Acronis

acronis.com

# Disaster Recovery Hybrid Cloud

25.03

Administrator Guide

REVISION: 3/21/2025

# Table of contents

# Copyright statement



# Acronis patented technologies

Technologies, used in this product, are covered and protected by one or more U.S. Patent Numbers: 7,047,380; 7,246,211; 7,275,139; 7,281,104; 7,318,135; 7,353,355; 7,366,859; 7,383,327; 7,475,282; 7,603,533; 7,636,824; 7,650,473; 7,721,138; 7,779,221; 7,831,789; 7,836,053; 7,886,120; 7,895,403; 7,934,064; 7,937,612; 7,941,510; 7,949,635; 7,953,948; 7,979,690; 8,005,797; 8,051,044; 8,069,320; 8,073,815; 8,074,035; 8,074,276; 8,145,607; 8,180,984; 8,225,133; 8,261,035; 8,296,264; 8,312,259; 8,347,137; 8,484,427; 8,645,748; 8,732,121; 8,850,060; 8,856,927; 8,996,830; 9,213,697; 9,400,886; 9,424,678; 9,436,558; 9,471,441; 9,501,234; and patent pending applications.

# Introduction

This guide describes the Disaster Recovery (DR) hybrid deployment case when Acronis Cyber Protect Cloud is deployed in the Acronis data center, while the backup storage and disaster recovery infrastructure are deployed on the partner's premises.

The guide describes how to deploy and configure the Acronis Cyber Protect Cloud with Advanced Disaster Recovery service on top of Acronis Cyber Infrastructure.

The Disaster Recovery infrastructure is a set of software components that you must install on Acronis Cyber Infrastructure, and register in Acronis Cyber Protect Cloud as an infrastructure entity with disaster recovery capabilities.

This document assumes that you are familiar with the following products:

- Linux
- Acronis Cyber Protect Cloud
- Acronis Cyber Infrastructure

With the DR hybrid deployment, partners can build their own Disaster Recovery infrastructure to achieve optimal configuration and control costs.

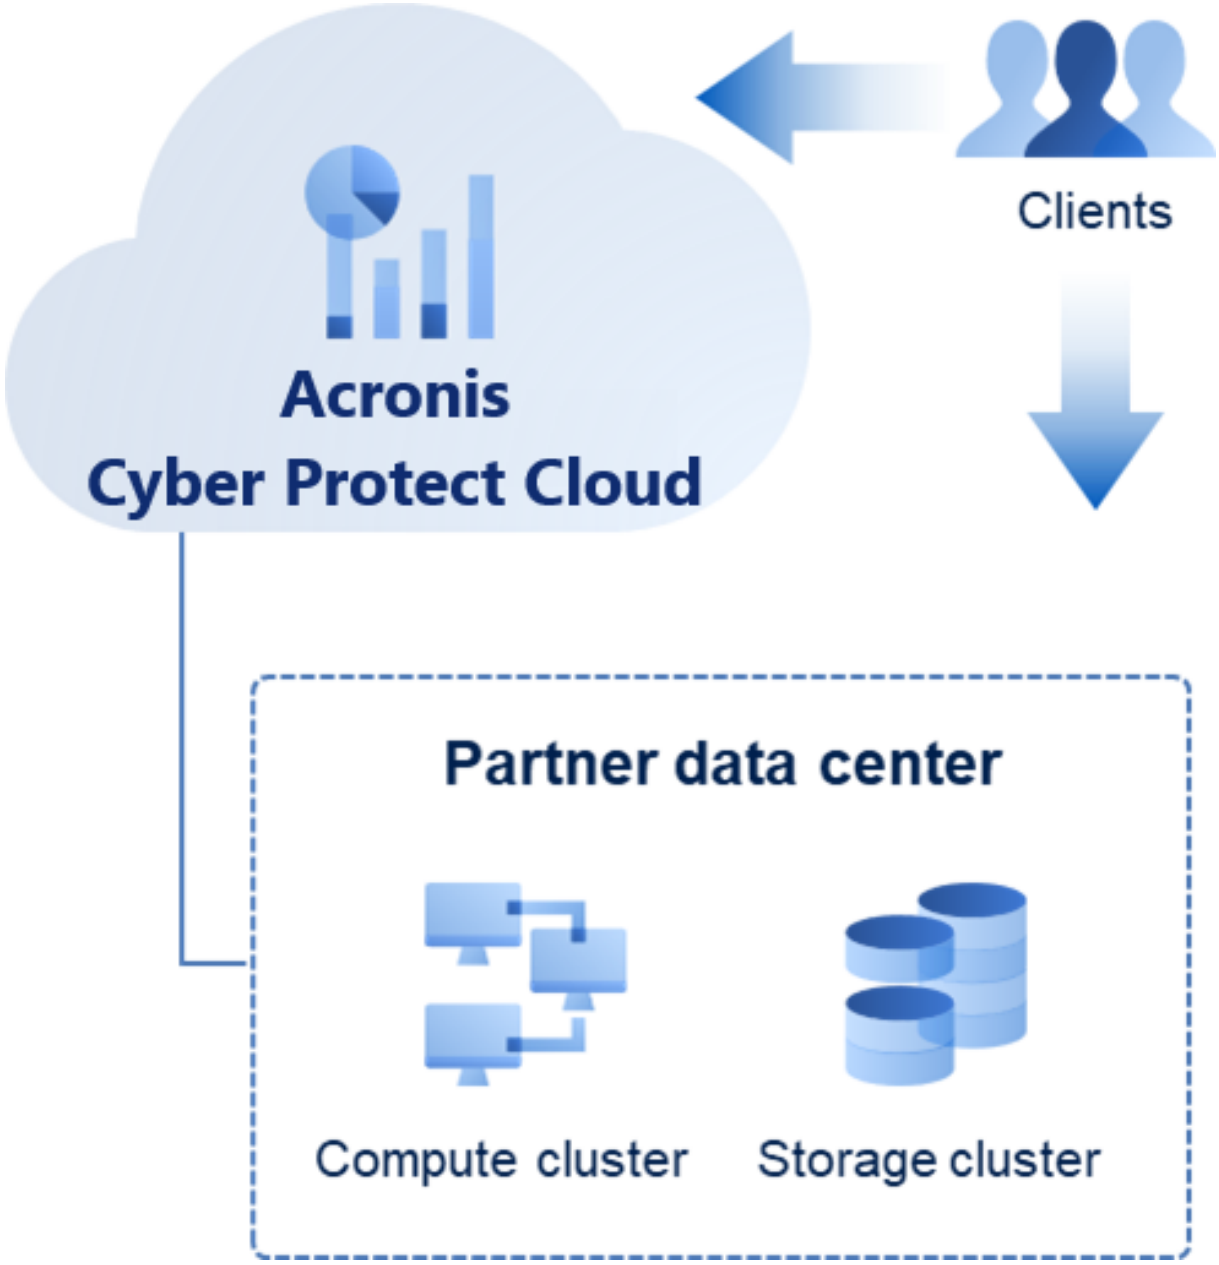


For more details about the Disaster Recovery architecture and main components, refer to "Appendix C".

# Infrastructure planning

This section describes different options of deployment of the backup storage and DR infrastructure.

## Planning capacity for new deployment of Backup and DR infrastructure

If you do not have backup storage deployed, follow the instructions in this section. It describes two options for deploying the backup storage and DR infrastructure.

### Planning capacity for a backup storage

The backup storage requires deployment of a storage cluster.

The Acronis Cyber Infrastructure storage cluster is a group of physical nodes connected to each other by network. Each of the nodes has one or several roles and runs the corresponding services according to its role.

The Acronis Cyber Infrastructure storage cluster consists of a single management node and a number of storage nodes.

To organize a backup storage, do the following:

1. Plan the infrastructure for the backup storage cluster. Consider the hardware requirements for the storage node depending on a usage scenario, the minimum configuration for a storage cluster, the recommended configuration for a storage cluster, and the network hardware recommendations.
   For more information, see Acronis Cyber Infrastructure Installation Guide.
2. Install and configure Acronis Cyber Infrastructure.
   For more details about the installation procedure, see Installing Acronis Cyber Infrastructure.
3. Configure the networks in Acronis Cyber Infrastructure.
   For more information, see Configuring networks.
4. Create the storage cluster.
   For more information, see Creating the storage cluster.

## Option 1. Compute (with disks for hot storage) cluster and storage (with disks for cold storage) cluster

If you want to have two separate clusters - a compute cluster with compute resources (without disks) only, and another cluster with hot and cold storage, use this deployment case.

### Minimum configuration

The infrastructure parameters listed below provide the following protected environment profile: 100 tenants, 500 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of HDD, 500 GB of

SSD. About 10% of the protected virtual machines (50 VMs) can run simultaneously on the specified below infrastructure.

The minimum required number of nodes for the DR infrastructure deployment:

- Three compute nodes
- Three storage nodes

### Compute resources

| Server name | Compute node |
|---|---|
| **CPU** | 32 cores |
| **RAM** | 256 GB |

### Disaster Recovery storage (hot) and Backup storage (cold)

| Server name | Storage node |
|---|---|
| **CPU** | 8 cores |
| **RAM** | 32 GB |
| **Disk** | 9 x 1,86 TB SSD, 4 x 6 TB HDD |

## Production configuration

The infrastructure parameters listed below provide the following protected environment profile: 100 tenants, 1000 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of HDD, 500 GB of SSD. About 10% of the protected virtual machines (100 VMs) can be run simultaneously on the specified below infrastructure.

The minimum required number of nodes for the DR infrastructure deployment:

- Five compute nodes
- Six storage nodes

### Compute resources

| Server name | Compute node |
|---|---|
| **CPU** | 24 cores |
| **RAM** | 256 GB |

### Disaster Recovery storage (hot) and Backup storage (cold)

| Server name | Storage node |
|---|---|

| | |
|---|---|
| **CPU** | 8 cores |
| **RAM** | 32 GB |
| **Disk** | 9 x 1,86 TB SSD, 4 x 6 TB HDD |

## Option 2. Hyperconverged cluster

A hyperconverged cluster implies that each cluster node has compute resources, hot and cold disks. Below you will find the minimum and production configurations.

### Minimum configuration

The infrastructure parameters listed below provide the following protected environment profile: 100 tenants, 500 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of HDD, 500 GB of SSD. About 10% of the protected virtual machines (50 VMs) can be run simultaneously on the specified below infrastructure.

To deploy the DR infrastructure, you need a minimum of three compute nodes.

#### Compute resources, hot and cold storage

| Server name | Compute node with disks for storage |
|---|---|
| **CPU** | 32 cores |
| **RAM** | 512 GB |
| **Disks** | 6 x 8 TB SSD (hot)<br>4 x 16 TB HDD (cold) |

### Production configuration

The infrastructure parameters listed below provide the following protected environment profile: 100 tenants, 1000 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of HDD, 500 GB of SSD. About 10% of the protected virtual machines (100 VMs) can be run simultaneously on the specified below infrastructure.

To deploy the DR infrastructure, you need a minimum of five cluster nodes.

#### Compute resources, hot and cold storage

| Server name | Compute node with disks for storage |
|---|---|
| **CPU** | 32 cores |
| **RAM** | 512 GB |
| **Disks** | 2 x 8 TB SSD (hot) |

| | |
|---|---|
| | 3 x 16 TB HDD (cold) |

# Planning capacity for DR infrastructure

If you already have a backup storage deployed, you may deploy only the DR infrastructure. This requires a deployment of a compute cluster with disks for hot storage.

## General hardware recommendations

Acronis Cyber Infrastructure works on the same hardware that is recommended for Red Hat Enterprise Linux 7 and Red Hat Enterprise Linux 9 (for the Acronis Cyber Infrastructure version 6.0 and above). For more information, see servers and hardware components.

General considerations:

- One of the strongest features of Acronis Cyber Infrastructure is scalability. The bigger the cluster, the better Acronis Cyber Infrastructure performs.
- It is not recommended for production to run Acronis Cyber Infrastructure on top of SAN/NAS hardware that has its own redundancy mechanisms.
- To achieve best performance, keep at least 20% of cluster capacity free.
- It is recommended to have the same CPU models on each node to avoid virtual machine live migration issues.

## Compute (with hot disks) cluster

If you already have a backup storage (cold) and you want to add the DR infrastructure, use this deployment case. Below, you will find the minimum and production configurations.

### Minimum configuration

The infrastructure parameters listed below provide the following protected environment profile: 100 tenants, 500 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of HDD, 500 GB of SSD. About 10% of the protected virtual machines (50 VMs) can be run simultaneously on the specified below infrastructure.

To deploy the DR infrastructure, you need a minimum of three compute nodes.

#### Compute resources and Disaster Recovery storage (hot)

| Server name | Compute node |
|---|---|
| **CPU** | 32 cores |
| **RAM** | 512 GB |
| **Disk** | 8 x 2 TB SSD |

### Production configuration

The infrastructure parameters listed below provide the following protected environment profile: 100 tenants, 1000 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of HDD, 500 GB of SSD. About 10% of the protected virtual machines (100 VMs) can be run simultaneously on the specified below infrastructure.

To deploy the DR infrastructure, you need a minimum of three compute nodes.

### Compute resources and Disaster Recovery storage (hot)

| Server name | Compute node |
|---|---|
| **CPU** | 32 cores |
| **RAM** | 512 GB |
| **Disk** | 8 x 2 TB SSD |

## Network infrastructure requirements

- Logical component requirements

  Public IP: 1 per Hybrid DR cluster (HA Proxy deployed) for the dmzvpn network.

  +n IP addresses for customer environments (optional). These IP addresses are for customers who want a public IP address for their DR workloads.
- Required minimum network infrastructure

  Enterprise-level network adapters, 2x 10 Gbit in the XOR or LACP bonding mode
- Recommended network infrastructure

  Enterprise-level network adapters, 6x 10 Gbit in the XOR or LACP bonding mode
- Recommended network configurations

  For more information about the networks that you must configure before deploying the DR infrastructure, see "Network and firewall requirements" (p. 16).

**Note**
Disaster Recovery Hybrid Cloud supports only the local Acronis Cyber Infrastructure cluster storage space for backups.

# Evaluation configuration for testing purposes

For testing and evaluation purposes, you may use this deployment option.

The infrastructure parameters listed below provide the following protected environment profile: 1 tenant, 50 protected virtual machines with 2 vCPU, 8 GB of RAM, 250 GB of SSD, 500 GB of HDD. About 10% of the protected virtual machines (5 VMs) can be run simultaneously on the specified below infrastructure.

For the evaluation purpose of the DR infrastructure, you may use one node with compute and storage services.

### Compute resources, hot and cold storages

| Server name | Compute node with storage |
|---|---|
| **CPU** | 16 cores |
| **RAM** | 128 GB |
| **Disks** | 3 x 1.86 TB SSD (hot)<br>5 x 6 TB HDD (cold) |

## Compute cluster limitations

Disaster Recovery Hybrid deployment does not support nested virtualization. This means that in Acronis Cyber Infrastructure, nested virtualization for the compute cluster is disabled by default and is not supported for servers in the DR Hybrid cloud.

For more information about nested virtualization, see Enabling nested virtualization.

# Deployment procedure

To deploy the DR infrastructure on Acronis Cyber Infrastructure, do the following:

1. Check all of the prerequisites.
2. Prepare nodes according to the hardware and software requirements and check the network and firewall requirements.
3. Check the "Deployment scenarios" (p. 13) to identify the deployment procedures that you must complete depending on the state of your environment.
4. Complete the deployment procedures.

## Deployment scenarios

The deployment and configuration of Disaster Recovery Hybrid Cloud consists of several procedures that you can perform sequentially.

1. "Installing Acronis Cyber Infrastructure" (p. 15) on a hardware node. This procedure is valid for the management node, and for the rest of the nodes in your environment.

   **Note**
   The recommended number of network adapters for the deployment of Disaster Recovery Hybrid Cloud is six, while the required minimum number of network adapters is two.
   If your node has less than six physical networks adapters, you must add the corresponding number of virtual network interfaces, so that the total number of network adapters and virtual network adapters becomes six. To do that, follow the procedure: "Creating VLAN adapters" (p. 21).
   If your node has six physical networks adapters, skip the procedure.

   a. "Creating VLAN adapters" (p. 21), if necessary. This procedure is valid for the management node, and for the rest of the nodes in your environment.
   b. "Creating infrastructure networks" (p. 22). This procedure is valid for the management node only.
   c. "Assigning network adapters to networks" (p. 22). This procedure is valid for the management node, and for the rest of the nodes in your environment.
   d. "Assigning traffic types to the infrastructure networks" (p. 23). This procedure is valid for the management node only.
2. "Creating a storage cluster" (p. 24)
3. "Creating a compute cluster" (p. 26)
   a. Configuring the compute cluster.
   b. Configuring the networks of the compute cluster, as described in "Creating compute networks" (p. 29)

**Note**
Disaster Recovery Hybrid Cloud does not support compute nodes that are included in a placement. Therefore, you must ensure that the nodes that you add to the compute clusters are not included in a placement. For more information about placements, see Managing placements in the Acronis Cyber Infrastructure Admin Guide.

4. "Adding nodes to the cluster" (p. 25)
5. "Deploying the Disaster Recovery infrastructure" (p. 34)

The procedure from which you should start depends on the current state of your Acronis Cyber Infrastructure. For more information, see the following table.

| Acronis Cyber Infrastructure is installed on a hardware node | Storage cluster is available | Compute cluster is available | What to do |
|---|---|---|---|
| No | No | No | Complete all procedures in the "Deployment procedure" (p. 13) section, starting from the first one. |
| Yes | No | No | Start from "Creating VLAN adapters" (p. 21) and complete all the procedures that follow. |
| Yes | Yes | No | Check if the existing configurations are correct, and then complete all the procedures starting from "Creating a compute cluster" (p. 26). |
| Yes | Yes | Yes | Check if the existing configurations are correct, and |

| Acronis Cyber Infrastructure is installed on a hardware node | Storage cluster is available | Compute cluster is available | What to do |
|---|---|---|---|
| | | | then complete all the procedures starting from "Deploying the Disaster Recovery infrastructure" (p. 34). |

# Prerequisites

The following prerequisites must be met prior to the DR infrastructure deployment:

- On the machine where the DR installer will be run, generate the RSA SSH key pair for the access to Acronis Cyber Infrastructure or use the existing SSH key pair. Upload the public key to the Acronis Cyber Infrastructure node by doing the following:
  1. In the Acronis Cyber Infrastructure admin panel, go to **Settings** > **Security** > **SSH** tab.
  2. Click **Add**.
  3. Copy your public key, paste it in the **Key** field, and click **Add key**.
     For more information about SSH keys, see Adding SSH keys for virtual machines.
- You must have a partner account that has the company administrator role in Acronis Cyber Protect Cloud.
- A stable connection must be available between the machine on which you will run the DR installer and the target data center where you will deploy the DR infrastructure.

# Installing Acronis Cyber Infrastructure

## Installing Acronis Cyber Infrastructure

The installation procedure of Acronis Cyber Infrastructure consists of several stages. Depending on your existing infrastructure, follow all the stages of the procedure, or follow specific stages that are suitable for your existing infrastructure.

**Note**
Disaster Recovery Hybrid Cloud supports Acronis Cyber Infrastructure version 4.7 Update 1 or later.

***To install Acronis Cyber Infrastructure***

1. Follow the installation instructions from the Acronis Cyber Infrastructure Installation Guide.
2. After the installation completes, activate the Acronis Cyber Infrastructure license. For more information, see Installing license keys.

# Configuring networks in Acronis Cyber Infrastructure

## Network and firewall requirements

Before deploying the DR infrastructure, you must create and configure the following networks in Acronis Cyber Infrastructure:

| Name | Description |
|---|---|
| **DMZ+VPN (dmzvpn)** | The public-facing VPN gateways and their proxies. It is used for public access to the servers over the Internet. |
| **Management (drmgmt)** | The network required for general communications between the RunVM components, for outbound access to the Acronis Cyber Protect Cloud, and for the Acronis Cyber Infrastructure OpenStack API access.<br><br>This network is created automatically when you deploy Acronis Cyber Infrastructure with the default name **Public**. If you are using an existing Acronis Cyber Infrastructure, the network might have been renamed. |
| **Content (content)** | The network required for transferring backed-up data by using the NBD protocol between the RunVM components and the Acronis Cyber Infrastructure nodes.<br><br>Note that this network is unprotected and not encrypted users' backup data are passed. |
| **DR Backup storage (drcoldstorage)** | The network providing fast, direct access to the cold storage backup gateway that contains customer backups. |
| **Private** | The network that is used for storage traffic.<br><br>This network is created automatically when you deploy Acronis Cyber Infrastructure. If you are using an existing Acronis Cyber Infrastructure, the network might have been renamed. |
| **acioverlay** | The network that is used for overlay network traffic between virtual machines. |

In the diagram below, you can find these networks.

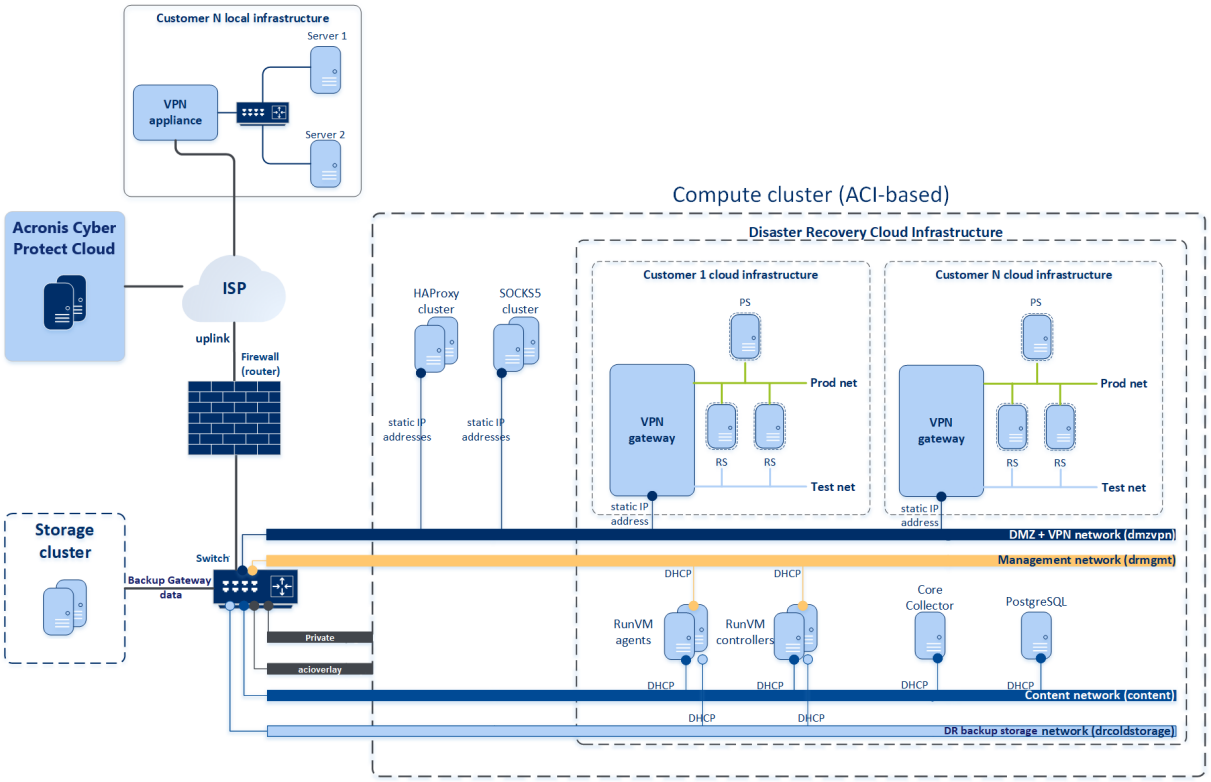


## General requirements to networks

All networks must meet the following requirements:

- The address assignment in all networks must be managed by Acronis Cyber Infrastructure. When creating the network in the UI, you must enable DHCP.
- There are no external DHCP servers in any of the networks.

## Firewall requirements: protocols and ports

You must set the following firewall rules to ensure that Disaster Recovery works properly.

| VLAN | Protocol | TCP ports | Connection direction | Comments |
|---|---|---|---|---|
| dmzvpn | VRRP | n/a | ACI ↔ ACI | VRRP traffic among highly available proxy nodes (to be deployed on different ACI nodes) |
| | TCP | 443 | Internet → ACI | VPN traffic from the local client environments to the VPN servers |
| | TCP/UDP | <any> | Internet → ACI | Any traffic directed to client servers must be |

| VLAN | Protocol | TCP ports | Connection direction | Comments |
|---|---|---|---|---|
| | | | | allowed (Further filtering must be configured on each VPN Server) |
| | <any> | <any> | ACI → Internet | Recovery/primary servers can freely connect to the Internet resources |
| content | TCP | 10809<br>49300-65535 | ACI ↔ ACI | Auxiliary VM → RunVM Controller (NBD server) :10809<br>ACI (internal NBD server) → Auxiliary VM :49300-65535 |
| drmgmt | TCP | 22<br>2650<br>5432<br>8080<br>8888<br>9090-9653<br>OpenStack API ports (see default OpenStack ports)<br>9090 | ACI ↔ ACI | ACI (a service workstation) → RunVM Agent, RunVM Controller, PostgreSQL, Core Collector :22 (ssh access for troubleshooting)<br>RunVM Agent → RunVM Controller :2650 (to manage Controller)<br>RunVM Agent, RunVM Controller → PostgreSQL :5432 (to acquire/release distributed locks)<br>RunVM Agent, RunVM Controller → Core Dump Collector :8080 (to post `core` files for future analysis)<br>ACI Admin panel:8888<br>ACI → RunVM Agent :9090-9653<br>• node_exporter default configuration (see Prometheus |

| VLAN | Protocol | TCP ports | Connection direction | Comments |
|---|---|---|---|---|
| | | | | default port allocations)<br>RunVM Agent → ACI : (see default OpenStack ports)<br>RunVM Agent → ACI :9090 (request from the agent to ACI's Prometheus the used disk space) |
| | UDP | 123 | Internet ↔ ACI | The NTP protocol for NTP clients running in RunVM Agent VMs |
| drcoldstorage | TCP | 44445<br>443 | ACI → ABGW | RunVM Agent, RunVM Controller → ABGW :44445, :443 |
| Deployment/update using Acronis DCO jenkins | TCP | 22 (ssh)<br>OpenStack API ports (see default OpenStack ports) | Acronis Cyber Protect Cloud ACC ↔ ACI (DCO jenkins job placed into ACC) | Access can be limited by jenkins container IP address |

## Networks used by the ACI cluster

| Network name | IP network | DHCP | Description |
|---|---|---|---|
| dmzvpn | vpn—100.64.0.0/10<br>dmz—data center specific (public IP pools)<br>The default gateway is to be assigned on the compute network | None | Customer VPN traffic from customer premises to their private cloud environment.<br>There are 2 IP networks configured over a single vlan.<br>The DR service assigns IP addresses to VPN servers. The range is configured during the DR infrastructure deployment as a subset of 100.64.0.0/10. |
| content | x.x.x.x/16 | yes, managed by | Data traffic: read/write of virtual disk data |

| | | | |
|---|---|---|---|
| | Not routable outside of the network<br>No default gateway | ACI | between RunVM Agent/Controller and ACI internal processes (possible location on different ACI nodes). The communication is done within the same VLAN. |
| drmgmt | y.y.y.y/24<br>The default gateway is to be assigned on the ACI host interfaces | 1. Some IP addresses in drmgmt network (for ACI hosts, etc) are assigned statically.<br>2. The rest of the range is managed by ACI. | Connection from RunVM Agent/Controller to Acronis Cyber Protect Cloud (ACPC) component.<br>Connection among RunVM Agent, RunVM Controller, and PostgreSQL VMs.<br>Connection from RunVM components to the Internet.<br>Public network in terms of ACI. |
| drcoldstorage | z.z.z.z/16<br>No default gateway | yes, managed by ACI | Data traffic:<br>RunVM Agents/Controllers read archives located on Backup (cold) storage.<br>RunVM Agents (backupers) write new backups to Backup (cold) storage.<br>Access to the Backup (cold) storage is done using public DNS names (IPs of Backup storage).<br>Maximum throughput and minimum latency over this RunVM ↔ Backup storage communication is critical for the main DR service operation. |
| Private | c.c.c.c/24<br>Not routable outside of the network<br>No default gateway | None | Inter ACI cluster communication: storage<br>Storage Internal network in terms of ACI, see Network requirements and recommendations. |
| acioverlay | d.d.d.d/24<br>Not routable outside of the network<br>No default gateway | None | Stretching the internal (virtual) networks across the ACI cluster (encapsulation of private vxlan traffic)<br>Overlay Networking network in terms of ACI, see Network requirements and recommendations. |

### ACI cluster network configuration

#### ACI infrastructure traffic type configuration

| | | acioverlay | Private | content | dmzvpn | drcoldstorage | drmgmt |
|---|---|---|---|---|---|---|---|
| **Exclusive traffic types** | Storage | | yes | | | | |
| | Internal management | | yes | | | | |
| | OSTOR private | | yes | | | | |
| | ABGW private | | yes | | | | |
| | VM private | yes | | | | | |
| | Compute API | | | | | | yes |
| | VM backups | | | yes | | | |
| **Regular traffic types** | S3 Public | | | | | | yes |
| | NFS | | | | | | yes |
| | Admin panel | | | | | | yes |
| | SSH | | | | | | yes |
| | VM public | yes | | yes | yes | yes | yes |
| | Self-service panel | | | | | | yes |
| **Custom traffic types** | Prometheus | | | | | | yes |

## Creating VLAN adapters

As a minimum configuration requirement, after deploying Acronis Cyber Infrastructure, you already have two network adapters assigned to the drmgmt (Public) and Private networks. You must now create the VLAN adapters for the acioverlay, content, drcoldstorage, and dmzvpn networks.

1. In the Acronis Cyber Infrastructure admin panel, go to **Infrastructure** > **Nodes**, and click the node.
2. Click **Network interfaces**.

   **Note**
   You need network adapters for six networks. If the number of physical network adapters on the node is less than 6, add the missing number of VLANs by repeating the steps below.

3. Click **Create**.
4. In the **Create network interface** wizard, select the **VLAN** interface type.
5. Click the network interface from which you want to create the VLAN, and click **Next**.
6. Type the VLAN ID of the adapter.
7. In **Select a network to assign to the interface**, select **Not assigned**.
8. In **Specify the network parameters**, select **Manually**.
9. In the **IP address** field, click **Add**, and type the IP address of the VLAN using the CIDR notation.
10. Configure the Gateway and the MTU, or leave the fields empty.
11. Click **Create**.

For more details, see Acronis Cyber Infrastructure Installation Guide.

## Creating infrastructure networks

1. In the Acronis Cyber Infrastructure admin panel, go to **Infrastructure**>**Network**.
   Note that the drmgmt (Public) and Private networks are already present, as they were created during the installation of the compute cluster. You must create four more networks with the following names:
   - acioverlay
   - content
   - dmzvpn
   - drcoldstorage
2. For each of the listed networks, repeat the steps:
   a. Click **Create network** and type the network name.
   b. Click **Create**.
3. Create a new traffic type: Prometheus with the 9090 port.
   a. Click **Create traffic type**.
   b. In the **Name**, type Prometheus.
   c. In the **Port**, type 9090.
   d. Click **Create**.

## Assigning network adapters to networks

***To assign network adapters to a network***

1. Go to **Infrastructure**>**Nodes**, and click the node.
2. Click **Network interfaces**.
3. For each physical adapter or VLAN that you created, repeat the steps below.
   a. Click the VLAN that you want to assign to a network (content, dmzvpn, drmgmt, or drcoldstorage), and click **Edit**.
   b. In the **Edit network interface** wizard, select a network to assign to the network interface.
   c. Click **Save**.

## Assigning traffic types to the infrastructure networks

After assigning network adapters to the networks, you must assign the traffic types, as shown in the table.

The following table shows the correct ACI Infrastructure traffic type configuration.

| | | acioverlay | Private | content | dmzvpn | drcoldstorage | drmgmt |
|---|---|---|---|---|---|---|---|
| **Exclusive traffic types** | Storage | | yes | | | | |
| | Internal management | | yes | | | | |
| | OSTOR private | | yes | | | | |
| | Backup (ABGW) private | | yes | | | | |
| | VM private | yes | | | | | |
| | Compute API | | | | | | yes |
| | VM backups | | | yes | | | |
| **Regular traffic types** | S3 Public | | | | | | yes |
| | NFS | | | | | | yes |
| | Admin panel | | | | | | yes |
| | SSH | | | | | | yes |
| | VM public | yes | | yes | yes | yes | yes |

| | | acioverlay | Private | content | dmzvpn | drcoldstorage | drmgmt |
|---|---|---|---|---|---|---|---|
| | Self-service panel | | | | | | yes |
| **Custom traffic types** | Prometheus | | | | | | yes |

***To assign an exclusive traffic type to a network***

1. Go to **Infrastructure**>**Network**
2. In the **Exclusive traffic types** section, in the **Assign to network** drop-down list, click the traffic type that you want to assign.
3. Select the radio button in the column of the network to which you want to assign the traffic type.
4. Click **Save**.

***To assign a regular traffic type to a network***

1. In the **Regular traffic types** section, click **Assign to networks**.
2. Select the checkbox in the column of each network to which you want to assign the corresponding traffic type.
3. Click **Save**.

***To assign a custom traffic type to a network***

1. In the **Custom traffic types** section, click **Assign to networks**.
2. Select the checkbox in the column of each network to which you want to assign the corresponding traffic type.
3. Click **Save**.

## Creating a storage cluster

***To create a storage cluster***

1. In the admin panel, go to **Infrastructure** > **Nodes** and click **Create storage cluster**.

   **Note**
   The storage cluster must have at least three nodes. All nodes which will be configured with high availability must have disks with **Tier 0** assigned.

2. Specify the cluster name.
3. Click **Create**.

For more information about creating the storage cluster, see the Acronis Cyber Infrastructure Administrator Guide.

## Adding nodes to the cluster

You may expand your existing cluster by adding a node.

***To add a node to the cluster***

1. In the admin panel, go to **Infrastructure** > **Nodes**.
2. Click **Connect node**.
3. You will see the management panel IP address (also known as management node IP address) and token. Copy them.

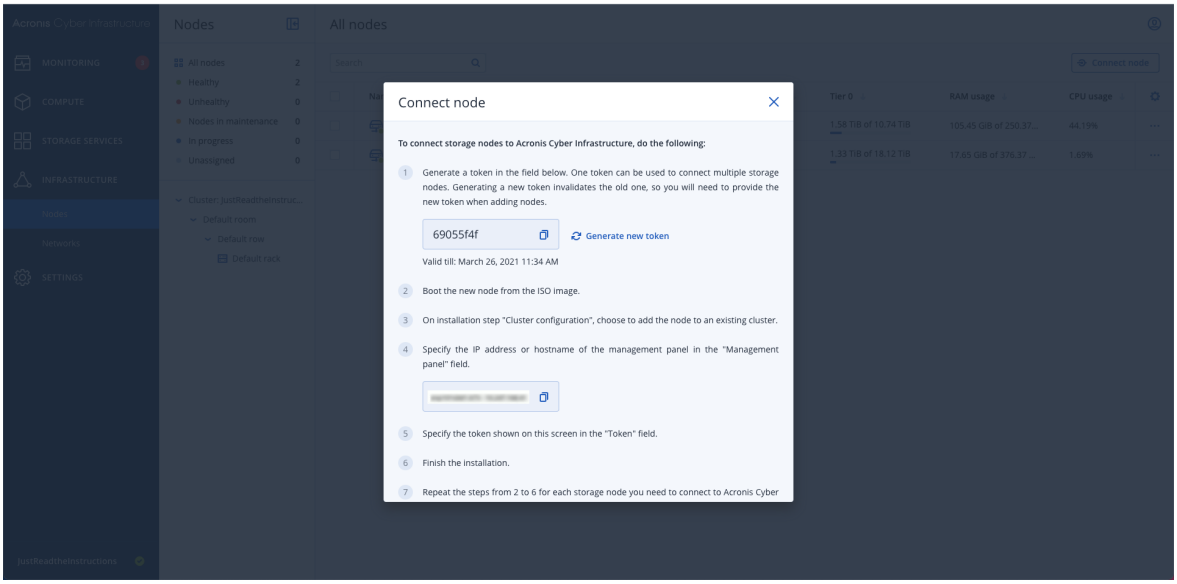


4. Install Acronis Cyber Infrastructure on the additional node.

> **Note**
> To install Acronis Cyber Infrastructure on the node, follow the "Installing Acronis Cyber Infrastructure" (p. 15) procedure. In the **Cluster configuration** step, select **No, add it to existing cluster**.

5. Enter the management node IP address and token that you copied in step 3.
6. In the ACI management portal, go to **Infrastructure** > **Nodes**.
7. Select the node, and click **Join to cluster**.

The rest of the steps are similar to the instructions for configuring the first node.

For more information about adding nodes to the cluster, see the Acronis Cyber Infrastructure Administrator Guide.

# Configuring high availability

To make your infrastructure more resilient and redundant, you can create a high availability configuration of three nodes. To configure high availability, follow the procedure described in Enabling high availability.

## Creating a compute cluster

We recommend that you complete this procedure after you complete the installation of Acronis Cyber Infrastructure on all nodes and add the nodes to the infrastructure cluster. This will simplify the configuration.

***To create a compute cluster***

1. In the admin panel, go to **Compute**.
2. Click **Create compute cluster**.

   You can adjust network roles, if needed, by clicking the gear icon. If the rules are assigned correctly, the network state will be **Configured**. When ready, click **Next**.

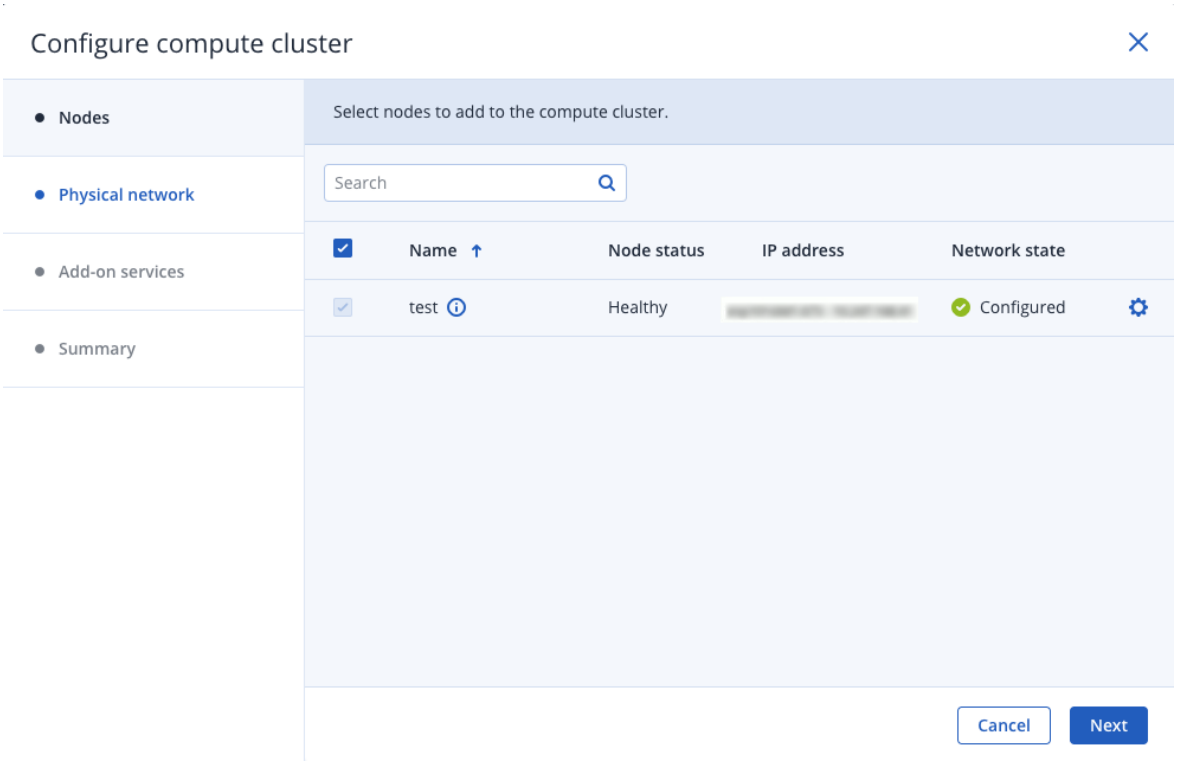


3. Select the drmgmt network and specify the gateway. When ready, click **Next**.

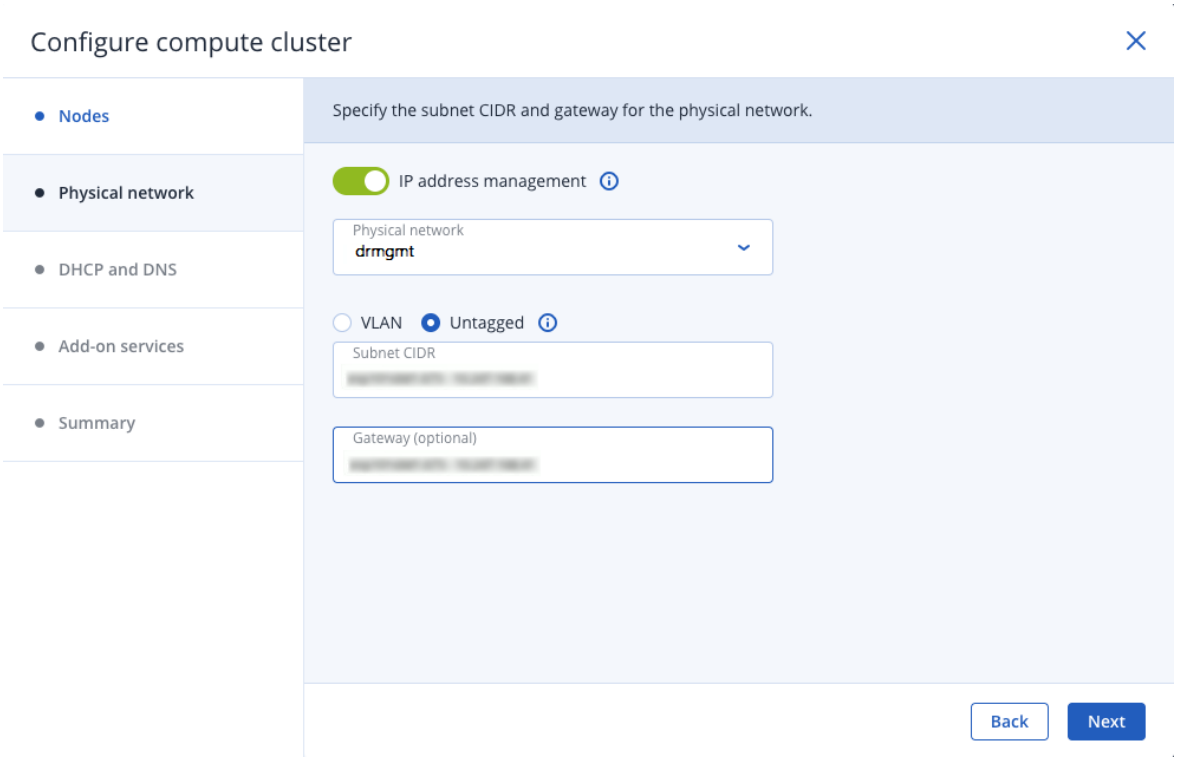


4. Enable DHCP, configure the DNS settings, and allocate pools for the public virtual network. When ready, click **Next**.

**Note**
Both DNS servers must be preliminary set for the dmzvpn and drmgmt networks in **Compute** > **Network** > **IP Management**. Do not set a default gateway for the dmzvpn and drmgmt networks.

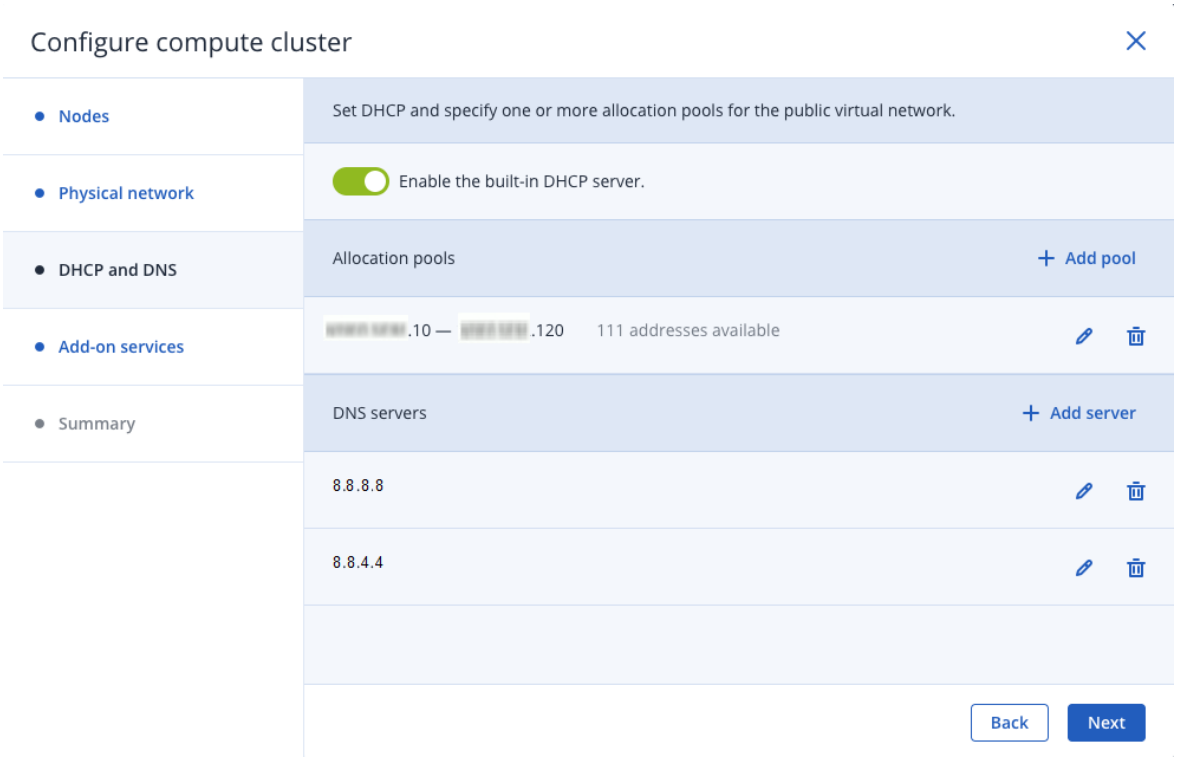


5. [Optional] To install an additional service on your compute cluster, enable the switch of the service, and click **Next**.

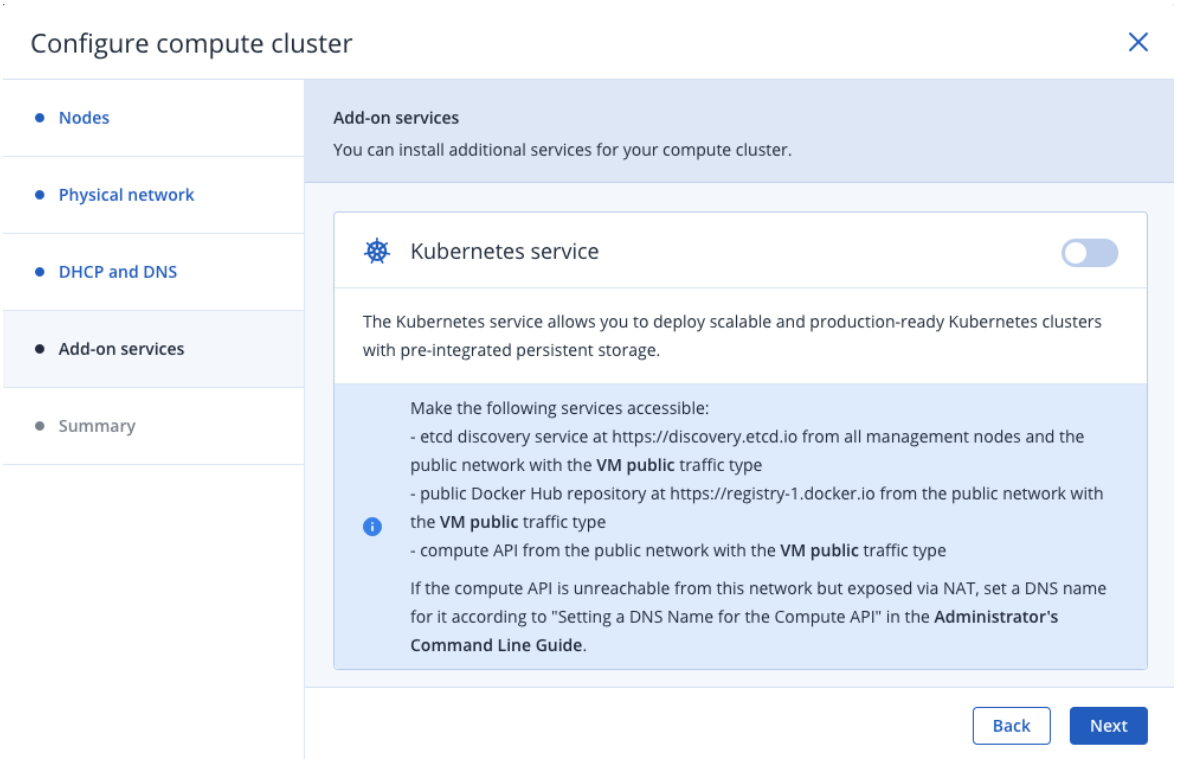


6. Review the settings and click **Create cluster**.

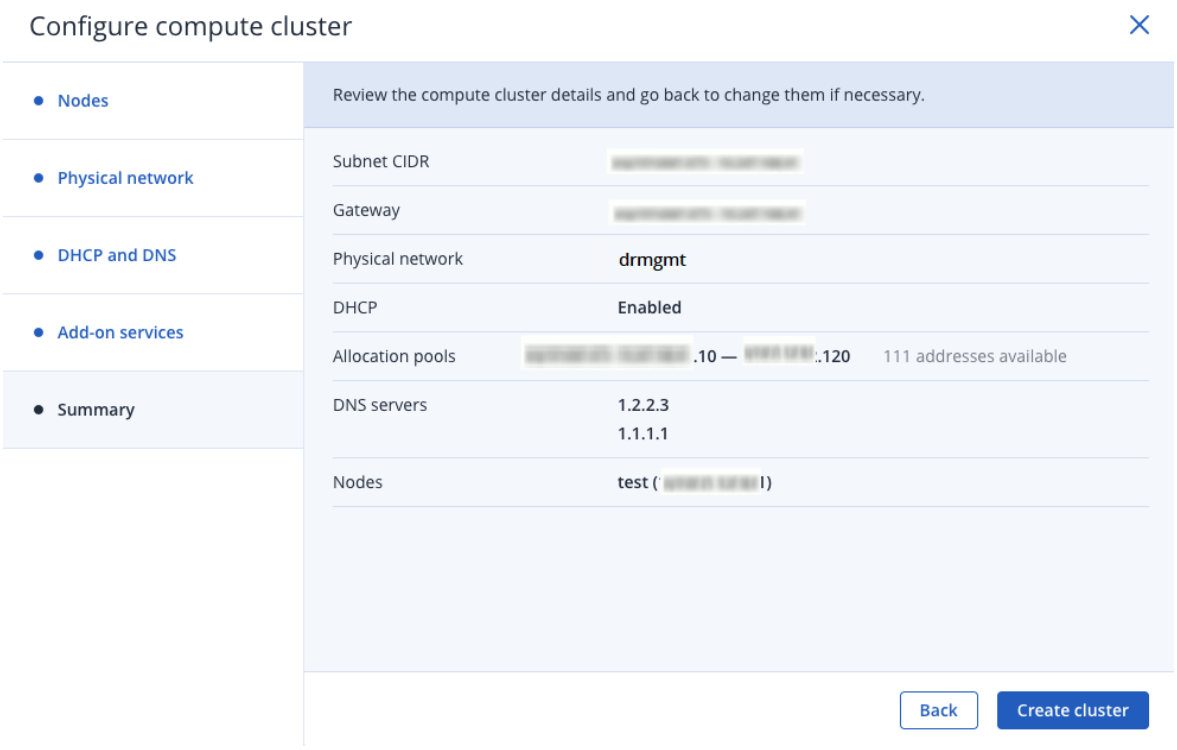


For more information about the compute cluster, see Creating the compute cluster.

To add nodes to the compute cluster, follow the procedure in Adding nodes to the compute cluster.

## Creating compute networks

You must create virtual networks for the Public, Private, drmgmt, dmzvpn, content, and drcoldstorage networks, and bind them to the existing networks on the Acronis Cyber Infrastructure management node. Thus, you give the compute cluster access to the networks that are connected to the Acronis Cyber Infrastructure cluster.

- The drmgmt, dmzvpn, content, and drcoldstorage networks must have the Physical type.
- Only the Private network must have the Virtual type.
- The gateway must be specified only for the drmgmt and dmzvpn networks.
- Note that the virtual dmzvpn network should have a different subnet CIDR than the physical dmzvpn network. This is needed for a VPN gateway to use internal addresses rather than real public IP addresses that will be used for customers' recovery and primary servers.

**Important**

When deploying the dmzvpn network (for example, 100.64.0.0/10), you must specify the gateway so that it takes the first address in this virtual network (for example, 100.64.0.1). Otherwise, proxying will not work.

## ACI compute networking configuration

| Network | IP address management | Type | Subnet range | Gateway | DNS | DHCP | Project |
|---|---|---|---|---|---|---|---|
| content | Enabled | Physical | /16 | - | - | Enabled | admin |
| dmzvpn | Enabled | Physical | /10 (or a subnet of /10) | yes | yes | Enabled | admin |
| drcoldstorage | Enabled | Physical | /16 | - | - | Enabled | admin |
| drmgmt | Enabled | Physical | /16 | yes | yes | Enabled | admin |

To learn more about configuring direct routing for a cold storage, refer to Appendix E. Advanced network configuration.

***To create the virtual networks***

1. In the admin panel, go to **Compute**>**Networks**.
2. Click **Create network**.
3. Specify the network type, network name, select the corresponding network from the drop-down list. The CIDR is filled in automatically.

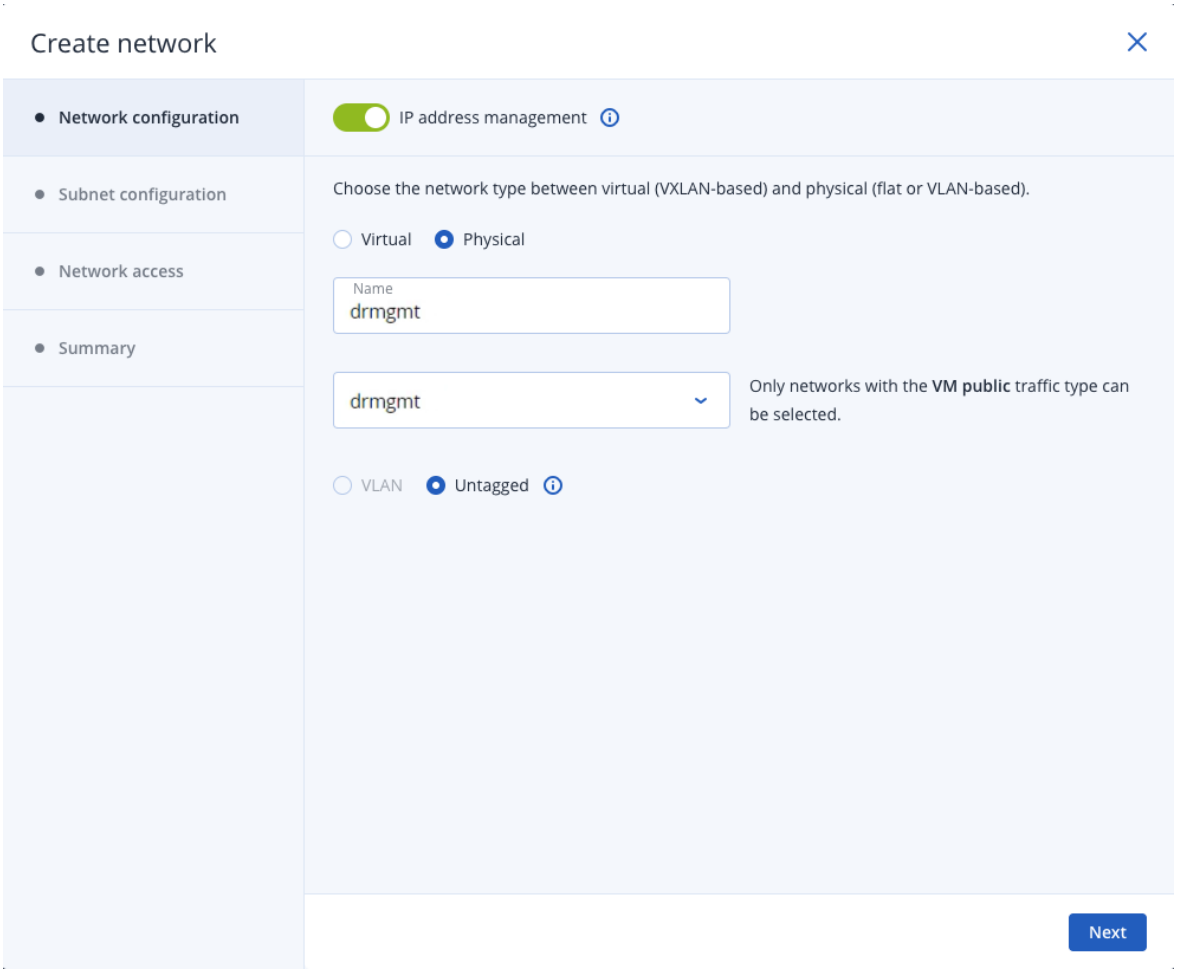


4. Enable DHCP, allocate the IP pool, and specify the DNS servers.

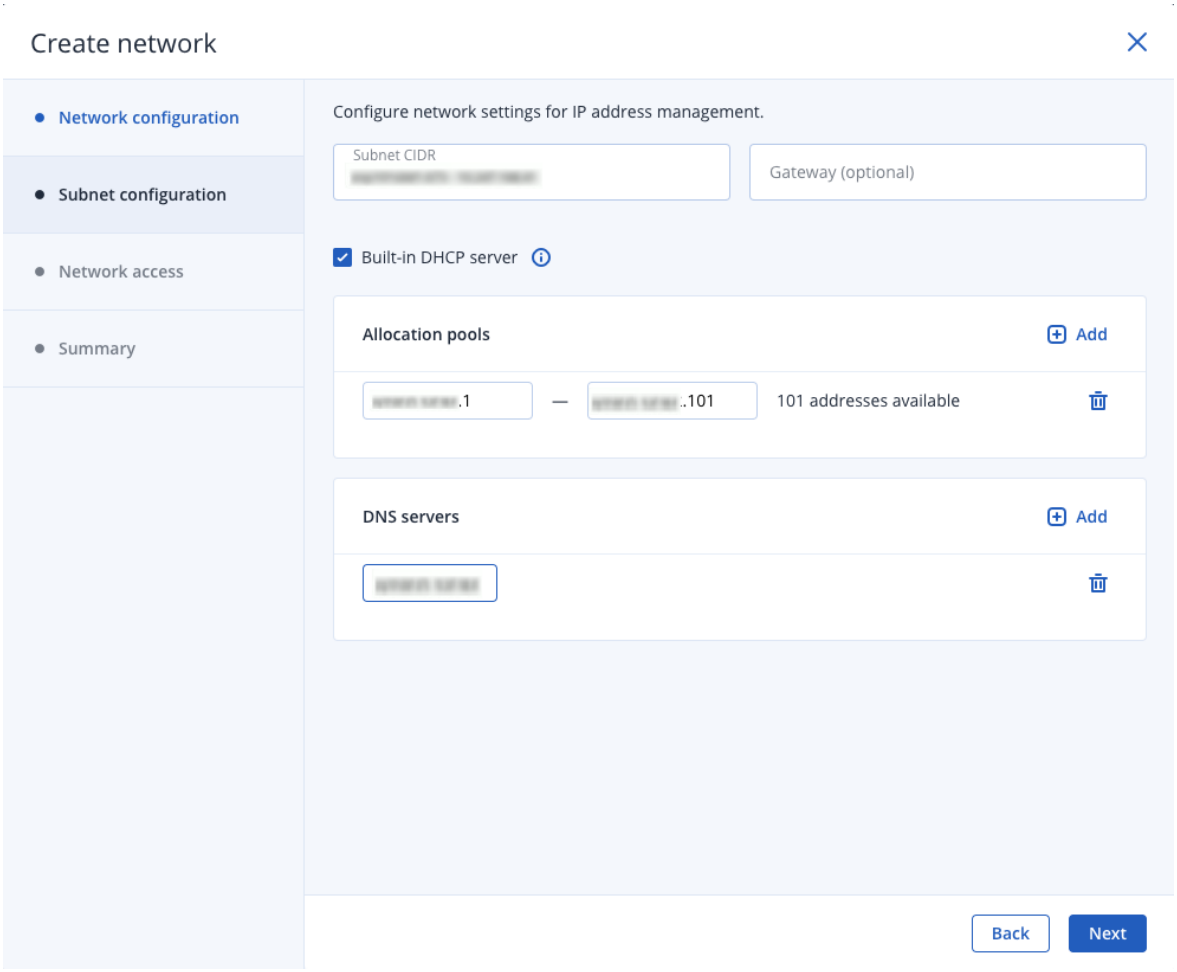


5. To configure the network access rights to the projects, select the checkbox of a project, and in the **Access options** field, select the access rights.
   All networks that you created should have **Full** access rights to the **Default** project.

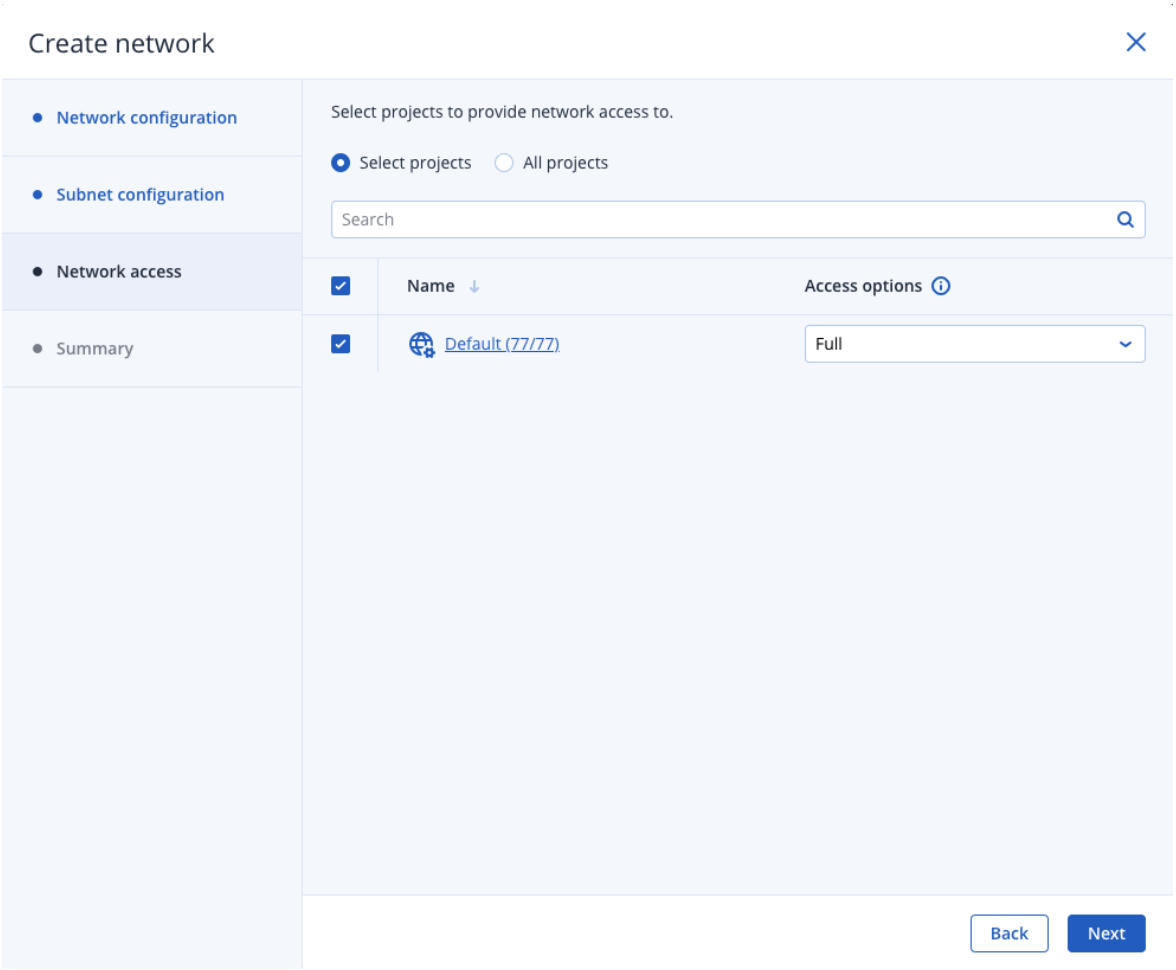


6. Review all of the settings and click **Create virtual network**.

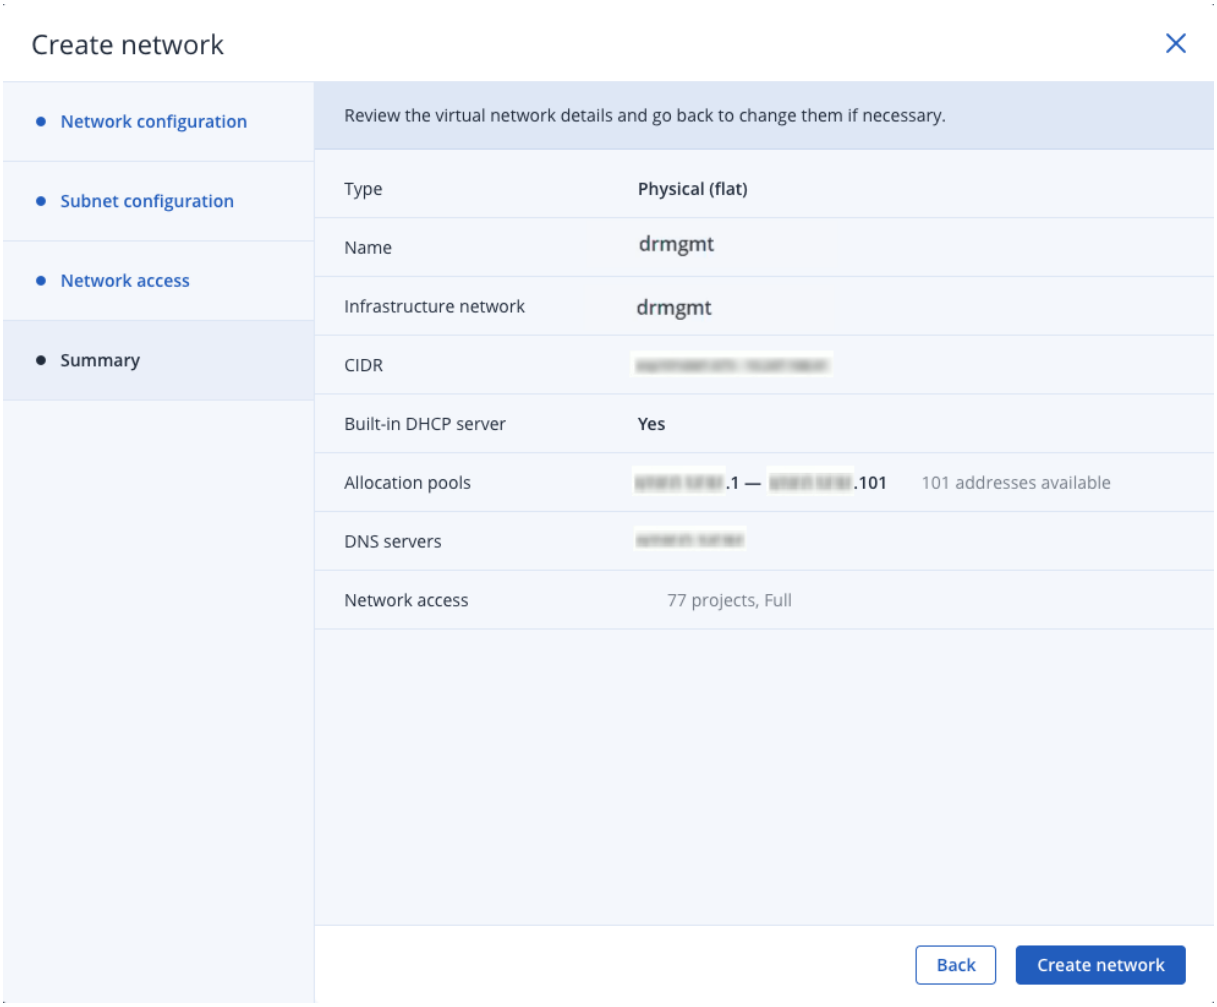


As a result, you will create virtual networks for the Private, drmgmt, dmzvpn, content, and drcoldstorage networks.

## Managing the compute storage

The compute storage is installed automatically on **Tier 0**. The DR Hybrid infrastructure is installed automatically on **Tier 0**.

You can manage the compute storage after the deployment of Acronis Cyber Infrastructure is complete.

For more information about the managing the compute storage, and creating volumes, see Managing the compute storage.

# Deploying the Disaster Recovery infrastructure

The deployment procedure consists of the following steps:

1. Download the Disaster Recovery installation archive to the machine from which you will run the installer. You can find the download links for the released versions of the DR installer in KB

article 69638.

You can run the DR installer on a Windows, Linux, or macOS machine. Ensure that the machine has a stable connection to the target data center where the DR infrastructure will be deployed.

2. Prepare the configuration file with parameters of deployment.
3. Run the DR installer.
4. Prepare templates for primary servers and upload them to Acronis Cyber Infrastructure.

## Step 1. Download the Disaster Recovery installation archive

The Disaster Recovery installation archive includes:

- The DR installers for Windows, macOS, and Linux
- The configuration file `config.yml`
- The templates for virtual machines located in the `images` folder

After you downloaded the archive, unpack it on your machine.

## Step 2. Prepare the configuration file

The configuration file defines the parameters of the DR infrastructure to be deployed. The default configuration file can be viewed in Appendix A.

1. On your machine in the unpacked archive folder, find the configuration file `config.yml`.
2. Define all of the mandatory parameters in it, according to your data center configuration, and specify the paths to the SSH keys to access Acronis Cyber Infrastructure.

   The file is split into several sections where you need to fill in the empty required parameters for deploying the DR infrastructure.

   - Parameters of the Acronis Cyber Infrastructure cluster
   - Parameters of Acronis Cyber Protect Cloud
   - Deployment parameters such as the project, key pair, flavors, and other
   - Deployment settings related to Acronis Cyber Protect Cloud (infrastructure registration, location)
   - Network settings for Acronis Disaster Recovery on Acronis Cyber Infrastructure
   - VPN settings
   - Acronis Disaster Recovery RunVM platform settings
   - PostgreSQL settings
   - Core-dump-server settings

## Step 3. Run the DR installer

The installer deploys the DR infrastructure components on Acronis Cyber Infrastructure and registers it as an IaaS in the proper partner account in Acronis Cyber Protect Cloud. For more details about all possible commands of the `dr-installer` tool, refer to Appendix B.

***To deploy the DR infrastructure***

1. Specify the public address of the SOCKS proxy server in `{path_to_unpacked_dr-installer}/config.yml` in the **HAProxy settings** section.
2. To perform a dry run and validate the correctness of the deployment settings, run the `validate` command

   For Windows:

   ```
   dr-installer-windows.exe --config {your-config-file-path}/config.yml validate
   ```

   For Linux:

   ```
   ./dr-installer-linux --config {your-config-file-path}/config.yml validate
   ```

   For macOS:

   ```
   ./dr-installer-macos --config {your-config-file-path}/config.yml validate
   ```

3. Install proxy servers (HAProxy, Dante) by running the command `vpn-proxy` on the machine. Depending on the operating system of the machine where you run the `dr-installer`, use one of the following commands:

   For Windows:

   ```
   dr-installer-windows.exe --config {your-config-file-path}/config.yml vpn-proxy
   install
   ```

   For Linux:

   ```
   ./dr-installer-linux --config {your-config-file-path}/config.yml vpn-proxy install
   ```

   For macOS:

   ```
   ./dr-installer-macos --config {your-config-file-path}/config.yml vpn-proxy install
   ```

4. To install the Disaster Recovery infrastructure, run the `install` command.

   For Windows:

   ```
   dr-installer-windows.exe --config {your-config-file-path}/config.yml install
   ```

   For Linux:

   ```
   ./dr-installer-linux --config {your-config-file-path}/config.yml install
   ```

   For macOS:

   ```
   ./dr-installer-macos --config {your-config-file-path}/config.yml install
   ```

Now you have the DR infrastructure deployed on top of Acronis Cyber Infrastructure.

As a result, a new project and security group in Acronis Cyber Infrastructure are created. The templates for the DR components are loaded to Acronis Cyber Infrastructure, then the respective

virtual machines are launched from these templates. To find these virtual machines (the DR components), log in to the Acronis Cyber Infrastructure admin panel, open the list of virtual machines, and filter them by the project name (the name that you defined in the configuration file).

**Important**
All the changes of a VM flavor, number of DR components, networks must be done only via `dr-installer`.

**Note**
If the installation of the DR infrastructure fails at any step, you must perform the deletion procedure before trying to install the DR infrastructure again. For more information about the deletion procedure, see "Deletion procedure" (p. 43).

## Step 4. Prepare templates for primary servers and upload them to Acronis Cyber Infrastructure

To learn how to create and upload templates for primary servers, refer to Creating templates for primary servers and Uploading templates for primary servers.

# Management

## Creating templates for primary servers

Templates for primary servers are qcow images of servers.

To create primary server images for Acronis Cyber Infrastructure (ACI):

1. In the ACI admin panel, upload the required ISO image of operating system. For more information, see Uploading images for virtual machines.
2. Create a virtual machine and install the desired operating system into this virtual machine from the ISO image. For more information, see Creating virtual machines.

   **Note**
   When you specify the **Flavor** of the virtual machine, select one of the options that you configured in the config.yml file and created using the `dr-installer` utility.

3. Install the ACI guest tools on the virtual machine. For more information, see Installing guest tools.
4. [For Linux] Install the `cloud-init` package on the virtual machine.
5. Configure the operating system in a special way (Sysprep is needed in case of Windows).
6. Shut down the virtual machine.
7. In the ACI admin panel, convert the volume into an image.

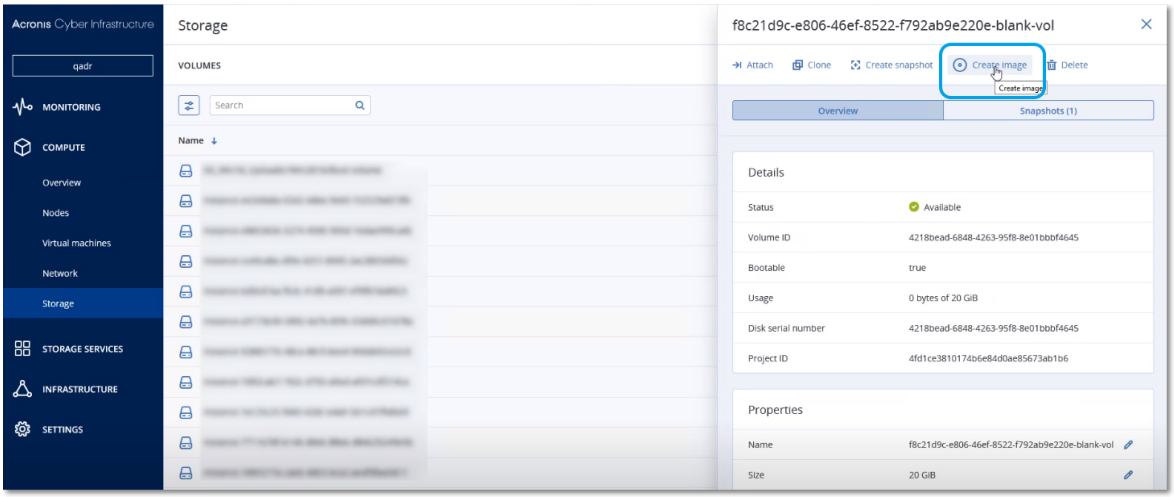

8. Download the qcow image file.

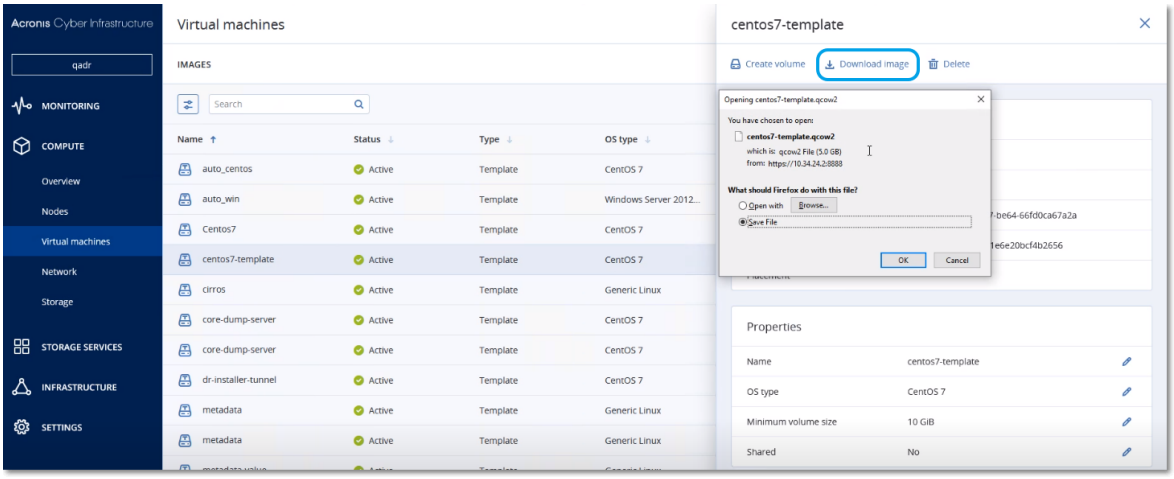


As a result, you can now upload this image file as a template for a primary server by using the `dr-installer` utility.

# Uploading templates for primary servers

The templates are used for creating primary servers. By default, there are no ready templates, and you need to create them first. For more information, see "Creating templates for primary servers" (p. 38).

Before you upload the templates to Acronis Cyber Infrastructure, you must create `image-batch.json` file an specify the following parameters:

- `location` - the full path to the image.
- `name` - a free text.
- `os-type` and `os-distro` - the values are fixed for each image of an operating system that is supported. Check their values in the table below.
- `sha256` - a hash of this image.

  You can find the `sha256` by running the following command:

  - For Windows:

    ```
    sha256sum .{path-to-image}\image.qcow2
    ```

  - For macOS/Linux:

    ```
    sha256sum .{path-to-image}/image.qcow2
    ```

The following example consists of two images and is run on a Linux operating system.

```
{
    "images": [
        {
            "location": "/root/templates/auto-cen7.qcow2",
```

```
            "name": "centos7",
            "os-type": "linux",
            "os-distro": "centos7",
            "sha256": "2fccd52a5031a95f232e756bf69e8ddee82eb5f648a0369c93ecf1495c6cfe8e"
        },
        {
            "location": "/root/templates/auto-win2012r2.qcow2",
            "name": "win12",
            "os-type": "windows",
            "os-distro": "win2k12",
            "sha256": "4face143f9510aea4ce6c5104efd77087c57a192efe6238358056609bde77215"
        }
    ]
}
```

To upload the templates to Acronis Cyber Infrastructure, run the following command:

- For Windows:

  ```
  dr-installer-windows.exe --config .{your-config-file-path}\config.yml image upload-batch {path-to-json-file}\image-batch.json
  ```

  For macOS/Linux:

  ```
  ./dr-installer-linux --config .{your-config-file-path}/config.yml image upload-batch {path-to-json-file}/image-batch.json
  ```

  For macOS/Linux:

  ```
  ./dr-installer-macos --config .{your-config-file-path}/config.yml image upload-batch {path-to-json-file}/image-batch.json
  ```

Currently, images for the following operating systems are supported:

| Name | os-distro | os-type |
|---|---|---|
| Generic Linux | linux | linux |
| CentOS 7 | centos7 | linux |
| CentOS 6 | centos6 | linux |
| Red Hat Enterprise Linux 8 | rhel8 | linux |
| Red Hat Enterprise Linux 7 | rhel7 | linux |
| Ubuntu 18.04 | ubuntu18.04 | linux |
| Ubuntu 16.04 | ubuntu16.04 | linux |
| Debian 9 | debian9 | linux |

| | | |
|---|---|---|
| Generic Windows | windows | windows |
| Windows Server 2019 | win2k19 | windows |
| Windows Server 2016 | win2k16 | windows |
| Windows Server 2012 R2 | win2k12r2 | windows |
| Windows Server 2012 | win2k12 | windows |
| Windows Server 2008 R2 | win2k8r2 | windows |
| Windows Server 2008 | win2k8 | windows |
| Windows 10 | win10 | windows |
| Windows 8.1 | win8.1 | windows |
| Windows 7 | win7 | windows |

# Maintenance

## Update procedure

When a Disaster Recovery update is available, Acronis Cyber Infrastructure might generate an alert.

For more information about alerts and configuring the alert settings, see Acronis Cyber Infrastructure Administrator Guide.

**Important**
We recommend that you install Disaster Recovery updates as soon as possible to prevent potential risks, such as a data loss.

The update procedure implies that the virtual servers (DR infrastructure components) are redeployed according to the settings specified in the config file. You can increase or decrease the number of RunVM agents or change flavors of virtual servers. Note that there is downtime during the update procedure.

**Note**
If several updates are available, before installing the latest update, you must sequentially install all previous updates following the procedure below.

***To update the existing DR infrastructure components***

1. Download the DR installer archive on your machine. You can find the download links for the released versions of the DR installer in KB article 69638.
2. Unpack the DR installer archive on your machine.
3. Copy your existing config file to the folder where you unpacked the new DR installer.

   ```
   cp {your-config-file-path} {path-to-extracted-dr-installer-archive}/
   ```

4. Compare your existing config file with the default one that you just downloaded. If there are any new parameters in the downloaded config file, add them to your existing config file.
5. Copy the SSH keys that were created prior to the initial DR infrastructure deployment. For more information about the SSH keys, see "Prerequisites" (p. 15).
6. Go to the folder where you unpacked the new DR installer:

   ```
   cd {path-to-extracted-dr-installer-archive}
   ```

7. Run `dr-installer` with the `update` option. Ensure that the config file is from the same project as it might be different files in case when the installer is used for DR infrastructure installation in different locations.

   For Windows:

   ```
   .\dr-installer-windows.exe --config {your-config-file-path}\config.yml update
   ```

For Linux:

```
./dr-installer-linux --config {your-config-file-path}/config.yml update
```

For macOS:

```
./dr-installer-macos --config {your-config-file-path}/config.yml update
```

As a result, your DR infrastructure components will be updated to the actual version.

8. Check your updated infrastructure.
   - Perform test failover of a test tenant.
   - Perform production failover.
   - Check the statuses of all the servers.

# Scaling procedure

When you add new nodes to the compute cluster, use the scaling procedure to increase the number of RunVM agents.

The command does not change the existing agents, therefore you cannot change the flavors as it requires recreation of a virtual machine. There is no downtime during the scaling procedure.

If you need to change the existing RunVM agents, use the `update` command.

***To scale out the existing DR infrastructure***

1. In the **config.yml** file, increase the value of the `count` parameter.
2. Run `dr-installer` with the `scale` option and specify the already existing config file:

   For Windows:

   ```
   dr-installer-windows.exe --config {your-config-file-path}/config.yml scale
   ```

   For Linux:

   ```
   ./dr-installer-linux --config {your-config-file-path}/config.yml scale
   ```

   For macOS:

   ```
   ./dr-installer--macos --config {your-config-file-path}/config.yml scale
   ```

# Deletion procedure

Before deleting the DR infrastructure, ensure that

- There is no active usage of the DR offering items.
- The DR services are disabled for all Acronis Cyber Protect Cloud customers.

***To delete the existing DR infrastructure***

1. Run the `dr-installer` command with the `remove` option to uninstall the DR infrastructure.

   For Windows:

   ```
   dr-installer-windows.exe --config {your-config-file-path}/config.yml remove
   ```

   For Linux:

   ```
   ./dr-installer-linux --config {your-config-file-path}/config.yml remove
   ```

   For macOS:

   ```
   ./dr-installer-macos --config {your-config-file-path}/config.yml remove
   ```

2. Run the `dr-installer` command with the `remove` option to uninstall the VPN proxy.

   For Windows:

   ```
   dr-installer-windows.exe --config {your-config-file-path}/config.yml vpn-proxy remove
   ```

   For Linux:

   ```
   ./dr-installer-linux --config {your-config-file-path}/config.yml vpn-proxy remove
   ```

   For macOS:

   ```
   ./dr-installer-macos --config {your-config-file-path}/config.yml vpn-proxy remove
   ```

As a result, the DR infrastructure is unregistered from Acronis Cyber Protect Cloud and the DR infrastructure components are removed from Acronis Cyber Infrastructure.

# Troubleshooting

All logs are written in the log file `dr-installer.log`. The log file is always located in the same folder from which the DR installer is launched.

The DR infrastructure installation can be divided in several stages. Depending on the stage when an issue happens, the investigation actions are different:

1. First, the installer checks that all the required parameters are specified correctly in `config.yml`. If there are any wrong or missing parameters in the configuration file, the installer will stop the installation and write the issue to the log file.
2. The next step is launching virtual machines for the DR infrastructure components in Acronis Cyber Infrastructure. If you have any issues at this step, you may either refer to `dr-installer.log` or investigate the issue on the Acronis Cyber Infrastructure side.
3. The last point to investigate the issue is Acronis Cyber Protect Cloud.

# How to collect logs for further investigation

## RunVM service logs

**agent-runner**

To collect the logs for a Agent Runner (RunVM Runner) instance, do the following:

1. Log in to a virtual machine with the Agent Runner (RunVM Runner) role by using the 'centos' login and private SSH key that was used during Acronis Cyber Infrastructure installation.

   **Note**
   To find the IP address of the virtual machine with the Agent Runner, in Acronis Cyber Infrastructure management portal, go to **Compute** > **Virtual machines**, and check the value of the **IP address** column.

2. Get root privileges by running the command:

   ```
   sudo -i
   ```

3. Navigate to `/var/log/runvm-agent/`. You can find the `runvm-agent.log` file there.
4. Copy the file to your logs directory (desktop machine, network share, etc.)

**agent-backuper**

1. Log in to a virtual machine with the Agent Backuper (RunVM Backup) role by using the 'centos' login and private SSH key that was used during Acronis Cyber Infrastructure installation.

   **Note**
   To find the IP address of the virtual machine with the Agent Backuper, in Acronis Cyber Infrastructure management portal, go to **Compute** > **Virtual machines**, and check the value of the **IP address** column.

2. Get root privileges by running the command:

   ```
   sudo -i
   ```

3. Navigate to `/var/log/runvm-agent/`. You can find the `runvm-agent.log` file there.
4. Copy the file to your logs directory (desktop machine, network share, etc.)

**agent-gateway**

1. Log in to a virtual machine with the Agent Gateway (RunVM Gateway) role by using the 'centos' login and private SSH key that was used during Acronis Cyber Infrastructure installation.

**Note**
To find the IP address of the virtual machine with the Agent Gateway, in Acronis Cyber Infrastructure management portal, go to **Compute** > **Virtual machines**, and check the value of the **IP address** column.

2. Get root privileges by running the command:

```
sudo -i
```

3. Navigate to `/var/log/runvm-agent/`. You can find the `runvm-agent.log` file there.
4. Copy the file to your logs directory (desktop machine, network share, etc.)

## Getting technical support for ACI

You can download a problem report for the Acronis Cyber Infrastructure cluster, and send it to the technical support team. For more information, see Getting technical support.

You can find more information about the ACI logs in Viewing cluster logs.

# Monitoring the Disaster Recovery Hybrid infrastructure

After you install the Disaster Recovery infrastructure, you can monitor the status of the virtual infrastructure, virtual machines and services. This might help you resolve issues proactively. If critical Disaster Recovery operations, such as failover, failback, power on, and power off, are not possible because of a problem with a component that is running on the Disaster Recovery Hybrid infrastructure, you will receive an alert in Acronis Cyber Infrastructure.

If the Disaster Recovery Hybrid infrastructure is unavailable, or the VPN tunnels are down, you will receive an email notification.

## Configuring alerts in Acronis Cyber Infrastructure

Alerts in Acronis Cyber Infrastructure are raised in the following cases.

- Hybrid Disaster Recovery agent (runvm-agent) is unavailable.
- Hybrid Disaster Recovery database is unavailable.
- Hybrid Disaster Recovery agent (runvm-agent) cannot access compute services.

For more information about alerts in Acronis Cyber Infrastructure, see Managing notifications.

**Important**
You can configure email notifications Acronis Cyber Infrastructure to send an email for each alert that is raised. For more information, see Sending email notifications.

## Configuring email notifications

Acronis Cyber Protect Cloud sends email notifications in the following cases.

- Disaster Recovery Hybrid infrastructure is unavailable.
- No VPN tunnels are available.

For more information about configuring notifications in Acronis Cyber Protect Cloud, see Changing the notification settings of a user.

## Troubleshooting email notifications

AcronisCyber Protect Cloud sends email notifications when the connection to the Disaster Recovery Hybrid infrastructure is unavailable, or when the VPN tunnels are unavailable.

For more information about the how to localize and fix the issues, depending on the email notification you get, see "Disaster Recovery Hybrid infrastructure is unavailable" (p. 48), and "No VPN tunnels are available" (p. 49).

## Disaster Recovery Hybrid infrastructure is unavailable

This notification means that the connection between Acronis Cyber Protect Cloud and the Disaster Recovery infrastructure virtual machines and services cannot be established.

***To resolve this issue***

1. Check the most common causes of the issue.
   a. Ensure that Acronis Cyber Protect Cloud is not under maintenance, and there are no reported incidents.
   b. Using Acronis Cloud Connection Verification Tool, check the connectivity from the agent gateway virtual machine (VM) to the cloud. For more information about locating the agent gateway VM, see step 2 a).
      i. If Acronis Cloud Connection Verification Tool is not available on the agent gateway VM, download and run **AcronisCloud_linux_checker32.zip** or **AcronisCloud_linux_checker32.zip**. You can find the files in KB article 47145. For more information about using the tool, see section **Usage in Linux**.
      ii. Ping the Acronis Cyber Protect Cloud public IP addresses, and try to reach them via HTTP.
2. Locate the agent gateway virtual machine and check its status.
   a. In Acronis Cyber Infrastructure, open **Compute** > **Virtual machines** and locate the virtual machine with a name starting with 'agent-gateway'.
   b. Ensure that the virtual machine is running.
   c. Click the virtual machine, and find its IP address in the **Network Interfaces** > **Public section**.

   **Note**
   If the virtual machine is missing, use the `dr-installer scale` command to redeploy it.

3. Log in to the virtual machine using SSH and the SSH private key specified in the dr-installer configuration file.

   **Note**
   If you cannot log in to the virtual machine due to network errors, download the virtual machine log file using the Acronis Cyber Infrastructure GUI, and contact the Support team.

4. Run the following command to check if the **runvm-agent** system service is active (running).

   ```
   systemctl status runvm-agent
   ```

   - If the service is not active (running), start it.
   - If you want to troubleshoot the issue, create a ticket to the Support team, and attach the service log files that are located in the `/var/log/runvm-agent` directory.
   - If the service exits immediately, send the service log files that are located in the `/var/log/runvm-agent` directory to the Support team, and then delete the virtual machine, and use the `dr-installer scale` command to redeploy the missing agent.

5. In your browser, enter `<agent IP address>:2661/hci/v1/healthcheck` to send an HTTP GET health check request. If the request fails, investigate for Acronis Cyber Infrastructure networking issues.

## No VPN tunnels are available

You receive this notification when one of the following issues has occured.

- At least one configured and established VPN tunnel is currently not available.
- All of the connected VPN tunnels for multiple tenants have suddenly become unavailable.

***To resolve this issue***

1. Check the most common causes of the issue.
   a. Ensure that Acronis Cyber Protect Cloud is not under maintenance, and there are no reported incidents.
   b. Using Acronis Cloud Connection Verification Tool, check the connectivity from the agent gateway virtual machine (VM) to the cloud. For more information about locating the agent gateway VM, see step 2 a).
      i. If Acronis Cloud Connection Verification Tool is not available on the agent gateway VM, download and run **AcronisCloud_linux_checker32.zip** or **AcronisCloud_linux_checker32.zip**. You can find the files in KB article 47145. For more information about using the tool, see section **Usage in Linux**.
      ii. Ping the Acronis cloud public IP addresses, and try to reach them via HTTP.
2. Check the status of the VPN appliance and the VPN gateway.
   a. In Acronis Cyber Protect Cloud, open **Disaster Recovery** > **Connectivity** and locate the VPN appliance and the VPN gateway.
   b. Ensure that the status of the VPN appliance and VPN gateway is **OK**.
   c. For the VPN appliance and VPN gateway, click the **Settings (gear)** icon, and then click **Download log**.
   d. Investigate the logs for network errors and other errors which prevent the VPN appliance and VPN gateway from establishing connections. If you cannot resolve the errors, create a ticket to the Support team and attach the logs.
3. Try to reinstall the VPN gateway.
   a. In Acronis Cyber Protect Cloud, open **Disaster Recovery** > **Connectivity**, and locate the VPN gateway.
   b. Click the **Settings (gear)** icon of the VPN gateway, and then click **Reinstall VPN gateway**.

# Troubleshooting alerts raised in Acronis Cyber Infrastructure

RunVM agents are crucial part of the failover, test failover, and backup operations. Therefore, the health of the RunVM agents is essential for the Disaster Recovery functionality.

If there is an issue with one of the RunVM agents or with their database, a notification appears in Acronis Cyber Infrastructure. For more information about the how to localize and fix the issues, depending on the notification you get, see "The Disaster Recovery Hybrid database is unavailable" (p. 50), "Disaster Recovery Hybrid agent (runvm-agent) is unavailable" (p. 50), "Disaster Recovery Hybrid agent (runvm-agent) cannot access compute services" (p. 51), and "Hybrid DR <available_ version> is now available" (p. 52).

## The Disaster Recovery Hybrid database is unavailable

***To resolve this issue***

1. Use the **UUID** from the alert to locate the affected virtual machine.
   a. In Acronis Cyber Infrastructure, open **Compute** > **Virtual machines**.
   b. Click a virtual machine from the list, and check the value of the **VM ID**. It should match the **UUID** from the alert.
   c. After you locate the virtual machine, check its IP address in **Compute** > **Virtual machines**.
   d. **Note**
   If the virtual machine is missing, use the `dr-installer scale` command to redeploy it.
2. Log in to the virtual machine using SSH and the SSH private key specified in the `dr-installer` configuration file.

   **Note**
   If you cannot log in to the virtual machine due to network errors, download the virtual machine log file using the Acronis Cyber Infrastructure GUI, and contact the Support team.
3. Run the following command to check if the `postgresql-10` system service is active (running).
   `systemctl status postgresql-10`
   - If the service is not active (running), start it.
   - If you want to troubleshoot the issue, create a ticket to the Support team, and attach the service log files that are located in the `/var/lib/pgsql/10/data/pg_log` directory.
4. In your browser, enter `<agent IP address>:9187/metrics` to send an HTTP GET request. If the request fails, investigate for Acronis Cyber Infrastructure networking issues.

## Disaster Recovery Hybrid agent (runvm-agent) is unavailable

***To resolve this issue***

1. Use the **UUID** from the alert to locate the affected virtual machine.
   a. In Acronis Cyber Infrastructure, open **Compute** > **Virtual machines**.
   b. Click a virtual machine from the list, and check the value of the **VM ID**. It should match the **UUID** from the alert.
   c. After you locate the virtual machine, check its IP address in **Compute** > **Virtual machines**.

**Note**
If the virtual machine is missing, use the `dr-installer scale` command to redeploy it.

2. Log in to the virtual machine using SSH and the SSH private key specified in the `dr-installer` configuration file.

**Note**
If you cannot log in to the virtual machine due to network errors, download the virtual machine log file using the Acronis Cyber Infrastructure GUI, and contact the Support team.

3. Run the following command to check if the **runvm-agent** system service is active (running).
   `systemctl status runvm-agent`
   - If the service is not active (running), start it.
   - If you want to troubleshoot the issue, create a ticket to the Support team, and attach the service log files that are located in the `/var/log/runvm-agent` directory.
   - If the service exits immediately, send the service log files that are located in the `/var/log/runvm-agent` directory to the support, and then delete the virtual machine, and use the `dr-installer scale` command to redeploy the missing agent.
4. In your browser, enter `<agent IP address>:2661` to send an HTTP GET request. If the request fails, investigate for Acronis Cyber Infrastructure networking issues.

## Disaster Recovery Hybrid agent (runvm-agent) cannot access compute services

***To resolve this issue***

1. Use the **UUID** from the alert to locate the affected virtual machine.
   a. In Acronis Cyber Infrastructure, open **Compute** > **Virtual machines**.
   b. Click a virtual machine from the list, and check the value of the **VM ID**. It should match the **UUID** from the alert.
   c. After you locate the virtual machine, check its IP address in **Compute** > **Virtual machines**.

**Note**
If the virtual machine is missing, use the `dr-installer scale` command to redeploy it.

2. Log in to the virtual machine using SSH and the SSH private key specified in the `dr-installer` configuration file.
3. Open the `/etc/runvm-agent/config.yml` file.
4. Find the section that has a line starting with `OpenStack:`.
5. In that section, check the following:
   a. If the credentials are correct, and if the specified user has Super user permission in the Acronis Cyber Infrastructure.

   b. If the specified Acronis Cyber Infrastructure address is correct.
6. Check if the management network that you used in dr-installer is the same as, or can reach, the Acronis Cyber Infrastructure network that is assigned the "Compute API" role.
7. If everything looks correct, create a ticket to the Support team and attach the agent logs from the `/var/log/runvm-agent` directory.

## Hybrid DR <available_version> is now available

When there is a new version of Disaster Recovery Hybrid, you will see the following alert:

Hybrid DR <available_version> is now available. Install the update as soon as possible. Otherwise your product functionality might be limited.

To resolve the issue, complete the update procedure. For more information, see "Update procedure" (p. 42).

# Working with Grafana dashboards

You can monitor the health of the disaster recovery RunVM system components on Grafana dashboards.

**Note**
Grafana dashboards are supported for Acronis Cyber Infrastructure version 5.4 Update 2 or later.

## Installing the Grafana dashboards

Install the Grafana dashboard on each Acronis Cyber Infrastructure (ACI) node that is a part of the high availability (HA) configuration. In that way, you will ensure that the dashboard will remain available even if the master node is changed due to a high availability event.

***To install the Grafana dashboards***

1. Log in to the command line interface of the node as a system administrator with root privileges.

   **Note**
   You can view which nodes belong to the HA configuration in the ACI admin panel, on the **Settings** > **System settings** > **Management node high availability** page. Alternatively, you can use the `vinfra cluster ha show` command.

2. Run the following command.

   ```
   # yum install vstorage-grafana-dashboards-dr.x86_64
   ```

## Uninstalling the Grafana dashboards

You can uninstall the Grafana dashboards from the nodes that are a part of the high availability (HA) configuration.

***To uninstall the Grafana dashboards***

1. Log in to Acronis Cyber Infrastructure (ACI) as a system administrator with root privileges.
2. Run the following command on each ACI node that is a part of the HA configuration.

   ```
   # yum remove vstorage-grafana-dashboards-dr.x86_64
   ```

## Monitoring the health of the RunVM system components

You can monitor the health of the RunVM system components.

### Prerequisites

Grafana dashboards are installed on each node that is part of the high availability (HA) configuration.

***To monitor the health of the RunVM system components***

1. Log in to the Acronis Cyber Infrastructure admin panel.
2. Navigate to **Monitoring** > **Grafana dashboards**.
3. Find the **Hybrid DR agents overview** dashboard.

# Appendix A. The default config.yml

```
##############################################################################################
##############

#

# Disclaimer for DCO:

#    You need to fill empty values only.

#    Please do not modify any defaults if you are not sure that you really need it.

#

##############################################################################################
##############



# Parameters of the Acronis Cyber Infrastructure (ACI) cluster.

infrastructure:

  #

  # Host name or IP address for the management node of the Acronis Cyber
Infrastructure cluster.

  # [Required parameter]

  # Example: name.company.com or 10.201.197.8

  address:

  #

  # Allow using insecure SSL connections to API if the authenticity of the host
cannot be validated.

  # If the ACI stand is accessible from the Internet, we recommend that you set the
"insecure" parameter to "false".

  insecure: false

  #

  # Authentication parameters in the Acronis Cyber Infrastructure cluster.

  # Please, use WebCP authentication parameters for Acronis Cyber Infrastructure.

  # If you use a user other than an administrator, ensure that the custom user has
the OpenStack API access rights.

  # [Required parameter]

  auth:

    username: admin
```

```
    password:

    #

    # OpenStack project name that Disaster Recovery will use to log in and start the
deployment.

    # It is highly recommended to use the default value. If you need to use a custom
value, contact the Support team.

    project: admin

    #

    # OpenStack domain name that Disaster Recovery will use to log in and start the
deployment.

    # It is highly recommended to use default value. If you need to use a custom
value, contact the Support team.

    domain: Default




# Parameters of the Acronis Cyber Protect Cloud (ACPC).

cloud:

  #

  # URL of the Acronis Cyber Protect Cloud API endpoint.

  # This should be the name of the current data center (us-cloud.acronis.com, etc).

  # [Required parameter]

  url: https://cloud.acronis.com

  #

  # Allow using insecure SSL connections to ACC if thw authenticity of the host
cannot be validated.

  # Please use the "false" value for the "insecure" parameter.

  insecure: false

  #

  # Authentication parameters in Acronis Cyber Protect Cloud.

  # This user will be used to register Acronis Disaster Recovery Infrastructure.

  # Attention for DCO: this should be a service user on the root level for each DC.

  # [Required parameters]

  auth:

    username:
```

```
    password:



# Deployment parameters
deployment:
  #
  # Acronis Cyber Infrastructure-related deployment settings.
  infrastructure:
    #
    # The name of the sub-project that will be created and used for the deployment.
    # All Disaster Recovery components will be created in ACI under this sub-
project.
    # The parent project is configured with the infrastructure.auth.project
parameter.
    project: Disaster-Recovery
    #
    # The name of a role that will be assigned to the current user in the created
sub-project.
    # It is recommended to use the default value.
    user-role: admin
    #
    # The default key pair is the key pair which is assigned to the new servers.
    default-key-pair: Disaster-Recovery
    #
    # A list of key pairs that are used for connecting to existing servers.
    #
    # One of them (default) is used for new servers deployed during the
installation.
    # NOTE: Private key file should be in PEM format. It will be used only during
    #       installation to access existing or new servers over SSH,
    #       and will not be preserved anywhere after installation.
    # NOTE: Public key should be in OpenSSH Public Key format. It will be used
    #       during installation to setup SSH authentication for centos user
```

```
    #       in newly deployed servers.
    #
    # !!! SSH keys should be generated manually. Please, keep them in a safe storage
with backup.
    # !!! If you lose the SSH keys, you won't be able to access Acronis Cyber
Infrastructure for maintenance.
    #
    key-pairs:
      - name: Disaster-Recovery
        private-key-file:
        public-key-file:
    #
    # Storage policies which should be deployed.
    storage-policies:
      #
      # System storage policy (for internal virtual machines).
      system:
        #
        # Enable storage policy.
        enabled: true
        #
        # Storage policy name template.
        # Storage policy name must be unique across domain in Acronis Cyber
Infrastructure.
        # Use `{{.Project}}` as a placeholder which will be replaced by the name of
the project.
        name: 'Disaster-Recovery-Project (System)'
        #
        # I/O operations per second limit.
        iops-limit: 1000
      #
      # Default storage policy (for customers' virtual machines).
      default:
```

```
        #
        # Enable storage policy.
        enabled: true
        #
        # Storage policy name template.
        # Storage policy name must be unique across domain in Acronis Cyber
Infrastructure.
        # Use `{{.Project}}` as a placeholder which will be replaced by the name of
the project.
        name: '{{.Project}} (Default)'
        #
        # I/O operations per second limit.
        iops-limit: 400
    #
    # Flavors which will be created during deployment.
    # These flavors will be available for a customer to create recovery and primary
servers.
    # You may add new flavors or adjust the existing ones as needed.
    flavors:
      #
      # Flavor name.
    - name: 1 vCore, 2 GB RAM
      #
      # Virtual CPU count.
      cpu: 1
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 2048
      #
      # Flavor name.
    - name: 1 vCore, 4 GB RAM
      #
      # Virtual CPU count.
```

```
      cpu: 1
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 4096
      #
      # Flavor name.
    - name: 2 vCores, 8 GB RAM
      #
      # Virtual CPU count.
      cpu: 2
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 8192
      #
      # Flavor name.
    - name: 4 vCores, 16 GB RAM
      #
      # Virtual CPU count.
      cpu: 4
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 16384
      #
      # Flavor name.
    - name: 8 vCores, 32 GB RAM
      #
      # Virtual CPU count.
      cpu: 8
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 32768
```

```
      #
      # Flavor name.
    - name: 16 vCores, 64 GB RAM
      #
      # Virtual CPU count.
      cpu: 16
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 65536
      #
      # Flavor name.
    - name: 16 vCores, 128 GB RAM
      #
      # Virtual CPU count.
      cpu: 16
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 131072
      #
      # Flavor name.
    - name: 16 vCores, 256 GB RAM
      #
      # Virtual CPU count.
      cpu: 16
      #
      # Memory amount, in MiB (1024 * 1024 bytes).
      ram: 262144
  #
  # Deployment settings related to Acronis Cyber Protect Cloud (infrastructure
registration, etc.).
  cloud:
    #
```

```
    # Parameters for registration of the infrastructure in the Acronis Cyber Protect
Cloud platform.

    infrastructure:

      # Optional owner of the newly created Disaster Recovery infrastructure.

      # If not specified, the owner will be the root tenant associated with the
username.

      # Example: 0d53a6ee-2922-408b-af59-9969851710ae

      # It is recommended to leave this parameter empty.

      owner-tenant-id:

      #

      # Infrastructure name that will be displayed in the Acronis Cyber Protect
Cloud management console.

      # You can change this parameter as needed.

      # [Required parameter]

      name: 'Acronis Disaster Recovery Cloud'

      #

      # The name of the infrastructure backend type.

      backend-type:

      #

      # The location is a group of infrastructures in the Acronis Cyber Protect
Cloud platform.

      location:

        #

        # Name of the Acronis Cyber Protect Cloud location in which the DR
infrastructure should be added.

        # You must define the location where your Backup storage (cold storage) is
already added.

        # Optional, if omitted, a new location is created automatically for each
infrastructure.

        name: 'Acronis Cloud'



# Network settings for Acronis Disaster Recovery Cloud on Acronis Cyber
Infrastructure

# Please, check the network diagram in the Network and firewall requirements topic
before modifying the following parameters.
```

```
# Network names below are Open Stack networks + VLANs.

networks:

  #

  # Disaster Recovery management network is used for services communications
(drmgmt).

  management:

    #

    # Name of the Disaster Recovery management network.

    # [Required parameter]

    # Example: drmgmt

    name:

  #

  # Content network is used for sharing archive data over NBD (NBD VXLAN).

  content:

    #

    # Name of the content network.

    # [Required parameter]

    # Example: content

    name:

  #

  # Storage network is used for communicating with the Acronis storage where the
archives are stored.

  storage:

    #

    # Name of the storage network.

    # [Required parameter]

    # Example: drcoldstorage

    name:

    #

    # Direct cold storage access settings.

    direct-routing:

      #
```

```
      # This parameter defines whether direct routing is enabled.
      # If direct routing is disabled, the network traffic between the DR components
and the Backup storage will pass through the Internet.
      # It is recommended to configure and enable direct routing. It requires
advanced network configuration.
      enabled: false
      #
      # Public IP subnet (in CIDR notation) assigned to the cold storage.
      # This parameter is required if direct routing is enabled.
      subnet:
      #
      # The IP address of the gateway, through which the cold storage will be
reached.
      # This parameter is required if direct routing is enabled.
      gateway:
  #
  # Demilitarized network is used for public access to the servers over the
Internet.
  dmz:
    #
    # Name of the dmz network.
    # [Required parameter]
    # Example: dmzvpn
    name: dmzvpn



# VPN settings.
vpn:
  #
  # Specification used for spawning VPN servers.
  gateway-specification:
    #
    # Settings used for cloud-only VPN.
```

```
    cloud-only:
      #
      # Flavor used for spawning VPN servers for cloud-only deployments.
      flavor:
        #
        # Flavor name.
        name: vpn-server-cloud-only
        #
        # Virtual CPU count.
        cpu: 1
        #
        # Memory amount, in MiB (1024 * 1024 bytes).
        ram: 512
    #
    # Settings used for ip-sec VPN.
    ip-sec:
      #
      # Flavor used for spawning VPN servers for ip-sec deployments.
      flavor:
        #
        # Flavor name.
        name: vpn-server-ip-sec
        #
        # Virtual CPU count.
        cpu: 2
        #
        # Memory amount, in MiB (1024 * 1024 bytes).
        ram: 2048
    #
    # Settings used for site-to-site VPN.
    site-to-site:
```

```
      #
      # Flavor used for spawning VPN servers for site-to-site deployments.
      flavor:
        #
        # Flavor name.
        name: vpn-server-site-to-site
        #
        # Virtual CPU count.
        cpu: 2
        #
        # Memory amount, in MiB (1024 * 1024 bytes).
        ram: 2048
  #
  # VPN proxying settings: SOCKS5 (Dante) and HAProxy.
  proxying:
    #
    # A name for a project that will be created for containing VPN proxies.
    #  It is recommended to use the default project name, but you may change it if
needed.
    project: vpn-proxy
    #
    # List of public IP pools for allocation to customers in this DR infrastructure.
    # [Required parameter]
    # Use exactly the same syntax as in the example entry:
    # - start: 102.44.32.5
    #   end: 102.44.32.70
    public-ip-pools:
        - start:
          end:


    #
    # SOCKS5 proxy settings.
```

```
    socks-proxy:
      #
      # Number of SOCKS5 proxy (Dante) servers. You can adjust this parameter
according to the workload in your data center.
      # [Required parameter]
      count: 2
      #
      # The name of a flavor for the server.
      # Flavor is an OpenStack parameter which defines the compute, memory, and
storage capacity of a virtual server.
      # It is recommended to use the default value "small".
      flavor: small
    #
    # HAProxy settings.
    haproxy:
      #
      # Number of HAProxy servers.
      # [Required parameter]
      count: 2
      #
      # The name of a flavor for the server.
      # Flavor is an OpenStack parameter which defines the compute, memory, and
storage capacity of a virtual server.
      # It is recommended to use the default value "small".
      flavor: small
      #
      # HAProxy network settings.
      networks:
        #
        # Demilitarized network settings for SOCKS proxy.
        dmz:
          #
          # Public IP address of the SOCKS proxy servers.
```

```
          # [Required parameter]
          # Example: 192.0.2.0/24
          address:
          #
          # Gateway is used for traffic from the demilitarized network to the
Internet.
          # [Required parameter]
          # Example: 10.248.81.1
          gateway:
      #
      # Virtual router ID for VRRP. We strongly recommend using the default value of
this parameter.
      virtual-router-id: 101
  #
  # Settings for the temporary tunnel server that is created during the VPN proxies
deployment.
  tunnel-server:
    #
    # Flavor name.
    # It is recommended to use the value "tiny".
    flavor: tiny



# Acronis Disaster Recovery RunVM platform settings.
runvm:
  #
  # RunVM Agent settings.
  agent:
    #
    # RunVM Agent system volume size, in gibibytes (1024*1024*1024 bytes).
    # It is recommended to use the default value.
    volume-size: 100
    #
```

```
    # Used as a start of the name for per-tenant ACI projects. Use a short DC name.
    # [Required parameter]
    # It is recommended to use the value from the deployment.infrastructure.project
parameter with a hyphen at the end.
    # Example: Disaster-Recovery-
    project-prefix:
    #
    # RunVM Agent roles configuration.
    # It is recommended to use the default RunVM Agent role settings, unless you
know exacly how to adjust each of the parameters below.
    roles:
      #
      # Runner role configuration, which is used for running (starting) virtual
machines from backups.
      runner:
        #
        # Number of instances of RunVM Agents with the runner role.
        # The minimum required value is 3.
        count: 3
        #
        # The name of a flavor for the server.
        # Flavor is an OpenStack parameter which defines the compute, memory, and
storage capacity of a virtual server.
        flavor: large
        #
        #
        limits:
          #
          # Specifies the maximum number of concurrently running tasks for VM
creation.
          vm-create-tasks: 5
          #
          # Specifies the maximum number of concurrently running tasks for VM
finalization.
```

```
          vm-finalize-tasks: 100

          #

          # Specifies the maximum number of concurrently running tasks for VM
deletion.

          vm-delete-tasks: 100

          #

          # Specifies maximum number of concurrently executing tasks for backup
optimization.

          backup-optimize-tasks: 10

          #

          # Specifies maximum number of concurrently executing tasks for backup
deoptimization.

          backup-deoptimize-tasks: 10

          #

          # Specifies the maximum number of concurrently running tasks for delta
creation.

          delta-create-tasks: 10

          #

          # Specifies the maximum number of concurrently running tasks for delta
deletion.

          delta-delete-tasks: 100

      #

      # Backuper role configuration, which is used for backing up virtual machines.

      backuper:

        #

        # Number of instances of RunVM Agents with the backuper role.

        # The minimum recommended value is 3.

        count: 3

        #

        # The name of a flavor for the server.

        # Flavor is an OpenStack parameter which defines the compute, memory, and
storage capacity of a virtual server.

        flavor: large

        #
```

```
        #
        limits:
          #
          # Specifies maximum number of concurrently executing backup tasks.
          backup-tasks: 7
          #
          # Specifies the maximum number of concurrently running archive-diff-size
tasks.
          archive-diff-size-tasks: 2
      #
      # Gateway role configuration, which is used for providing access to the
hypervisor for the Disaster Recovery service.
      gateway:
        #
        # Number of instances of RunVM Agents with the gateway role.
        # The minimum recommended value is 3.
        count: 3
        #
        # The name of a flavor for the server.
        # Flavor is an OpenStack parameter which defines the compute, memory, and
storage capacity of a virtual server.
        flavor: small
        #
        #
        limits:
          #
          # Specifies the maximum number of concurrently running hypervisor
management tasks.
          gateway-tasks: 100
          #
          # Specifies the maximum number of concurrently running cold data retention
tasks.
          cleanup-tasks: 3
    # RunVM Agent timeouts.
```

```
    timeouts:
      # Maximum allowed time for RunVM Agent deactivation initiated by the update
procedure.
      # In case an RunVM Agent instance is found that did not manage to deactivate
within the specified timeout,
      # it will be deleted forcibly.
      deactivation: 168h
  #
  # RunVM Controller settings.
  # It is recommended to use the default RunVM Controller settings, unless you know
how exactly to adjust each of the parameters below.
  controller:
    #
    # RunVM Controller system volume size, gibibytes.
    volume-size: 100
    #
    # The name of a flavor for the server.
    # Flavor is an OpenStack parameter which defines the compute, memory, and
storage capacity of a virtual server.
    flavor: small
    #
    # Options for a pool of RunVM Controller instances.
    pool:
      #
      # Size of the pool of pre-created RunVM Controller instances.
      size: 4
      #
      # Enable or disable the pool of RunVM Controller instances.
      disabled: false
      #
    # On-configuration hook allows specifying an external command that will be
executed while configuring the controller image.
    on-configuration-hook:
      # Command is a list of strings containing a command, optionally followed by a
```

```
list of arguments.

      # The following environment variables will be passed into the command:

      # DR_INSTALLER_HOOK_TARGET_SERVER_ADDR - IP address of virtual machine that
currently configures controller image.

      # DR_INSTALLER_HOOK_POSTGRES_ADDR - IP address of current infra PostgreSQL
server.

      # DR_INSTALLER_HOOK_COREDUMP_SERVER_ADDR - IP address of current infra core-
dump server.

      # DR_INSTALLER_HOOK_SSH_KEY - path to ssh key.

      command: []



# PostgreSQL settings.

# It is recommended to use the default PostgreSQL settings, unless you know how
exactly to adjust each of the parameters below.

postgresql:

  #

  # The name of a flavor for the server.

  # Flavor is an OpenStack parameter which defines the compute, memory, and storage
capacity of a virtual server.

  flavor: medium


# Core-dump-server settings.

# It is recommended to use the default Core-dump-server settings unless you know how
exactly to adjust each of the parameters below.

core-dump-server:

  #

  # Core Dump Server system volume size, gibibytes.

  volume-size: 200

  #

  # The name of a flavor for the server.

  # Flavor is an OpenStack parameter which defines the compute, memory, and storage
capacity of a virtual server.

  flavor: tiny
```

```
  # Storage settings
  storage:
    # The maximum disk space that can be used by the core dump files.
    max-space-usage: '130GiB'
    # The maximum disk used by the core dump files of a single client.
    max-space-usage-per-client: '30GiB'
    # The maximum log file size.
    max-log-file-size: '1GiB'
    # The maximum log file count.
    max-log-file-count: 300
    # This parameter defines whether log compression is enabled.
    enable-log-compression: true



# Timeouts.
timeouts:
  # Server operation timeouts
  server:
    # Server creation timeout
    creation: 15m
    # Server deletion timeout
    deletion: 10m
    # Server shutoff timeout
    shutoff: 15m
    # Server shelving timeout
    shelving: 15m
  # Image operation timeouts
  image:
    # Image deletion timeout
    deletion: 5m
    # Image uploading completion timeout
```

```
    uploading-completion: 10m
  # Volume operation timeouts
  volume:
    # Volume readiness for deletion timeout
    deletion-readiness: 5m
    # Volume deletion timeout
    deletion: 5m
```

# Appendix B. Working with the dr-installer tool

**NAME**:

`dr-installer` – Acronis Disaster Recovery installer for the Acronis Cyber Infrastructure environment

**USAGE**:

First, you must specify the global options, then the commands.

```
dr-installer [global options] command [command options] [arguments...]
```

**GLOBAL OPTIONS**:

`--force` – to force deletion of the Disaster Recovery infrastructure

`--config <value>` – a path to the deployment configuration file

`--components-config <value>` – a path to the component configuration file

`--verbose` – a detailed output in the Debug mode, showing both [INFO] and [DEBUG] messages

`--json-log` – a json log output

`--log-file <value>` – a log file name (default: "`dr-installer.log`")

`--log-to-stderr` – output console logs to stderr instead of stdout

`--skip-cleanup` – skip cleanup on exit

**COMMANDS**:

`help` – show a list of available commands or help for one command

`version` – show the version and exit

`install` – install a new DR infrastructure

`update` – update the existing DR infrastructure (the config file must be the same as for the initial installation)

`scale` – scale the existing DR infrastructure

`remove` – remove the existing DR infrastructure

`validate` – validate the configuration file settings

`info` – show the installation information

`image` – manage images for primary servers

`vpn-proxy` – manage VPN proxying infrastructure for the Acronis Cyber Infrastructure cluster

# Appendix C. Disaster Recovery architecture and components

In the diagram below, you can find the DR infrastructure components.

There are three infrastructural layers. From the bottom to the top:

1. **Physical layer**: hardware nodes combined in a highly available cluster.
2. **Software-defined infrastructure layer**: Acronis Cyber Infrastructure installed on the hardware nodes.
3. **Service infrastructure layer**: virtual machines for each DR service component and a customer's cloud infrastructure running on top of Acronis Cyber Infrastructure.

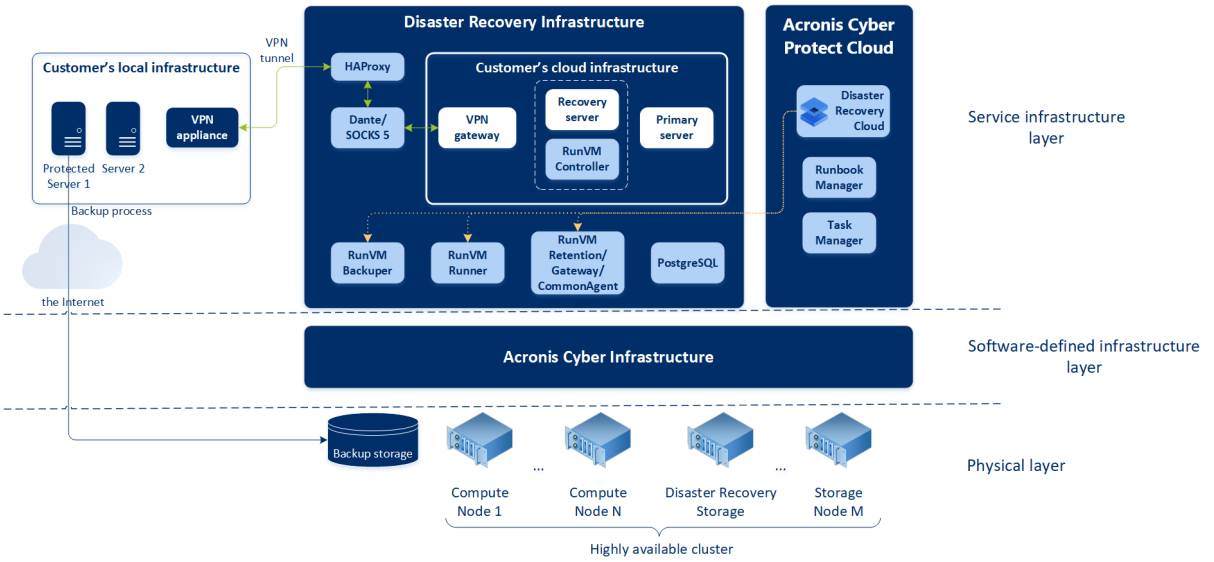


The main components of the Disaster Recovery infrastructure are listed in the table below.

| Component name | Description |
|---|---|
| **RunVM Agent** | An agent that can have one of the following three roles:<br>• **RunVM Backuper** – an agent role responsible for backing up cloud servers.<br>• **RunVM Runner** – an agent role responsible for creating deltas (service files that optimize a virtual machine start from a backup, including bootability and AUR fixes).<br>• **RunVM Gateway/Retention/CommonAgent** – an agent role that acts as a hypervisor gateway and performs backup retention tasks.<br>One RunVM Agent can process concurrently the following number of tasks:<br>• vm-create-tasks: 5<br>• vm-finalize-tasks: 100<br>• vm-delete-tasks: 100 |

| | |
|---|---|
| | • delta-create-tasks: 10<br>• delta-delete-tasks: 100<br><br>The number of RunVM Agents must be equal to the number of nodes in the ACI cluster. Thus, for two nodes you need two RunVM Agents on the two ACI nodes. |
| **RunVM Controller** | A component responsible for attaching a backup as a disk to a virtual machine. The RunVM Controller is created per recovery server in failover mode and exists as long as the respective recovery server exists. If the recovery server is deleted, the RunVM Controller is deleted automatically. |
| **HAProxy/Dante (SOCKS5)** | These virtual machines are mainly used to eliminate the need to provide a dedicated public IP address for each client. It is recommended to have at least two proxy servers of each type for high availability. |
| **VPN gateway** | A special virtual machine providing a connection between the customer's local network and the cloud recovery site via a secure VPN tunnel. The VPN gateway is deployed on the cloud recovery site. |
| **VPN appliance** | 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. |
| **Primary server templates** | Server templates from which the primary servers are launched. |
| **PostgreSQL RDBMS** | This database is currently used only for distributed locks between agents. |
| **Virtual router ID for Virtual Router Redundancy Protocol (VRRP)** | It is used for selecting the active instance between HAProxy instances. The parameter is needed for using the other ID if the current one is already used by any device or application in the dmzvpn network. |
| **Tunnel-server** | The dr-installer needs access to the dmz network from your local machine where it is run. The tunnel-server is a temporary auxiliary server that is used for connection to the management and dmz networks. After the deployment of proxy servers, the tunnel-server is deleted. |
| **Core Collector** | It is an auxiliary server for collecting logs and core files from RunVM Agents and RunVM Controllers to troubleshoot issues. |

# Appendix D. Calculating hardware needs

## Calculating hardware needs for a compute cluster

To calculate your hardware needs for a compute cluster based on the desired number of virtual machines and their profile, follow the steps described below.

**Step 1. Define the number of virtual machines and their configuration.**

Define the following parameters:

- VM count – the number of machines that can be run simultaneously in a failover case.
- VM CPU, cores – the number of CPU cores for one virtual machine.
- VM RAM, GB – the amount of RAM for one virtual machine.
- VM disk (hot storage), TB – the disk amount for one virtual machine.

**Step 2. Define the desired overcommit parameter.**

Define the following parameter:

- CPU overcommit ratio

**Step 3. Define the redundancy parameter.**

Define the following parameter:

- Redundancy (replicas=3), Hot

**Step 4. Define the possible hardware node configurations.**

Define the following parameters for hardware nodes that you are going to deploy in your data center:

- HW CPU, cores – the number of CPU cores for a hardware node
- HW RAM, GB – the amount of RAM for a hardware node

**Step 5. Calculate how much resources of hardware a VM will consume taking into account the defined overcommit ratio.**

A VM will consume the following amount of CPU taking into account the overcommit ratio:

VM CPU (overcommit) = VM CPU * CPU overcommit ratio

**Step 6. Calculate how much CPU and RAM are actually available on the nodes.**

For each node you can calculate the amount of CPU and RAM actually available for VMs considering that some amount of these resources is consumed by different services. See the tablebellow.

CPU available = HW CPU – 9 (value is taken from the table below)

RAM available = HW RAM – 32 (value is taken from the table below)

**Step 7. Calculate how many VMs with the desired profile can be provisioned on a hardware node.**

VM count (CPU) = CPU available / VM CPU (overcommit)

VM count (RAM) = RAM available / VM RAM

VM count available = MIN(VM count (CPU);VM count(RAM))

**Step 8. Calculate how many hardware nodes you need for a compute cluster.**

**Node count** = VM count / VM count available

**Step 9. Calculate hot disk space required for VMs per node.**

**Disk space (hot) per node, TB** = VM count available * VM disk (hot storage) * Redundancy (replicas=3)

## Predefined parameters for hardware needs calculations

The following table lists the recommended amount of RAM and CPU cores for one node according to the services you will use on a compute node:

| Service | RAM, GB | CPU, cores |
|---|---|---|
| System | 6 | 2 |
| Compute | 8 | 3 |
| Disk HOT (10 disks per node) | 10 | 2 |
| DR components | 8 | 2 |
| **Total** | 32 | 9 |

# Calculating hardware needs for a storage cluster

To calculate your hardware needs for a cold storage cluster, follow the steps described below.

**Step 1. Define the amount of disk on cold storage for a VM**

VM disk (cold storage), TB – the disk amount for one virtual machine.

**Step 2. Define the redundancy parameter.**

Redundancy (EC=3+2)

**Step 3. Calculate the cold disk space required for VMs per node.**

**Disk space (cold) per node, TB** = VM count available * VM disk (cold storage) * Redundancy (EC=3+2)

**Step 4. Calculate how many hardware nodes you need for a storage cluster.**

**Node count** = VM count * VM disk (cold storage) / Disk space (cold) per node

## Predefined parameters for hardware needs calculations

The following table lists the recommended amount of RAM and CPU cores for one node according to the services you will use on a storage node:

| Service | RAM, GB | CPU, cores |
|---|---|---|
| System | 6 | 2 |
| Backup storage COLD (10 disks per node) | 10 | 2 |
| Backup gateway | 1 | 2 |
| **Total** | 17 | 6 |

# Example of hardware needs calculation

Now, let us consider an example how to calculate the number of hardware nodes needed for a DR infrastructure if you have the desired number of VMs to be run in case of failover per node and their profile. We will consider a case when you are going to deploy a compute cluster (compute resources and a hot storage) and a storage cluster (a cold storage).

**Step 1. Define the number of virtual machines and their configuration.**

You plan to run 100 virtual machines (VM count = 100) in case of failover with the following parameters:

- VM CPU, cores = 2
- VM RAM, GB = 8
- VM disk (hot storage), TB = 0,25
- VM disk (cold storage), TB = 0,5

**Step 2. Define the desired overcommit ratio.**

- CPU overcommit ratio = 0,25

**Step 3. Define the redundancy parameters.**

You plan to use the following parameters:

- Redundancy (replicas=3), hot = 3
- Redundancy (EC=3+2), cold = 5/3 = 1,67

**Step 4. Define the possible hardware node configurations.**

Define the following parameters for hardware nodes that you are going to deploy in your data center:

- HW CPU, cores = 20
- HW RAM, GB = 192

- HW disk count = 10 disks per node

**Step 5. Calculate how much CPU of hardware a VM will consume taking into account the defined overcommit ratio.**

A VM will consume the following amount of CPU taking into account the overcommit ratio:

VM CPU (overcommit), cores = 2 * 0,25 = 0,5

**Step 6. Calculate how much CPU and RAM are actually available on the nodes.**

For each node you can calculate the amount of CPU and RAM actually available for VMs considering that some amount of these resources is consumed by different services.

CPU available, cores = 20 – 9 = 11

RAM available, GB = 192 – 32 = 160

**Step 7. Calculate how many VMs with the desired profile can be provisioned on a hardware node.**

VM count (CPU), cores = 11 / 0,5 = 22

VM count (RAM), GB = 160 / 8 = 20

VM count available = MIN(22;20) = 20

**Step 8. Calculate the disk space required for VMs per node.**

**Disk space (hot) per node, TB** = 20 * 0,25 * 3 = 15

**Disk space (cold) per node, TB** = 20 * 0,5 * 1,67 = 16

**Step 9. Calculate how many hardware nodes you need for a compute and storage clusters.**

**Node count (for compute)** = 100 / 20 = 5

**Node count (for storage)** = 100 * 0,5 / 16 = 4

As a result, you will need five compute nodes and four storage nodes to protect 1000 VMs from a disaster and run 100 VMs simultaneously in case of failover.

# Appendix F. Direct routing to the Backup storage

It is recommended to configure and enable direct routing to the backup storage. If direct routing is disabled, the network traffic between the DR components and backup storage will go through the Internet.

# Helpful links

Disaster Recovery Administrator Guide:
https://www.acronis.com/support/documentation/DisasterRecovery/#43224.html

Disaster Recovery Quick Start Guide: https://dl.managed-protection.com/u/pdf/DisasterRecovery_quickstart_en-US.pdf

Disaster Recovery video tutorials:
https://www.youtube.com/watch?v=p3RDUB5UQSY&list=PLJbh8iM59BMeWCLCErGmwO4QxOBkg6bDi

Acronis Cyber Infrastructure documentation see at
https://www.acronis.com/support/documentation/

# Index

**T**



**U**



**W**