleration can speed up the Adaptive codec, which is responsible for the Balanced quality mode, by running some computations on the graphics processing unit (GPU). This requires an installation of an OpenCL driver on the remote Linux. The Adaptive Codec is using OpenCL on macOS and Windows, if it is available in your graphics drivers.	Linux
Use hardware H.264 encoding	NEAR supports three quality modes: Smooth , Balanced , and Sharp . Smooth mode uses hardware H.264 encoding to encode the desktop picture. Balanced mode uses Adaptive codec, which provides full picture quality in 32 bit, compared to the 'video' mode used by H.264. The picture quality is automatically adjusted according to your current network conditions and retains the current framerate. Sharp mode uses Adaptive codec, which provides full picture quality in 32 bit, compared to the 'video' mode used by H.264. The picture quality is always full, but it might be with reduced frames per seconds, if your	Windows, macOS

Setting	Description	Available on
	network or processor/video card are overloaded.	

- If you want the information about the users who last logged in to the workloads to be visible in the workload's details, click **Toolbox**, select **Show last logged-in users**, and then click **Done**.
For more information about the last logged-in users, see "Find the last logged in user" (p. 374).
- Click **Create**.

Adding a workload to a remote management plan

Depending on your needs, you can add workloads to a remote management plan after the plan was created.

Prerequisites

2FA is enabled for your user account.

To add a workload to a remote management plan

From Remote management plans

- In the Cyber Protect console, go to **Management > Remote management plans**.
- Click the remote management plan.
- Depending on whether or not the plan was already applied to any workload, do the following:
 - Click **Add workloads**, if the plan was not applied to any workloads yet.
 - Click **Manage workloads**, if the plan was applied to any workloads.
- Select a workload from the list, and then click **Add**.
- Click **Save**.
- Click **Confirm** to apply the required service quota to the workload.

From All devices

- In the Cyber Protect console, go to **Devices > All devices**.
- Click the workload to which you want to apply a remote management plan.
- Click **Protect**, and then click **Add plan**.
- In **Select a plan from the list below**, select **Remote management** to view only the remote management plans.
- Click **Apply**.
- Click **Confirm** to apply the required service quota to the workload.

Removing workloads from a remote management plan

Depending on your needs, you can remove workloads from a remote management plan.

Prerequisites

2FA is enabled for your user account.

To remove workloads from a remote management plan

1. In the Cyber Protect console, go to **Management > Remote management plans**.
2. Click the remote management plan.
3. Click **Manage workloads**.
4. Select one or several workloads that you want to remove from the remote management plan, and then click **Remove**.
5. Click **Done**.
6. Click **Save**.

Additional operations with existing remote management plans

From the **Remote management plans** screen, you can perform the following additional operations with remote management plans: view details, edit, view the activities, view the alerts, rename, enable, disable, clone, export, and delete.

View details

Prerequisites

2FA is enabled for your user account.

To view the details of a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **View details**.

Edit

Prerequisites

2FA is enabled for your user account.

To edit a plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Edit**.

Activities

To view the activities related to a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.

2. Click **Activities**.
3. Click an activity to view more details about it.

Alerts

To view the alerts

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Alerts**.

Rename

Prerequisites

2FA is enabled for your user account.

To rename a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Rename**.
3. Enter the new name of the plan, and then click **Proceed**.

Enable

Prerequisites

2FA is enabled for your user account.

To enable a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Enable**.

Disable

Prerequisites

2FA is enabled for your user account.

To disable a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Disable**.

Clone

Prerequisites

2FA is enabled for your user account.

To clone a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Clone**.
3. Click **Create**.

Export

Prerequisites

2FA is enabled for your user account.

To export a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Export**.
The plan configuration is exported in a JSON format to the local machine.

Delete

Prerequisites

2FA is enabled for your user account.

To delete a remote management plan

1. In the **Remote management plans** screen, click the **More actions** icon of the remote management plan.
2. Click **Delete**.
3. Select **I confirm**, and then click **Delete**.

Compatibility issues with remote management plans

In some cases, applying a remote management plan on a workload might cause compatibility issues. You might observe the following compatibility issues:

- **Conflicting plans** - this issue appears when another remote management plan is already applied on the workload, as only one remote management plan can be applied on a workload.
- **Incompatible operating system**- this issue appears when the workload's operating system is not supported.
- **Unsupported agent** - this issue appears when the version of the protection agent on the workload is outdated and does not support the remote desktop functionality.
- **Insufficient quota** - this issue appears when there is not enough service quota in the tenant to assign to the selected workloads.

If the remote management plan is applied to up to 150 individually selected workloads, you will be prompted to resolve the existing conflicts before saving the plan. To resolve a conflict, remove the

root cause for it or remove the affected workloads from the plan. For more information, see "Resolving compatibility issues with remote management plans" (p. 938). If you save the plan without resolving the conflicts, it will be automatically disabled for the incompatible workloads, and alerts will be shown.

If the remote management plan is applied to more than 150 workloads or to device groups, first it will be saved, and then checked for compatibility. The plan will be automatically disabled for the incompatible workloads, and alerts will be shown.

Resolving compatibility issues with remote management plans

Depending on the cause of the compatibility issues, you can perform different actions to resolve the compatibility issues as a part of the process of creating a new remote management plan.

Note

When resolving a compatibility issue by removing workloads from a plan, you cannot remove workloads that are part of a device group.

To resolve the compatibility issues

1. Click **Review issues**.
2. [To resolve compatibility issues with existing remote management plans by removing workloads from the new plan]
 - a. On the **Conflicting plans** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
3. [To resolve compatibility issues with remote management plans by disabling the plans that are already applied on the workloads]
 - a. Click **Disable applied plans**.
 - b. Click **Disable**, and then click **Close**.
4. [To resolve compatibility issues with incompatible operating systems]
 - a. On the **Incompatible operating system** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
5. [To resolve compatibility issues with unsupported agents by removing workloads from the plan]
 - a. On the **Unsupported agents** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
6. [To resolve compatibility issues with unsupported agents by updating the agent version] Click **Go to the Agents list**.

Note

This option is available only for customer administrators.

7. [To resolve compatibility issues with insufficient quota by removing workloads from the plan]
 - a. On the **Insufficient quota** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
8. [To resolve compatibility issues with insufficient quota by increasing the quota of the tenant]

Note

This option is available only for partner administrators.

- a. On the **Insufficient quota** tab, click **Go to the Management portal**.
- b. Increase the service quota for the customer.

Workload credentials

You can add administrator or non-administrator credentials of the remote workloads (username and password, or VNC password), save them in the cloud credentials store, and then use them for automatic authentication when connecting to the workloads that you manage. In that way, instead of entering these credentials manually every time during the authentication step of the connection, you can save them in the credentials store once, assign them to multiple workloads, and then the Connect Client will use these credentials every time you want to connect to the workloads remotely.

Note

The credentials that are stored in the credentials store are not shared between different tenant levels. They are shared only on the same tenant level for the same customer tenant or partner tenant.

This means that if a customer tenant has several administrators, they will see and share the credentials in the credentials store, while any other partner administrators, or customer administrators of other tenants will not be able to view or use these credentials.

Adding credentials

You can add credentials and then use them for remote connections to multiple workloads.

To add credentials to a workload and save them in the Credentials store

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click the workload for which you want to add credentials.
3. Access the **Settings** menu in one of the following ways:
 - Click **Remote desktop**, and then click **Settings**.
 - Click **Manage**, and then click **Settings**.
4. Click **Add credentials**.
5. In the **Credentials store**, click **Add credentials**.

6. Enter the credentials.

Field	Description
Credentials name	Identifier of the credentials that will be visible in the credentials store.
Username	Username that will be used for remote connections to the target workload.
Password	Password that will be used for remote connections to the target workload.
VNC password	This field is available for Apple Screen Sharing only.

7. Click **Save**.

Assigning credentials to a workload

After you add credentials, you can use them to authenticate automatically when you connect to a workload that you manage.

To assign saved credentials for automatic authentication to a workload

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Access the **Settings** menu in one of the following ways:
 - Click **Remote desktop**, and then click **Settings**.
 - Click **Manage**, and then click **Settings**.
3. On the tab of the supported protocol (NEAR, RDP, or Apple Screen Sharing), click **Add credentials**.
4. In the **Credentials store**, select the credentials from the list, and then click **Select credentials**.

Deleting credentials

You can delete credentials that are not needed anymore.

To delete credentials from the Credentials store

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Access the **Settings** menu in one of the following ways:
 - Click **Remote desktop**, and then click **Settings**.
 - Click **Manage**, and then click **Settings**.
3. On the tab of the supported protocol (NEAR, RDP, or Apple Screen Sharing), click **Delete**.
4. In the confirmation window, click **Delete**.

Unassigning credentials from a workload

You can unassign credentials from a workload, but still keep them in the Credentials store.

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Access the **Settings** menu in one of the following ways:
 - Click **Remote desktop**, and then click **Settings**.
 - Click **Manage**, and then click **Settings**.
3. On the tab of the supported protocol (NEAR, RDP, or Apple Screen Sharing), click **Unassign**.
4. In the confirmation window, click **Unassign**.

Working with managed workloads

Managed workloads are workloads on which the Protection agent is installed.

You can perform the following actions on the remote managed workloads:

- connect for remote assistance or remote desktop by using NEAR in control or view-only mode
- connect for remote desktop by using RDP in control mode
- connect for remote assistance or remote desktop by using Apple Screen Sharing in control, view-only, or curtain mode
- connect for remote desktop via a web client
- restart, shut down, sleep, empty recycle bin, log out remote user from the remote workloads
- transfer files between your workload and the remote workloads
- monitor them by taking screenshots

Note

The remote desktop connections to managed workloads require installing a Protection agent and applying a remote management plan on the workload.

Configuring RDP settings

You can configure the settings that will be applied automatically for remote control RDP connections to the managed workload.

To configure the RDP settings of a workload

1. In the Cyber Protect console, go to **Devices > Machines with agents**.
2. Access the **Settings** menu in one of the following ways:
 - Click **Remote desktop**, and then click **Settings**.
 - Click **Manage**, and then click **Settings**.

- On the **RDP** tab, configure the settings.

Setting	Description
Audio playback	This setting enables or disables the redirection of the remote workload sound on your local workload.
Audio recording	This setting determines whether audio recording (speaking to the microphone) will be transferred to the remote workload.
Redirect printers	If you select this setting, the printers from your workload will be available on the remote workload.
Redirect files	This setting defines whether files from your local workload will be shared to remote workload.
Color depth	This setting determines the number of colors in the picture that RDP will transfer. Higher value requires higher bandwidth. High color: 16 bit True color: <ul style="list-style-type: none"> 24 bit for RDP connections via the web client 32 bit for RDP connections via Connect Client

- Click the close button.

Connecting to managed workloads for remote desktop or remote assistance

Note

The availability of the connection protocols that you can use for remote connections depends on the remote management plan configuration and on the remote workload's operating system.

Prerequisites

- A remote management plan with the corresponding connection protocol enabled is applied on the managed workload.
- The required service quota is assigned to workload. (The service quota is automatically acquired when you apply a remote management plan to workload.)
- [For Apple Screen Sharing connections] Apple Screen Sharing is enabled on the macOS workload.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

To remotely connect to a managed workload for remote desktop or remote assistance

- In the Cyber Protect console, go to **Devices > Machines with agents**.
- Click the workload to which you want to connect.
- Click **Remote desktop**.

By default, NEAR is selected as a connection protocol.

4. [Optional] In the **Connection protocol** drop-down list, select the connection protocol that you want to use.
5. Click the view mode that you want to use.

Protocol	Remote connections to	View mode	Supported remote action
NEAR	Windows Linux macOS	<p>Control- In this mode, you will be able to observe and perform operations on the remote workload.</p> <p>View-only- In this mode, you will be only able to observe the remote workload.</p>	Remote desktop Remote assistance
RDP	Windows	<p>Control- In this mode, you will be able to view and perform operations on the remote workload.</p> <hr/> <p>Note If RDP is disabled in the OS settings of the workload, a pop-up window will appear. Use this window to enable RDP for the workload for the current session or in general:</p> <ul style="list-style-type: none"> • If you want to enable RDP for this workload only for the current session, select Disable it after the session is over, and then click Allow. • If you want to enable RDP for this workload, click Allow. <hr/>	Remote desktop
Apple Screen Sharing	macOS	<p>Control- In this mode, you will be able to observe and perform operations on the remote workload.</p> <p>View-only- In this mode, you will be only able to observe the remote workload.</p> <p>Curtain - available only for macOS workloads. If you connect to the remote workload in curtain mode, the display of the remote workload will be dimmed, and the remote user will not be able to see your actions on the workload.</p>	Remote desktop Remote assistance

6. Depending on whether or not Connect Client is installed on your workload, do one of the following:
 - If Connect Client is not installed, download it, install it, and then in the confirmation pop-up that appears, select **Allow**.

- If Connect Client is already installed, in the confirmation pop-up that appears, click **Open Connect Client**.
7. In the **Authentication** window, select an authentication option, and then provide the required credentials.

Note

If you have assigned credentials to the workload, authentication will be done automatically and this step will be skipped. For more information, see "Assigning credentials to a workload" (p. 940).

Authentication option	Description
With remote workload credentials	<p>You will be allowed to establish the remote connection after you provide username and password of an administrator user of the remote workload.</p> <p>This option is available for NEAR, RDP, and Apple Screen Sharing.</p> <p>You can use this option to authenticate for remote desktop and remote assistance.</p>
Ask for permission to observe	<p>You will be allowed to establish the remote connection in observe mode after the user who is logged in on the remote workload allows it.</p> <p>This option is available for NEAR and Apple Screen Sharing.</p> <p>You can use this option to authenticate for remote assistance.</p>
Ask for permission to control	<p>You will be allowed to establish the remote connection in control mode after the user who is logged in on the remote workload allows it.</p> <p>This option is available for NEAR and Apple Screen Sharing.</p> <p>You can use this option to authenticate for remote assistance.</p>

8. Click **Connect**, and then click the session to display (if more than one user session is available on the workload).

Connect Client will open a new viewer window on which you will be able to see the remote workload's desktop. The viewer has a toolbar with additional actions that you can perform on the remote workload after the remote connection is established. For more information, see "Using the toolbar in the Viewer window" (p. 952).

Connecting to a managed workload via a web client

You can establish a remote desktop connection to a managed workload via a web client.

Prerequisites

- Standard service quota is assigned to the workload.
- A remote management plan that has RDP enabled is applied to the managed workload.
- RDP is enabled on the managed workload.

- Your browser supports HTML5.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

To remotely connect to a workload via a web client

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click the workload which you want to connect remotely, and then click **Remote desktop>Connect via web client**.
3. Enter the login and password to access the workload, and then click **Connect**.

Note

If you have assigned credentials to the workload, authentication will be done automatically and this step will be skipped. For more information, see "Assigning credentials to a workload" (p. 940).

Transferring files

You can easily transfer files between the local workload and a managed workload.

Prerequisites

- A remote management plan that has the NEAR protocol and File transfer enabled is applied on the workload.
- Advanced Management quota is applied on the workload.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

To remotely transfer files between your workload and a managed workload

1. In the Cyber Protect console, go to **Devices > Machines with agents**.
2. Click the workload with which you want to transfer files.
3. Click **Manage**, and then **Transfer files**.
4. Depending on whether or not Connect Client is installed on your workload, do one of the following:
 - If Connect Client is not installed, download it, install it, and then in the confirmation pop-up that appears, click **Allow**.
 - If Connect Client is already installed, in the confirmation pop-up that appears, click **Open Connect Client**.
5. In the **Authentication** window, select an authentication option, and then provide the required credentials.

Authentication option	Description
With remote workload credentials	You will be allowed to establish the remote connection after you provide username and password of an administrator user of the remote workload.

Authentication option	Description
Ask for permission to transfer files	You will be allowed to transfer files after the user who is logged in on the remote workload allows it.

- In the **File transfer** window, browse files and drag and drop them to the desired destination.

Note

The files of the local workload are listed in the left pane, and the files of the remote workload are listed in the right pane.

When a file transfer begins, it is listed in the **Tasks** pane.

- [Optional] If you want to remove the completed tasks from the **Tasks** pane, click **Clear finished**.
- When all transfers complete, close the window.

Performing control actions on managed workloads

You can manage a remote workload by performing the following basic control actions on it: empty recycle bin, sleep, restart, shut down, and log out remote user.

Prerequisites

- Standard service quota is applied to the workload.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

Empty recycle bin

To empty the recycle bin on the remote workload

- In the Cyber Protect console, go to **Devices > Machines with agents**.
- Click the workload on which you want to perform this action.
- Click **Manage**, and then click **Empty recycle bin**.
- Select the user session for which you want to perform this action, and then click **Empty recycle bin**.

Sleep

To put a remote workload to sleep

- In the Cyber Protect console, go to **Devices > Machines with agents**.
- Click the workload on which you want to perform this action.
- Click **Manage**, and then click **Sleep**.

Restart

To restart a remote workload

1. In the Cyber Protect console, go to **Devices > Machines with agents**.
2. Click the workload on which you want to perform this action.
3. Click **Manage**, and then click **Restart**.
 - For Windows workloads, select if you want to allow the user who is currently logged in locally to the workload to save the changes before the workload is restarted, select the user, and then click **Restart** again.
 - For macOS workloads, select if you want to allow the user who is currently logged in locally to the workload to save the changes before the workload is restarted, and then click **Restart** again.
 - For Linux workloads, click **Restart**.

Shut down

To shut down a remote workload

1. In the Cyber Protect console, go to **Devices > Machines with agents**.
2. Click the workload on which you want to perform this action.
3. Click **Manage**, and then click **Shut down**.
 - For Windows workloads, select if you want to allow the user who is currently logged in locally to the workload to save the changes before the workload is shut down, select the user, and then click **Shut down** again.
 - For macOS workloads, select if you want to allow the user who is currently logged in locally to the workload to save the changes before the workload is shut down, and then click **Shut down** again.
 - For Linux workloads, click **Shut down** again.

Log out remote user

To log out the user of a remote workload

1. In the Cyber Protect console, go to **Devices > Machines with agents**.
2. Click the workload on which you want to perform this action.
3. Click **Manage**, and then click **Log out remote user**.
4. Select the user that you want to log out, and then click **Log out**.

Monitoring workloads via screenshot transmission

You can monitor the status of a workload by using the Screenshot transmission feature.

Prerequisites

- A remote management plan with the Screenshot transmission feature enabled is applied on the workload.
- The Protection agent version is up-to-date and supports the Screenshot transmission feature.
- Advanced Management service quota is applied on the workload.

- The workload is online.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

Monitoring a workload via Screenshot transmission

To monitor a workload via Screenshot transmission

1. In the Cyber Protect console, go to **Devices>Screenshot transmission**.
2. Click the workload that you want to monitor.
3. Select the user session.
4. Select the display.
5. Select refresh rate at which to take a new screenshot of the desktop.
6. Select the image quality.
7. To download the screenshot, click the download icon.

Taking a screenshot of a workload

To take a screenshot of a managed workload

1. In the Cyber Protect console, go to **Devices > Machines with agents**.
2. Click the workload from which you want to take a screenshot.
3. Click **Manage**, and then click **Take desktop screenshot**.

The **Screenshot transmission** screen will open, with the workload preselected. Depending on the settings of the remote management plan that is applied on the workload, you will see the screenshot or you will see the screenshot after the user of the remote workload approves the request.

Observing multiple managed workloads simultaneously

You can observe the desktops of multiple remote workloads simultaneously in a single window.

Note

The number of desktops that you can see simultaneously in the window depends on the size of your monitor.

Prerequisites

- NEAR / Apple Screen Sharing is enabled in the remote management plans that are applied to the workloads.
- Advanced Management service quota is applied on the workload.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

To observe multiple workloads simultaneously

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Select the workloads which you want to observe.
3. Click **Multi view**.

4. Depending on whether or not Connect Client is installed on your workload, do one of the following:
 - If Connect Client is not installed, download it, install it, and then in the confirmation pop-up that appears, select **Allow**.
 - If Connect Client is already installed, in the confirmation pop-up that appears, click **Open Connect Client**.
5. In the **Authentication** window, select an authentication option, and then provide the required credentials.

Authentication option	Description
With remote workload credentials	You will be allowed to establish the remote connection after you provide username and password of an administrator user on the remote workload.
Ask for permission to observe	You will be allowed to establish the remote connection in observe mode after the user who is logged in on the remote workload allows it.

6. If you want to use the same authentication method and credentials when connecting to all the remote workloads that you selected in step 2, select **Use on other computers**.
7. Click **Connect**.

In toolbar of the multi-view window, you can select a view mode in which to connect to a workload. This action will open a separate Viewer window for that workload.

Note

If any of the selected workloads is offline, or has an outdated version of the agent installed, it will not be shown in the multi-view window.

All multi-view connections to remote workloads are in **View-only** mode.

Working with unmanaged workloads

Unmanaged workloads are workloads on which the Protection agent is not installed.

You can perform the following actions on the unmanaged remote workloads:

- connect for remote assistance by using Acronis Quick Assist
- connect for remote desktop or remote assistance by using an IP address
- transfer files between your workload and the remote workload by using Quick Assist

Note

The remotely connect to unmanaged workloads by using Quick Assist, ensure that:

- The Advanced Management pack is activated for your customer tenant.
 - The Quick Assist application is running on the remote workload to which you want to connect.
-

Connecting to unmanaged workloads via Acronis Quick Assist

You can use the Quick assist feature to connect remotely on demand to unmanaged workloads and provide one-time assistance.

Prerequisites

- The Advanced Management pack is assigned to your customer tenant.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.
- The remote user has provided the workload ID and access code from Quick Assist.
- The remote user has downloaded and run Acronis Quick Assist.

To connect to a workload for remote assistance by using Quick Assist

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click **Quick Assist**.
3. In the **Quick Assist** window, enter the workload ID that the end user gave you, and then select **Connect**.
4. Click **Connect**.
5. Depending on whether or not Connect Client is installed on your workload, do one of the following:
 - If Connect Client is not installed, download it, install it, and then in the confirmation pop-up that appears, select **Allow**.
 - If Connect Client is already installed, in the confirmation pop-up that appears, click **Open Connect Client**.
6. In the **Authentication** window, enter the access code.
7. Connect Client will open a new viewer window on which you will be able to see the remote workload's desktop. The viewer has a toolbar with additional actions that you can perform on the remote workload after the remote connection is established. For more information, see "Using the toolbar in the Viewer window" (p. 952).

Connecting to unmanaged workloads via IP address

If an unmanaged workload is in your LAN, you can connect to it for remote control or remote assistance by using its IP address. This connection does not require Internet access.

Prerequisites

- The Advanced Management pack is assigned to your customer tenant.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.

To connect to a workload for remote desktop or remote assistance by using its IP address

1. In the Cyber Protect console, go to **All devices**.
2. Click **Quick Assist**.

3. Click the **Via IP address** tab.
4. Enter the workload's IP address and port.
5. Select a connection protocol - RDP (Windows workloads) or Apple Screen Sharing (for macOS workloads), depending on the remote workload's operating system.

Note

Connections via RDP support the remote desktop action, and connections via Apple Screen Sharing support both the remote desktop and remote assistance actions.

6. Click **Connect**.
7. In the **Authentication** window, provide the required credentials.

For Apple Screen Sharing connections, Connect Client will open a new viewer window on which you will be able to see the remote workload's desktop. The viewer has a toolbar with additional actions that you can perform on the remote workload after the remote connection is established. For more information, see "Using the toolbar in the Viewer window" (p. 952).

Transferring files via Acronis Quick Assist

You can use the Quick assist feature to transfer files between your workload and unmanaged workloads.

Prerequisites

- The Advanced Management pack is assigned to your customer tenant.
- 2FA is enabled for your user account in Acronis Cyber Protect Cloud.
- The remote user has downloaded and run Acronis Quick Assist.
- The remote user has provided the computer ID and access code from Quick Assist.

To transfer files to a workload by using Quick Assist

1. In the Cyber Protect console, go to **Devices > All devices**.
2. Click **Quick assist**.
3. In the **Quick assist** window, enter the workload ID that the end user gave you, and then select **File Transfer**.
4. Click **Connect**.
5. Depending on whether or not Connect Client is installed on your workload, do one of the following:
 - If Connect Client is not installed, download it, install it, and then in the confirmation pop-up that appears, select **Allow**.
 - If Connect Client is already installed, in the confirmation pop-up that appears, click **Open Connect Client**.
6. In the **Authentication** window, enter the access code.
7. In the **File transfer** window, browse files and drag and drop them to the desired destination.

Note








The files of the local workload are listed in the left pane, and the files of the remote workload are listed in the right pane.






When a file transfer begins, it is listed in the **Tasks** pane.

- [Optional] If you want to remove the completed tasks from the **Tasks** pane, click **Clear finished**.
- When all transfers complete, close the window.

Using the toolbar in the Viewer window

After you connect to a remote workload, you can use the toolbar of the viewer window to quickly perform the different actions.

Icon	Description
	Actual size Scales the remote workload's desktop so that one pixel of the remote desktop corresponds to one pixel on the viewer window.
	Zoom to fit Scales the remote workload desktop to fit the viewer window.
	Lock and Unlock screen Shows a placeholder on the remote workload's display so that the remote user will not see your actions.
	Take screenshot Save the remote server's desktop image to a local file.
	Select display Select the remote workload display that you want to view and the desired resolution. Available for Apple Screen Sharing connections to macOS, and NEAR connections to any operating system.
	Image quality Adjusts the remote screen image quality from black and white to the highest possible on Apple Screen Sharing connections.
	NEAR image quality Adjusts the quality/performance ratio on NEAR connections. The left side of the slider (Smooth) prioritizes performance over image quality, the right one (Sharp) means the best quality of remote desktop screen, but probably worse performance.

Icon	Description
	<p>Send Ctrl+Alt+Del</p> <p>Sends a Ctrl + Alt + Delete sequence to the remote workload.</p> <p>Available for Windows and Linux workloads.</p>
	<p>File Transfer</p> <p>Opens the File Manager window to exchange files between remote and local workload. Available for NEAR connections.</p>
	<p>Pin toolbar</p> <p>Turns off automatic hiding of viewer toolbar.</p> <p>Available for Windows workloads.</p>
	<p>Full screen</p> <p>Switches to the full screen mode and scales the remote workload to completely fill your local screen.</p> <p>Available for Windows workloads.</p>
	<p>Close</p> <p>Closes the Viewer window and ends the remote control session.</p> <p>Available for Windows workloads.</p>

Depending on connection type, additional options might be available when you click the **Other** icon.

Option	Description
<p>Start recording / Stop recording</p>	<p>Record the current remote desktop session.</p> <p>Session recordings are saved as .crec files on the local workload. You can open .crec files with Acronis Connect Client.</p> <p>Available for NEAR connections</p>
<p>Clipboard auto sync</p>	<p>When this option is on, the client will automatically synchronize your local clipboard and the clipboard of the remote computer.</p> <p>Available for NEAR and Apple Screen Sharing connections</p>
<p>Send clipboard Get clipboard</p>	<p>Send Clipboard replaces the remote computer clipboard contents with the contents of the local clipboard.</p> <p>Get Clipboard transfers the contents of the remote computer clipboard to the local clipboard.</p>
<p>Smart keyboard / Raw keys / Raw keys with all shortcuts</p>	<p>Changes the keyboard input mode for the current connection.</p> <p>Smart keyboard- the client transmits Unicode codes of the locally typed symbols to the remote computer</p>

Option	Description
	<p>Raw keys- the client uses the raw codes of the keyboard buttons you press.</p> <p>Raw keys with all shortcuts- the client disables local system shortcuts so that they are also transmitted to the remote operating system.</p>
Keyboard focus on mouse hover	<p>When enabled, the client only captures the keyboard input while your local mouse cursor is placed over the Viewer window.</p> <p>When disabled, the client captures your keyboard whenever its window is active.</p>
Show connection info / Hide connection info	<p>When Show connection info is selected, a small information panel will appear over the remote desktop screen, showing the most essential information about current connection.</p>
Remote sound	<p>Enables the client to redirect the sound from the remote computer to the local one.</p> <p>Available for NEAR connections</p>
Preferences	<p>Configure the settings of Connect Client. For more information, see "Configuring the Connect Client settings" (p. 955).</p>

Recording and playing remote connection sessions

You can record a remote connection session via NEAR in Acronis Connect Client.

To record a remote connection session

1. On the Viewer toolbar in Connect Client, click **Other**, and select **Start Recording**.
2. Select a name and location for the record.
By default, the file will be named with the current date and time and located in the **Documents** folder in the current user home directory. While the recording is active, in the **Viewer** toolbar you will see a flashing red circle over the top right corner of the remote screen and the recording timer.
3. To stop the recording, click **Other**, and then click **Stop Recording**. On a Mac, you can also click **Stop** on the toolbar.
All .crec files made by Acronis Connect Client will be opened with Acronis Connect Client by default.

To play a recording

1. Locate a recording file.
2. Open it.

The recording player of Acronis Connect Client opens. Note that it is not possible to navigate through the recording. To find a certain moment in the recording, wait until the player reaches it.

- [Optional] To adjust playback speed, use the << and >> icons in the playback controls section. The recording is stored as a sequence of events that have been transmitted to and from the remote server during a connection. This ensures the best quality of the recording at the minimum file size. However, this also means that it is not possible to navigate through the recording. At the moment it is also not possible to convert the recordings to a video format.

Configuring the Connect Client settings

After you install Connect Client on your workload, you can configure its settings according to your preferences.

To configure the settings of Connect Client

- In the start menu, find **Connect Client**, and start it.
- Configure the settings on the **General** tab.

Option	Description
Write verbose logs	Select this option to allow Connect Client to write verbose logs. If disabled, the client will only write general information to the log file.
Proxy settings	Select whether to use the default System proxy, or configure a Custom SOCKSS proxy.

- Configure the settings on the **Viewer** tab.

Option	Description
Ask for confirmation when closing a viewer	Select this option if you want Connect Client to display a confirmation message when you attempt closing the Viewer window in order to prevent accidental closing.
When minimized	Select whether to suspend the Viewer activity when minimized, in order to reduce the CPU load.
When maximized	Select whether to enable the full screen mode when maximized.
Clipboard transfer	Enable showing the Clipboard transfer indicator in the Viewer window whenever you copy or paste text and images.
Keyboard Mode	Enable showing the Input mode indicator in the Viewer window title whenever mouse and keyboard events are being sent to the remote machine.
Clipboard	Select Automatically synchronize clipboard to enable automatic

Option	Description
	clipboard synchronization, when available.
Send keyboard events	Choose whether to grab your local keyboard input whenever the Connect Client window is active or only when your local mouse pointer is over it.
Viewer background color	Change the Viewer window background color.
Reconnect automatically	Select Enable to reconnect automatically , if you want Connect Client to automatically re-establish the connection if it has been interrupted.
H.264	You can disable hardware decoders.
Close when idle	Select the time interval of being idle after which to close the Viewer window.

4. Configure the settings on the **Keyboard** tab.


Option	Description
Modifier mappings	Change the behavior of modifier keys with a pop-up menu. These settings are stored separately for NEAR, Apple Screen Sharing, and RDP connections.
Input mode	For each type of connection (selected in the header of pane), select the default keyboard input mode.

5. Click **OK**.


The remote desktop notifiers

The Connect Agent displays action dialogs (notifiers) on the remote workload's desktop in the following cases:


- when you try to connect to the workload remotely by asking for permission to observe. The user who is logged in to the remote workload locally can allow or deny the request.



- when you try to connect to the workload remotely by asking for permission to control. The user who is logged in to the remote workload locally can allow or deny the request.




- when you try to exchange files between your workload and the remote workload by asking for permission to transfer files. The user who is logged in to the remote workload locally can allow or deny the request.



When you establish a remote desktop connection to a workload, the user who is logged in to the workload will view a different connection notifier that contains the following information:

- name of the user who is connected remotely
- connection protocol that is used to establish the remote connection
- duration of the remote connection

The user who is logged in to the remote workload locally can end the connection at any time by clicking the **Disconnect** icon or the **Close** icon.



Monitoring the health and performance of workloads

You can monitor the system parameters and the health of the workloads in your organization. If a parameter is out of norm, you will be notified immediately and will be able to quickly resolve the issue. You can also configure custom alerts and automatic response actions. These are actions that will be performed automatically to resolve anomalies in workload behavior.

Note

The monitoring functionality requires an installation of Protection agent version 15.0.35324 or later on the workloads.

Monitoring plans

To start monitoring the performance, hardware, software, system, and security parameters of your managed workloads, apply a monitoring plan to them. The monitoring plans consist of different monitors that you can enable and configure. Some monitors support the anomaly-based monitoring type. For more information about monitoring plans, see "Monitoring plans" (p. 990). For more information about the available monitors that you can configure in the monitoring plans, see "Configurable monitors" (p. 959).

If the agent cannot collect data from a workload for some reason, the system will generate an alert.

Monitoring types

You must configure the monitoring type for each monitor that you enable in the plan. The monitoring type determines the algorithm that the monitor will use to estimate the normal behavior and deviation of the workload. There are two monitoring types: threshold-based and anomaly-based. Some monitors support only the threshold-based monitoring type.

Threshold-based monitoring tracks if the values of the parameters are above or below a threshold value that you configure. With this monitoring type, you are responsible for defining the correct threshold values for the workloads. The system determines the normal behavior based on these static threshold values and without considering other specific conditions that might cause the behavior. For this reason, threshold-based monitoring might be less accurate as compared to anomaly-based monitoring.

Anomaly-based monitoring uses machine learning to create the normal behavior patterns for a workload and to detect abnormal behavior. For more information, see "Anomaly-based monitoring" (p. 958).

Anomaly-based monitoring

Anomaly-based monitoring uses machine learning models to create the normal behavior patterns for a workload and to detect anomalies (unexpected spikes in the time-series data) in the workload's

behavior. When you activate this monitoring type, the system creates a model and starts training itself and adjusting the model for the specific workload based on the data that it collects from the workload. This means that at the beginning of the training period, the data might not be fully accurate. Creating a reliable model requires at least three weeks of training of the model. As the system collects more data and analyzes historical data sets, it gradually refines the model and creates the dynamic upper and lower thresholds for each metric of the workload. This monitoring type is more flexible as compared to the threshold-based monitoring because the system monitors the values of the parameters and their context. For example, it might be normal for a specific workload to have bigger load in certain hours of the day. A threshold-based monitoring type would falsely interpret this as an abnormal behavior and would trigger an alert.

You can reset the machine learning models of a workload. In this case, the system will delete all data and models for the monitors that are applied to the workload. For more information, see "Resetting the machine learning models" (p. 999).

Supported platforms for monitoring

The monitoring functionality is supported for the following operating systems.

Supported Windows versions	Supported macOS versions
<ul style="list-style-type: none"> Windows 7 SP1 Windows 8, 8.1 Windows 10 Windows 11 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 Windows Server 2019 Windows Server 2022 	<ul style="list-style-type: none"> macOS 10.14 (Mojave) macOS 10.15 (Catalina) macOS 11.x (Big Sur) macOS 12.x (Monterey) macOS 13.x (Ventura)

Configurable monitors

The monitoring functionality supports the following monitors, divided into six categories: Hardware, Performance, Software, System, Security, and Custom.

Monitor	Description	Supported operating systems	Frequency of data collection	Support of anomaly-based monitoring	Availability in Standard protection or Advanced Management
Hardware					
Disk space	Monitors the	Windows	1 minute	Yes	Standard

Monitor	Description	Supported operating systems	Frequency of data collection	Support of anomaly-based monitoring	Availability in Standard protection or Advanced Management
	free space on a specific drive of the workload.	macOS			protection
CPU temperature	Monitors the CPU temperature.	Windows macOS	30 sec	Yes	Advanced Management
GPU temperature	Monitors the GPU temperature.	Windows macOS	30 sec	Yes	Advanced Management
Hardware changes	Monitors the hardware changes, such as adding, removing, or replacing hardware on a workload	Windows macOS	24 hours	No	Standard protection
Performance					
CPU usage	Monitors the overall CPU usage (by all CPUs on the workload).	Windows macOS	30 sec	Yes	Advanced Management
Memory usage	Monitors the overall memory usage (by all memory slots on the workload).	Windows macOS	30 sec	Yes	Advanced Management
Disk transfer rate	Monitors the read and write speed of each physical disk on the workload.	Windows macOS	30 sec	Yes	Advanced Management

Monitor	Description	Supported operating systems	Frequency of data collection	Support of anomaly-based monitoring	Availability in Standard protection or Advanced Management
Network usage	Monitors the incoming and outgoing traffic for each network adapter of the workload.	Windows macOS	30 sec	Yes	Advanced Management
CPU usage by process	Monitors the CPU usage by certain process.	Windows macOS	30 sec	No	Advanced Management
Memory usage by process	Monitors the memory usage by the selected process.	Windows macOS	30 sec	No	Advanced Management
Disk transfer rate by process	Monitors the read and write speed of the selected process.	Windows macOS	30 sec	No	Advanced Management
Network usage by process	Monitors the incoming and outgoing traffic of the selected process.	Windows macOS	30 sec	No	Advanced Management
Software					
Windows service status	Monitors the status of the selected Windows service (Running or Stopped).	Windows	30 sec	No	Advanced Management
Process status	Monitors the status of the	Windows	30 sec	No	Advanced Management

Monitor	Description	Supported operating systems	Frequency of data collection	Support of anomaly-based monitoring	Availability in Standard protection or Advanced Management
	selected process (Running or Stopped).	macOS			
Installed software	Monitors the installation, update, or deletion of software applications.	Windows macOS	24 hours	No	Advanced Management
System					
Last system restart	Monitors when the workload was restarted.	Windows macOS	1 hour	No	Standard protection
Windows event log	Monitors specific business-critical events in the Windows event logs.	Windows	10 min	No	Advanced Management
Files and folders size	Monitors the total size of the selected files or folders.	Windows macOS	10 min	No	Standard protection
Security					
Windows Update status	Monitors the Windows Update status of the workload and whether the latest updates are installed.	Windows	15 min	No	Advanced Management

Monitor	Description	Supported operating systems	Frequency of data collection	Support of anomaly-based monitoring	Availability in Standard protection or Advanced Management
Firewall status	Monitors the status of the built-in or third-party firewall that is installed on the workload.	Windows macOS	5 min	No	Advanced Management
Antimalware software status	Monitors the status of the built-in or third-party antimalware software that is installed on the workload.	Windows macOS	5 min	No	Advanced Management
Failed logins	Monitors unsuccessful login attempts on the workload.	Windows	1 hour	No	Advanced Management
AutoRun status	Monitors if the AutoRun feature for removable storage media is enabled.	Windows	1 hour	No	Advanced Management
Custom					
Custom	Monitors custom objects via running scripts.	Windows macOS	custom	No	Advanced Management

Settings of the Disk space monitor

Disk space monitors the free space on a specific drive of the workload.

Note

When calculating the space, the monitor uses binary bytes (1024 bytes per KB, 1024 KB per MB, and 1024 MB per GB) for both Windows and macOS workloads.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
Drive	The drive that you want to monitor. The following values are available. <ul style="list-style-type: none">• System drive —This is the default value.• Any drive
Operator	The operator is a conditional function that defines how to measure the performance on the metric. The following values are available. <ul style="list-style-type: none">• Less than —This is the default value.• Less than or equal to
Disk free space threshold	The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert. Enter an integer value in the range 1-100 (%). The default value is 20.
Include removable drives	This setting is available if the Drive value is Any drive . Select this setting if you want to add removable drives, like USB flash drives, for monitoring. By default, the setting is disabled.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 1-60 (min). The default value is 30.
Anomaly-based monitoring	
Drive	The drive that you want to monitor. The following values are available. <ul style="list-style-type: none">• System drive —This is the default value.• Any drive
Model training period	The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training period, the more precise the long-term behavior pattern that the system will

Setting	Description
	<p>create. We recommend that the minimum model training period is twenty-one days.</p> <p>Enter an integer value (days). The default value is 21.</p>
<p>Receive anomaly alerts during the training period</p>	<p>If you select this setting, you will receive alerts about anomalies during the model training period. These alerts might be false, because the models are still being trained and might not be accurate enough.</p> <p>By default, the setting is selected.</p>
<p>Sensitivity level</p>	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none"> 1. The algorithm is trained using the data that is collected during the training. 2. The algorithm performs anomaly detection on the training data. 3. A filtering process based on mean and standard deviation is applied. 4. Any anomalies that are in the specified interval are filtered. 5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model. <p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard deviation, according to the sensitivity level. 3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior. <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
<p>Anomaly duration</p>	<p>The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.</p> <p>The default value is 30 minutes.</p>

Settings of the CPU temperature monitor

CPU temperature monitors the CPU temperature of the workload.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
CPU temperature has exceeded (C°)	<p>The maximum value of the monitored metric. If the value is exceeded, the system generates an alert.</p> <p>Enter an integer value (C°). The default value is 80.</p>
Time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Anomaly-based monitoring	
Model training period	<p>The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training period, the more precise the long-term behavior pattern that the system will create. We recommend that the minimum model training period is twenty-one days.</p> <p>Enter an integer value (days). The default value is 21.</p>
Sensitivity level	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none">1. The algorithm is trained using the data that is collected during the training.2. The algorithm performs anomaly detection on the training data.3. A filtering process based on mean and standard deviation is applied.4. Any anomalies that are in the specified interval are filtered.5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model.

Setting	Description
	<p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard deviation, according to the sensitivity level. 3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior. <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
Anomaly duration	<p>The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 15.</p>

Settings of the GPU temperature monitor

GPU temperature monitors the GPU temperature of the workload.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
GPU temperature has exceeded	<p>The maximum value of the monitored metric. If the value is exceeded, the system detects an anomaly.</p> <p>Enter an integer value (C°). The default value is 80.</p>
Time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Anomaly-based monitoring	
Model training period	<p>The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training</p>

Setting	Description
	<p>period, the more precise the long-term behavior pattern that the system will create. We recommend that the minimum model training period is twenty-one days.</p> <p>Enter an integer value (days). The default value is 21.</p>
<p>Sensitivity level</p>	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none"> 1. The algorithm is trained using the data that is collected during the training. 2. The algorithm performs anomaly detection on the training data. 3. A filtering process based on mean and standard deviation is applied. 4. Any anomalies that are in the specified interval are filtered. 5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model. <p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard deviation, according to the sensitivity level. 3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior. <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
<p>Anomaly duration</p>	<p>The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 15.</p>

Settings of the Hardware changes monitor

Hardware changes monitors the hardware changes, such as adding, removing, or replacing hardware on a workload.

You can configure the following settings for the monitor.

Setting	Description
Hardware components	<p>Select one or multiple hardware components that you want to monitor for changes.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• All — This is the default value.• Motherboard• CPU• RAM• Disk• GPU• Network adapter
What to monitor	<p>Specify the changes for which you want to monitor the selected hardware components. You can select multiple items from the list.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• Any change — This is the default value.• Newly added components• Replaced components• Removed components

Settings of the CPU usage monitor

CPU usage monitors the total CPU usage (processor utilization) of the workload. If the workload has multiple CPUs, the total CPU usage will be the sum of the CPU usage of each CPU.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
Operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• More than — This is the default value.• More than or equal to

Setting	Description
	<ul style="list-style-type: none"> • Less than • Less than or equal to
CPU usage threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value in the range 1-100 (%). The default value is 90.</p>
Time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Anomaly-based monitoring	
Model training period	<p>The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training period, the more precise the long-term behavior pattern that the system will create. We recommend that the minimum model training period is twenty-one days.</p> <p>Enter an integer value (days). The default value is 21.</p>
Receive anomaly alerts during the training period	<p>If you select this setting, you will receive alerts about anomalies during the model training period. These alerts might be false, because the models are still being trained and might not be accurate enough.</p> <p>By default, the setting is selected.</p>
Sensitivity level	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none"> 1. The algorithm is trained using the data that is collected during the training. 2. The algorithm performs anomaly detection on the training data. 3. A filtering process based on mean and standard deviation is applied. 4. Any anomalies that are in the specified interval are filtered. 5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model. <p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard

Setting	Description
	<p>deviation, according to the sensitivity level.</p> <p>3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
Anomaly duration	<p>The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 15.</p>

Settings of the Memory usage monitor

Memory usage monitors the total memory usage by all memory modules of the workload.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
Operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than — This is the default value. • More than or equal to • Less than • Less than or equal to
Memory usage threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value in the range 1-100 (%). The default value is 90.</p>
Time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Anomaly-based monitoring	

Setting	Description
Model training period	<p>The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training period, the more precise the long-term behavior pattern that the system will create. We recommend that the minimum model training period is twenty-one days.</p> <p>Enter an integer value (days). The default value is 21.</p>
Receive anomaly alerts during the training period	<p>If you select this setting, you will receive alerts about anomalies during the model training period. These alerts might be false, because the models are still being trained and might not be accurate enough.</p> <p>By default, the setting is selected.</p>
Sensitivity level	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none"> 1. The algorithm is trained using the data that is collected during the training. 2. The algorithm performs anomaly detection on the training data. 3. A filtering process based on mean and standard deviation is applied. 4. Any anomalies that are in the specified interval are filtered. 5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model. <p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard deviation, according to the sensitivity level. 3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior. <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
Anomaly	<p>The system will generate an alert for a detected anomaly only if the abnormal</p>

Setting	Description
duration	behavior persists for the specified period. Enter an integer value in the range 1-60 (min). The default value is 30 minutes.

Settings of the Disk transfer rate monitor

Disk transfer rate monitors the read and write speed of each physical disk on the workload.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
What to monitor	Select the speed that you want to monitor. The following values are available. <ul style="list-style-type: none"> • Read speed and Write speed. This is the default value. • Read speed • Write speed
Read speed operator	The operator is a conditional function that defines how to measure the performance on the metric. The following values are available. <ul style="list-style-type: none"> • More than. This is the default value. • More than or equal to • Less than • Less than or equal to
Read speed threshold	The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert. Enter an integer value (kb/s). The default value is 0 kb/s.
Read speed time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 1-60 (min). The default value is 5.
Write speed operator	The operator is a conditional function that defines how to measure the performance on the metric. The following values are available. <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to

Setting	Description
Write speed threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Write speed time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Anomaly-based monitoring	
Model training period	<p>The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training period, the more precise the long-term behavior pattern that the system will create. We recommend that the minimum model training period is twenty-one days.</p> <p>Enter an integer value (days). The default value is 21.</p>
Receive anomaly alerts during the training period	<p>If you select this setting, you will receive alerts about anomalies during the model training period. These alerts might be false, because the models are still being trained and might not be accurate enough.</p> <p>By default, the setting is selected.</p>
What to monitor	<p>Select the speed that you want to monitor.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • Read speed and Write speed. This is the default value. • Read speed • Write speed
Sensitivity level	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none"> 1. The algorithm is trained using the data that is collected during the training. 2. The algorithm performs anomaly detection on the training data. 3. A filtering process based on mean and standard deviation is applied. 4. Any anomalies that are in the specified interval are filtered. 5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model.

Setting	Description
	<p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard deviation, according to the sensitivity level. 3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior. <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
Anomaly duration (Read speed)	<p>The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.</p> <p>Enter an integer value in the range 1--60 (min).</p> <p>The default value it 25.</p>
Anomaly duration (Write speed)	<p>The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.</p> <p>Enter an integer value in the range 1--60 (min).</p> <p>The default value it 25.</p>

Settings of the Network usage monitor

Network usage monitors the incoming and outgoing traffic for each network adapter of the workload.

You can configure the following settings for the monitor.

Setting	Description
Threshold-based monitoring	
Traffic direction	<p>The traffic direction that you want to monitor.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • Incoming and Outgoing traffic. This is the default value. • Incoming traffic • Outgoing traffic

Setting	Description
Incoming traffic operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Incoming traffic threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Incoming traffic time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Outgoing traffic operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Outgoing traffic threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Outgoing traffic time period	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Anomaly-based monitoring	
Model training period	<p>The period during which the system will train the machine learning models based on the data that is collected from the agents, and will then create the normal behavior pattern of the workload. The longer the model training period, the more precise the long-term behavior pattern that the system will create. We recommend that the minimum model training period is twenty-one days.</p>

Setting	Description
	Enter an integer value (days). The default value is 21.
Receive anomaly alerts during the training period	<p>If you select this setting, you will receive alerts about anomalies during the model training period. These alerts might be false, because the models are still being trained and might not be accurate enough.</p> <p>By default, the setting is selected.</p>
Traffic direction	<ul style="list-style-type: none"> • Incoming and Outgoing traffic. This is the default value. • Incoming traffic • Outgoing traffic
Sensitivity level	<p>The sensitivity level acts as a preliminary filter for anomalies if their values are within a specific range. This filter operates independently from the anomaly detection algorithm. Its purpose is to stop the anomalies that are in the specified range from being processed by the anomaly detection algorithm.</p> <p>During the training period:</p> <ol style="list-style-type: none"> 1. The algorithm is trained using the data that is collected during the training. 2. The algorithm performs anomaly detection on the training data. 3. A filtering process based on mean and standard deviation is applied. 4. Any anomalies that are in the specified interval are filtered. 5. From the remaining anomalous data points, the anomaly with the lowest anomaly level is selected. This level (a float number between 0 and 1) is recorded in the model. <p>During the prediction:</p> <ol style="list-style-type: none"> 1. The algorithm predicts anomalies on the inference data. 2. The predicted anomalies are filtered based on the mean and standard deviation, according to the sensitivity level. 3. The remaining anomalies are further filtered based on the following principle: values above the threshold level are considered an anomaly, while values below the threshold level are considered normal behavior. <p>The following values are available.</p> <ul style="list-style-type: none"> • Low — The low level equals the mean value and the standard deviation value. • Normal — This is the default value. The normal level equals the mean value and two times the standard deviation value. • High — The high level equals the mean value and three times the standard deviation value.
Anomaly duration (Incoming)	The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period.

Setting	Description
	Enter an integer value in the range 1--60 (min). The default value it 25.
Anomaly duration (Outgoing)	The system will generate an alert for a detected anomaly only if the abnormal behavior persists for the specified period. Enter an integer value in the range 1--60 (min). The default value it 25.

Settings of the CPU usage by process monitor

CPU usage by process monitors the CPU usage of the selected process. If there are multiple instances of the same process, the system will monitor the total usage by all process instances and will generate an alert when the conditions are met.

You can configure the following settings for the monitor.

Setting	Description
Process name	Name of the process that you want to monitor. Enter the process name without the extension.
Operator	The operator is a conditional function that defines how to measure the performance on the metric. The following values are available. <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Threshold	The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert. Enter an integer value in the range 1-100 (%). The default value is 90.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 1-60 (min). The default value is 5.

Settings of the Memory usage by process monitor

Memory usage by process monitors the memory usage of the selected process. If there are multiple instances of the same process, the system will monitor the total usage by all process instances and will generate an alert when the conditions are met.

Note

The agents use the total process working set (private and shared) to estimate the size of the memory usage by process. That is why the size that the widget shows might be different from the size of the memory usage that is shown in Windows Task Manager (private working set).

You can configure the following settings for the monitor.

Setting	Description
Process name	Name of the process that you want to monitor. Enter the process name without the extension.
Operator	The operator is a conditional function that defines how to measure the performance on the metric. The following values are available. <ul style="list-style-type: none">• More than —This is the default value.• More than or equal to• Less than• Less than or equal to
Threshold	The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert. Enter an integer value (kb). The default value is 1.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 1-60 (min). The default value is 5.

Settings of the Disk transfer rate by process monitor

Disk transfer rate by process monitors the read and write speed of the selected process. If there are multiple instances of the same process, the system will monitor the total usage by all process instances and will generate an alert when the conditions are met.

You can configure the following settings for the monitor.

Setting	Description
Process name	The name of the process that you want to monitor. Enter the process name without the extension.
What to monitor	The speed that you want to monitor. The following values are available. <ul style="list-style-type: none">• Read speed and Write speed. This is the default value.

Setting	Description
	<ul style="list-style-type: none"> • Read speed • Write speed
Read speed operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Read speed threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Read speed time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Write speed operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Write speed threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Write speed time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>

Settings of the Network usage by process monitor

Network usage by process monitors the incoming and outgoing traffic of the selected process. If there are multiple instances of the same process, the system will monitor the total usage by all process instances and will generate an alert when the conditions are met for all instances.

You can configure the following settings for the monitor.

Setting	Description
Process name	Name of the process that you want to monitor. Enter the process name without the extension.
Traffic direction	<p>The traffic direction that you want to monitor.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • Incoming traffic and Outgoing traffic. This is the default value. • Incoming traffic • Outgoing traffic
Incoming traffic operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Incoming traffic threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Incoming traffic time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>
Outgoing traffic operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • More than or equal to • Less than • Less than or equal to
Outgoing traffic threshold	<p>The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert.</p> <p>Enter an integer value (kb/s). The default value is 0 kb/s.</p>
Outgoing traffic time period	<p>The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.</p> <p>Enter an integer value in the range 1-60 (min). The default value is 5.</p>

Settings of the Windows service status monitor

Windows service status monitors whether the selected Windows service is running or stopped.

You can configure the following settings for the monitor.

Setting	Description
Service name	The name of the Windows service that you want to monitor. You can select a service name from the list of Windows services. The list is populated by all agents of the tenant after software inventory scan completes successfully on the workloads. You can also add a service name that is not in the list. This is the only available option if software inventory scan was not performed on the workloads.
Service status	If the service is in the selected status, the system will generate an event. The following values are available. <ul style="list-style-type: none">• Running• Stopped—This is the default value.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 1-60 (min). The default value is 1.

Settings of the Process status monitor

Process status monitors whether the selected process is running or stopped. If there are multiple instances of the same process, the system will monitor each instance of the process and will generate the alert when the conditions are met for all instances of the process.

You can configure the following settings for the monitor.

Setting	Description
Process name	The name of the process that you want to monitor. Enter the name of the executable file without the extension.
Process status	If the process is in the selected status, the system will generate an event. The following values are available. <ul style="list-style-type: none">• Running• Stopped—This is the default value.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 1-60 (min). The default value is 1.

Settings of the Installed software monitor

Installed software monitors the installation, updates, or deletion of software applications on the workload.

You can configure the following settings for the monitor.

Setting	Description
What software to monitor	<p>Specify the software that you want to monitor.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• Any software —This is the default value.• Specific software
Software names	<p>This setting becomes available if you select the Specific software value for What software to monitor.</p> <p>Enter the name of one or several software applications.</p> <p>You can select a software application name from the list of Windows services. The list is populated by all agents of the tenant after software inventory scan completes successfully on the workloads. You can also add a software application name that is not in the list. This is the only available option if software inventory scan was not performed on the workloads.</p>
Installation status	<p>Specify if you want to monitor installed, not installed, or updated software.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• Installed - This is the default value. If you select this value, the monitor will generate an alert when a new software application is installed on the workload.• Updated - If you select this value, the monitor will generate an alert when a software application is updated.• Not installed - If you select this value, the monitor will generate an alert when a software application is uninstalled or not available on the workload.

Settings of the Last system restart monitor

Last system restart when the workload was last restarted.

You can configure the following setting for the monitor.

Setting	Description
The workload has not been restarted for	<p>The period (number of days) since the last restart of the workload. If the workload has not been restarted for a longer period than the period you specify, the system will generate an alert.</p> <p>Enter an integer value in the range 1-180 (days). The default value is 30.</p>

Settings of the Windows event log monitor

Windows event log monitors specific business-critical events in the Windows event logs.

You can configure the following settings for the monitor.

Setting	Description
Event log name	<p>Select a certain event log from a list of Windows event logs that are available in Windows Event Viewer.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• Any —This is the default value.• Application• Security• System
Event source	<p>Event source name</p> <p>You can select the value from a list of event sources that are collected from all agents of the tenant or enter a new source name manually.</p> <p>If the software inventory scan is disabled for the tenant, the event source list will be empty.</p>
Matching mode	<p>In this field, you can specify whether to connect the Event IDs, Event type, and Event description settings by using the Any or the All operator.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• Any —This is the default value. An alert will be generated only if any of the selected criteria is matched.• All — An alert will be generated if all the selected criteria are matched.
Event IDs	<p>Enter one or multiple event IDs separated by comma. If the system finds in the event log any of the event codes that you entered in this field, it generates an alert.</p>
Event type	<p>Select one or multiple event types that you want to monitor.</p> <p>The following values are available.</p> <ul style="list-style-type: none">• Any —This is the default value.• Error• Warning• Information• Success-audit• Failure-audit
Event description	<p>Specific keywords or phrases in the event description for which you want to search. Each keyword or phrase that you enter must be enclosed in quotation</p>

Setting	Description
	marks and must be separated by comma. If the system finds any of the keywords or phrases that you entered, it will generate an alert.
Number of occurrences	The minimum number of occurrences in the log that an event must have during the time period for the system to generate an alert. Enter an integer value in the range 1-1000.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value and then select the unit: minutes or hours. The default value is 60 minutes.

Settings of the Files and folders size monitor

Files and folders size monitors the total size of the selected files and folders.

You can configure the following settings for the monitor.

Setting	Description
Files or folders to monitor	<p>The paths to the files or folders that you want to monitor. You can also specify files or folders that you want to exclude from monitoring.</p> <p>You can use the following wildcard characters.</p> <ul style="list-style-type: none"> • * — for zero or more characters in a file or folder name • ? — for exactly one character in a file or folder name <p>For Windows workloads:</p> <ul style="list-style-type: none"> • The full path should start from the drive letter followed by the :\ separator. • You can use slash or backslash as a path separator character. • The file or folder name must not end with a space or a period. <p>For macOS workloads:</p> <ul style="list-style-type: none"> • The full path should start from the root directory. • You can use slash as a path separator character. • The file or folder name must not end with a space or a period. <p>Specifying a specific location is not mandatory for exclusion filters. The files entered without a specific location will be excluded in the monitored folders.</p>
Operator	<p>The operator is a conditional function that defines how to measure the performance on the metric.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • More than —This is the default value. • Less than

Setting	Description
Threshold value	The threshold value and the Operator value determine the normal performance of the monitored metric. When the value of the monitored metric is out of the norm, the system generates an alert. Enter an integer value (MB).
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period. Enter an integer value in the range 10-60 (min). The default value is 10.

Settings of the Windows Update status monitor

Windows Update status monitors the Windows Update status of the workload and whether the latest updates are installed.

If you enable this monitor, the system will generate an alert in the following cases.

- Windows Update is disabled on the workload.
- Windows Update is enabled on the workload, but the latest updates are not installed.

Settings of the Firewall status monitor

Firewall status monitors the built-in or third-party firewall that is installed on the workload.

If you enable this monitor, the system will generate an alert in the following cases.

- The built-in OS firewall (Windows Defender Firewall or macOS firewall) is disabled and no third-party firewall is running.
- Windows Defender Firewall is disabled for public networks.
- Windows Defender Firewall is disabled for private networks.
- Windows Defender Firewall is disabled for domain networks.

Settings of the Failed logins monitor

Failed logins monitors the unsuccessful login attempts on the workload.

You can configure the following settings for the monitor.

Setting	Description
Failed login attempts threshold	The threshold value determines the boundaries for the normal performance of the monitored metric. When the threshold value is exceeded, the value is out of norm. Enter an integer value. The default value is 60.
Time period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.

Setting	Description
	Enter an integer value in the range 1-24 and select a unit: hours or days. The default value is 12.

Settings of the Antimalware software status monitor

Antimalware software status monitors the built-in or third-party antimalware software that is installed on the workload.

If you enable this monitor, the system will generate an alert when it identifies one of the following conditions.

- Antimalware software is not installed on the workload.
- Antimalware software is installed, but not running.
- Antimalware software is installed and running, but the malware definitions are not up to date.

Note

This condition is checked for Windows and Windows Server operating systems.

Operating system	Supported antimalware software
Windows	<ul style="list-style-type: none"> • Acronis Cyber Protect • Windows Defender • Symantec Endpoint Security • Norton 360 • Norton antivirus • SentinelOne • Trend Micro Endpoint Security with Apex One • Trend Micro Worry-Free Business • McAfee Endpoint Security • McAfee Endpoint Protection for SMB • FireEye Endpoint Security • F-Secure SAFE • F-Secure Client Security • CrowdStrike Falcon • Kaspersky Endpoint Security Cloud • BitDefender Antivirus • Sophos Intercept X Endpoint • Avast Business Antivirus • AVG Antivirus Business Edition • AVG Internet Security Business Edition • Panda Endpoint Protection • Tencent PC Manager

Operating system	Supported antimalware software
	<ul style="list-style-type: none"> • Webroot Business Endpoint Protection • ESET Endpoint Security • Avira Antivirus • Comodo Internet Security • Comodo Business Antivirus • K7 Business Security • K7 Total Security • Vipre Endpoint Protection • Total AV
Windows Server	<ul style="list-style-type: none"> • Acronis Cyber Protect • Windows Defender • ESET Endpoint Security <hr/> <p>Note The monitor might work with other antimalware applications, but this is not guaranteed.</p> <hr/>
macOS	<ul style="list-style-type: none"> • Acronis Cyber Protect • F-Secure Safe • BitDefender Anti-virus for Mac • Sophos Home • Sophos Endpoint Protection • Avast Security for Mac • AVG AntiVirus for Mac • Webroot SecureAnywhere • ESET Cybersecurity • Avira Antivirus for Mac • Comodo Antivirus for Mac • K7 Antivirus for Mac • Vipre Advanced Security • Total AV for Mac <hr/> <p>Note The monitor might work with other antimalware applications, but this is not guaranteed.</p> <hr/>

Settings of the AutoRun feature status monitor

AutoRun feature status monitors if the AutoRun feature for removable media is enabled.

For security reasons, we recommend that you disable the AutoRun feature for removable media on the workload. If the feature is enabled, the system will generate an alert.

Settings of the Custom monitor

Custom monitors custom objects via running a script.

You can configure the following settings for the monitor.

Setting	Description
Script to run	List of predefined scripts from the script repository.
Schedule	<p>The time when the script is run and, optionally, additional conditions that should be met to run the script.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • Schedule by time — The script will run in the exact time, days, weeks, or months that you specify. This is the default value. <p>Schedule type — Hourly, Daily, or Monthly</p> <p>Run within a date range — A time range in which to run the script.</p> <ul style="list-style-type: none"> • When user logs in to the system — The script will run when a user logs in to the workload. • When user logs off the system — The script will run when a user logs out of the workload. • On the system startup — The script will run when the operating system of the workload starts. • When system is shut down — The script will run when the workload is shut down. • When system goes online — The script will run when the workload becomes available online. <p>Start conditions — The task will be performed at a specified time or event only if the condition is met. With multiple conditions are selected, all of them must be met simultaneously to start the task.</p> <p>By default, the Prevent the sleep or hibernate mode to start a scheduled task condition is selected.</p> <p>If start conditions are not met, run the task anyway after — By default, this condition is enabled. The default value is 1 hour.</p>
Account to execute the script	<p>The account on which the script will be run.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • System account — This is the default value. • Currently logged in account
Maximum duration	<p>The maximum period during which the script can run on the workload.</p> <p>If the script does not complete during this period, the operation will fail.</p> <p>Enter an integer value in the range 1-1440 (minutes). The default value is 3</p>

Setting	Description
	minutes.
PowerShell execution policy	<p>The PowerShell execution policy.</p> <p>The following values are available.</p> <ul style="list-style-type: none"> • Undefined • AllSigned • Bypass — This is the default value. • RemoteSigned • Restricted • Unrestricted <p>For more information about these values, see the Microsoft documentation.</p>

Monitoring plans

Monitoring plans are plans that you apply on your managed workloads to enable and configure the monitoring functionality.

If no monitoring plan is applied on a workload, the monitoring features will not be available for the workload.

Note

The availability of the settings that you can configure in the monitoring plan depends on the service pack that is applied on the tenant. To access all settings, activate the Advanced Management pack.

Creating a monitoring plan

You can create a monitoring plan, and then add workloads to it to configure the monitoring functionality on the managed workloads.

Prerequisites

The version of the agent that is installed on the workload supports the monitoring functionality.

To create a monitoring plan

From Monitoring plans

1. In the Protection console, go to **Management > Monitoring plans**.
2. Create a monitoring plan by using one of the two options.
 - If there are no monitoring plans in the list, click **Create**.
 - If there are monitoring plans in the list, click **Create plan**.
3. In the **Create monitoring plan** window, depending on whether the Advanced Management pack is enabled for your tenant, do the following:

- If your tenant is using Standard protection, the following four monitors are automatically added to the monitoring plan: Disk space, Hardware changes, Last system restart, and Files and folders size.
- If the Advanced Management pack is enabled for your tenant, select one of the template options, and then click **Next**.

Option	Description
Recommended	Select this option to create a monitoring plan with the default monitoring configuration.
Custom	Use this option to create a monitoring plan from scratch.

4. [Optional] To change the default name of the plan, click the pencil icon, enter the name of the plan, and then click **OK**.
5. [Optional] To add a monitor to the plan, click **Add monitor**, click the monitor in the list, and then click **Add**.

Note

The settings of the monitor will be populated automatically with the default values. You can add maximum three monitors of the same type and up to 30 monitors in total to a monitoring plan.

6. [Optional] In the monitor parameters screen, change the default settings of the monitor and alerts, and then click **Done**.

Note

You can configure different settings for each monitor. For more information, see "Configurable monitors" (p. 959) and "Configuring monitoring alerts" (p. 1000).

7. [Optional] To delete a monitor, click the bin icon, and then click **Delete**.
8. [Optional] To add workloads to the plan:
 - a. Click **Add workloads**.
 - b. Select the workloads, and then click **Add**.
 - c. If there are compatibility issues that you want to resolve, follow the procedure as described in "Resolving compatibility issues with monitoring plans" (p. 998).
9. Click **Create**.

From All devices

1. In the Protection console, go to **Devices > All devices**.
2. Click the workload to which you want to apply a monitoring plan.
3. Click **Protect**.
4. Depending on whether a monitoring plan is applied to the workload, do the following:
 - If a monitoring plan is already applied on the workload, click **Create plan**, and then select **Monitoring**.

- If no monitoring plan is applied on the workload, click **Add plan**, and then **Create plan**, and select **Monitoring**.

5. In the **Create monitoring plan** window, select one of the template options, and then click **Next**.

Option	Description
Recommended	Select this option to create a monitoring plan with the default monitoring configuration.
Custom	Use this option to create a monitoring plan from scratch.

6. [Optional] To change the default name of the plan, click the pencil icon, enter the name of the plan, and then click **OK**.
7. [Optional] If you want to change the default settings of the monitor and alerts, configure the new values, and then click **Done**.

Note

You can add maximum three monitors of the same type and up to 30 monitors in total to a monitoring plan.

8. [Optional] In the monitor parameters screen, change the default settings of the monitor and alerts, and then click **Done**.

Note

You can configure different settings for each monitor. For more information, see "Configurable monitors" (p. 959) and "Configuring monitoring alerts" (p. 1000).

9. [Optional] To delete a monitor, click the bin icon, and then click **Delete**.
10. Click **Create**.

Adding workloads to monitoring plans

Depending on your needs, you can add workloads to a monitoring plan after the plan was created.

Prerequisites

- 2FA is enabled for your user account.
- The version of the agent that is installed on the workload supports the monitoring functionality.
- At least one monitoring plan is available.

To add a workload to a monitoring plan

From Monitoring plans

1. In the Protection console, go to **Management > Monitoring plans**.
2. Click the monitoring plan.
3. Depending on whether the plan was already applied to any workload, do the following:

- Click **Add workloads**, if the plan was not applied to any workloads yet.
 - Click **Manage workloads**, if the plan was applied to any workload.
4. Select a workload from the list, and then click **Add**.
 5. Click **Save**.
 6. If necessary, click **Confirm** to apply the required service quota to the workload.

From All devices

1. In the Protection console, go to **Devices > All devices**.
2. Click the workload to which you want to apply a monitoring plan.
3. Click **Protect**.
4. Find the monitoring plan to which you want to add the workload, and click **Apply**.
5. If necessary, click **Confirm** to apply the required service quota to the workload.

Revoking monitoring plans

You can revoke a monitoring plan from a workload to which the plan was applied.

Prerequisites

At least one monitoring plan is applied to the workload.

To revoke a monitoring plan

1. In the Protection console, go to **Devices > All devices**.
2. Click the workload, and then click **Protect**.
3. Click the **More actions** icon of the monitoring plan that you want to revoke, and then click **Revoke**.

Configuring automatic response actions

Automatic response actions on the alerted events are predefined actions or measures that are triggered automatically in response to detected events or incidents. These actions are designed to mitigate potential threats and to minimize damage.

You can configure one or several automatic response actions on the alerted events. The maximum number of automatic response actions per monitor can be 20.

To configure automatic response actions

1. In the Protection console, go to **Management > Monitoring plans**.
2. Select the monitoring plan for which you want to configure automatic response actions.
3. Select the monitor, to which you want to configure automatic response actions, or, if you have not added monitors yet, click **Add monitor**, click the monitor in the list, click **Add**, and then select the monitor.
4. Click the link next to **Automatic response actions**.

5. In the **Automatic response actions** window, add one or several response actions that will be performed automatically when an alert is triggered.
6. Configure each response action. For example, if you have added the response action **Start a Windows service**, do the following:
 - a. Next to **Windows service**, click **Specify**.
 - b. In the **Service** field, select a service to start as a response action.
 - c. Click **Done**.
7. In the list of all added response actions use the up and down arrows or drag and drop to set the sequence of the response actions.
8. Configure how to handle successive response actions if a previous response action fails. Select one of the following:
 - a. **Continue with the next response action.**
 - b. **Do not continue with the next response action.**
9. Click **Done**.
 You will see the number of configured actions next to the **Automatic response actions** setting of your monitoring plan. You can edit or delete these actions, as well as add the new ones at any time later.

The following table lists and describes all the automatic response actions available in the monitor settings.

Automatic response action	Description	Supported OS
Run a script	If you add this action, you can: <ol style="list-style-type: none"> 1. Select a certain script to run on the workload. 2. Specify the account under which you want to execute the script. 3. Specify maximum duration of the operation. 4. Specify PowerShell execution policy. 5. Run a script. To perform this action, you need an Advanced Management pack license for the workload (if not assigned yet). The system will run the selected remote script with specified parameters when the conditions are met.	Windows, macOS
Restart the workload	If you add this action, the system will restart the workload remotely when the conditions are met.	Windows, macOS

Automatic response action	Description	Supported OS
Stop the process	If you add this action, you can specify the process to stop via manual input of process name. The system will stop the process when the conditions are met.	Windows, macOS
Start the Windows service	If you add this action, you can select which Windows service to start from the dynamic list of services populated from the agents. The system will start the service when the conditions are met.	Windows
Stop the Windows service	If you add this action, you can select which Windows service to stop from the dynamic list of services populated from the agents. The system will stop the service when the conditions are met.	Windows
Enable Windows Update	If you add this action, the system will enable Windows Update when the conditions are met. This action is available only for Windows Update status monitor.	Windows
Disable AutoRun on removable drives	If you add this action, the system will disable the AutoRun feature on removable storage media for the workload when the conditions are met. This action is available only for Autorun feature status monitor.	Windows

Additional operations with monitoring plans

From the **Monitoring plans** screen, you can perform the following additional operations with monitoring plans: view details, edit, view the activities, view the alerts, rename, enable, disable, clone, export, and delete.

View details

To view the details of a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **View details**.
3. [Optional] If you want to view the details of a monitor that is enabled in the plan, click the monitor name.

Edit

Prerequisites

2FA is enabled for your user account.

To edit a plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Edit**.
3. [Optional] To delete a monitor from the plan, click the recycle bin icon that is situated to the right of the monitor name.
4. [Optional] To enable or disable a monitor in the plan, use the toggle next to the monitor name.
5. [Optional] To edit the monitor parameters, do the following.
 - a. Click the monitor name.
 - b. Click the overview of the monitor parameters.
 - c. In the **Monitor parameters** screen, configure the parameters, and then click **Done**.

Note

You can configure different settings for each monitor. For more information, see ["Configurable monitors"](#) (p. 959) and ["Configuring monitoring alerts"](#) (p. 1000).

- d. Close the screen and confirm the changes.
6. [Optional] To add a monitor, click **Add monitor**, and then, if necessary, edit the parameters as explained in the previous step.
 7. Click **Save**.

Activities

To view the activities related to a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Activities**.
3. Click an activity to view more details about it.

Alerts

To view the alerts

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Alerts**.

Rename

Prerequisites

2FA is enabled for your user account.

To rename a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Rename**.
3. Enter the new name of the plan, and then click **OK**.

Enable

Prerequisites

- 2FA is enabled for your user account.
- The monitoring plan is applied to at least one workload.

To enable a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Enable**.

Disable

Prerequisites

2FA is enabled for your user account.

To disable a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Disable**.

Clone

Prerequisites

2FA is enabled for your user account.

To clone a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Clone**.
3. Click **Create**.

Export

Prerequisites

2FA is enabled for your user account.

To export a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Export**.

The plan configuration is exported in a JSON format to the local machine.

Delete

Prerequisites

2FA is enabled for your user account.

To delete a monitoring plan

1. In the **Monitoring plans** screen, click the **More actions** icon of the monitoring plan.
2. Click **Delete**.
3. Select **I confirm**, and then click **Delete**.

Compatibility issues with monitoring plans

In some cases, applying a monitoring plan on a workload might cause compatibility issues. You might observe the following compatibility issues:

- Incompatible operating system- this issue appears when the workload's operating system is not supported.
- Unsupported agent - this issue appears when the version of the protection agent on the workload is outdated and does not support the monitoring functionality.
- Insufficient quota - this issue appears when there is not enough service quota in the tenant to assign to the selected workloads.

If the monitoring plan is applied to up to 150 individually selected workloads, you will be prompted to resolve the existing conflicts before saving the plan. To resolve a conflict, remove the root cause for it or remove the affected workloads from the plan. For more information, see "Resolving compatibility issues with monitoring plans" (p. 998). If you save the plan without resolving the conflicts, it will be automatically disabled for the incompatible workloads, and alerts will be shown.

If the monitoring plan is applied to more than 150 workloads or to device groups, first it will be saved, and then checked for compatibility. The plan will be automatically disabled for the incompatible workloads, and alerts will be shown.

Resolving compatibility issues with monitoring plans

Depending on the cause of the compatibility issues, you can perform different actions to resolve the compatibility issues as part of the process of creating a new monitoring plan.

To resolve the compatibility issues

1. Click **Review issues**.
2. [Optional] To resolve compatibility issues with incompatible operating systems by removing workloads from the plan:
 - a. On the **Incompatible operating system** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.

3. [Optional] To resolve compatibility issues with incompatible operating systems by disabling a monitor in the plan:
 - a. On the **Incompatible operating system** tab, select the monitors that you want to remove.
 - b. Click **Disable monitor**.
 - c. Click **Disable**, and then click **Close**.
4. [Optional] To resolve compatibility issues with unsupported agents by removing workloads from the plan:
 - a. On the **Unsupported agents** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
5. [Optional] To resolve compatibility issues with unsupported agents by updating the agent version, click **Go to Agents list**.

Note

This option is available only for customer administrators.

6. [Optional] To resolve compatibility issues with insufficient quota by removing workloads from the plan:
 - a. On the **Insufficient quota** tab, select the workloads that you want to remove.
 - b. Click **Remove workloads from plan**.
 - c. Click **Remove**, and then click **Close**.
7. [Optional] To resolve compatibility issues with insufficient quota by increasing the quota of the tenant:
 - a. On the **Insufficient quota** tab, click **Go to Management portal**.
 - b. Increase the service quota for the customer.

Note

This option is available only for partner administrators.

Resetting the machine learning models

You can reset the models of a workload when they become outdated or invalid for some reason. This action will delete the created models and the data that was collected for the workload by the monitors with anomaly-based monitoring type, and then will start training the machine learning models for the workload from scratch.

To reset the machine learning models for a workload

1. In the Protection console, go to **Devices > All devices**.
2. Click a workload from the list, and then click the **Details** tab.
3. In the **Reset machine learning models** section, click **Reset**.
4. In the confirmation window, click **Reset** again.

Monitoring alerts

Monitoring alerts are displayed in the Protection console and are sent via email when the monitored behavior of workloads is out of norm. The alerts ensure that the stakeholders are informed as soon as possible when there are any issues in the IT environment of the organization.

Note

To enable monitoring alerts via email, you must configure at least one email notification policy for the corresponding alert type. For more information, see "Configuring email notification policies" (p. 1007).

Configuring monitoring alerts

You can configure the monitor's alert settings when you add a monitor to a monitoring plan, or when you edit a monitor that is already available in a monitoring plan.

To configure monitoring alerts

1. In the **Monitor parameters** window, go to the **Generate alerts** section.
2. In **Alert severity**, select the severity that corresponds to the priority of the alert.

Option	Description
Critical	These alerts have the highest priority and are related to issues that are critical for the operation of the workload. Resolve these issues as soon as possible.
Error	An error alert is less severe and indicates that something is wrong or is not behaving normally. Resolve the issues on time to prevent them causing more severe issues.
Warning	A warning alert indicates that there is some condition of which you should be aware, but it might not be causing a problem yet. Resolve these issues after you fix the issues that are causing critical and error alerts. This is the default value.
Informational	These alerts have the lowest priority. The Informational severity does not indicate a problem. Such alerts provide information about actions that are related to a monitored object.

3. In **Alert frequency**, select how often the system should generate an alert when the condition is met.

Option	Description
Once until the check passes	The system will generate an alert one time until the check completes successfully.

Option	Description
	This is the default value.
After X consecutive failures	The system will generate an alert after X consecutive failed checks, where X is an integer value.

- In **Alert message**, click the pencil icon to edit the default alert message that will be used when the system generates an alert. You can specify a custom alert message that contains variables. For more information about the variables that you can use, see "Monitoring alert variables" (p. 1001).

Note

You can configure more than one alert message for some of the monitors.

- Enable **Alert auto-resolution**, if you want the system to automatically resolve the alert when the monitored metric returns to normal state and the behavior is normal again. By default, the setting is enabled.

Monitoring alert variables

You can configure different alert variables for different monitors. To use a variable, it must be enclosed in {{}}.

The following table provides more information about the available variables.

Variable	Description	Available for monitor
plan_name	The name of the policy	All monitors
monitor_name	The name of the sub policy in the monitoring plan	All monitors
workload_name	The name of the workload	All monitors
threshold_value	Specific monitoring conditions or thresholds for generating an alert	All monitors that support threshold-based monitoring.
threshold_unit	The unit that is associated with the threshold value. For example, %, MB, or mb/s.	All monitors that support threshold-based monitoring.
time_period	The system will generate an alert for a detected issue only if the metric value is out of the norm during the specified period.	All monitors that support threshold-based monitoring.
time_unit	The unit that will be associated with the time period (sec/min/hours/day).	All monitors that support threshold-

Variable	Description	Available for monitor
		based monitoring.
anomaly_value	The anomaly value	All monitors that support anomaly-based monitoring.
anomaly_unit	The unit that will be associated with the anomaly value	All monitors that support anomaly-based monitoring.
deviation_value	The deviation value	All monitors that support anomaly-based monitoring.
deviation_unit	The unit that will be associated with the deviation value	All monitors that support anomaly-based monitoring.
drive_name	The drive for Windows, or partition for macOS	Disk space,
CPU_model	The model of the monitored CPU	CPU temperature
GPU_model	The model of the monitored GPU	GPU temperature
hardware_model	The model of the monitored component	Hardware changes
hardware_component	The type of monitored hardware	Hardware changes
hardware_model_old	The model of the monitored component that was replaced	Hardware changes
hardware_model_new	The model of the new monitored component that was added	Hardware changes
disk_model	The model of the disk	Disk transfer rate
network_adapter_model	The model of the network adapter	Network usage
process_name	The name of the process	CPU usage by process Memory usage by process Disk transfer rate by process Network usage by process

Variable	Description	Available for monitor
		Process status
service_name	The name of the service	Windows service status
software_name	The name of the software application	Installed software
software_version	The version of the software application	Installed software
software_version_old	The version of the software application before the update	Installed software
software_version_new	The version of the new or updated software application	Installed software
number_of_occurrences	The number of times an event appears in the log	Windows event log
event_types	The type of the event	Windows event log
event_source	The source of the event	Windows event log
event_log_name	The name of the event	Windows event log
firewall_software_name	The name of the firewall software	Firewall status
antimalware_software_name	The name of the antimalware software	Antimalware software status
user_name	The name of the user	AutoRun feature status
script_name	The name of the script	Custom

Manual response actions

When you see an alert, you can select a response action that you want to perform on the alerted events.

To perform a manual response action

1. In the Protection console, go to **Alerts**.
2. Open the alert that you want to view.
3. Click **Response action**, and then select a response action from the drop-down list.

The list of response actions available for a particular alert depends on the alert type, availability of features for a particular tenant and the workload operating system.

The following table lists and describes all the manual response actions for your reference.

Manual response action	Description	Supported OS
Browse disk space usage trend	<p>Opens a window with Disk space usage graph, where you can:</p> <ul style="list-style-type: none"> • Browse how the disk space usage changed over time (for the last 1 day / 7 days / 1 month). • Browse the delta for disk space usage in relative value (%) for the selected period. 	Windows, macOS
Browse files size growth trend	<p>Opens a window with File size growth graph, where you can:</p> <ul style="list-style-type: none"> • Browse how the total size of the monitored files and folders changed over time (for the last 1 day / 7 days / 1 month). • Browse the delta for total size of files in relative value (%) for the selected period. 	Windows, macOS
Run a script	<p>Opens a window, where you can:</p> <ol style="list-style-type: none"> 1. Select a certain script to run on the workload. 2. Specify the account under which you want to execute the script. 3. Specify maximum duration of the operation. 4. Specify PowerShell execution policy. 5. Run a script. <p>To perform this action, you need Advanced Management Pack license for the workload (if not assigned yet).</p>	Windows, macOS
Connect via NEAR	Acronis Connect Client establishes a remote connection.	Windows, macOS
Connect via RDP	Acronis Connect Client establishes a remote connection.	Windows
Open hardware inventory	You are redirected to Hardware inventory tab for the current workload.	Windows, macOS
Browse top 10 processes that loaded CPU	Opens a window with top 10 processes that have loaded the CPU and may have caused	Windows, macOS

Manual response action	Description	Supported OS
	its overheating (The system snapshot at the moment of alert generation).	
Browse top 10 processes that loaded GPU	Opens a window with top 10 processes that have loaded the GPU and may have caused its overheating (The system snapshot at the moment of alert generation).	Windows, macOS
Browse top 10 processes that loaded memory	Opens a window with top 10 processes that have loaded the memory (The system snapshot at the moment of alert generation).	Windows, macOS
Browse top 10 processes that loaded disk	Opens a window with top 10 processes that have loaded the disk (The system snapshot at the moment of alert generation).	Windows, macOS
Browse top 10 processes that loaded network	Opens a window with top 10 processes that have loaded the network interface adapter (The system snapshot at the moment of alert generation).	Windows, macOS
Browse resource usage by process	Opens a window with detailed information about the usage of hardware resources by the related process: CPU usage, memory usage, disk I/O, network usage.	Windows, macOS
Restart workload	Opens a confirmation window. Restarts the workload after the confirmation.	Windows, macOS
Start Windows service	Opens a confirmation window. Starts the Windows service after the confirmation.	Windows
Stop Windows service	Opens a confirmation window. Stops the Windows service after the confirmation.	Windows
Stop process	Opens a confirmation window. Stops the process to which the alert refers to after the confirmation.	Windows, macOS
Enable Windows Update	Opens a confirmation window. Enables Windows Update after the confirmation.	Windows
Disable AutoRun feature on removable drives	Opens a confirmation window. Disables AutoRun feature on the system level of the workload after the confirmation.	Windows

Important

For security reasons, [two-factor authentication](#) is required to perform the following manual response actions:

- Run a script
 - Connect via NEAR
 - Connect via RDP
 - Restart workload
 - Start Windows service
 - Stop Windows service
 - Stop process
 - Enable Windows Update
 - Disable AutoRun feature on removable drives
-

Viewing the monitoring alerts for a workload

On the **Alerts** tab, you can view the monitoring alerts of a specific workload and perform different alert actions.

To view the monitoring alerts for a workload

1. In the Protection console, go to **All devices**.
2. Click a workload, and then select the **Alerts** tab.
3. [Optional] In the monitoring alert pane, perform one of the following actions:
 - To clear the alert, click **Clear**.
 - To take a response action, click **Response action**, and then click the action.
 - To contact the Support team, click **Get support**.
4. [Optional] To clear all monitoring alerts for the workload, click **Clear all**.

Viewing the alert log of monitoring alerts

You can see all events that are related to a monitoring alert in a chronological order: the response actions (both automatic or manual) that were performed, and the email notifications that were sent.

To view the audit log of a monitoring alert

1. In the Protection console, go to **Alerts**.
2. Open the **Table view**.
3. In the list of alerts, click the monitoring alert that you want to view.
4. Click **Details**, and then click **Alert log**.

Configuring email notification policies

Email notification policies specify which users will receive email notifications from different monitors.

From the **Email notifications** screen, you can perform the following actions with email notification policies: add, edit, enable, disable, and delete.

Add

To add a new email notification policy

1. In the Protection console, go to **Settings > Email notifications**.
2. Click **Add policy**.
3. Click **Select recipients**.
4. In the **Select recipients** screen, select the users that you want to receive email alerts, and then click **Select**.
5. In **Alert types**, select the monitors for which you want the system to send email alerts.
6. Click **Add**.

Edit

To edit an email notification policy

1. In the Protection console, go to **Settings > Email notifications**.
2. Click the ellipsis icon of the notification policy, and then click **Edit**.
3. [Optional] To change the recipients, click **Edit recipients**, add or remove users from the list, and then click **Select**.
4. [Optional] In **Alert types**, select the types of monitoring alerts that you want to be sent to the selected recipients.
5. Click **Save**.

Enable

To enable an email notification policy

1. In the Protection console, go to **Settings > Email notifications**.
2. In the **Email notifications** screen, click the ... icon of the email notification policy.
3. Click **Enable**.

Disable

To disable an email notification policy

1. In the Protection console, go to **Settings > Email notifications**.
2. In the **Email notifications** screen, click the ... icon of the email notification policy.
3. Click **Disable**.

Delete

To delete an email notification policy

1. In the Protection console, go to **Settings > Email notifications**.
2. In the **Email notifications** screen, click the ... icon of the email notification policy.
3. Click **Delete**, and then click **Confirm**.

Viewing monitor data

For each workload, you can view the list of applied monitors, the current status of the monitors, and the historical performance details in a graphical view. You can use this information to analyze the state of the workload and how the state changed in time.

Prerequisites

- A monitoring plan is applied on the workload.
- The workload is online and has data for the corresponding monitor.
- The version of the agent that is installed on the workload supports the monitoring plans.

To view the monitors that are applied to a workload and the monitor data

1. In the Protection console, go to **Devices > All devices**.
2. Click a workload, and then click the **Monitoring** tab.

The **Monitoring** tab displays a widget for each monitor that is enabled for the workload. Each widget displays the following information.

Displayed information	Description
Monitor name	The monitor name
Last result	The latest value of the monitored metric or the latest state of the event
Last check	The date and time when the monitor collected the last data
Alerts	The number of alerts that were generated by the monitor and are still unresolved. If there is at least one unresolved alert generated by this monitor, clicking the number will open the Alerts tab. The alerts will be filtered, and only the alerts for this monitor will be listed.

Note

The widgets become visible on the tab 15 minutes (or the minimum monitor frequency that is set for a monitor) after you apply a monitoring plan to the workload.

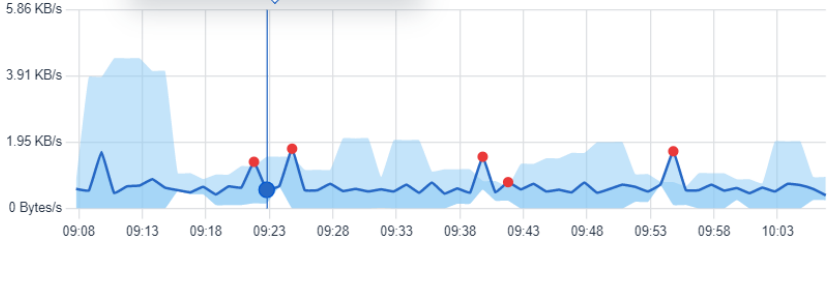
3. [Optional] To view more details about the monitor, and if applicable, the historical data that was collected for the monitored metric, in the monitor's widget, click the ellipsis icon, and then click **Details**.

For more information about the monitor details that you can see in the widgets, see "Monitor widgets" (p. 1009).

Monitor widgets

In the monitor widget, you can see the following details about the monitor.

Detail	Description
Monitoring plan	The name of the monitoring plan that contains the monitor. The name of the monitoring plan is a link that opens the monitoring plan in view mode.
Monitor frequency	The time interval at which the monitor collects data from the workload
Last result	The latest value of the monitored metric or the latest state of the event
Last check	The date and time when the monitor collected the last data
Last alert	The date and time when the last alert was generated. The field is displayed only if at least one alert has been generated for the monitor.
Historical graph	<p>For monitors that collect time-series data, the widget displays historical data for a selected period (1 hour, 6 hours, 12 hours, 1 day, 1 week, or 1 month) in a graphical view.</p> <p>The graph displays the actual values of the metrics during a period that you select. If for some reason the agent did not send the collected data to the cloud, the missing values are displayed as a dotted line that connects the data points with actual values that precede and follow the missing value.</p> <p>For monitors that are using Anomaly-based monitoring, the graph displays the baselines area, a line that shows the actual values of the metric, and the anomalies. The anomalies are the spikes or values that are out of the baselines and are displayed as red dots on the graph.</p> <p>If you hover the mouse over the graph, you can see the actual value and the threshold values for a specific time.</p>

Detail	Description								
	<div data-bbox="389 277 1268 1086"> <p>Monitor details</p> <table border="1"> <tr> <td>Monitoring plan</td> <td>Monitoring plan</td> </tr> <tr> <td>Monitor frequency</td> <td>Every 25 minutes</td> </tr> <tr> <td>Last result</td> <td>Incoming traffic: 0.39 Kb/s</td> </tr> <tr> <td>Last check</td> <td>a few seconds ago</td> </tr> </table> <p> 16 May 2023 09:22:48 Incoming : 563 Bytes/s Lower threshold : 157 Bytes/s Upper threshold : 1.52 KB/s </p> <p> Network usage 1 hour ▾ </p>  <p>The graph displays network usage from 09:08 to 10:03. The y-axis ranges from 0 Bytes/s to 5.06 KB/s. A blue line represents 'Normal behavior', and a light blue area represents the range of usage. A tooltip highlights a data point at 09:22:48 with a value of 563 Bytes/s, which is above the 157 Bytes/s lower threshold.</p> </div> <hr/> <p>Note The data on the graphs is displayed in the time zone of the local system. That is the time zone of the browser of the workload from which you access the Protection console.</p>	Monitoring plan	Monitoring plan	Monitor frequency	Every 25 minutes	Last result	Incoming traffic: 0.39 Kb/s	Last check	a few seconds ago
Monitoring plan	Monitoring plan								
Monitor frequency	Every 25 minutes								
Last result	Incoming traffic: 0.39 Kb/s								
Last check	a few seconds ago								

Additional Cyber Protection tools

Enhanced security mode

The Enhanced security mode is designed for clients with higher security demands. This mode requires mandatory encryption for all backups and allows only locally set encryption passwords.

With the Enhanced security mode, all backups created in a customer tenant and its units are automatically encrypted with the AES algorithm and a 256-bit key. Users can set the encryption passwords only on the protected devices, and cannot set them in the protection plans.

Important

The Enhanced security mode cannot be disabled.

Limitations

- The Enhanced security mode is compatible only with agents version 15.0.26390 or higher.
- The Enhanced security mode is not available for devices running Red Hat Enterprise Linux 4.x or 5.x, and their derivatives.
- Cloud services cannot access encryption passwords. Due to this limitation, some features are not available for tenants in the Enhanced security mode.

Unsupported features

The following features are not available for tenants in the Enhanced security mode:

- Recovery through the Cyber Protect console
- File-level browsing of backups through the Cyber Protect console
- Cloud-to-cloud backup
- Website backup
- Application backup
- Backup of mobile devices
- Antimalware scan of backups
- Safe recovery
- Automatic creation of corporate whitelists
- Data protection map
- Disaster recovery
- Reports and dashboards related to the unavailable features

Setting the encryption password

You must set the encryption password locally, on the protected device. You cannot set the encryption password in the protection plan. Without a password, creating backups will fail.

Warning!

There is no way to recover encrypted backups if you lose or forget the password.

You can set the encryption password in the following ways:

1. During the installation of a protection agent (for Windows, macOS, and Linux).
2. By using the command line (for Windows and Linux).

This is the only way to set an encryption password on a virtual appliance.

For more information on how to set an encryption password with the **Acropsh** tool, refer to "Encryption" (p. 417).

3. In the Cyber Protect Monitor (for Windows and macOS).

To set the encryption password in the Cyber Protect Monitor

1. On the protected device, log on as an administrator.
2. Click the Cyber Protect Monitor icon in the notification area (in Windows) or the menu bar (in macOS).
3. Click the gear icon.
4. Click **Encryption**.
5. Set the encryption password.
6. Click **OK**.

Changing the encryption password

You can change the encryption password before a protection plan creates any backups.

We recommend that you do not change the encryption password after backups are created, because subsequent backups will fail. To continue protecting the same machine, you must create a new protection plan for it. Changing both the encryption password and the protection plan will result in creating new backups that are encrypted with the changed password. The backups that were created before these changes will not be affected.

Alternatively, you can keep the applied protection plan, and change only the backup file name in it. This will also result in creating new backups that are encrypted with the changed password. To learn more about the backup file name, refer to "Backup file name" (p. 425).

You can change the encryption password in the following ways:

1. In the Cyber Protect Monitor (for Windows and macOS).
2. By using the command line (for Windows and Linux).

For more information on how to set an encryption password with the **Acropsh** tool, refer to "Encryption" (p. 417).

Recovering backups for tenants in the Enhanced security mode

With the Enhanced security mode, you cannot recover backups in the Cyber Protect console.

The following options are available:

- Recovering an entire machine, its disks, or files, by using a bootable media.
- Extracting files from local backups of Windows machines with installed agent, by using Windows File Explorer.

Immutable storage

With immutable storage, you can access deleted backups during a specified retention period. You can recover content from these backups, but you cannot change, move, or delete them. When the retention period ends, the deleted backups are permanently deleted.

The immutable storage contains the following backups:

- Backups that are deleted manually.
- Backups that are deleted automatically, according to the settings in the **How long to keep** section in a protection plan or the **Retention rules** section in a cleanup plan.

Deleted backups in the immutable storage still use storage space and are charged accordingly.

Deleted tenants are not charged for any storage, including immutable storage.

Immutable storage modes

For customer tenants, immutable storage is available in the following modes:

Immutable storage is available in the following modes:

- **Governance mode**
You can disable and re-enable the immutable storage. You can change the retention period or switch to Compliance mode.
- **Compliance mode**

Warning!

Selecting Compliance mode is irreversible.

You cannot disable the immutable storage. You cannot change the retention period and cannot switch back to Governance mode.

Supported storages and agents

- Immutable storage is supported only on the cloud storage. Immutable storage is available for Acronis-hosted and partner-hosted storages that use Acronis Cyber Infrastructure (version 4.7.1 or later).

Immutable storage requires that TCP port 40440 is open for the Backup Gateway service in Acronis Cyber Infrastructure. In version 4.7.1 and later, TCP port 40440 is automatically opened with the **Backup (ABGW) public** traffic type. For more information about the traffic types, see [Acronis Cyber Infrastructure documentation](#).

- Immutable storage requires a protection agent version 21.12 (build 15.0.28532) or later.
- Only TIBX (Version 12) backups are supported.

Enabling immutable storage

You can configure the immutable storage settings in the Cyber Protect console or in the management portal. They both provide access to the same settings. The procedure below uses the Cyber Protect console. To learn how to configure the immutable storage settings in the management portal, see [Configuring immutable storage](#) in the administrator guide.

Configuring the immutable storage settings requires two-factor authentication in the tenant to which the administrator account belongs.

To enable immutable storage

1. Log in to the Cyber Protect console as an administrator.
2. Go to **Settings > System settings**.
3. Scroll through the list of default backups options, and then click **Immutable storage**.
4. Enable the **Immutable storage** switch.
5. Specify a retention period between 14 and 3650 days.
The default retention period is 14 days. A longer retention period will result in increased storage usage.
6. Select the immutable storage mode, and then confirm your choice, if prompted.
In the Governance mode, you can enable or disable immutable storage, and change the retention period. You can switch from Governance mode to Compliance mode.

Warning!

Switching to Compliance mode is irreversible. After you select Compliance mode, you cannot disable the immutable storage, or change its mode or retention period.

7. Click **Save**.
8. To make an existing archive support the immutable storage, create a new backup in that archive.
To create a new backup, run the protection plan manually or on a schedule.

Warning!

If you delete a backup before making the archive support the immutable storage, the backup is deleted permanently.

Disabling immutable storage

Note

You can disable the immutable storage only in the Governance mode.

To disable immutable storage

1. Log in to the Cyber Protect console as an administrator.
2. In the navigation menu, click **Settings** > **System settings**.
3. Scroll through the list of default backups options, and then click **Immutable storage**.
4. Disable the **Immutable storage** switch.
5. Confirm your choice by clicking **Disable**.

Warning!

Disabling the immutable storage does not come into effect immediately. During a grace period of 14 days, the immutable storage is still active and you can access the deleted backups according to their original retention period. When the grace period ends, all backups in the immutable storage are permanently deleted.

Accessing deleted backups in immutable storage

During the retention period, you can access deleted backups and recover data from them.

Note

To allow access to deleted backups, port 40440 on the backup storage should be enabled for incoming connections.

To access a deleted backup

1. On the **Backup storage** tab, select the cloud storage that contains the deleted backup.
2. [Only for deleted archives] To see the deleted archives, click **Show deleted**.
3. Select the archive that contains the backup that you want to recover.
4. Click **Show backups**, and then click **Show deleted**.
5. Select the backup that you want to recover.
6. Proceed with the recovery operation, as described in "Recovery" (p. 469).

Geo-redundant storage

Geo-redundant storage ensures data durability by asynchronously copying it to a secondary location that is geographically distant to the primary location. With geo-redundancy, your data is accessible even if the primary location is unavailable.

Important

The replicated data takes up the same storage space as the original data.

Enabling and disabling geo-redundant storage

Prerequisites

- The geo-redundant storage becomes available in the Cyber Protect console only after a partner administrator enables in the Management portal or via API.

- Only administrators can enable or disable the geo-redundant storage in the Cyber Protect console. Make sure you have the administrator rights.

To enable geo-redundant storage

1. [Only if the geo-redundant storage was enabled via API] In the alert on the top "Geo-redundancy is available for all your data in the cloud", click **Enable Geo-redundant Cloud Storage**.
2. In the Cyber Protect console, go to **Settings > System settings**.
3. Scroll through the list of default backups options, and then click **Geo-redundant Cloud Storage**.
4. Enable the **Geo-redundant Cloud Storage** switch.
5. Click **Save**.

Now, your data will be replicated to a secondary location and will stay available even if the primary location fails.

To disable geo-redundant storage

Warning!

The replicated data is deleted within one day after you disable the geo-redundancy.

1. In the Cyber Protect console, go to **Settings > System settings**.
2. Scroll through the list of backups options, and then click **Geo-redundant Cloud Storage**.
3. Disable the **Geo-redundant Cloud Storage** switch.
4. Confirm your choice by typing **Disable**, and then click **Disable**.

Geo-replication status

Geo-redundancy implies that data is replicated to a secondary location. Geo-replication status shows the stages of this process. The following statuses are possible:

- **In sync**—The data has been replicated to the secondary location.
- **Syncing**—The data is being replicated to the secondary location. The duration of this operation depends on the size of the data.
- **On hold**—Data replication is temporarily suspended.
- **Disabled**—Data replication is disabled.

To check the replication status in the Cyber Protect console

1. In the Cyber Protect console, go to **Backup storage**.
2. Select the location and the backup set.
3. Click **Details**, and then check the status in **Geo-replication status**.

Limitations

- Currently, secondary locations for replicated data are only available in the United States and Canada.

- For information about the Disaster Recovery service limitations when using geo-redundancy, see the Disaster Recovery documentation.

Glossary

B

Backup set

A group of backups to which an individual retention rule can be applied. For the Custom backup scheme, the backup sets correspond to the backup methods (Full, Differential, and Incremental). In all other cases, the backup sets are Monthly, Daily, Weekly, and Hourly. A monthly backup is the first backup created after a month starts. A weekly backup is the first backup created on the day of the week selected in the Weekly backup option (click the gear icon, then Backup options > Weekly backup). If a weekly backup is the first backup created after a month starts, this backup is considered monthly. In this case, a weekly backup will be created on the selected day of the next week. A daily backup is the first backup created after a day starts, unless this backup falls within the definition of a monthly or weekly backup. An hourly backup is the first backup created after an hour starts, unless this backup falls within the definition of a monthly, weekly, or daily backup.

C

Cloud server

[Disaster Recovery] General reference to a recovery or a primary server.

Cloud site (or DR site)

[Disaster Recovery] Remote site hosted in the cloud and used for running recovery infrastructure, in case of a disaster.

D

Data loss prevention (formerly, data leak prevention)

A system of integrated technologies and organizational measures aimed at detecting and preventing accidental or intentional disclosure / access to confidential, protected, or sensitive data by unauthorized entities outside or inside the organization, or the transfer of such data to untrusted environments.

Data loss prevention agent

A data loss prevention system's client component that protects its host computer from unauthorized use, transmission, and storage of confidential, protected, or sensitive data by applying a combination of context and content analysis techniques and enforcing centrally managed data loss prevention policies. Cyber Protection provides a fully featured data loss prevention agent. However, the functionality of the agent on a protected computer is limited to the set of data loss prevention features available for licensing in Cyber Protection, and depends upon the protection plan applied to that computer.

Device control module

As part of a protection plan, the device control module leverages a functional subset of the data loss prevention agent on each protected computer to detect and prevent unauthorized access and transmission of data over local computer channels. These include user access to peripheral devices and ports, document printing, clipboard copy/paste operations, media format and eject operations, as well as synchronizations with locally connected mobile

devices. The device control module provides granular, contextual control over the types of devices and ports that users are allowed to access on the protected computer and the actions that users can take on those devices.

Differential backup

A differential backup stores changes to the data against the latest full backup. You need access to the corresponding full backup to recover the data from a differential backup.

F

Failback

Switching a workload from a spare server (such as a virtual machine replica or a recovery server running in the cloud) back to the production server.

Failover

Switching a workload from a production server to a spare server (such as a virtual machine replica or a recovery server running in the cloud).

Finalization

The operation that makes a temporary virtual machine that is running from a backup into a permanent virtual machine. Physically, this means recovering all of the virtual machine disks, along with the changes that occurred while the machine was running, to the datastore that stores these changes.

Full backup

A self-sufficient backup containing all data chosen for backup. You do not need access to any other backup to recover the data from a full backup.

I

Incremental backup

A backup that stores changes to the data against the latest backup. You need access to other backups to recover data from an incremental backup.

L

Local site

[Disaster Recovery] The local infrastructure deployed on your company's premises.

M

Module

Module is a part of protection plan providing a particular data protection functionality, for example, the backup module, the Antivirus & Antimalware protection module, and so on.

O

Orphaned backup

An orphaned backup is a backup that is not associated to a protection plan anymore.

P

Physical machine

A machine that is backed up by an agent installed in the operating system.

Point-to-site (P2S) connection

[Disaster Recovery] A secure VPN connection from outside to the cloud and local sites by using your endpoint devices (such as a computer or laptop).

Primary server

[Disaster Recovery] A virtual machine that does not have a linked machine on the local site (such as a recovery server). Primary servers are used for protecting an application or running various auxiliary services (such as a web server).

Production network

[Disaster Recovery] The internal network extended by means of a VPN tunneling and covering both local and cloud sites. Local servers and cloud servers can communicate with each other in the production network.

Protection agent

Protection agent is the agent to be installed on machines for data protection.

Protection plan

Protection plan is a plan that combines the data protection modules including Backup, Antivirus & Antimalware protection, URL filtering, Windows Defender Antivirus, Microsoft Security Essentials, Vulnerability assessment, Patch management, Data protection map, Device control.

Public IP address

[Disaster Recovery] An IP address that is needed to make cloud servers available from the Internet.

R

Recovery point objective (RPO)

[Disaster Recovery] Amount of data lost from outage, measured as the amount of time from a planned outage or disaster event. RPO

threshold defines the maximum time interval allowed between the last suitable recovery point for a failover and the current time.

Recovery server

[Disaster Recovery] A VM replica of the original machine, based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers, in case of a disaster.

Runbook

[Disaster Recovery] Planned scenario consisting of configurable steps that automate disaster recovery actions.

S

Single-file backup format

A backup format, in which the initial full and subsequent incremental backups are saved to a single .tibx file. This format leverages the speed of the incremental backup method, while avoiding its main disadvantage—difficult deletion of outdated backups. The software marks the blocks used by outdated backups as "free" and writes new backups to these blocks. This results in extremely fast cleanup, with minimal resource consumption. The single-file backup format is not available when backing up to locations that do not support random-access reads and writes.

Site-to-site (S2S) connection

[Disaster Recovery] Connection extending the local network to the cloud, via a secure VPN tunnel.

T

Test IP address

[Disaster Recovery] An IP address that is needed in case of a test failover, to prevent duplication of the production IP address.

Test network

[Disaster Recovery] Isolated virtual network that is used to test the failover process.

U

USB devices database

[Device control] The device control module maintains a database of USB devices from which they can be added to the list of exclusions from device access control. The database registers USB devices by device ID, which can be entered by hand or selected from known devices in the Cyber Protect console.

V

Validation

An operation that checks the possibility of data recovery from a backup. Validation of a file backup imitates recovery of all files from the backup to a dummy destination. Validation of a disk backup calculates a checksum for every data block saved in the backup. Both procedures are resource-intensive. While the successful validation means a high probability of successful recovery, it does not check all factors that influence the recovery process.

Virtual machine

A virtual machine that is backed up at a hypervisor level by an external agent such as Agent for VMware or Agent for Hyper-V. A

virtual machine with an agent inside is treated as physical from the backup standpoint.

VPN appliance

[Disaster Recovery] A special virtual machine that enables connection between the local network and the cloud site via a secure VPN tunnel. The VPN appliance is deployed on the local site.

VPN gateway (formerly, VPN server or connectivity gateway)

[Disaster Recovery] A special virtual machine providing a connection between the local site and the cloud site networks via a secure VPN tunnel. The VPN gateway is deployed on the cloud site.

Index

#

- #CyberFit Score by machine 275
- #CyberFit Score for machines 221
- #CyberFit scoring mechanism 222

3

- 32-bit or 64-bit? 658

A

- About Cyber Disaster Recovery Cloud 679
- About Secure Zone 19, 392
- About the backup schedule 594
- About the Physical Data Shipping service 457
- Access settings 351
- Accessing a virtual appliance via an SSH client 171
- Accessing deleted backups in immutable storage 1015
- Accessing the Cyber Protection service 26
- Action field values 364
- Action on detection 779
- Actions 812
- Actions with protection plans 210
- Activating Startup Recovery Manager 677
- Activating the account 23
- Active Directory Domain Controller for L2 Open VPN connectivity 705
- Active Directory Domain Controller for L3 IPsec VPN connectivity 705
- Active point-to-site connections 716

- Active Protection 763
- Active Protection in the Cyber Backup Standard edition 778
- Active Protection settings in Cyber Backup Standard 779
- Activities tab 306
- Adaptive codec 923
- Add or remove a process, file or network in the protection plan blocklist or allowlist 880
- Add or remove USB devices from the database 349
- Adding a Google Workspace organization 595
- Adding a Microsoft 365 organization 558, 562
- Adding a workload to a remote management plan 934
- Adding access to a Microsoft Azure subscription 514
- Adding credentials 939
- Adding quarantined files to the whitelist 801
- Adding VLANs 672
- Adding workloads to a static group 322
- Adding workloads to monitoring plans 992
- Adding workloads to the Cyber Protect console 310
- Additional Cyber Protection tools 1011
- Additional operations with existing remote management plans 935
- Additional operations with monitoring plans 995
- Additional options 399
- Additional parameters 108

Additional requirement for virtual machines 528

Additional requirements for application-aware backups 519

Additional requirements for machines running Windows 528

Additional scheduling options 410

Adjusting the permissions in data flow policy rules 810

Administering Microsoft 365 organizations added on different levels 563

Advanced 796

Advanced Antimalware 764

Advanced Data Loss Prevention 806

Advanced Data Loss Prevention widgets on the Overview dashboard 825

Advanced settings 817

Advanced storage option 391

Agent-based and agentless backup 67

Agent for Advanced Data Loss Prevention 29

Agent for Data Loss Prevention 28

Agent for Exchange (for mailbox backup) 29

Agent for File Sync & Share 29

Agent for Hyper-V 32

Agent for Linux 30

Agent for Mac 31

Agent for Microsoft 365 29

Agent for MySQL/MariaDB 30

Agent for Oracle 30

Agent for oVirt 33

Agent for oVirt – required roles and ports 156

Agent for Scale Computing HC3 33

Agent for Scale Computing HC3 – required roles 142

Agent for SQL, Agent for Active Directory, Agent for Exchange (for database backup and application-aware backup) 28

Agent for Synology 33

Agent for Virtuozzo 32

Agent for Virtuozzo Hybrid Infrastructure 32

Agent for VMware - LAN-free backup 638

Agent for VMware – necessary privileges 647

Agent for VMware (Virtual Appliance) 32

Agent for VMware (Windows) 32

Agent for Windows 27

Aggregated workloads 371

Alert types 251

Alert widgets 269

Alerts 424

Alerts tab 306

Allowing DHCP traffic over L2 VPN 715

Amazon 46

Analyze incident details 841

Anomaly-based monitoring 958

Antimalware features 762

Antimalware protection alerts 261

Antimalware scan of backups 801

Antivirus and antimalware protection 761

Antivirus and antimalware protection settings 763

Apple Screen Sharing 924

Application-aware backup 526

Applying a default protection plan 220

Applying a plan to a group 340

Applying a protection plan to a workload 211

Approving patches manually 906

Are the required packages already installed? 71

Assessing vulnerabilities and managing patches 885

Assigning credentials to a workload 940

Attaching SQL Server databases 537

Autodiscovery and manual discovery 127

Autodiscovery of machines 124

Automated detection of destination 818

Automated test failover 726, 728

Automatic adding to the whitelist 800

Automatic deletion of unused customer environments on the cloud site 697

Automatic driver search 482

Automatic patch approval 901

Automatic updates for components 179

Availability of the backup options 422

Availability of the recovery options 491

B

Backing up a website 625

Backing up clustered Hyper-V machines 650

Backing up databases included in an AAG 523

Backing up the cloud servers 752

Backing up the Exchange cluster data 525

Backing up workloads to public clouds 510

Backup 375

Backup alerts 251

Backup consolidation 425

Backup file name 425

Backup format 430

Backup format and backup files 430

Backup format compatibility across different product versions 431

Backup options 422

Backup plans for cloud applications 192

Backup replication 194

Backup scanning details 284

Backup scanning plans 192

Backup schedule 395

Backup schemes 395

Backup types 397

Backup validation 431, 492

Backup window 454

Basic parameters 106

Before you start 134, 138, 143, 151, 157

Behavior-engine 768

Boot mode 492

Bootable Media Builder 658

Browsing the hardware inventory 913

Browsing the software inventory 909

Built-in groups 319

Built-in groups and custom groups 319

C

Cache storage 180

calculate hash 445

Capturing network packets 719

Categories to filter 788

Changed block tracking (CBT) 432

Changed Block Tracking (CBT) 636

Changing the backup format to version 12 (TIBX) 430
 Changing the encryption password 1012
 Changing the logon account on Windows machines 85
 Changing the Microsoft 365 access credentials 560
 Changing the ports used by the protection agent 63
 Changing the registration of a workload 124
 Changing the script status 233
 Changing the service quota of machines 181
 Changing the SQL Server or Exchange Server access credentials 546
 Changing the timeout for VM heartbeat and screenshot validation 201
 Check access to the drivers in bootable environment 482
 Check device IP address 409
 Check for indicators of compromise (IOCs) from publicly known attacks on your workloads 854
 Check for publicly disclosed attacks on your workloads using threat feeds 834
 Checking the cloud firewall activities 751
 Checking the size of a search index 615
 Checking the validation status of a backup 203
 Checksum verification 200
 Citrix 41
 Cleanup 203
 Cloning a script 231
 Cloud-only mode 689, 708
 Cloud-to-cloud groups and non-cloud-to-cloud groups 320
 Cloud agent and local agent 554
 Cloud applications 285
 Cloud network infrastructure 687
 Cluster-aware backup 524
 Cluster backup mode 432
 Combining data flow policy rules 811
 Common backup rule 47
 Common installation rule 47
 Common requirements 518
 Comparing script versions 234
 Comparison of the default protection plans 216
 Compatibility issues with monitoring plans 998
 Compatibility issues with remote management plans 937
 Compatibility issues with scripting plans 242
 Compatibility with Dell EMC Data Domain storages 48
 Compatibility with encryption software 47
 Components for unattended installation (EXE) 95
 Components for unattended installation (MSI) 103
 Compression level 433
 Compute points 683
 Configurable monitors 959
 Configuring a CDP backup 389
 Configuring a Site-to-site Open VPN connection 698
 Configuring an application-aware backup 619
 Configuring automated test failover 729
 Configuring automatic patch approval 902

Configuring automatic response actions 993

Configuring Cloud-only mode 698

Configuring custom DNS servers 713

Configuring email notification policies 1007

Configuring encryption as a machine property 418

Configuring encryption in the protection plan 418

Configuring local routing 714

Configuring monitoring alerts 1000

Configuring Multi-site IPsec VPN 700

Configuring network settings 672

Configuring networks in Virtuozzo Hybrid Infrastructure 144

Configuring Point-to-site remote VPN access 705

Configuring proxy server settings 74

Configuring proxy server settings in Cyber Protect Monitor 297

Configuring RDP settings 941

Configuring retention rules 415

Configuring Site-to-site Open VPN 698

Configuring the Connect Client settings 955

Configuring the Multi-site IPsec VPN settings 700

Configuring the number of retries in case of an error 202

Configuring the patch lifetime in the list 901

Configuring the Production patching protection plan 904

Configuring the Test patching protection plan 903

Configuring the virtual appliance 135, 140, 148, 153

Configuring user accounts in Virtuozzo Hybrid Infrastructure 145

Configuring your antivirus and antimalware protection 758

Conflict between a new and existing plan 214

Conflict between an individual and group plan 215

Connecting to a machine booted from bootable media 672

Connecting to a managed workload via a web client 944

Connecting to managed workloads for remote desktop or remote assistance 942

Connecting to unmanaged workloads via Acronis Quick Assist 950

Connecting to unmanaged workloads via IP address 950

Connecting to workloads for remote desktop or remote assistance 918

Connections to remote workloads for remote desktop or remote assistance 925

Continuous data protection (CDP) 386

Control type 665

Conversion to a virtual machine 204

Copying Microsoft Exchange Server libraries 546

Corporate whitelist 800

CPU priority 455

Create a disaster recovery protection plan 684

Creating a backup replication plan 194

Creating a dynamic group 322

Creating a monitoring plan 990

Creating a personal Google Cloud project 596

Creating a primary server 744

Creating a protection plan 209
Creating a recovery server 722
Creating a remote management plan 926
Creating a replication plan 634
Creating a runbook 753
Creating a script 230
Creating a scripting plan 236
Creating a static group 321
Creating a validation plan 198
Creating backups in an existing backup archive 429
Creating bootable media to recover operating systems 656
Creating physical bootable media 657
Creating the data flow policy and policy rules 806
Creating the transform file and extracting the installation packages 168
Creating WinPE or WinRE bootable media 668
Cross-platform recovery 470
Cryptomining process detection 767
Custom groups 319
Custom or ready-made bootable media? 656
Custom scripts 663
Custom sensitivity categories 826
Cyber Disaster Recovery Cloud trial version 682
Cyber Protect console – partner level view 305
Cyber Protect Monitor 33, 296
Cyber Protection 270
Cyber Protection services installed in your environment 182

Cyber Scripting 228
CyberApp workloads 371

D

Data considered PCI DSS 822
Data considered Personally Identifiable Information (PII) 820
Data considered Protected Health Information 819
Data Deduplication 60
Data flow policy renewal 812
Data flow policy structure 808
Data Loss Prevention events 824
Data protection map 280, 292
Data protection map settings 293
Database backup 520
Date and time for files 494
Deactivating Startup Recovery Manager 678
Default actions 795
Default backup file name 427
Default backup options 421
Default protection plans 215
Define response actions for a suspicious file 878
Define response actions for a suspicious process 874
Define response actions for a suspicious registry entry 879
Define response actions for an affected workload 864
Define threat feed settings 855
Defining a backup location in Microsoft Azure 510

Defining how and what to protect 191

Deleting a group 340

Deleting a Microsoft 365 organization 564

Deleting a protection plan 214

Deleting all alerts 292

Deleting backups 504

Deleting backups outside the Cyber Protect console 505

Deleting credentials 940

Deleting custom DNS servers 714

Deleting the machine 631

Deploying Agent for oVirt (Virtual Appliance) 151

Deploying Agent for Scale Computing HC3 (Virtual Appliance) 138

Deploying Agent for Synology 157

Deploying Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance) 143

Deploying Agent for VMware (Virtual Appliance) 134

Deploying agents through Group Policy 165

Deploying the OVA template 152

Deploying the OVF template 135

Deploying the QCOW2 template 139, 147

Description 794

Detection by tactics 274

Device control alerts 363

Device Control alerts 267

Device groups 318

Device types allowlist 356

Devices tab 306

Different login options 923

Disable automatic DRS for the agent 135

Disabling automated test failover 730

Disabling automatic assignment for an agent 643

Disabling full-text search for Gmail backups 617

Disabling immutable storage 1014

Disabling One-click recovery 451

Disaster recovery alerts 255

Disaster Recovery compatibility with encryption software 682

Disaster Recovery failover 873

Discovered machines 271

Disk health monitoring 276

Disk health status alerts 280

Disk health widgets 277

Disk provisioning 636

Distribution algorithm 642

Do not show messages and dialogs while processing (silent mode) 434, 494

Do not start when connected to the following Wi-Fi networks 408

Do not start when on metered connection 407

Download configuration for OpenVPN 716

Downloading data for recently affected workloads 285

Downloading files from the cloud storage 485

Downloading protection agents 78

Downloading the IPsec VPN log files 721

Downloading the logs of the VPN appliance 718

Downloading the logs of the VPN gateway 718

Downloading the output of a scripting operation 234

Downloading the setup program 158

Dynamic groups 320

Dynamic installation and uninstallation of components 87

E

Easy to understand visualization of the attack storyline 833

Editing a default protection plan 221

Editing a dynamic group 339

Editing a protection plan 212

Editing or deleting a script 232

Editing the Recovery server default parameters 686

EDR alerts 267

Enable or disable device control 344

Enable or disable OS notification and service alerts 348

Enable VSS full backup 467

Enabling Advanced Data Loss Prevention in protection plans 815

Enabling and disabling firewall management 798

Enabling and disabling geo-redundant storage 1015

Enabling and disabling the Site-to-site connection 711

Enabling Endpoint Detection and Response (EDR) functionality 834

Enabling enhanced search in encrypted backups 616

Enabling immutable storage 1014

Enabling monitoring mode for Endpoint Detection and Response (EDR) 882

Enabling One-click recovery 449

Enabling or disabling a protection plan 213

Enabling or disabling enhanced search in existing plans 617

Enabling the hardware inventory scanning 912

Enabling the software inventory scanning 908

Enabling the use of the device control module on macOS 345

Encryption 417

Endpoint Detection and Response (EDR) 831

Endpoint Detection and Response (EDR) widgets 271

Enhanced security mode 1011

Error handling 434, 494, 637

Event parameters 403

Example 89, 98, 111, 145-147, 405-409, 414

- Emergency backup in case of bad blocks on the hard disk 403
- Installing the packages manually in Fedora 14 73

Examples 88, 90, 97, 99, 110

Exchange Server clusters overview 524

Exclude device subclasses from access control 348

Exclude individual USB devices from access control 348

Excluding processes from access control 361

Exclusions 797

Executing a runbook 757

Existing vulnerabilities 282

Exploit prevention 769

Exporting backups 503
Extensions and exception rules 295
Extracting files from local backups 488
Extracting the MSI, MST, and CAB files 96

F

Failback options 637
Failback to a target physical machine 738
Failback to a target virtual machine 733
Failing back 636
Failing over to a replica 635
Fast incremental/differential backup 435
Features 833
File-level backup snapshot 437
File-level security 495
File exclusions 494
File filters (Inclusions/Exclusions) 435
Files of a script 663
Filter criteria 436
Finalization of machines running from cloud backups 632
Finalization vs. regular recovery 632
Finalizing the machine 631
Find the last logged in user 374
Firewall management 797
Firewall rules for cloud servers 748
Fits the time interval 406
Flashback 495
Forensic backup process 439
Forensic data 438
Full-text search 614

Full path recovery 495

G

General recommendations for local sites 702
Generating a registration token 166
Geo-redundant storage 1015
Geo-replication status 1016
get content 444
Getting started with Cyber Protection 23
Getting the certificate for backups with forensic data 441
Granting the required system permissions to the Connect Agent 84

H

H.264 923
Hardware inventory 912
Hardware inventory widgets 287
High Availability of a recovered machine 650
How autodiscovery works 125
How creating Secure Zone transforms the disk 19, 393
How do files get into the quarantine folder? 798
How does it work? 828
How failback works 733
How failover works 725
How it works 222, 277, 289, 292, 387, 421, 441, 612, 778, 786
How many agents are required for cluster-aware backup and recovery? 525
How many agents are required for cluster data backup and recovery? 523

How many agents do I need? 135, 139, 144, 152

How remote installation of agents works 127

How routing works 689, 691, 696

How the regular conversion to a virtual machine works 207

How to analyze which security incidents need immediate attention 838

How to assign the user rights 86

How to create Secure Zone 20, 393

How to delete Secure Zone 21, 394

How to get forensic data from a backup? 439

How to investigate incidents in the cyber kill chain 844

How to navigate attack stages 848

How to perform failover of a DHCP server 732

How to perform failover of servers using local DNS 732

How to recover data to a mobile device 549

How to reduce bottlenecks 507

How to review data via the Cyber Protect console 549

How to start backing up your data 548

How to test if Endpoint Detection and Response (EDR) is working correctly 883

How to use Endpoint Detection and Response (EDR) 836

How to use notarization 421, 612

I

If you choose to create the virtual machine on a virtualization server 208

If you choose to save the virtual machine as a set of files 207

Ignore bad sectors 434

Ignore failed VSS writers 466

Immutable storage 1013

Immutable storage modes 1013

Implementing disaster recovery 679

Important tips 412

In 556

In-archive deduplication 431

In Cyber Protection 594

In Google Workspace 594

In Microsoft 365 556

Incident severity history 273

Inclusion and exclusion filters 436

Individual protection plans for hosting control panel integrations 221

Information for partner administrators 314

Information parameters 109

Initial connectivity configuration 698

Installation 81

Installation parameters 106

Installing Agent for Synology 158

Installing agents and components (MSI and MST combination) 96

Installing and deploying Cyber Protection agents 61

Installing and uninstalling agents and components (EXE) 88

Installing and uninstalling agents and components (MSI and direct selection) 97

Installing patches on demand 906

Installing protection agents 78

- Installing protection agents in Linux 80
- Installing protection agents in macOS 83
- Installing protection agents in Windows 78
- Installing the packages from the repository 72
- Installing the packages manually 73
- Integrations for DirectAdmin, cPanel, and Plesk 628
- Interaction with other backup options 460
- Intermediate snapshots 208
- Investigate individual nodes in the cyber kill chain 850
- Investigate the attack stages of an incident 847
- Investigating incidents 843
- IP address reconfiguration 708
- IPsec/IKE security settings 702
- Isolating a workload from the network 368

K

- Kernel parameters 659
- Keyword groups 824
- Known issues 619
- Known issues and limitations 831

L

- License issue 215
- License management for on-premises management servers 190
- Licensing alerts 265
- Limitations 19, 36, 39-40, 42-43, 45-46, 144, 152, 158, 206, 228, 276, 379-380, 384, 386, 392, 472, 493, 557, 575, 580, 583, 595, 601, 605, 609, 619, 625, 639, 677, 681, 802, 1011

- Limitations and known issues 592
- Limitations for backup file names 427
- Limitations for recovering files in the Cyber Protect console 489
- Limitations when using Geo-redundant Cloud Storage 682
- Limiting the total number of simultaneously backed-up virtual machines 650
- Linking workloads to specific users 373
- Linux 381
- Linux-based 657
- Linux-based bootable media 659
- Linux-based or WinPE/WinRE-based bootable media? 657
- Linux packages 70
- list backups 443
- list content 443
- List of USB devices on a computer 361
- Local connection 672
- Local operations with bootable media 672
- Log truncation 446
- Logical expression for all supported languages except Japanese 821
- Logical expression for Japanese 822
- Logical expression used for content detection 819, 821, 823
- LVM snapshotting 447

M

- Mac 382
- Machine migration 652
- Mailbox backup 528
- Malicious website access 788

Manage the network isolation of a workload 865

Manage your incidents in the Incident page 833

Managing access to Microsoft Azure subscriptions 513

Managing discovered machines 132

Managing found vulnerabilities 892

Managing network exclusions 369

Managing networks 706

Managing point-to-site connection settings 715

Managing public cloud account access 513

Managing quarantined files 799

Managing the backup and recovery of workloads and files 375

Managing the cloud servers 747

Managing the detected unprotected files 292

Managing the isolation of workloads 367

Managing the target workloads for a plan 240

Managing the VPN appliance settings 710

Managing virtualization environments 645

Managing workloads in the Cyber Protect console 303

Managing your software and hardware inventory 908

Manual adding to the whitelist 800

Manual binding 643

Manual failback 742

Manual response actions 1003

Marked as Confidential 823

Mass storage drivers to install anyway 482

McAfee Endpoint Encryption and PGP Whole Disk Encryption 48

Microsoft 38

Microsoft 365 seats licensing report 558

Microsoft Azure 46

Microsoft Azure and Amazon EC2 virtual machines 655

Microsoft BitLocker Drive Encryption 47

Microsoft Defender Antivirus 794

Microsoft Defender Antivirus and Microsoft Security Essentials 794

Microsoft Exchange Server 433

Microsoft products 895

Microsoft Security Essentials 794

Microsoft SQL Server 432

Migration via a bootable media 655

Missing updates by categories 283

Monitor widgets 1009

Monitoring 248

Monitoring alert variables 1001

Monitoring alerts 1000

Monitoring plans 958, 990

Monitoring the health and performance of workloads 958

Monitoring types 958

Monitoring workloads via screenshot transmission 947

Mount points 447, 495

Mounting Exchange Server databases 540

Mounting volumes from a backup 501

Multi-site IPsec VPN connection 695

Multi-site IPsec VPN log files 722

Multi-volume snapshot 448

Multitenancy support 307

N

Names without variables 428
NEAR 923
Network folder protection 765
Network management 706
Network requirements for the Agent for
Virtuozzo Hybrid Infrastructure (Virtual
Appliance) 144
Network settings 671
Networking concepts 688
No successful backups for a specified number
of consecutive days 424
Notarization 420, 612
Notarization of backups with forensic data 440
Note for Mac users 471
Nutanix 44

O

Observing multiple managed workloads
simultaneously 948
Obtaining application ID and application
secret 559
Off-host data processing 193
On what workloads, agents, and backup
locations are bottlenecks shown? 509
On Windows Event Log event 402
One-click recovery 448
Operations with a primary server 747
Operations with backups 499
Operations with Microsoft Azure virtual
machines 744
Operations with runbooks 756

Options description 445
Oracle 43
Orchestration (runbooks) 752
Organization map 828
OS notification and service alerts 355
Output speed during backup 456
Overview of the physical data shipping
process 457
oVirt/Red Hat Virtualization 4.2 and 4.3/Oracle
Virtualization Manager 4.3 156
oVirt/Red Hat Virtualization 4.4, 4.5 156

P

Parallels 42
Parameters 660
Parameters for legacy features 109
Parameters for unattended installation
(EXE) 90
Parameters for unattended installation
(MSI) 99
Password requirements 23
Passwords with special characters or blank
spaces 123
Patch a workload 868
Patch installation history 283
Patch installation status 282
Patch installation summary 283
Patch installation widgets 282
Patch management 893
Patch management settings in the protection
plan 894
Payment Card Industry Data Security Standard
(PCI DSS) 822

Performance 496, 637

Performance and backup window 453

Performing a failover 730

Performing a permanent failover 636

Performing a test failover 726

Performing control actions on managed workloads 946

Performing failback to a physical machine 739

Performing failback to a virtual machine 735

Performing manual failback 742

Permissions 811

Personally Identifiable Information (PII) 820

Physical Data Shipping 457

Physical machine to virtual 475

Plan statuses 191

Plans on different administration levels 241

Point-to-site remote VPN access 696

Policy review and management 812

Policy rules for disks and volumes 382

Policy rules for files and folders 384

Ports 698

Ports required by the Downloader component 63

Post-backup command 460

Post-data capture command 462

Post-recovery command 497

Power off target virtual machines when starting recovery 498

Power on the target virtual machine when recovery is complete 498

Pre-backup command 459

Pre-data capture command 461

Pre-recovery command 496

Pre-update backup 899

Pre/Post commands 458, 496, 637

Pre/Post data capture commands 460

Preconfiguring multiple network connections 671

Predefined scripts 662

Preparation 61, 80, 481

 WinPE 2.x and 3.x 669

 WinPE 4.0 and later 670

Prepare drivers 481

Preparing a machine for remote installation 130

Prerequisites 125, 159, 161, 163, 165, 171-172, 228, 372-373, 386, 451, 489, 518, 619, 629, 644, 700, 705, 710, 713-714, 721-722, 735, 740, 744, 906, 909, 911, 913-914, 916, 927, 934-937, 942, 944-948, 950-951, 990, 992-993, 996-998, 1008

Preventing unauthorized uninstallation or modification of agents 176

Primary servers 694

Prioritize which incidents need immediate attention 838

Privacy settings 25

Privileges required for the logon account 86

Production failover 725

Protected Health Information (PHI) 819

Protecting a domain controller 517

Protecting Always On Availability Groups (AAG) 522

Protecting Database Availability Groups (DAG) 524

Protecting Exchange Online data 566

Protecting Exchange Online mailboxes 560

Protecting Gmail data 601

Protecting Google Drive files 604

Protecting Google Workspace data 593

Protecting Hosted Exchange data 550

Protecting Microsoft 365 collaboration app seats 593

Protecting Microsoft 365 data 553

Protecting Microsoft 365 Teams 583

Protecting Microsoft applications 517

Protecting Microsoft SharePoint 517

Protecting Microsoft SQL Server and Microsoft Exchange Server 517

Protecting mobile devices 547

Protecting MySQL and MariaDB data 618

Protecting OneDrive files 575

Protecting OneNote notebooks 592

Protecting Oracle Database 618

Protecting SAP HANA 618

Protecting Shared drive files 608

Protecting SharePoint Online sites 579

Protecting web hosting servers 627

Protecting websites 624

Protecting websites and hosting servers 624

Protection exclusions 775

Protection of collaboration and communication applications 247

Protection plan cheat sheet 377

Protection plans 192

Protection plans and modules 208

Protection settings 179

Protection status 270

Public and test IP address 693

Q

Quarantine 768, 798

Quarantine location on machines 799

Quick glance overview in the dashboard 834

Quotas 627

R

RDP 924

Re-attempt, if an error occurs 434, 494

Re-attempt, if an error occurs during VM snapshot creation 435

Re-generate configuration 716

Real-time protection 762, 770, 796

Reassigning IP addresses 712

Receive alert notifications when a breach happens 833

Recently affected 284

Recommendations 493

Recommendations and remediation steps 834

Recommendations for the Active Directory Domain Services availability 705

Recording and playing remote connection sessions 954

Recovering a machine 472

Recovering a machine with One-click recovery 451

Recovering a team mailbox 588

Recovering a team site or specific items of a site 591

Recovering a virtual machine 477

Recovering a website 626

Recovering an entire Google Drive 606
 Recovering an entire OneDrive 576
 Recovering an entire Shared drive 610
 Recovering an entire team 585
 Recovering applications 517
 Recovering backed-up OneNote notebooks 592
 Recovering backups for tenants in the Enhanced security mode 1012
 Recovering data from an application-aware backup 620
 Recovering databases 622
 Recovering disks by using bootable media 479
 Recovering email messages and meetings 590
 Recovering entire mailboxes to PST data files 571
 Recovering ESXi configuration 490
 Recovering Exchange databases 538
 Recovering Exchange mailboxes and mailbox items 540
 Recovering files 484
 Recovering files by using bootable media 487
 Recovering files in the Cyber Protect console 484
 Recovering Google Drive and Google Drive files 606
 Recovering Google Drive files 607
 Recovering instances 621
 Recovering mailbox items 543, 552, 562, 569, 603
 Recovering mailbox items to PST files 572
 Recovering mailboxes 542, 551, 561, 568, 602
 Recovering mailboxes and mailbox items 551, 561, 568, 602
 Recovering OneDrive and OneDrive files 576
 Recovering OneDrive files 578
 Recovering physical machines 472
 Recovering public folders and folder items 573
 Recovering Shared drive and Shared drive files 610
 Recovering Shared drive files 611
 Recovering SharePoint Online data 581
 Recovering SQL databases 530
 Recovering SQL databases as files 534
 Recovering SQL databases to a non-original machine 532
 Recovering SQL databases to the original machine 530
 Recovering stored routines 624
 Recovering system databases 537
 Recovering system state 490
 Recovering tables 623
 Recovering team channels or files in team channels 586
 Recovering team mailbox items to PST files 589
 Recovering the entire server 621
 Recovering the Exchange cluster data 525
 Recovering the master database 537
 Recovery 469
 Recovery cheat sheet 469
 Recovery from a network share 662
 Recovery from backup 872
 Recovery from the cloud storage 662
 Recovery of databases included in an AAG 523

- Recovery options 491
- Recovery servers 692
- Recovery to Virtuozzo containers or Virtuozzo virtual machines 489
- Recovery with bootable media on-premises 673
- Recovery with restart 473
- Red Hat and Linux 41
- Redirecting sound from a remote Linux workload 925
- Redirecting sound from a remote macOS workload 924
- Redirecting sound from a remote Windows workload 924
- Redistribution 642
- Registering and unregistering workloads manually 119
- Registering the bootable media 670
- Registration parameters 107
- Regular conversion to virtual machine vs. running a virtual machine from a backup 207
- Reinstalling the VPN gateway 710
- Remediate a false positive incident 861
- Remediate an entire incident 857
- Remediating incidents 857
- Remote connection protocols 923
- Remote connection to a workload 871
- Remote management plans 926
- Remote operations with bootable media 674
- Remote sessions widget 288
- Remote sound redirection 924
- Removing access to a Microsoft Azure subscription 516
- Removing workloads from a remote management plan 934
- Removing workloads from the Cyber Protect console 315
- Renewing access to a Microsoft Azure subscription 515
- Renewing the policy for a company or unit 813
- Renewing the policy for one or more users in the company or unit 813
- Replication 416
- Replication of virtual machines 632
- Replication options 636
- Replication vs. backing up 633
- Reported data according to widget type 301
- Reports 297
- Required permissions for unattended installation in macOS 112
- Required ports 156
- Required roles 156
- Required user rights 530, 556, 594
- Required user rights for application-aware backups 527
- Requirements 488, 501
- Requirements for ESXi virtual machines 519
- Requirements for Hyper-V virtual machines 520
- Requirements for the VPN appliance 698
- Requirements on User Account Control (UAC) 130
- Requirements on user accounts 541
- Resetting the machine learning models 999
- Resolving compatibility issues with monitoring

- plans 998
- Resolving compatibility issues with remote management plans 938
- Resolving compatibility issues with scripting plans 243
- Resolving plan conflicts 214
- Response actions for individual cyber kill chain nodes 862
- Restart a workload 869
- Restrictions 633
- Retention rules 411
- Retention rules according to the backup scheme 412
- Reverting to the original initial RAM disk 483
- Review and analyze discovered IOCs 856
- Review and mitigate IOCs on affected workloads 855
- Reviewing incidents 837
- Revoking a plan from a group 341
- Revoking a protection plan 213
- Revoking monitoring plans 993
- Rule structure 808
- Run an on-demand forensic backup on a workload 870
- Run as virtual machine 200
- Runbook parameters 755
- Running a #CyberFit Score scan 226
- Running a backup manually 410
- Running a backup on a schedule 397
- Running a hardware inventory scan manually 913
- Running a software inventory scan manually 909

- Running a virtual machine from a backup (Instant Restore) 628
- Running cloud-to-cloud backups manually 192
- Running pre-freeze and post-thaw scripts automatically 644
- Running the machine 629
- Running the Test patching protection plan and decline unsafe patches 905

S

- Safe recovery 471
- Save battery power 407
- Save system information if a recovery with reboot fails 494
- Saving an agent log file 183
- Scale Computing 40
- Scanning types 762
- Schedule 238, 293, 888, 897
- Schedule and start conditions 238
- Schedule by events 400
- Schedule by time 398
- Schedule scan 771, 795
- Scheduled scan 763
- Scheduling 463
- Screenshot validation 201
- Script quick run 244
- Script repository 234
- Script versions 232
- Scripting plans 235
- Scripts 229
- Scripts in bootable media 662
- Search attributes for cloud-to-cloud

- workloads 325
- Search attributes for non-cloud-to-cloud workloads 326
- Search in cloud-to-cloud backups 613
- Search indexes 614
- Search operators 337
- Sector-by-sector backup 463
- Security 924
- Security incident burndown 274
- Security incident MTTR 273
- Seeding an initial replica 637
- Select the snapshot provider 466
- Selecting a destination 390
- Selecting components for installation 131
- Selecting data to back up 379
- Selecting disks or volumes 380
- Selecting entire machine 379
- Selecting ESXi configuration 386
- Selecting Exchange Online mailboxes 551
- Selecting Exchange Server data 521
- Selecting Exchange Server mailboxes 529
- Selecting files or folders 383
- Selecting Gmail mailboxes 602
- Selecting Google Drive files 605
- Selecting mailboxes 567
- Selecting Microsoft 365 mailboxes 561
- Selecting OneDrive files 576
- Selecting public folders 568
- Selecting Shared drive files 609
- Selecting SharePoint Online data 580
- Selecting SQL databases 521
- Selecting system state 385
- Selecting teams 584
- Self-protection 766
- Self-service custom folder on-demand 799
- Sensitive data definitions 818
- Server-side protection 765
- Services installed in macOS 183
- Services installed in Windows 182
- Setting firewall rules for cloud servers 749
- Setting the encryption password 1011
- Setting the frequency of Google Workspace backups 600
- Setting the frequency of Microsoft 365 backups 565
- Setting the root password on a virtual appliance 170
- Setting up a display mode 673
- Setting up connectivity 687
- Setting up primary servers 744
- Setting up recovery servers 722
- Setting up the disaster recovery functionality 684
- Setting up the Group Policy object 169
- Settings of the Antimalware software status monitor 987
- Settings of the AutoRun feature status monitor 988
- Settings of the CPU temperature monitor 966
- Settings of the CPU usage by process monitor 978
- Settings of the CPU usage monitor 969
- Settings of the Custom monitor 989
- Settings of the Disk space monitor 963

Settings of the Disk transfer rate by process monitor 979

Settings of the Disk transfer rate monitor 973

Settings of the Failed logins monitor 986

Settings of the Files and folders size monitor 985

Settings of the Firewall status monitor 986

Settings of the GPU temperature monitor 967

Settings of the Hardware changes monitor 969

Settings of the Installed software monitor 983

Settings of the Last system restart monitor 983

Settings of the Memory usage by process monitor 978

Settings of the Memory usage monitor 971

Settings of the Network usage by process monitor 980

Settings of the Network usage monitor 975

Settings of the Process status monitor 982

Settings of the Windows event log monitor 984

Settings of the Windows service status monitor 982

Settings of the Windows Update status monitor 986

SID changing 498

Signing a file with ASign 486

Site-to-site Open VPN - Additional information 183

Site-to-site Open VPN connection 690, 706

Skip the task execution 465

Smart protection 288

Software-specific recovery procedures 47

Software inventory 908

Software inventory widgets 286

Software management tab 307

Software requirements 27, 680, 834

Sound transfer 923

Special operations with virtual machines 628

Splitting 464

SQL Server high-availability solutions overview 522

SSH connections to a virtual appliance 170

Start conditions 238, 404

Starting the Secure Shell daemon 170

Startup Recovery Manager 676

Static groups 319

Static groups and dynamic groups 319

Step 1 61

Step 2 61

Step 3 61

Step 4 61

Step 5 62

Step 6 63

Stopping a runbook execution 757

Stopping failover 635

Store security events for 180 days 834

Structure of autostart.json 664

Support for virtual machine migration 645

Supported features per platform 759

Supported Apple and third-party products 887

Supported Apple products 887

Supported cluster configurations 523-524

Supported data sources 388

Supported destinations 389

Supported file systems 58

Supported languages 819-820, 822-823

Supported Linux products 888

Supported locations 196-197, 203, 416

Supported MariaDB versions 34

Supported Microsoft and third-party products 885

Supported Microsoft Exchange Server versions 33

Supported Microsoft products 886

Supported Microsoft SharePoint versions 34

Supported Microsoft SQL Server versions 33

Supported mobile devices 547

Supported MySQL versions 34

Supported operating systems 680

Supported operating systems and environments 27

Supported operating systems and versions 49

Supported Oracle Database versions 34

Supported plans for device groups 320

Supported platforms 228, 758, 922

Supported platforms for monitoring 959

Supported protection features by operating system 49

Supported remote desktop and assistance features 919

Supported SAP HANA versions 34

Supported third-party products for macOS 887

Supported third-party products for Windows OS 887

Supported versions 592

Supported virtual machine types 206

Supported virtualization platforms 35, 680

Supported web browsers 27

Supported Windows operating systems 797

Switching the Site-to-site connection type 711

System alerts 269

System requirements 698

System requirements for agents 68

System requirements for the agent 134, 138, 143, 151

T

Task failure handling 464

Task start conditions 465

TCP ports required for backup and replication of VMware virtual machines 62

Tenants in the Enhanced security mode 489

Test failover 726

Testing a replica 634

The Activities dashboard 249

The Activities tab 295

The Alerts dashboard 250

The backup location's host is available 405

The Backup storage tab 499

The Cyber Protect console 303

The key functionality 679

The Management tab 191

The Overview dashboard 248

The patch management workflow 894

The remote desktop notifiers 956

The tool "tibxread" for getting the backed-up data 442

The way of using Secure Zone 47

Threat feed 288

Threat status 272
Top-level object 664
Top incident distribution per workload 272
Transferring files 945
Transferring files via Acronis Quick Assist 951
Troubleshooting 133
Troubleshooting IPsec VPN configuration issues 719
Troubleshooting the IPsec VPN configuration 719
Two-factor authentication 23

U

Unassigning credentials from a workload 940
Unattended installation and uninstallation in macOS 110
Unattended installation and uninstallation with an EXE file 88
Unattended installation and uninstallation with an MSI file 95
Unattended installation or uninstallation 87
Unattended installation or uninstallation in Linux 104
Unattended installation or uninstallation in Windows 87
Unattended installation or uninstallation parameters 106
Understand the actions taken to mitigate an incident 852
Understanding and customizing the cyber kill chain view 846
Understanding the detection of bottlenecks 506
Understanding the scope and impact of incidents 840

Understanding your current level of protection 248
Uninstallation parameters 109
Uninstalling agents 177
Universal Restore in Linux 483
Universal Restore in Windows 481
Universal Restore process 482
Universal Restore settings 482
Unsupported features 1011
Updating Agent for Synology 163
Updating agents 171
Updating agents automatically 174
Updating agents manually 172
Updating agents on BitLocker-protected workloads 176
Updating the Cyber Protection definitions by schedule 180
Updating the Cyber Protection definitions on-demand 180
Updating, rebuilding, or deleting indexes 615
URL exclusions 793
URL filtering 785
URL Filtering alerts 266
URL filtering configuration workflow 788
URL filtering settings 788
Usage examples 416, 628, 633, 644
Usage scenarios 501
USB devices allowlist 357
USB devices database 358
USB devices database management page 359
Use case for automatic patch approval and testing 902

Use case for automatic patch approval without testing 905

Useful tips 563, 596

User is idle 405

User roles and Cyber Scripting rights 245

Users logged off 406

Using a locally attached storage 641

Using device control 344

Using the Adaptive enforcement mode for renewing a user policy 814

Using the cloud Agent for Microsoft 365 562

Using the Cyber Protect console as a partner administrator 304

Using the locally installed Agent for Office 365 558

Using the Observation mode for renewing a user policy 813

Using the toolbar in the Viewer window 952

Using Universal Restore 481

Using variables 428

v

Validating backups 502

Validation 196

Validation methods 200

Validation status 197

Variable object 664

Verifying file authenticity with Notary Service 486, 613

View device control alerts 351

View or change access settings 347

Viewing and updating Microsoft Azure backup locations 512

Viewing backup status in vSphere Client 646

Viewing bottleneck details 507

Viewing details about items in the whitelist 801

Viewing monitor data 1008

Viewing the alert log of monitoring alerts 1006

Viewing the automated test failover status 729

Viewing the distribution result 643

Viewing the execution history 757

Viewing the hardware of a single device 916

Viewing the list of available patches 899

Viewing the monitoring alerts for a workload 1006

Viewing the software inventory of a single device 911

Viewing which incidents are currently not mitigated 839

Viewing workloads managed by RMM integrations 370

Virtual machine binding 642

Virtuozzo 44

Vituozzo Hybrid Infrastructure 45

VM heartbeat 201

VM power management 498, 637

VMware 35

Volume Shadow Copy Service (VSS) 465

Volume Shadow Copy Service (VSS) for virtual machines 467

Volume Shadow Copy Service VSS for virtual machines 637

VPN access to local site 716

VPN appliance 692

VPN gateway 691, 696

VPN gateway network configuration 692
Vulnerability assessment 885
Vulnerability assessment for Linux machines 891
Vulnerability assessment for macOS devices 891
Vulnerability assessment for Windows machines 890
Vulnerability assessment settings 888
Vulnerability assessment widgets 281
Vulnerable machines 281

W

Wait until the conditions from the schedule are met 465
Weekly backup 469
What's new in the Cyber Protect console 304
What do I need to back up a website? 624
What do I need to use application-aware backup? 526
What does a disk or volume backup store? 381
What does Google Workspace protection mean? 593
What exactly are incidents? 837
What information is included in an attack stage? 848
What is a backup file? 426
What is a bottleneck? 506
What items can be backed up? 550, 560, 566, 575, 579, 583, 601, 604, 608, 624
What items can be recovered? 550, 560, 566, 575, 580, 583, 601, 605, 609
What items cannot be recovered? 580
What to do next 685
What to replicate 195
What to scan 888
What triggers a policy rule? 810
What you can back up 547
What you can do with a replica 633
What you need to know 547
What you need to know about conversion 206
What you need to know about finalization 632
Where can I see backup file names? 426
Where to get the Cyber Protect app 548
Which agent do I need? 63
Which backup type do I need? 67
Whitelist settings 801
Why are there monthly backups with an hourly scheme? 413
Why back up Microsoft 365 data? 553
Why use application-aware backup? 526
Why use Bootable Media Builder? 658
Why use runbooks? 753
Why use Secure Zone? 19, 392
Why you need Endpoint Detection and Response (EDR) 832
Windows 381
Windows event log 469, 498
Windows third-party products 895
WinPE-based and WinRE-based bootable media 667
WinPE images 667
WinPE/WinRE-based 657
WinRE images 667
Wiping data from a managed workload 366
Working in VMware vSphere 632

Working with Advanced protection features 804

Working with aggregated workloads 372

Working with CyberApp workloads 371

Working with encrypted backups 743

Working with logs 717

Working with managed workloads 941

Working with the Device control module 342

Working with unmanaged workloads 949

Workload credentials 939

Workload network status 275

Workloads 308