Cyber Disaster Recovery Cloud

25.06

Administrator Guide

REVISION: 6/24/2025

Table of contents

About Acronis Disaster Recovery

Acronis Disaster Recovery is a part of Cyber Protection that provides disaster recovery as a service (DRaaS). Disaster Recovery is a fast and stable solution that, in the event of man-made or natural disaster, launches recovery servers (exact copies of your local servers) on the disaster recovery (DR) site.

The DR site is a secondary location that ensures that IT operations are restored and continue if the primary site fails. In such a case, DR switches the workloads from the corrupted original machines to the recovery servers at the DR site. The DR site includes the cloud servers and the cloud connectivity (cloud networks).

The DR site can be located in Acronis Cyber Protect Cloud or Microsoft Azure. The availability of the DR site locations depends on the offering items that are enabled for your tenant.

Note
Only one location is supported per customer tenant. If you want to change the location of your DR site, you must first remove the existing configuration, and then create a new one with the new location.

For more information about Disaster Recovery to Acronis Cyber Protect Cloud, see "Disaster Recovery to Cyber Protect Cloud" (p. 7).

For more information about Disaster Recovery to Microsoft Azure, see "Disaster Recovery to Microsoft Azure" (p. 93).

Disaster Recovery to Cyber Protect Cloud

Disaster Recovery to Cyber Protect Cloud has the following characteristics:

- The location of the DR site is Cyber Protect Cloud.
- The backups (recovery points) of the protected servers are stored in Cyber Protect Cloud.

The key functionality

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

- Manage the Disaster Recovery service from a single console
- Extend up to 23 local networks to the cloud, by using a secure VPN tunnel
- Establish the connection to the cloud site without any VPN appliance[1] deployment (the cloud-only mode)
- Establish the point-to-site connection[2] to your local[3] and cloud sites[4]
- Protect your machines by using recovery servers[5] in the cloud
- Protect applications and appliances by using primary servers[6] in the cloud
- Perform automatic disaster recovery operations for encrypted backups
- Perform a test failover in the isolated network
- Use runbooks[7] to automate the deployment to the production environment in the cloud

Software requirements for Disaster Recovery to Cyber Protect Cloud

Supported operating systems

Protection with a recovery server has been tested for the following operating systems:

- CentOS 6.6, 7.x, 8.x
- Debian 9.x, 10.x, 11.x

---

[1] 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.
[2] A secure VPN connection from outside to the cloud and local sites by using your endpoint devices (such as a computer or laptop).
[3] The local infrastructure deployed on your company's premises.
[4] Remote site hosted in the cloud and used for running recovery infrastructure, in case of a disaster.
[5] A VM replica of the original machine, based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers, in case of a disaster.
[6] A virtual machine that does not have a linked machine on the local site (such as a recovery server). Primary servers are used for protecting an application or running various auxiliary services (such as a web server).
[7] Planned scenario consisting of configurable steps that automate disaster recovery actions.

- Red Hat Enterprise Linux 6.6, 7.x, 8.x
- Ubuntu 18.04, 20.x, 21.x
- Oracle Linux 7.3 and 7.9 with Unbreakable Enterprise Kernel
- Windows Server 2008 R2
- Windows Server 2012/2012 R2
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows Server 2022 – all installation options, except for Nano Server
- AlmaLinux 8.x, 9.0, 9.1, 9.2, 9.3, 9.4, 9.5
- Rocky Linux 9.x, 8.x

The software may work with other Windows operating systems and Linux distributions, but this is not guaranteed.

Note
Protection with a recovery server has been tested for Microsoft Azure VM with the following operating systems.

- Windows Server 2008 R2
- Windows Server 2012/2012 R2
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows Server 2022 – all installation options, except for Nano Server
- Ubuntu Server 20.04 LTS - Gen2 (Canonical). For more information about accessing the recovery server console, see https://kb.acronis.com/content/71616.

Supported virtualization platforms

Protection of virtual machines with a recovery server has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7, 7.0
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2022 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016
- Kernel-based Virtual Machines (KVM) — fully virtualized guests (HVM) only. Paravirtualized guests (PV) are not supported.
- Red Hat Enterprise Virtualization (RHEV) 3.6

- Red Hat Virtualization (RHV) 4.0
- Citrix XenServer: 6.5, 7.0, 7.1, 7.2

The VPN appliance has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2022 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016

Linux workloads that have agentless backups from a guest OS and have volumes with the Logical Volume Manager (LVM) configurations are supported.

Windows workloads that have agentless backups from a guest OS and have dynamic disks (LDM) configurations are supported.

The software might work with other virtualization platforms and versions, but this is not guaranteed.

Limitations

The following platforms and configurations are not supported in Disaster Recovery:

1. Unsupported platforms:
   - Agents for Virtuozzo
   - macOS
   - Windows desktop operating systems are not supported due to Microsoft product terms.
   - Windows Server Azure Edition

     Azure Edition is a special version of Windows Server that was built specifically to run either as an Azure IaaS virtual machine (VM) in Azure or as a VM on an Azure Stack HCI cluster. Unlike the Standard and Datacenter editions, Azure Edition is not licensed to run on bare metal hardware, Windows client Hyper-V, Windows Server Hyper-V, third-party hypervisors, or in third-party clouds.
2. Unsupported configurations:

   Microsoft Windows
   - Windows desktop operating systems are not supported (due to Microsoft product terms).
   - Active Directory service with FRS replication is not supported.
   - Removable media without either GPT or MBR formatting (so-called "superfloppy") are not supported.

   Linux

- File systems without a partition table.
- Linux workloads that are backed up with an agent from a guest OS and have volumes with the following advanced Logical Volume Manager (LVM) configurations: Striped volumes, Mirrored volumes, RAID 0, RAID 4, RAID 5, RAID 6, or RAID 10 volumes.

Note
Workloads with multiple operating systems installed are not supported.

3. Unsupported tenant modes:
   - Disaster recovery is not available when Compliance mode is enabled for the tenant.
4. Unsupported backup types:
   - Continuous data protection (CDP) recovery points are incompatible.

     Important
     If you create a recovery server from a backup having a CDP recovery point, then during the failback or creating backup of a recovery server, you will loose the data contained in the CDP recovery point.

   - Forensic backups cannot be used for creating recovery servers.

A recovery server has one network interface. If the original machine has several network interfaces, only one is emulated.

Cloud servers are not encrypted.

Disaster Recovery trial version

You can use a trial version of Acronis Disaster Recovery for a period of 30 days. In this case, Disaster Recovery has the following limitations for partner tenants:

- No access to public internet for recovery and primary servers. You cannot assign public IP addresses to the servers.
- IPsec Multi-site VPN is not available.

Operations with Microsoft Azure virtual machines

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

You can perform failover of Microsoft Azure virtual machines to Acronis Cyber Protect Cloud. For more information, see "Performing a failover" (p. 73).

After that, you can perform failback from Acronis Cyber Protect Cloud back to Azure virtual machines. For more information, see "Failback from Cyber Protect Cloud to an Azure virtual machine" (p. 86).

You can configure a Multisite IPsec VPN connectivity between Acronis Cyber Protect Cloud and the Azure VPN gateway. For more information, see "Configuring Multi-site IPsec VPN" (p. 31).

Limitations when using Geo-redundant cloud storage

Geo-redundant cloud storage provides a secondary location for your backup data. The secondary location is in a region that is geographically distinct from the primary storage location. Geographical separation of regions ensures that - if there is a disaster that affects one of the regions and makes the backup data unrecoverable - the other region will not be affected, and operations will continue.

Important
The Disaster Recovery service is not supported if the backup storage location is switched from the primary location to the geo-redundant secondary location.

Automatic deletion of unused customer environments on the cloud site

The Disaster Recovery service tracks the usage of the customer environments created for disaster recovery purposes and automatically deletes them if they are unused.

The following criteria are used to define if the customer tenant is active:

- Currently, there is at least one cloud server or there were cloud server(s) in the last seven days.
  OR
- The **VPN access to local site** option is enabled and either the Site-to-site Open VPN tunnel is established or there are data reported from the VPN appliance for the last 7 days.

All the rest of the tenants are considered as inactive tenants. For such tenants the system performs the following:

- Deletes the VPN gateway and all cloud resources related to the tenant.
- Unregisters the VPN appliance.

The inactive tenants are rolled back to their state before the connectivity was configured.

Working with Disaster Recovery Cloud

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

The basic workflow for using disaster recovery is the following:

1. Create a recovery server of the workload that you want to protect in one of the following ways:
   - a. Create a protection plan that includes the **Disaster Recovery** module and the **Backup** module with the **What to backup** setting set to **Entire machine** or **System and boot**

volumes.

b. Apply the plan to your devices. This will automatically set up the default disaster recovery infrastructure. For more information, see Create a disaster recovery protection plan.

Note
Unit administrators cannot create, modify, or apply disaster recovery protection plans.

- Set up the disaster recovery cloud infrastructure manually and control each step. See "Creating a recovery server" (p. 54).

2. Configure the connectivity to the cloud site.
   - Cloud-only mode
   - Site-to-site OpenVPN connection
   - Multi-site IPsec VPN connection
   - Point-to-site connection
3. Configure automated test failover.
4. Perform a test failover.
5. [When a disaster occurs] Perform a production failover.
6. [After the disaster] Perform a failback to the local site.
7. [Optional] Configure runbooks.

Creating a disaster recovery protection plan

The disaster recovery protection plan is a protection plan in which the **Disaster Recovery** module is enabled.

After you enable the disaster recovery functionality and apply the plan to your devices, the cloud network infrastructure (cloud site) is created automatically. For more information, see "Default cloud network infrastructure" (p. 15).

Note

- Applying a disaster recovery protection plan creates recovery cloud network infrastructure only if it does not exist. Existing cloud networks are not changed or recreated.
- After you configure disaster recovery, you will be able to perform a test or production failover from any of the recovery points generated after the recovery server was created for the device. Recovery points (backups) that were generated before the device was protected with disaster recovery (before the recovery server was created) cannot be used for failover.
- A disaster recovery protection plan cannot be enabled if the IP address of a device cannot be detected. For example, when virtual machines are backed up agentless and are not assigned an IP address.
- When you apply a protection plan, the same networks and IP addresses are assigned in the cloud site. The IPsec VPN connectivity requires that network segments of the cloud and local sites do not overlap. If a Multi-site IPsec VPN connectivity is configured, and you apply a protection plan to one or several devices later, you must additionally update the cloud networks and reassign the IP addresses of the cloud servers. For more information, see "Reassigning IP addresses" (p. 47).

To create a disaster recovery protection plan

1. In the Cyber Protect console, go to **Devices** > **All devices**.
2. Select the machines that you want to protect.
3. Click **Protect**, and then click **Create plan**.
   The protection plan default settings open.
4. Configure the backup options.
   To use the disaster recovery functionality, the plan must back up the entire machine, or only the disks, required for booting up and providing the necessary services, to a cloud storage.
5. Enable the **Disaster recovery** module by turning on the switch next to the module name.
6. In the **Location** field, select where to create the disaster recovery infrastructure.
7. Click **Create**.
   The plan is created and applied to the selected machines. The default network infrastructure and the recovery servers with default parameters are created. For more information, see "Editing the default settings of a recovery server" (p. 13) and "Default cloud network infrastructure" (p. 15).

What to do next

- You can edit the default configuration of the recovery server. For more information, see "Editing the default settings of a recovery server" (p. 13).
- You can edit the default networking configuration. For more information, see "Connectivity and networks" (p. 16).

Editing the default settings of a recovery server

When you create and apply a disaster recovery protection plan, a recovery server is created with default settings. You can edit these default settings when necessary.

The following procedure applies to recovery servers that are located in Cyber Protect Cloud. If you want to configure the settings of a recovery server that is located in Microsoft Azure, follow the steps that are described in "Creating recovery servers in Microsoft Azure" (p. 110).

Note
A recovery server is created only if it does not exist. Existing recovery servers are not changed or recreated.

To edit the default settings of the recovery server

1. Go to **Devices** > **All devices**.
2. Select a device, and click **Disaster recovery**.
3. Edit the default settings of the recovery server.
   The recovery server settings are described in the following table.

Setting	Default value	Description
CPU and RAM	auto	The number of virtual CPUs and the amount of RAM for the recovery server. The default settings will be automatically determined based on the original device CPU and RAM configuration.
Cloud network	auto	Cloud network to which the server will be connected. For details on how cloud networks are configured, see Cloud network infrastructure.
IP address in production network	auto	The IP address that the server will have in the production network. By default, the IP address of the original machine is set.
Test IP address	disabled	Test IP address gives you the capability to test a failover in the isolated test network and to connect to the recovery server via RDP or SSH during a test failover. In the test failover mode, the VPN gateway will replace the test IP address with the production IP address by using the NAT protocol. If a test IP address is not specified, the console will be the only way to access the server during a test failover.
Internet Access	enabled	Enable the recovery server to access the Internet during a real or test failover. By default, TCP port 25 is denied for outbound connections.
Use Public address	disabled	Having a public IP address makes the recovery server available from the Internet during a failover or test failover. If you do not use a public IP address, the server will be available only in your production network. To use a public IP address, you must enable internet access. The public IP

		address will be shown after you complete the configuration. By default, TCP port 443 is open for inbound connections.
Set RPO threshold	disabled	RPO threshold defines the maximum allowable time interval between the last recovery point and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

Default cloud network infrastructure

The cloud network infrastructure that is created automatically when you apply a disaster recovery protection plan to your workloads consists of the following components:

- A recovery server for each protected device.

  The recovery server is a virtual machine in the cloud that is a copy of the selected device.

  For each of the selected devices, a recovery server with default settings is created in the **Standby** state (virtual machine not running). The recovery server is sized automatically depending on the CPU and RAM of the protected device.
- VPN gateway on the cloud site.
- Cloud networks to which the recovery servers are connected.

The system checks devices IP addresses and if there are no existing cloud networks where an IP address fits, it automatically creates suitable cloud networks. If you already have existing cloud networks where the recovery servers IP addresses fit, the existing cloud networks will not be changed or recreated.

- If you do not have existing cloud networks or you setup disaster recovery configuration for the first time, the cloud networks will be created with maximum ranges recommended by IANA for private use (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16) based on your devices IP address range. You can narrow your network by editing the network mask.
- If you have devices on multiple local networks, the network on the cloud site may become a superset of the local networks. You may reconfigure networks in the **Connectivity** section. See "Managing networks for Site-to-site Open VPN" (p. 24).
- If you need to set up Site-to-site Open VPN connectivity, download the VPN appliance and configure it. See "Configuring Site-to-site Open VPN" (p. 21). Make sure your cloud network ranges match your local network ranges connected to the VPN appliance.
- To change the default network configuration, navigate to **Disaster Recovery** > **Connectivity**, or on the **Disaster Recovery** module of the protection plan, click **Go to connectivity**.

If you revoke, delete, or switch off the **Disaster Recovery** module of a protection plan, the recovery servers and cloud networks will not be deleted automatically. You can remove the disaster recovery infrastructure manually, if necessary.

Connectivity and networks

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

With Disaster Recovery, you can define the following connectivity types to the cloud site:

- **Cloud-only mode**

  This type of connection does not require a VPN appliance deployment on the local site.

  The local and cloud networks are independent networks. This type of connection implies either the failover of all the local site's protected servers or partial failover of independent servers that do not need to communicate with the local site.

  Cloud servers on the cloud site are accessible through the point-to-site VPN, and public IP addresses (if assigned).
- **Site-to-site Open VPN connection**

  This type of connection requires a VPN appliance deployment on the local site.

  The Site-to-site Open VPN connection allows to extend your networks to the cloud and retain the IP addresses.

  Your local site is connected to the cloud site by means of a secure VPN tunnel. This type of connection is suitable in case you have tightly dependent servers on the local site, such as a web server and a database server. In case of partial failover, when one of these servers is recreated on the cloud site while the other stays on the local site, they will still be able to communicate with each other via a VPN tunnel.

  Cloud servers on the cloud site are accessible through the local network, point-to-site VPN, and public IP addresses (if assigned).
- **Multi-site IPsec VPN connection**

  This type of connection requires a local VPN device that supports IPsec IKE v2.

  When you start configuring the Multi-site IPsec VPN connection, Disaster Recovery automatically creates a cloud VPN gateway with a public IP address.

  With Multi-site IPsec VPN, your local sites are connected to the cloud site by means of a secure IPsec VPN tunnel.

  This type of connection is suitable for Disaster Recovery scenarios when you have one or several local sites hosting critical workloads or tightly dependent services.

  In case of partial failover of one of the servers, the server is recreated on the cloud site while the others remain on the local site, and they are still able to communicate with each other through an IPsec VPN tunnel.

  In case of partial failover of one of the local sites, the rest of the local sites remain operational, and will still be able to communicate with each other through an IPsec VPN tunnel.
- **Point-to-site remote VPN access**

  A secure Point-to-site remote VPN access to your cloud and local site workloads from outside by using your endpoint device.

For a local site access, this type of connection requires a VPN appliance deployment on the local site.

Cloud-only mode

The cloud-only mode does not require a VPN appliance deployment on the local site. It implies that you have two independent networks: one on the local site, another on the cloud site. Routing is performed with the router on the cloud site.

How routing works

In case the cloud-only mode is established, routing is performed with the router on the cloud site so that servers from different cloud networks can communicate with each other.

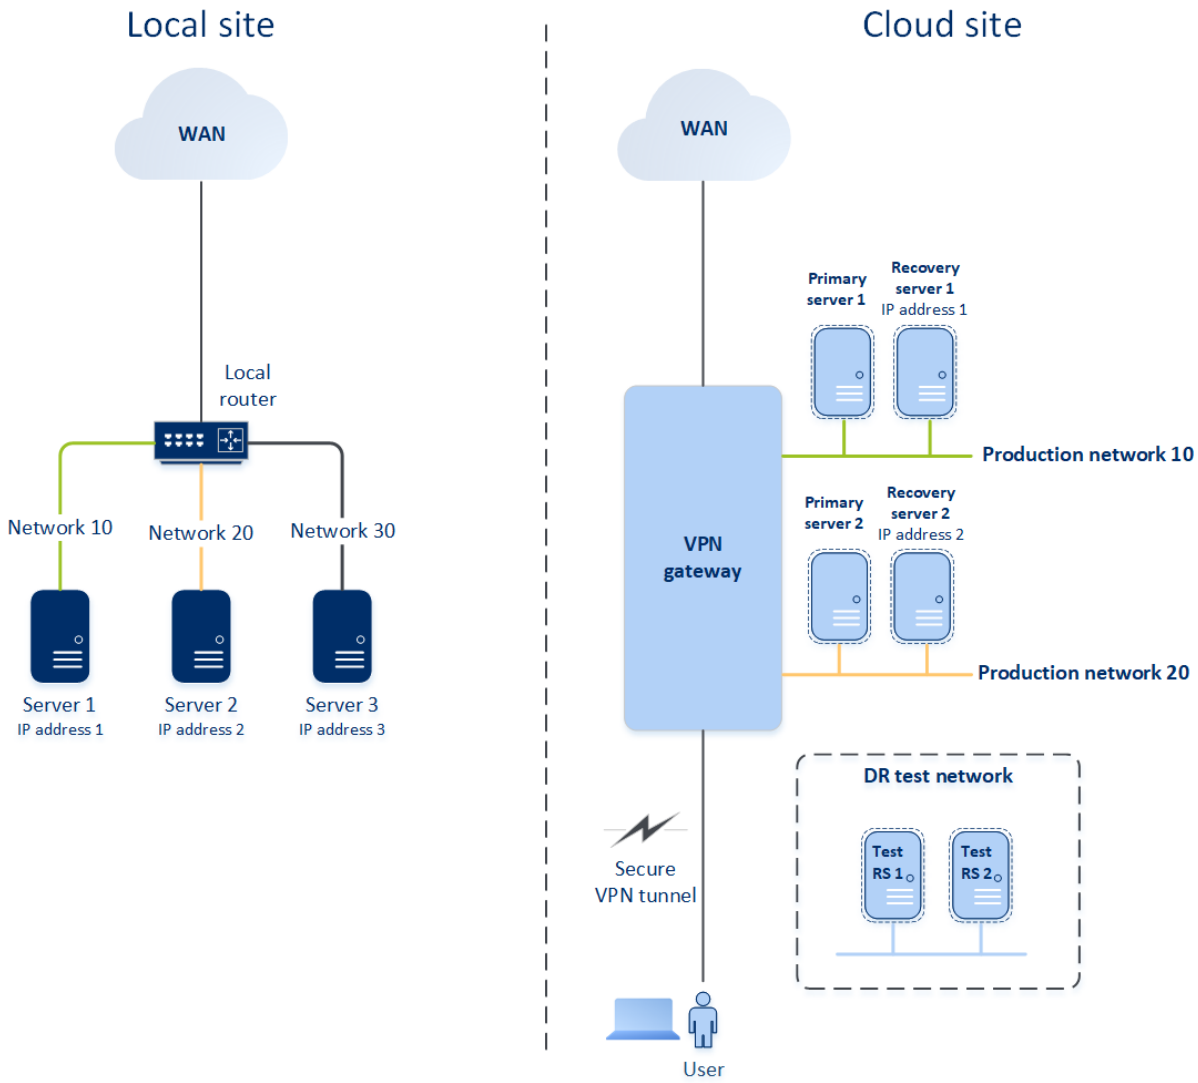

Configuring Cloud-only mode

Cloud-only mode is the default connectivity type that is created automatically when you apply a disaster recovery plan to a workload.

To configure a connection in the Cloud-only mode

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Select **Cloud-only** and click **Configure**.

   As a result, the VPN gateway and cloud network with the defined address and mask are deployed on the cloud site.

Managing networks in Cloud-only mode

You can add and manage up to 23 networks in the cloud.

Add network

To add a new cloud network

1. Go to **Disaster Recovery** > **Connectivity**.
2. On **Cloud site**, click **Add cloud network**.
3. Define the cloud network parameters: the network address and mask, and then click **Done**.

As a result, the additional cloud network with the defined address and mask is created on the cloud site.

Delete network

Prerequisites

All cloud servers are deleted from the network that you want to delete.

To delete a cloud network

1. Go to **Disaster Recovery** > **Connectivity**.
2. On **Cloud site**, click the network address that you want to delete.
3. Click **Delete** and confirm the operation.

Change parameters

To change cloud network parameters

1. Go to **Disaster Recovery** > **Connectivity**.
2. On **Cloud site**, click the network address that you want to edit.
3. Click **Edit.**
4. Define the network address and mask, and click **Done**.

Site-to-site Open VPN connection

Note
The availability of this feature depends on the service quotas that are enabled for your account.

To understand how networking works in Disaster Recovery, we will consider a case when you have three networks with one machine each in the local site. You are going to configure the protection from a disaster for the two networks – Network 10 and Network 20.

On the diagram below, you can see the local site where your machines are hosted, and the cloud site where the cloud servers are launched in case of a disaster.

With Disaster Recovery, you can fail over all the workload from the corrupted machines in the local site to the cloud servers in the cloud.

You can add and manage up to 23 networks in the cloud.

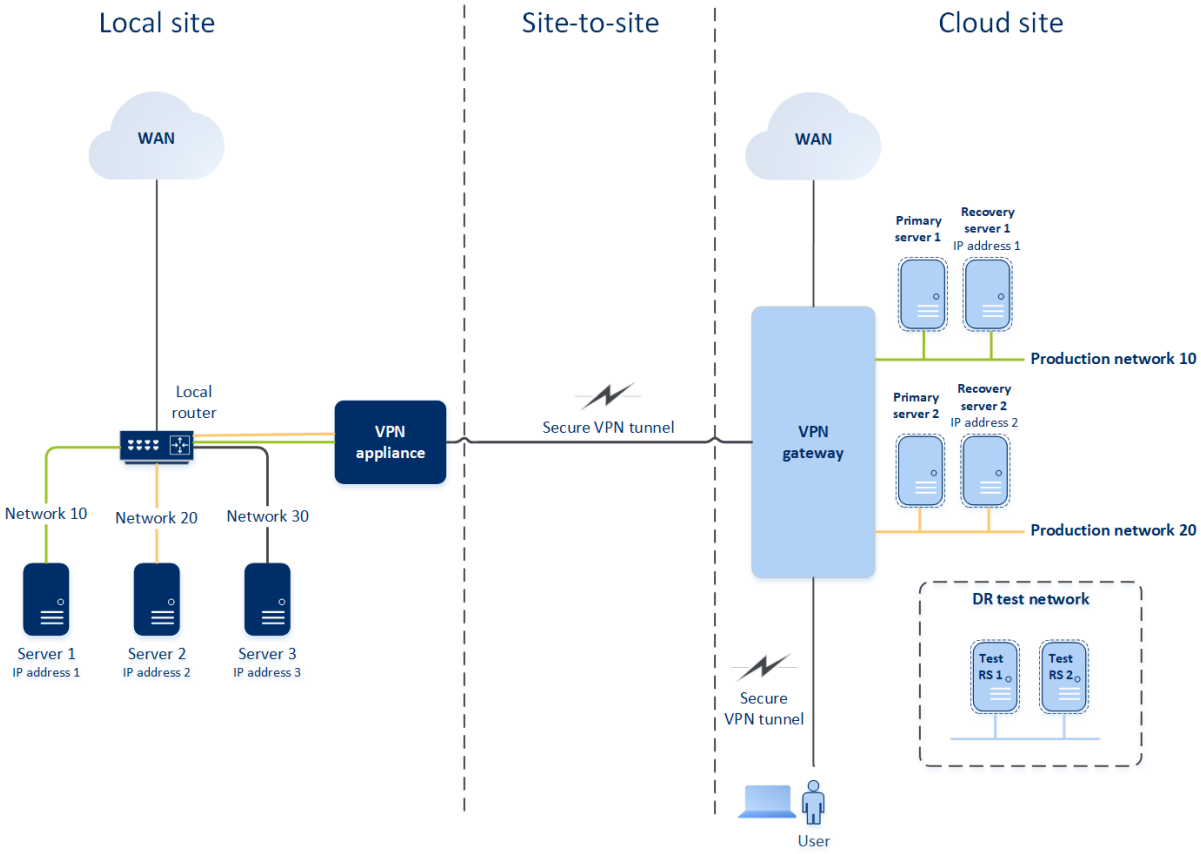

To establish a Site-to-site Open VPN connectivity between the local and cloud sites, a **VPN appliance** and a **VPN gateway** are used.

When you start configuring the Site-to-site Open VPN connectivity in the Cyber Protect console, the VPN gateway is automatically deployed in the cloud site.

After the VPN gateway is deployed, you must do the following:

- Deploy the VPN appliance on your local site.
- Add the networks that you want to be protected.
- Register the VPN appliance in the cloud.

Disaster Recovery will create a replica of your local network in the cloud. A secure VPN tunnel will be established between the VPN appliance and the VPN gateway. This VPN tunnel will provide your local network extension to the cloud. The production networks in the cloud will be bridged with your local networks. The local and cloud servers will communicate through this VPN tunnel as if they are all in the same Ethernet segment. Routing will be performed by your local router.

For each source machine that you want to protect, you must create a recovery server on the cloud site. It will stay in the **Standby** state until a failover event happens. If a disaster happens and you start a failover process (in the **production mode**), the recovery server representing the exact copy of your protected machine will be started in the cloud. It may be assigned the same IP address as the source machine, and it can be started in the same Ethernet segment. Your clients will continue working with the server, not noticing any background changes.

You can also start a failover process in the **test mode**. This means that the source machine will be working and, at the same time, the respective recovery server with the same IP address will be started in the cloud. To prevent IP address conflicts, a special virtual network will be created in the cloud – **test network**. The test network will be isolated to prevent duplication of the source machine IP address in one Ethernet segment. To access the recovery server in the test failover mode, when you create a recovery server, you must assign a **Test IP address** to it. There are other parameters for the recovery server that you can configure, too.

How routing works

When a Site-to-site connection is established, routing between cloud networks is performed by your local router. The VPN server does not perform routing between cloud servers located in different cloud networks. If a cloud server from one network wants to communicate with a server from another cloud network, the traffic goes through the VPN tunnel to the local router on the local site, then the local router routes it to another network, and it goes back through the tunnel to the destination server on the cloud site.

VPN gateway

The VPN gateway is the major component that enables communication between the local and cloud sites. It is a virtual machine in the cloud on which the special software is installed, and the network is specifically configured. The VPN gateway has the following functions:

- Connects the Ethernet segments of your local network and production network in the cloud in the L2 mode.
- Provides iptables and ebtables rules.
- Works as a default router and NAT for the machines in the test and production networks.
- Works as a DHCP server. All machines in the production and test networks get the network configuration (IP addresses, DNS settings) via DHCP. Every time a cloud server will get the same IP

address from the DHCP server. If you need to set up the custom DNS configuration, you should contact the support team.
- Works as a caching DNS.

VPN gateway network configuration

The VPN gateway has several network interfaces:

- External interface, connected to the Internet
- Production interfaces, connected to the production networks
- Test interface, connected to the test network

In addition, two virtual interfaces are added for Point-to-site and Site-to-site connections.

When the VPN gateway is deployed and initialized, the bridges are created – one for the external interface, and one for the client and production interfaces. Though the client-production bridge and the test interface use the same IP addresses, the VPN gateway can route packages correctly by using a specific technique.

VPN appliance

The **VPN appliance** is a virtual machine on the local site with Linux that has special software installed, and a special network configuration. It enables the communication between the local and cloud sites.

Enabling the Site-to-site connectivity

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can enable the Site-to-site connectivity in the following cases:

- If you need the cloud servers on the cloud site to communicate with servers on the local site.
- After a failover to the cloud, the local infrastructure is recovered, and you want to fail back your servers to the local site.

To enable the site-to-site connectivity

1. Go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**, and then enable the **Site-to-site connection** option.

As a result, the site-to-site VPN connection is enabled between the local and cloud sites. The Disaster Recovery service gets the network settings from the VPN appliance and extends the local networks to the cloud site.

Configuring Site-to-site Open VPN

Note
The availability of this feature depends on the service quotas that are enabled for your account.

Requirements for the VPN appliance

System requirements

- 1 CPU
- 1 GB RAM
- 8 GB disk space

Ports

- TCP 443 (outbound) – for VPN connection
- TCP 80 (outbound) – for automatic update of the appliance

Ensure that your firewalls and other components of your network security system allow connections through these ports to any IP address.

Configuring a Site-to-site Open VPN connection

The VPN appliance extends your local network to the cloud through a secure VPN tunnel. This kind of connection is often referred to as a "Site-to-site" (S2S) connection. You can follow the procedure below or watch the video tutorial.

To configure a connection through the VPN appliance

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Select **Site-to-site Open VPN connection**, and click **Configure**.
   The system starts deploying the VPN gateway in the cloud. This will take some time. Meanwhile, you can proceed to the next step.

   Note
   The VPN gateway is provided without additional charge. It will be deleted if the Disaster Recovery functionality is not used, i.e. no primary or recovery server is present in the cloud for seven days.

3. In the **VPN appliance** block, click **Download and deploy**. Depending on the virtualization platform you are using, download the VPN appliance for VMware vSphere or Microsoft Hyper-V.
4. Deploy the appliance and connect it to the production networks.
   In vSphere, ensure that **Promiscuous mode** and **Forged transmits** are enabled and set to **Accept** for all virtual switches that connect the VPN appliance to the production networks. To access these settings, in vSphere Client, select the host > **Summary** > **Network**, and then select the switch > **Edit settings...** > **Security**.
   In Hyper-V, create a **Generation 1** virtual machine with 1024 MB of memory. Also, we recommend that you enable **Dynamic Memory** for the machine. Once the machine is created, go to **Settings** > **Hardware** > **Network Adapter** > **Advanced Features** and select the **Enable MAC address spoofing** check box.
5. Power on the appliance.

6. Open the appliance console and log in with the "admin"/"admin" user name and password.
7. [Optional] Change the password.
8. [Optional] Change the network settings if needed. Define which interface will be used as the WAN for Internet connection.
9. Register the appliance in the Cyber Protection service by using the credentials of the company administrator.

   These credentials are only used once to retrieve the certificate. The data center URL is predefined.

   Note
   If two-factor authentication is configured for your account, you will also be prompted to enter the TOTP code. If two-factor authentication is enabled but not configured for your account, you cannot register the VPN appliance. First, you must go to the Cyber Protect console login page and complete the two-factor authentication configuration for your account. For more details on two-factor authentication, go to the Management Portal Administrator's Guide.

Once the configuration is complete, the appliance will have the **Online** status. The appliance connects to the VPN gateway and starts to report information about networks from all active interfaces to the Disaster Recovery service. The Cyber Protect console shows the interfaces, based on the information from the VPN appliance.

Managing the VPN appliance settings

Note
The availability of this feature depends on the service quotas that are enabled for your account.

On the **Disaster Recovery** > **Connectivity** tab, you can:

- Download log files.
- Unregister the appliance (if you need to reset the VPN appliance settings or switch to the Cloud-only mode).

To access these settings, click the **i** icon in the **VPN appliance** block.

In the VPN appliance console, you can:

- Change the password for the appliance.
- View/change the network settings and define which interface to use as the WAN for the Internet connection.
- Register/change the registration account (by repeating the registration).
- Restart the VPN service.
- Reboot the VPN appliance.
- Run the Linux shell command (only for advanced troubleshooting cases).

Managing networks for Site-to-site Open VPN

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

You can add and manage up to 23 networks in the cloud.

Adding networks

Prerequisites

Site-to-site Open VPN connectivity is configured, as described in "Configuring Site-to-site Open VPN" (p. 21).

To add networks on the local site and extend them to the cloud

1. On the VPN appliance, set up the new network interface with the local network that you want to extend in the cloud.
2. [Optional] If you want to add one or more networks, for each additional network, add one virtual network interface (network adapter) to the virtual machine on which the virtual appliance is running.
   The following example demonstrates the step for a virtual machine that is running on a Hyper-V hypervisor.

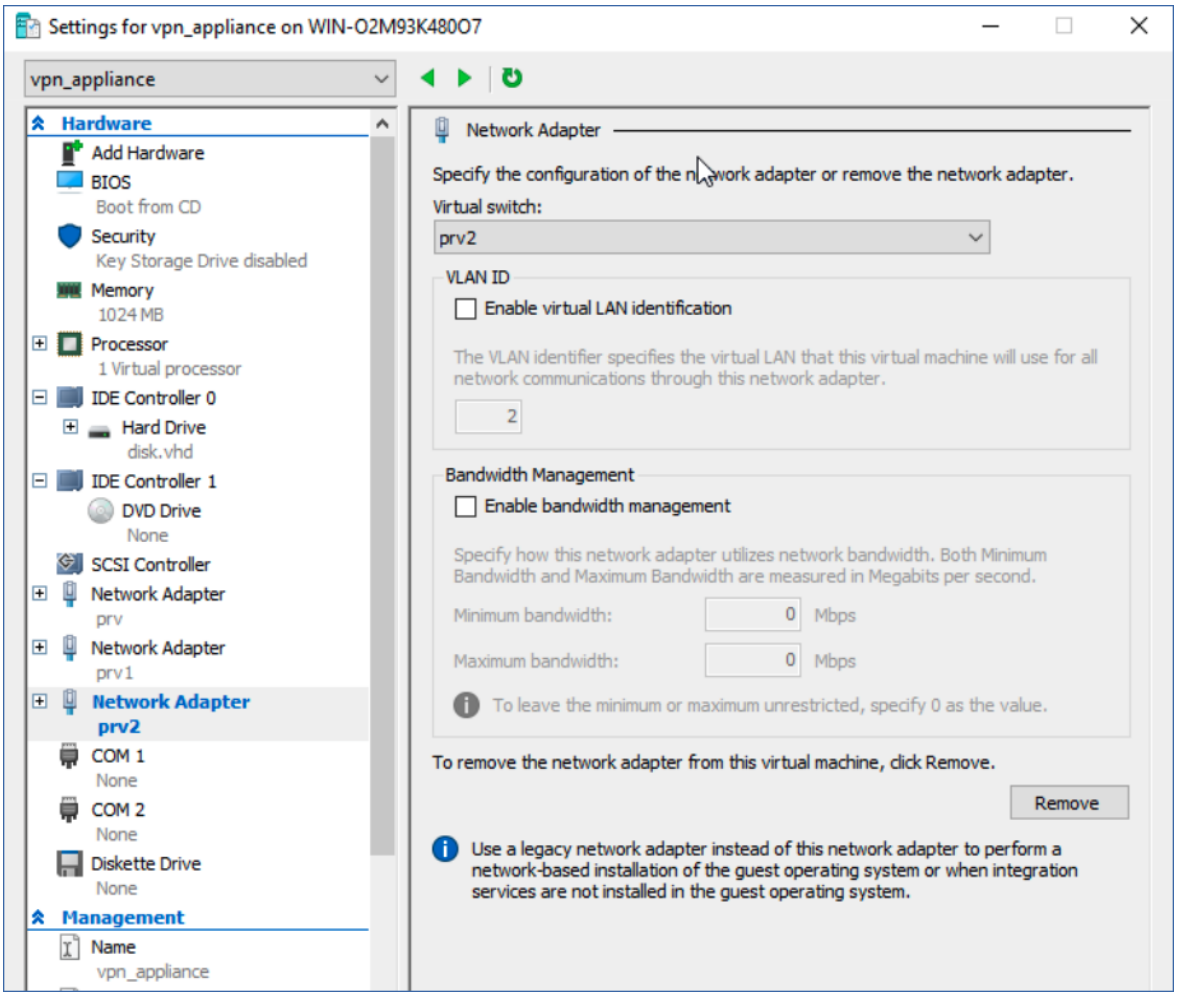

Note

The new virtual network adapters must be configured with the local virtual network that you want to extend to the cloud.

3. Log in to the console of the VPN appliance, and then in the **Networking** section, configure the network settings for one of the interfaces (adapters).

Note

- The IP address configuration is mandatory for only one of the virtual network interfaces, to enable Internet access. You can skip the IP configuration for the other network interfaces.
- Promiscuous mode and Forged transmits or MAC address spoofing must be enabled for each adapter. For more information, see this knowledge base article.

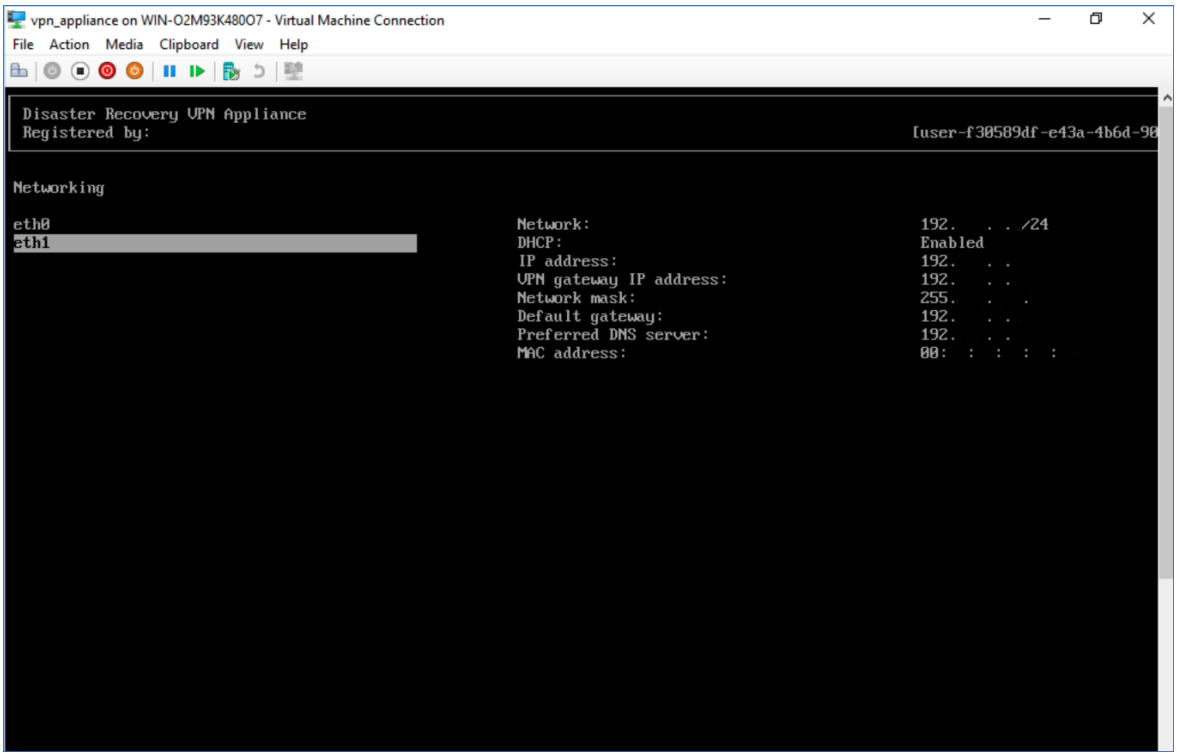

The VPN appliance starts automatically to report information about the networks from all active interfaces to Disaster Recovery.

4. Log in to the Cyber Protect console, and then go to **Disaster recovery** > **Connectivity**.

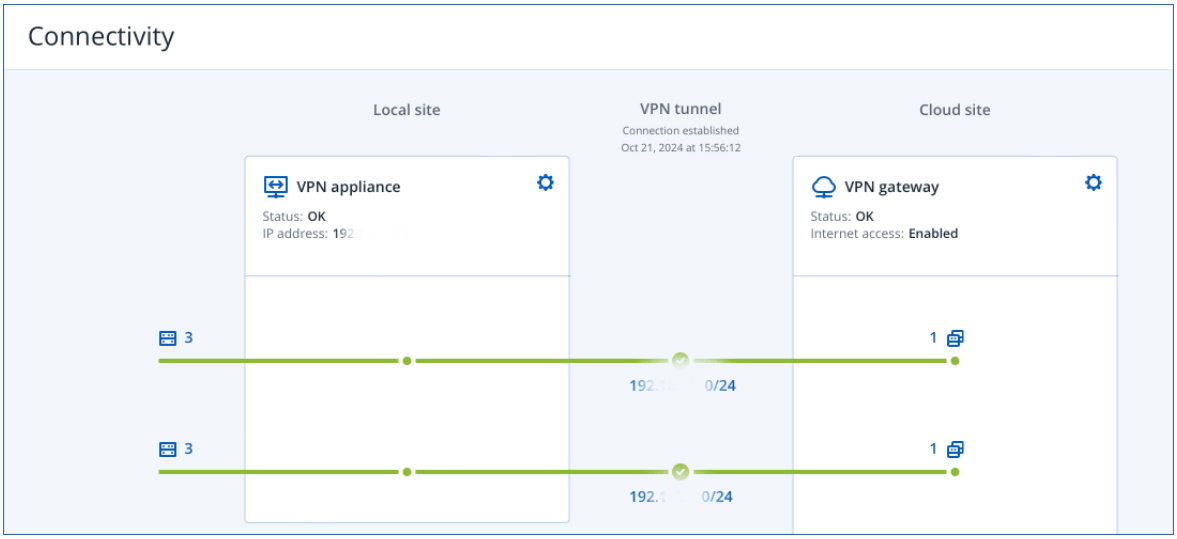

All local networks are automatically extended to the cloud site.

Delete network

To delete a network extended to the cloud

1. Log in to the VPN appliance console.
2. In the **Networking** section, select the interface that you want to delete, and then click **Clear network settings**.
3. Confirm the operation.

As a result, the local network extension to the cloud via a secure VPN tunnel will be stopped. This network will operate as an independent cloud segment. If this interface is used to pass the traffic to or from the cloud site, all of your network connections to or from the cloud site will be disconnected.

Change parameters

To change the network parameters

1. Log in to the VPN appliance console.
2. In the **Networking** section, select the interface that you want to edit.
3. Click **Edit network settings**.
4. Select one of the options:
   - For automatic network configuration via DHCP, click **Use DHCP**, and then confirm the operation.
   - For manual network configuration, click **Set static IP address**, configure the settings, and then click **Enter**.

Setting	Description
IP address	IP address of the interface in the local network.
VPN gateway IP address	Special IP address which is reserved for the cloud segment of network for the proper Disaster Recovery service work.
Network mask	Network mask of the local network.
Default gateway	Default gateway on the local site.
Preferred DNS server	Primary DNS server on the local site.
Alternate DNS server	Secondary DNS server on the local site.

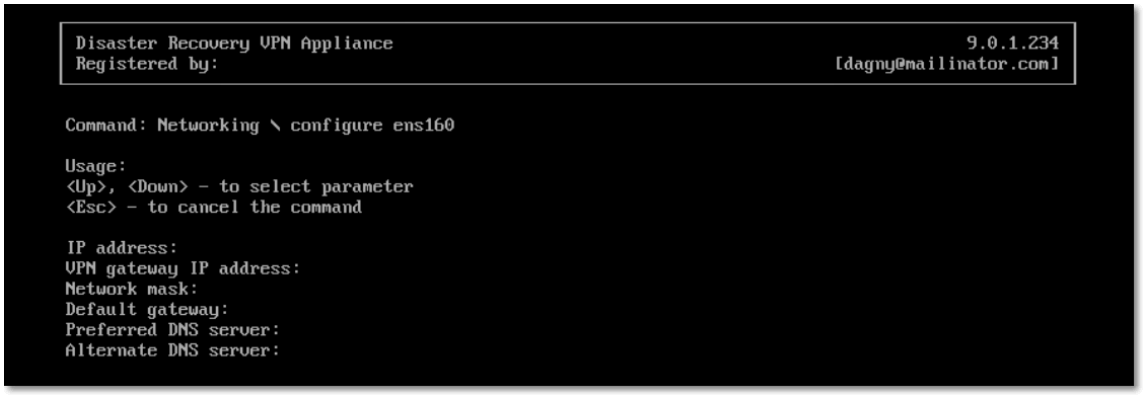

Allowing DHCP traffic over L2 VPN

If devices on your local site get their IP address from a DHCP server, you can protect the DHCP server with Disaster Recovery, fail it over to the cloud, and then allow the DHCP traffic to run over L2

VPN. Thus, your DHCP server will be running in the cloud, but will continue assigning IP addresses to your local devices.

Prerequisites

A Site-to-site L2 VPN connectivity type to the cloud site must be set.

To allow the DHCP traffic via the L2 VPN connection

1. Go to **Disaster Recovery** > **Connectivity** tab.
2. Click **Show Properties**.
3. Enable the **Allow DHCP traffic via L2 VPN** switch.

Switching from Site-to-site Open VPN to Multi-site IPsec VPN

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can easily switch form a Site-to-site Open VPN connection to a Multi-site IPsec VPN connection, and from a Multi-site IPsec VPN connection to a Site-to-site Open VPN connection.

When you switch the connectivity type, the active VPN connections are deleted, but the cloud servers and network configurations are preserved. However, you will still need to reassign the IP addresses of the cloud networks and servers.

The following table compares the basic characteristics of the Site-to-site Open VPN connection and the Multi-site IPsec VPN connection.

	Site-to-site Open VPN	Multi-site IPsec VPN
Local site support	Single	Single, Multiple
VPN Gateway mode	L2 Open VPN	L3 IPsec VPN
Network segments	Extends the local network to the cloud network	Local networks and cloud network segments should not overlap
Supports Point-to-Site access to local site	Yes	No
Supports Point-to-Site access to cloud site	Yes	Yes
Requires a public IP offering item	No	Yes

To switch form a Site-to-site Open VPN connection to a Multi-site IPsec VPN connection

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.

3. Click **Switch to multi-site IPsec VPN**.
4. Click **Reconfigure**.
5. Reassign the IP addresses of the cloud network and cloud servers.
6. Configure the Multi-site IPsec connection settings.

Disabling the Site-to-site connectivity

Note
The availability of this feature depends on the service quotas that are enabled for your account.

If you do not need cloud servers on the cloud site to communicate with servers on the local site, you can disable the Site-to-site connectivity.

To disable the Site-to-site connectivity

1. Go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**, and then disable the **Site-to-site connection** option.

As a result, the local site is disconnected from the cloud site.

Troubelshooting the Site-to-site Open VPN connectivity

When you experience connectivity issues with your Site-to-Site Open VPN site, you can troubleshoot it, and then fix the reported errors or send the information to the Support team for further analysis and assistance.

To troubleshoot the Site-to-site Open VPN connectivity

1. Open the GUI of the VPN appliance.
2. Under **Commands**, select **Troubleshooting**, and then press **Enter**.
3. In the command-line interface, on the **Do you want to run the diagnostic for the site-to-site connection [Y/N]** line, type **Y**, and then press **Enter**.
   The diagnostic tool starts. If an issue is detected, you will see a corresponding error and a detailed description. You can copy the text from the screen or take its screenshot, and then send it to the Support team for assistance.

Multi-site IPsec VPN connection

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can use the Multi-site IPsec VPN connectivity to connect a single local site, or multiple local sites to the Disaster Recovery through a secure L3 IPsec VPN connection.

This connectivity type is useful for Disaster Recovery scenarios if you have one of the following use cases:

- you have one local site hosting critical workloads.
- you have multiple local sites hosting critical workloads. For example, offices in different locations.
- you use third-party software sites or managed service providers sites, and are connected to them through an IPsec VPN tunnel.

To establish a Multi-site IPsec VPN communication between the local sites and the cloud site, a **VPN gateway** is used. When you start configuring the Multi-site IPsec VPN connection in the Cyber Protect console, the VPN gateway is automatically deployed in the cloud site. You should configure the cloud network segments and make sure that they do not overlap with the local network segments. A secure VPN tunnel is established between local sites and the cloud site. The local and cloud servers can communicate through this VPN tunnel as if they are all in the same Ethernet segment.

Note
When using a Multi-site IPsec VPN connection, the VPN gateway is automatically assigned a public IP address.

For each source machine to be protected, you must create a recovery server on the cloud site. It stays in the **Standby** state until a failover event happens. If a disaster happens and you start a failover process (in the **production mode**), the recovery server representing the exact copy of your protected machine is launched in the cloud. Your clients can continue working with the server, without noticing any background changes.

You can also launch a failover process in the **test mode**. This means that the source machine is still working and at the same time the respective recovery server is launched in the cloud in a special virtual network that is created in the cloud – **test network**. The test network is isolated to prevent duplication of IP addresses in the other cloud network segments.

VPN gateway

The major component that allows communication between the local sites and the cloud site is the **VPN gateway**. It is a virtual machine in the cloud on which the special software is installed, and the network is specifically configured. The VPN gateway serves the following functions:

- Connects the Ethernet segments of your local network and production network in the cloud in the L3 IPsec mode.
- Works as a default router and NAT for the machines in the test and production networks.
- Works as a DHCP server. All machines in the production and test networks get the network configuration (IP addresses, DNS settings) via DHCP. Every time a cloud server will get the same IP address from the DHCP server.

  If you prefer, you can set up a custom DNS configuration. For more information, see "Configuring custom DNS servers" (p. 49).
- Works as a caching DNS.

How routing works

Routing between the cloud networks is performed with the router on the cloud site so that servers from different cloud networks can communicate with each other.

Configuring Multi-site IPsec VPN

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can configure a Multi-site IPsec VPN connection in the following two ways:

- from the **Disaster Recovery** > **Connectivity** tab.
- by applying a protection plan on one or several devices, and then manually switching from the automatically created Site-to-site Open VPN connection to a Multi-site IPsec VPN connection, configuring the Multi-site IPsec VPN settings, and reassigning IP addresses.

Connectivity tab

To configure a Multi-site IPsec VPN connection from the Connectivity tab

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. In the **Multi-site VPN connection** section, click **Configure**.
   A VPN gateway is deployed on the cloud site.
3. Configure the Multi-site IPsec VPN settings.

Protection plan

To configure a Multi-site IPsec VPN connection from a protection plan

1. In the Cyber Protect console, go to **Devices**.
2. Apply a protection plan to one or multiple devices from the list.
   The recovery server and the cloud infrastructure settings are automatically configured for Site-to-site Open VPN connectivity.
3. Go to **Disaster Recovery** > **Connectivity**.
4. Click **Show properties**.
5. Click **Switch to Multi-site IPsec VPN**.
6. Configure the Multi-site IPsec VPN settings.
7. Reassign the IP addresses of the cloud network and cloud servers.

Configuring the Multi-site IPsec VPN settings

Note
The availability of this feature depends on the service quotas that are enabled for your account.

After you configure a Multi-site IPsec VPN, you must configure the cloud site and the local sites settings on the **Disaster Recovery** > **Connectivity** tab.

Prerequisites

- Multi-site IPsec VPN connectivity is configured. For more information about configuring the Multi-site IPsec VPN connectivity, see "Configuring Multi-site IPsec VPN" (p. 31).
- Each local IPsec VPN gateway has a public IP address.
- Your cloud network has enough IP addresses for the cloud servers that are copies of your protected machines (in the production network), and for the recovery servers (with one or two IP addresses, depending on your needs).
- [If you use a firewall between the local sites and the cloud site] The following IP protocols and UDP ports are allowed on the local sites: IP Protocol ID 50 (ESP), UDP Port 500 (IKE), and UDP Port 4500.
- The NAT-T configuration on the local sites is disabled.

To configure a Multi-site IPsec VPN connection

1. Add one or more networks to the cloud site.
   a. Click **Add Network**.

      Note
      When you add a cloud network, a corresponding test network is added automatically with the same network address and mask for performing test failovers. The cloud servers in the test network have the same IP addresses as the ones in the cloud production network. If you need to access a cloud server from the production network during a test failover, when you create a recovery server, assign it a second test IP address.

   b. In the **Network address** field, type the IP address of the network.

      Note
      Ensure that the cloud networks do not overlap with any local network in your environment. Otherwise, a tunnel cannot be established.

   c. In the **Network mask** field, type the mask of the network.
   d. Click **Add**.
2. Configure the settings for each local site that you want to connect to the cloud site, following the recommendations for the local sites. For more information about these recommendations, see "General recommendations for local sites" (p. 33).
   a. Click **Add Connection**.
   b. Enter a name for the of the local VPN gateway.
   c. Enter the public IP address of the local VPN gateway.
   d. [Optional] Enter a description of the local VPN gateway.
   e. Click **Next**.

f. In the **Pre-shared key** field, type the pre-shared key, or click **Generate a new pre-shared key** to use an automatically generated value.

Note
You must use the same pre-shared key for the local and the cloud VPN gateways.

g. Click **IPsec/IKE security settings** to configure the settings. For more information about the settings that you can configure, see "IPsec/IKE security settings" (p. 34).

Note
You can use the default settings, which are populated automatically, or use custom values. Only IKEv2 protocol connections are supported. The default **Startup action** when establishing the VPN is **Add** (your local VPN gateway initiates the connection), but you can change it to **Start** (the cloud VPN gateway initiates the connection) or **Route** (suitable for firewalls that support the route options).

h. Configure the **Network policies**.
The network policies specify the networks to which the IPsec VPN connects. Type the IP address and mask of the network using the CIDR format. The local and cloud network segments should not overlap.

i. Click **Save**.

General recommendations for local sites

Note
The availability of this feature depends on the service quotas that are enabled for your account.

When you configure the local sites for your Multi-site IPsec VPN connectivity, consider the following recommendations:

- For each IKE Phase, set at least one of the values that are configured in the cloud site for the following parameters: Encryption algorithm, Hash algorithm, and Diffie-Hellman group numbers.
- Enable Perfect forward secrecy with at least one of the values for Diffie-Hellman group numbers that is configured in the cloud site for IKE Phase 2.
- Configure the same **Lifetime** value for IKE Phase 1 and IKE Phase 2 as in the cloud site.
- Configurations with NAT traversal (NAT-T) are not supported. Disable the NAT-T configuration on the local site. Otherwise, the additional UDP encapsulation cannot be negotiated.
- The **Startup action** configuration defines which side initiates the connection. The default value **Add** means that the local site initiates the connection, and cloud site is waiting for the connection initiation. Change the value to **Start** if you want the cloud site to initiate the connection, or to **Route** if you want both sides to be able to initiate the connection (suitable for firewalls that support the route option).

For more information and configuration examples for different solutions, see:

- This series of knowledge base articles
- This video example

IPsec/IKE security settings

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The following table provides more information about the Psec/IKE security parameters.

Parameter	Description
Encryption algorithm	The encryption algorithm that will be used to ensure that data is not viewable while in transit. By default, all algorithms are selected. You must configure at least one of the selected algorithms on your local gateway device for each IKE phase.
Hash algorithm	The hash algorithm that will be used to verify the data integrity and authenticity. By default, all algorithms are selected. You must configure at least one of the selected algorithms on your local gateway device for each IKE phase.
Diffie-Hellman group numbers	The Diffie-Hellman group numbers define the strength of the key that is used in the Internet Key Exchange (IKE) process. Higher group numbers are more secure but require additional time for the key to compute. By default, all groups are selected. You must configure at least one of the selected groups on your local gateway device for each IKE phase.
Lifetime (seconds)	The lifetime value determines the duration of a connection instance with a set of encryption/authentication keys for user packets, from successful negotiation to expiry. Range for Phase 1: 900-28800 seconds with default 28800. Range for Phase 2: 900-3600 seconds with default 3600. The lifetime for Phase 2 must be less than the lifetime for Phase 1. The connection is re-negotiated through the keying channel before it expires, see **Rekey margin time**.

Parameter	Description
	If the local and the remote side do not agree on the lifetime, a clutter of superseded connections will occur on the side with the longer lifetime. See also **Rekey margin time** and **Rekey fuzz**.
Rekey margin time (seconds)	The margin time before connection expiration or keying-channel expiration, during which the local side of the VPN connection attempts to negotiate a replacement. The exact time of the rekey is randomly selected based on the value of **Rekey fuzz**. Relevant only locally, the remote side does not need to agree on it. Range: 900-3600 seconds. The default value is 3600.
Replay window size (packet)	The IPsec replay window size for this connection. The default -1 uses the value configured with charon.replay_window in the strongswan.conf file. Values larger than 32 are supported only when using the Netlink backend. A value of 0 disables the IPsec replay protection.
Rekey fuzz (%)	The maximum percentage by which marginbytes, marginpackets and margintime are randomly increased to randomize rekeying intervals (important for hosts with many connections). The Rekey fuzz value can exceed 100%. The value of marginTYPE, after the random increase, must not exceed lifeTYPE, where TYPE is one of bytes, packets or time. The value 0% disables randomization. Relevant only locally, the remote side does not need to agree on it.
DPD timeout (seconds)	Time after which a dead peer detection (DPD) timeout occurs. You can specify value 30 or higher. The default value is 30.
Dead peer detection (DPD) timeout action	The action to take after a dead peer detection (DPD) timeout occurs. **Restart** - Restart the session when DPD timeout occurs. **Clear** - End the session when DPD timeout occurs. **None** - Take no action when DPD timeout occurs.

Parameter	Description
Startup action	Determines which side initiates the connection and establishes the tunnel for the VPN connection. **Add** - your local VPN gateway initiates the connection. **Start** - the cloud VPN gateway initiates the connection. **Route** - suitable for VPN gateways that support the route option. The tunnel is up only when there is traffic initiated from either the local VPN gateway, or the cloud VPN gateway.

Switching from Multi-site IPsec VPN to Site-to-site Open VPN

You can easily switch form a Multi-site IPsec VPN connection to a Site-to-site Open VPN connection.

When you switch the connectivity type, the active VPN connections are deleted, but the cloud servers and network configurations are preserved. However, you will still need to reassign the IP addresses of the cloud networks and servers.

The following table compares the basic characteristics of the Site-to-site Open VPN connection and the Multi-site IPsec VPN connection.

	Site-to-site Open VPN	Multi-site IPsec VPN
Local site support	Single	Single, Multiple
VPN Gateway mode	L2 Open VPN	L3 IPsec VPN
Network segments	Extends the local network to the cloud network	Local networks and cloud network segments should not overlap
Supports Point-to-Site access to local site	Yes	No
Supports Point-to-Site access to cloud site	Yes	Yes
Requires a public IP offering item	No	Yes

To switch from a Multi-site IPsec VPN connection to a Site-to-site Open VPN connection

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.
3. Click **Switch to site-to-site Open VPN**.
4. Click **Reconfigure**.

5. Reassign the IP addresses of the cloud network and cloud servers.
6. Configure the Site-to-site connection settings.

Troubleshooting the IPsec VPN configuration

Note
The availability of this feature depends on the service quotas that are enabled for your account.

When you configure or use the IPsec VPN connection, you might experience problems.

You can learn more about the problems that you encountered in the IPsec log files, and check the Troubleshooting IPsec VPN configuration issues topic for possible solutions of some of the common problems that might occur.

Troubleshooting IPsec VPN configuration issues

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The following table describes the IPsec VPN configuration problems that occur most often, and explains how to troubleshoot them.

Problem	Possible solution
I see the following error message: **IKE phase 1 negotiation error. Check the IPsec IKE settings on the Cloud and the Local sites.**	Click **Retry** and check if a more specific error message appears. For example, a more specific error message may be an error message about an algorithm mismatch or an incorrect Pre-shared key. **Note** For security reasons, the following restrictions apply to the IPsec VPN connectivity: • IKEv1 is called for deprecation in RFC8247 and is not supported due to security risks. Only IKEv2 protocol connections are supported. • The following Encryption algorithms are not considered secure and are not supported: DES, and 3DES. • The following Hash algorithms are not considered secure and are not supported: SHA1, and MD5. • Diffie-Hellman group number 2 is not considered secure and is not supported.
The connection between my local site	Check:

Problem	Possible solution
and the cloud site stays in status **Connecting**.	• If the UDP port 500 is open (when you use a firewall). • The connectivity between the local site and the cloud site. • If the IP address of the local site is correct.
The connection between my local site and the cloud site stays in status **Waiting for a connection**.	You see this status when the **Startup action** for cloud site is set to **Add**, which means that the cloud site is waiting for the local site to initiate the connection. Initiate connection from the local site.
The connection between my local site and the cloud site stays in status **Waiting for traffic**.	You see this status when the **Startup action** for cloud site is set to **Route**. If you are expecting a connection from the local site, do the following: • From the local site, try to ping the virtual machine in the cloud site. This is a standard behavior necessary for establishing a tunnel for some devices, for example Cisco ASA. (Route mode) • Ensure that the local site established a tunnel by setting the **Startup action** of the local site to **Start**.
The connection between my local site and the cloud site is established, but I can see that one or more of the network policies are down.	This issue may be due to the following reasons: • Network mapping in the cloud IPsec site is different from the network mapping in the local site. Ensure that the network mappings and the sequence of the network policies in the local and cloud sites match exactly. • This state is correct when the **Startup action** of the local site and/or of the cloud site is set to **Route** (for example, on Cisco ASA devices), and currently there is no traffic. You can try to ping to make sure that the tunnel is established. If the ping is not working, check the network mapping on the local and the cloud site.
I want restart a specific IPsec connection.	To restart a specific IPsec connection: 1. In the **Disaster recovery** > **Connectivity** screen, click the IPsec connection. 2. Click **Disable connection**.

Problem	Possible solution
	3. Click the IPsec connection again. 4. Click **Enable connection**.

Downloading the IPsec VPN log files

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can find additional information about the IPsec connectivity in the log files on the VPN server. The log files are compressed in a .zip archive that you can download and extract.

Prerequisites

Multi-site IPsec VPN connectivity is configured.

To download the .zip archive with the log files

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click the gear icon next to the VPN gateway of the cloud site.
3. Click **Download log**.
4. Click **Done**.
5. When the .zip archive is ready for download, click **Download log**, and save it locally.

Multi-site IPSec VPN log files

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The following list describes the IPsec VPN log files that are part of the zip archive, and the information that they contain.

- `ip.txt` - The file contains the logs from the configuration of the network interfaces. You must see two IP addresses - a public IP address, and a local IP address. If you do not see these IP addresses in the log, there is a problem. Contact the Support team.

  Note
  The mask for the public IP address must be 32.

- `swanctl-list-loaded-config.txt` - The file contains information about all IPsec sites.
  If you do not see a site in the file, then the IPsec configuration was not applied. Try to update the configuration and save it, or contact the Support team.
- `swanctl-list-active-sas.txt` - The file contains connections and policies that are in status active or a connecting.

Point-to-site remote VPN access

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The Point-to-site connection is a secure connection from the outside by using your endpoint devices (such as computer or laptop) to the cloud and local sites through a VPN. It is available after you establish a Site-to-site Open VPN connection to the Disaster Recovery site. This type of connection is useful in the following cases:

- In many companies, the corporate services and web resources are available only from the corporate network. You can use the Point-to-site connection to securely connect to the local site.
- In case of a disaster, when a workload is switched to the cloud site and your local network is down, you may need direct access to your cloud servers. This is possible through the Point-to-site connection to the cloud site.

For the Point-to-site connection to the local site, you need to install the VPN appliance on the local site, configure the Site-to-site connection, and then the Point-to-site connection to the local site. Thus, your remote employees will have access to the corporate network through L2 VPN.

The scheme below shows the local site, cloud site, and communications between servers highlighted in green. The L2 VPN tunnel connects your local and cloud sites. When a user establishes a Point-to-site connection, the communications to the local site are performed through the cloud site.

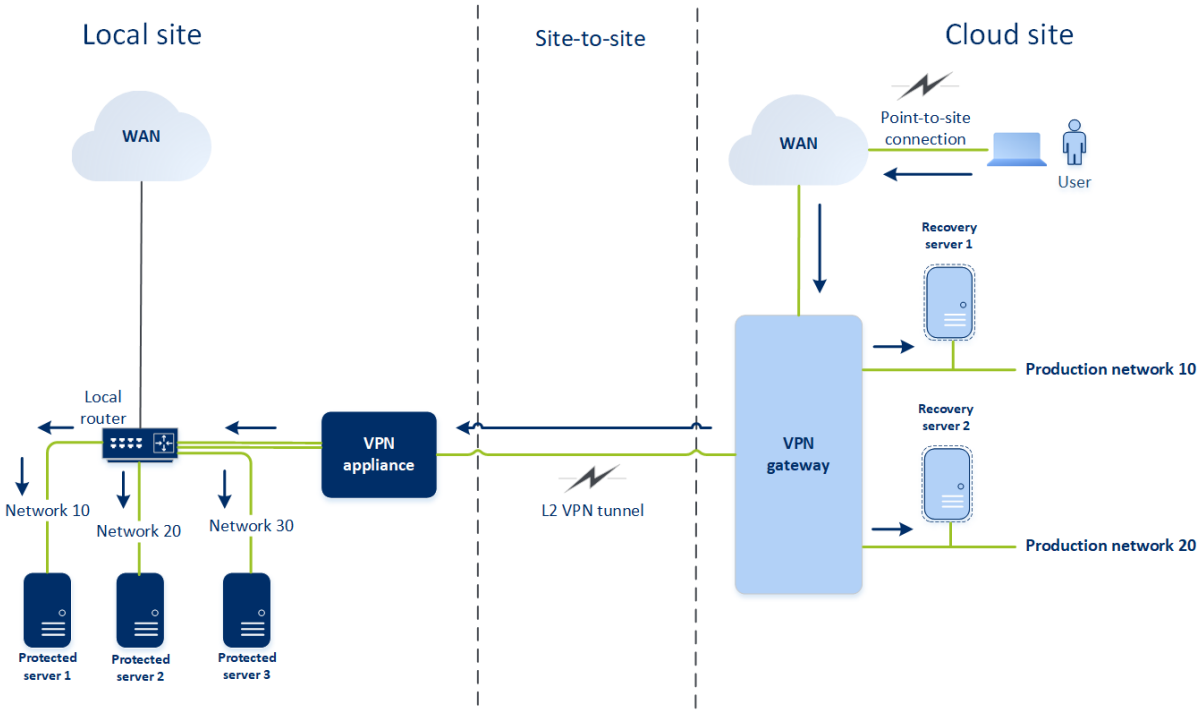

The Point-to-site configuration uses certificates to authenticate to the VPN client. Additionally user credentials are used for authentication. Note the following about the Point-to-site connection to the local site:

- Users should use their Cyber Protect Cloud credentials to authenticate in the VPN client. They must have either a "Company Administrator" or a "Cyber Protection" user role.
- If you re-generated the OpenVPN configuration, you need to provide the updated configuration to all of the users using the Point-to-site connection to the cloud site.

Configuring Point-to-site remote VPN access

Note
The availability of this feature depends on the service quotas that are enabled for your account.

If you need to connect to your local site remotely, you can configure the Point-to-site connection to the local site. You can follow the procedure below or watch the video tutorial.

Prerequisites

- Site-to-site Open VPN connectivity is configured.
- The VPN appliance is installed on the local site.

To configure the Point-to-site connection to the local site

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.
3. Enable the **VPN access to local site** option.
4. Ensure that your user who needs to establish the Point-to-site connection to the local site has:
   - a user account in Cyber Protect Cloud. These credentials are used for authentication in the VPN client. Otherwise, create a user account in Cyber Protect Cloud.
   - a "Company Administrator" or "Cyber Protection" user role.
5. Configure the OpenVPN client:
   a. Download the OpenVPN client version 2.4.0 or later from the following location https://openvpn.net/community-downloads/.

      Note
      OpenVPN Connect client is not supported.

   b. Install the OpenVPN client on the machine from which you want to connect to the local site.
   c. Click **Download configuration for OpenVPN**. The configuration file is valid for users in your organization with the "Company Administrator" or "Cyber Protection" user role.
   d. Import the downloaded configuration to the OpenVPN client.
   e. Log in to the OpenVPN client with your Cyber Protect Cloud user credentials (see step 4 above).

f. [Optional] If two-factor authentication is enabled for your organization, then you should provide the one-time generated TOTP code.

Important
If you enabled two-factor authentication for your account, you need to re-generate the configuration file and renew it for your existing OpenVPN clients. Users must re-log in to Cyber Protect Cloud to set up two-factor authentication for their accounts.

As a result, you will be able to connect to machines on the local site.

Managing point-to-site connection settings

Note
The availability of this feature depends on the service quotas that are enabled for your account.

In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity** and then click **Show properties** in the upper right corner.

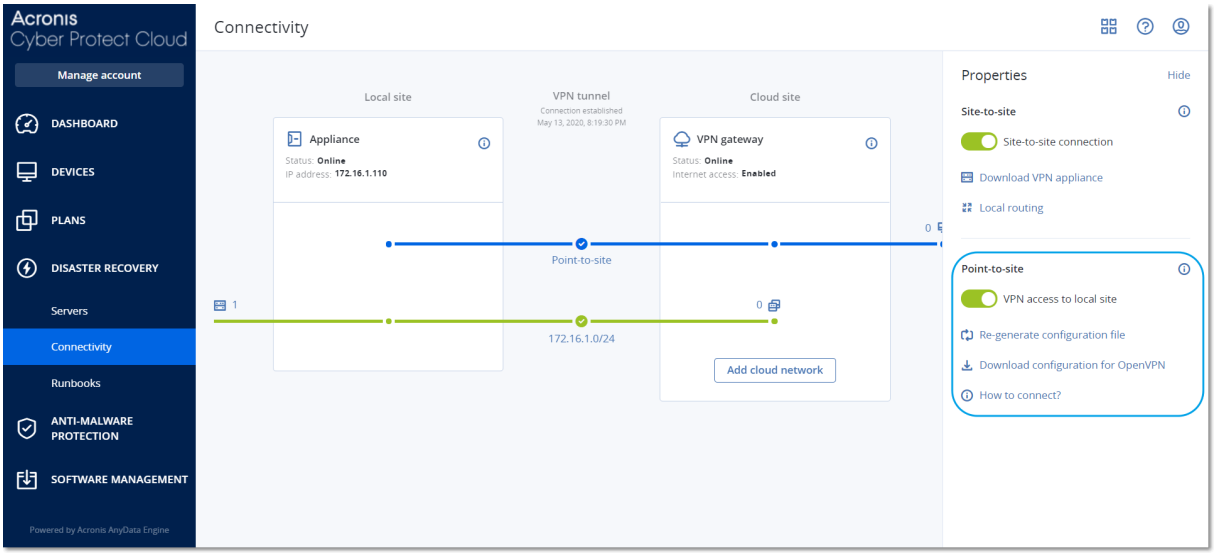

VPN access to local site

This option is used for managing VPN access to the local site. By default it is enabled. If it is disabled, then the Point-to-site access to the local site will be not allowed.

Download configuration for OpenVPN

This will download the configuration file for the OpenVPN client. The file is required to establish a Point-to-site connection to the cloud site.

Re-generate configuration

You can re-generate the configuration file for the OpenVPN client.

This is required in the following cases:

- If you suspect that the configuration file is compromised.
- If two-factor authentication was enabled for your account.

As soon as the configuration file is updated, connecting by means of the old configuration file becomes not possible. Make sure to distribute the new file among the users who are allowed to use the Point-to-site connection.

Active point-to-site connections

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can view all active point-to-site connections in **Disaster recovery** > **Connectivity**. Click the machine icon on the blue **Point-to-site** line and you will see the detailed information about active point-to-site connections grouped by the user name.

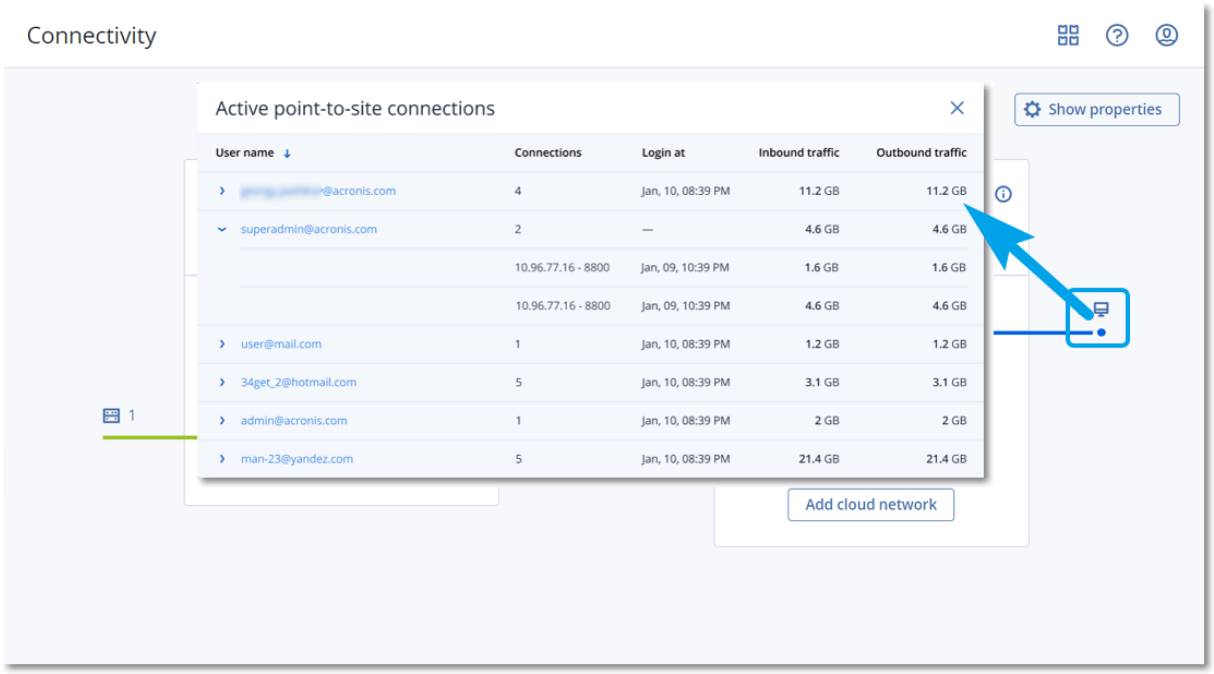

Recommendations for the Active Directory Domain Services availability

If your protected workloads need to authenticate in a domain controller, we recommend that you have an Active Directory Domain Controller (AD DC) instance at the Disaster Recovery site.

Active Directory Domain Controller for L2 Open VPN connectivity

With the L2 Open VPN connectivity, the IP addresses of the protected workloads are retained in the cloud site during a test failover or a production failover. Therefore, the AD DC during a test failover or a production failover has the same IP address as in the local site.

With custom DNS you can set your own custom DNS server for all cloud servers. For more information, see "Configuring custom DNS servers" (p. 49).

Active Directory Domain Controller for L3 IPsec VPN connectivity

With L3 IPsec VPN connectivity, the IP addresses of the protected workloads are not retained in the cloud site. Therefore, we recommend that you have an additional dedicated AD DC instance as a primary server in the cloud site before you perform a production failover.

The recommendations for a dedicated AD DC instance that is configured as a primary server in the cloud site are the following:

- Turn off Windows firewall.
- Join the primary server to the Active Directory service.
- Ensure that the primary server has Internet access.
- Add the Active Directory feature.

With custom DNS you can set your own custom DNS server for all cloud servers. For more information, see "Configuring custom DNS servers" (p. 49).

Network management

This section describes network management scenarios.

Public and test IP addresses

If you assign the public IP address when creating a recovery server, the recovery server becomes available from the Internet through this IP address. When a packet comes from the Internet with the destination public IP address, the VPN gateway remaps it to the respective production IP address by using NAT, and then sends it to the corresponding recovery server.

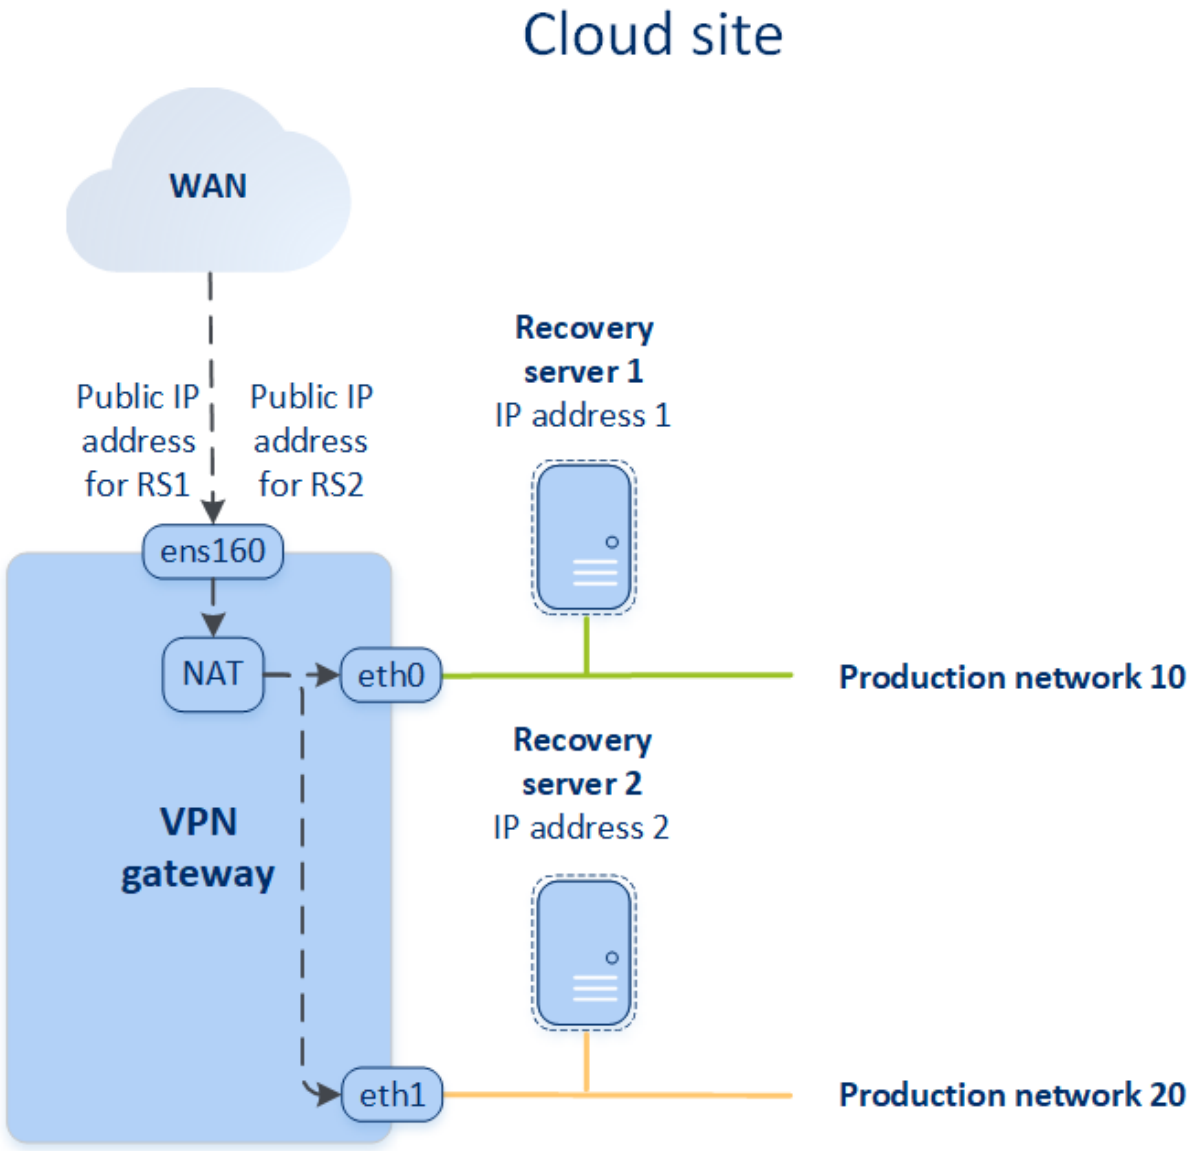

If you assign the test IP address when creating a recovery server, the recovery server becomes available in the test network through this IP address. When you perform the test failover, the original machine is still running while the recovery server with the same IP address is launched in the test network in the cloud. There is no IP address conflict, as the test network is isolated. The recovery servers in the test network are reachable by their test IP addresses, which are remapped to the production IP addresses through NAT.

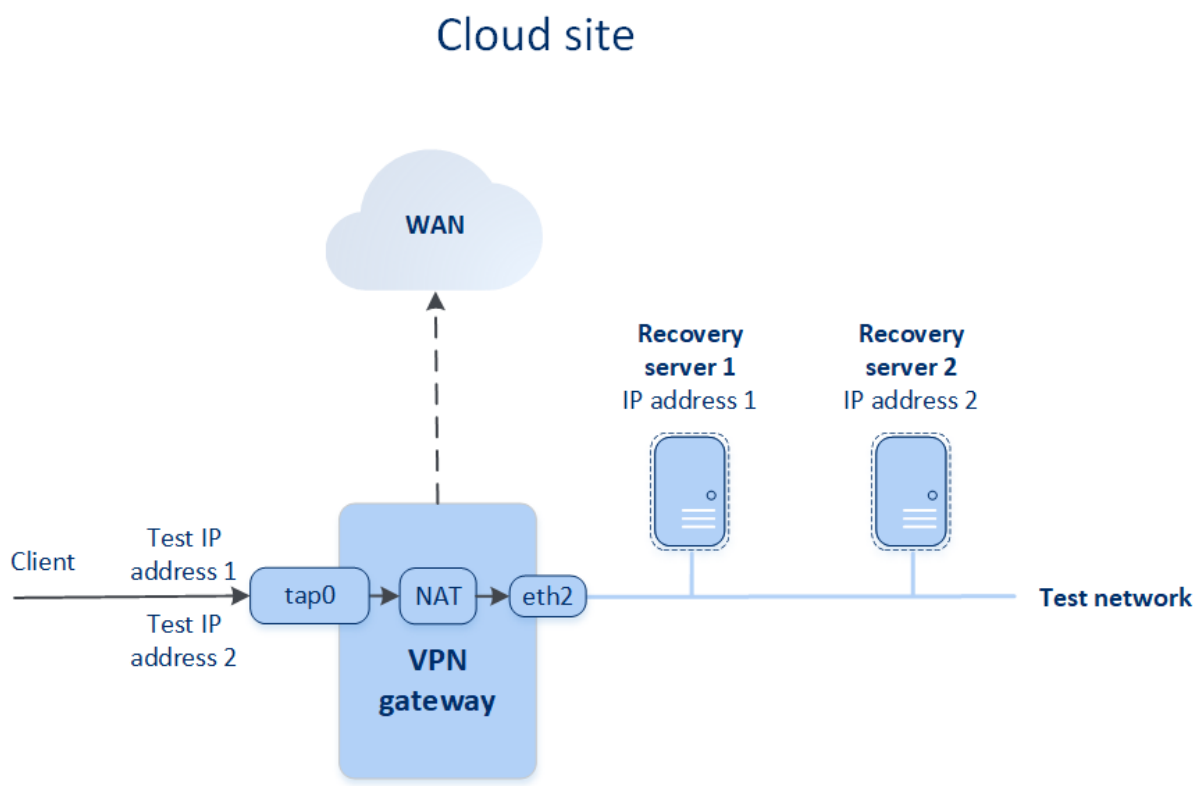

For more information about Site-to-site Open VPN, see "Site-to-site Open VPN - Additional information" (p. 143).

IP address reconfiguration

For proper disaster recovery performance, the IP addresses assigned to the local and cloud servers must be consistent. If there is any inconsistency or mismatch in IP addresses, you will see the exclamation mark next to the corresponding network in **Disaster Recovery** > **Connectivity**.

Some of the commonly known reasons for IP address inconsistency are listed below:

1. A recovery server was migrated from one network to another, or the network mask of the cloud network was changed. As a result, cloud servers have the IP addresses from networks to which they are not connected.
2. The connectivity type was switched from one without Site-to-site connection to a Site-to-site connection. As a result, a local server is placed in the network different from the one that was created for the recovery server on the cloud site.
3. The connectivity type was switched from Site-to-site Open VPN to Multi-site IPsec VPN, or from Multi-site IPsec VPN to Site-to-site Open VPN. For more information about this scenarios, see Switching connections, "Switching from Multi-site IPsec VPN to Site-to-site Open VPN" (p. 36), and Reassigning IP addresses.
4. Editing the following network parameters on the VPN appliance site:
   - Adding an interface via the network settings
   - Editing the network mask manually via the interface settings

- Editing the network mask via DHCP
- Editing the network address and mask manually via the interface settings
- Editing the network mask and address via DHCP

As a result of the actions listed above, the network on the cloud site may become a subset or superset of the local network, or the VPN appliance interface may report the same network settings for different interfaces.

To resolve the issue with network settings

1. Click the network that requires IP address reconfiguration.
   You will see a list of servers in the selected network, their status, and IP addresses. The servers whose network settings are inconsistent are marked with an exclamation mark.
2. To change network settings for a server, click **Go to server**. To change network settings for all servers at once, click **Change** in the notification block.
3. Change the IP addresses as needed by defining them in the **New IP** and **New test IP** fields.
4. When ready, click **Confirm**.

Move servers to a suitable network

When you create a disaster recovery protection plan and apply it on selected devices, the system checks device IP addresses and automatically creates cloud networks if there are not existing cloud networks where IP address fits. By default, the cloud networks are configured with maximum ranges recommended by IANA for private use (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16). You can narrow your network by editing the network mask.

In the case that the selected devices were on multiple local networks, the network on the cloud site might become a superset of the local networks. In this case, to reconfigure cloud networks:

1. Click the cloud network that requires network size reconfiguration and then click **Edit**.
2. Reconfigure the network size with the correct settings.
3. Create other required networks.
4. Click the notification icon next to the number of devices connected to the network.
5. Click **Move to a suitable network**.
6. Select the servers that you want to move to suitable networks and then click **Move**.

Reassigning IP addresses

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You must reassign the IP addresses of the cloud networks and the cloud servers in order to complete the configuration in the following cases:

- After you switch from Site-to-site Open VPN to Multi-site IPsec VPN, or the opposite.
- After you apply a protection plan (if the Multi-site IPsec VPN connectivity is configured).

Cloud network

To reassign the IP address of a cloud network

1. In the **Connectivity** tab, click the IP address of the cloud network.
2. In the **Network** pop-up, click **Edit**.
3. Type the new the network address and network mask.
4. Click **Done**.

After you reassign the IP address of a cloud network, you must reassign the cloud servers that belong to the reassigned cloud network.

Cloud server

To reassign the IP address of a server

1. In the **Connectivity** tab, click the IP address of the server in the cloud network.
2. In the **Servers** pop-up, click **Change IP address**.
3. In the **Change IP address** pop-up, type the new IP address of the server, or use the automatically generated IP address which is part of the reassigned cloud network.

   Note
   Disaster Recovery automatically assigns IP addresses from the cloud network to all cloud servers that were part of the cloud network before the reassignment of the network IP address. You can use the suggested IP addresses to reassign the IP addresses of all the cloud servers at once.

4. Click **Confirm**.

Reinstalling the VPN gateway

If there is an issue with the VPN gateway which you cannot resolve, you might want to reinstall the VPN gateway. Possible issues include the following:

- The VPN gateway is in **Error** status.
- The VPN gateway is in **Pending** status for a long time.
- The VPN gateway status is undetermined for a long time.

Reinstalling the VPN gateway process includes the following automatic actions: deleting the existing VPN gateway virtual machine completely, installing a new virtual machine from the template, and applying the settings of the previous VPN gateway on the new virtual machine.

Prerequisites:

One of the connectivity types to the cloud site must be set.

To reinstall the VPN gateway

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click the gear icon of the VPN gateway, and select **Reinstall VPN gateway**.
3. In the **Reinstall VPN gateway** dialog, enter your login.
4. Click **Reinstall**.

Configuring custom DNS servers

Note
The availability of this feature depends on the service quotas that are enabled for your account.

When you configure a connectivity, Disaster Recovery creates your cloud network infrastructure. The cloud DHCP server automatically assigns default DNS servers to the recovery servers and primary servers, but you can change the default settings and configure custom DNS servers. The new DNS settings will be applied at the time of the next request to the DHCP server.

Prerequisites

One of the connectivity types to the cloud site must be set.

To configure a custom DNS server

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.
3. Click **Default (Provided by Cloud Site)**.
4. Select **Custom servers**.
5. Type the IP address of the DNS server.
6. [Optional] If you want to add another DNS server, click **Add**, and type the DNS server IP address.

   Note
   After you add the custom DNS servers, you can also add the default DNS servers. In that way, if the custom DNS servers are unavailable, Disaster Recovery will use the default DNS servers.

7. Click **Done**.

Deleting custom DNS servers

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can delete DNS servers from the custom DNS list.

Prerequisites:

Custom DNS servers are configured.

To delete a custom DNS server

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.
3. Click **Custom servers**.
4. Click the delete icon next to the DNS server.

   Note
   The delete operation is disabled when only one custom DNS server is available. If you want to delete all custom DNS servers, select **Default (provided by Cloud Site)**.

5. Click **Done**.

Configuring local routing

In addition to your local networks that are extended to the cloud through the VPN appliance, you may have other local networks that are not registered in the VPN appliance but have servers which need to communicate with cloud servers. To establish the connectivity between such local servers and cloud servers, you need to configure the local routing settings.

To configure local routing

1. Go to **Disaster Recovery**>**Connectivity**.
2. Click **Show properties**, and then click **Local routing**.
3. Specify the local networks in the CIDR notation.
4. Click **Save**.

As a result, the servers from the specified local networks can communicate with the cloud servers.

Downloading MAC addresses

You can download a list of MAC addresses, and then extract them and import them in the configuration of your custom DHCP server.

Prerequisites:

- One of the connectivity types to the cloud site must be set.
- At least one primary or recovery server with a MAC address must be configured.

To download the list of MAC addresses

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.
3. Click **Download the list of MAC addresses**, and then save the CSV file.

Working with logs

Disaster Recovery collects logs for the VPN appliance and the VPN gateway. The logs are saved as .txt files, which are compressed in a .zip archive. You can download and extract the archive, and use the information for troubleshooting or monitoring purposes.

The following list describes the log files that are part of the .zip archive, and the information that they contain.

`dnsmasq.config.txt` - The file contains information about the configuration of the service that provides DNS and DHCP addresses.

`dnsmasq.leases.txt` - The file contains information about the current DHCP address leases.

`dnsmasq_log.txt` - The file contains logs of the dnsmasq service.

`ebtables.txt` - The file contains information about the firewall tables.

`free.txt` - The file contains information about the free memory.

`ip.txt` - The file contains the logs from the configuration of the network interfaces, including their names which can be used in the configuration of the **Capturing network packets** settings.

`NetworkManager_log.txt` - The file contains logs from the NetworkManager service.

`NetworkManager_status.txt` - The file contains information about the status of the NetworkManager service.

`openvpn@p2s_log.txt` - The file contains logs from the OpenVPN service.

`openvpn@p2s_status.txt` - The file contains information about the status of the VPN tunnels.

`ps.txt` - The file contains information about the currently running processes on the VPN gateway or VPN appliance.

`resolf.conf.txt` - The file contains information about the configuration of the DNS servers.

`routes.txt` - The file contains information about the networking routes.

`uname.txt` - The file contains information about the current version of the kernel of the operating system.

`uptime.txt` - The file contains information about the length of period for which the operating system has not been restarted.

`vpnserver_log.txt` - The file contains logs from the VPN service.

`vpnserver_status.txt` - The file contains information about the status of the VPN server.

For more information about log files that are specific to the IPsec VPN connectivity, see "Multi-site IPSec VPN log files" (p. 39).

Downloading the logs of the VPN appliance

You can download and extract the archive that contains the logs of the VPN appliance, and use the information for troubleshooting or monitoring purposes.

To download the logs of the VPN appliance

1. On the **Connectivity** page, click the gear icon next to the VPN appliance.
2. Click the **Download log**.

3. [Optional] Select **Capture network packets**, and configure the settings. For more information, see "Capturing network packets" (p. 52).
4. Click **Done**.
5. When the .zip archive is ready for download, click **Download log**, and save it locally.

Downloading the logs of the VPN gateway

You can download and extract the archive that contains the logs of the VPN gateway, and use the information for troubleshooting or monitoring purposes.

To download the logs of the VPN gateway

1. On the **Connectivity** page, click the gear icon next to the VPN gateway.
2. Click the **Download log**.
3. [Optional] Select **Capture network packets**, and then configure the settings. For more information, see "Capturing network packets" (p. 52).
4. Click **Done**.
5. When the .zip archive is ready for download, click **Download log**, and save it locally.

Capturing network packets

To troubleshoot and analyze the communication between the local production site and a primary or recovery server, you can choose to collect network packets on the VPN gateway or VPN appliance.

After collecting 32,000 network packets, or reaching time limit, capturing network packets stops, and the results are written in a .libpcap file that is added to the logs .zip archive.

The following table provides more information about the **Capture network packets** settings that you can configure.

Setting	Description
Network interface name	The network interface on which to capture network packets. If you want to capture network packets on all network interfaces, select **Any**.
Time limit (seconds)	The time limit for capturing network packets. The maximum value you can set is 1800.
Filtering	An extra filter to apply on the captured network packets. You can enter a string containing protocols, ports, directions, and their combinations, separated by space, such as: "and", "or", "not", " (", ") ", "src", "dst", "net", "host", "port", "ip", "tcp", "udp", "icmp", "arp", "esp". If you want to use brackets, surround them with spaces. You can also enter IP addresses and network addresses, for example: "icmp or arp" and "port 67 or 68". For more information about the values that you can enter, see the Linux tpcdump help.

Cloud servers

With Disaster recovery, you can use two types of cloud servers: primary and recovery.

A primary server is a virtual machine that is not linked to a machine on the local site. You can use primary servers to protect a specific application or run various auxiliary services (such as a web server).

A recovery server is a virtual machine that is a replica of the original machine (protected server). The recovery server is based on the protected server backups that are stored in the cloud. In case of a disaster, recovery servers are used for switching workloads from the original servers.

Configuring recovery servers

A **recovery server** – a replica of the original machine based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers in case of a disaster.

When creating a recovery server, you must specify the following network parameters:

Parameter	Description
Cloud network	(required) The cloud network to which a recovery server will be connected.
IP address in production network	(required) The IP address with which a virtual machine for a recovery server will be launched. This address is used in both the production and test networks. Before launching, the virtual machine is configured for getting the IP address via DHCP.
Test IP address	(optional) The IP address to access a recovery server from the client-production network during a test failover, to prevent the production IP address from being duplicated in the same network. This IP address is different from the IP address in the production network. Servers in the local site can reach the recovery server during the test failover via the test IP address, while access in the reverse direction is not available. Internet access from the recovery server in the test network will be available if the **Internet access** option is selected during the recovery server creation.
Public IP address	(optional) The IP address to access a recovery server from the Internet. If a server has no public IP address, it can be reached only from the local network.
Internet access	(optional) Enables the recovery server to access the Internet (in both the production and test failover cases).

Creating a recovery server

To create a recovery server that will be a copy of your workload, follow the procedure below. You can also watch the video tutorial that demonstrates the process.

Important
When you perform a failover, you can select only recovery points that were created after the creation of the recovery server.

Prerequisites

- A protection plan must be applied to the original machine that you want to protect. This plan must back up the entire machine, or only the disks, required for booting up and providing the necessary services, to a cloud storage.
- One of the connectivity types to the cloud site must be set.

To create a recovery server

1. On the **All devices** tab, select the machine that you want to protect.
2. Click **Disaster recovery**, and then click **Create recovery server**.
3. In the **Create recovery server** wizard, on the **Server configuration** tab, do the following:
   a. Select the number of virtual cores and the size of RAM.

      Note
      You can see the compute points for every option. The number of compute points reflects the cost of running the recovery server per hour. For more information, see "Compute points" (p. 67).

   b. [Optional] Change the default name of the recovery server.
   c. [Optional] Add a description.
4. On the **Network** tab, do the following:
   a. Specify the cloud network to which the server will be connected.
   b. Select the **DHCP** option.

DHCP option	Description
Provided by cloud site	This is the default setting. The IP address of the server will be provided by an automatically configured DHCP server in the cloud.
Custom	The IP address of the server will be provided by your own DHCP server in the cloud.

   c. Specify the **MAC address**.
      The MAC address is a unique identifier that is assigned to the network adapter of the server. If you use custom DHCP, you can configure it to always assign a specific IP address to a

specific MAC address. Thus you will ensure that the recovery server always gets the same IP address. You can run applications that have licenses that are registered with the MAC address.

d. Specify the IP address that the server will have in the production network. By default, the IP address of the original machine is set.

   Note
   If you use a DHCP server, add this IP address to the server exclusion list in order to avoid IP address conflicts.
   If you use a custom DHCP server, you must specify the same IP address in **IP address in production network** as the one configured in the DHCP server. Otherwise, test failover will not work properly, and the server will not be reachable via a public IP address.

e. [Optional] Select the **Test IP address** check box, and then specify the IP address.
   If you select this setting, you can test a failover in the isolated test network and connect to the recovery server via RDP or SSH during a test failover. During a test failover, the VPN gateway will replace the test IP address with the production IP address by using the NAT protocol.
   If you do not select the setting, the console will be the only way to access the server during a test failover.

   Note
   If you use a DHCP server, add this IP address to the server exclusion list to avoid IP address conflicts.

   You can select one of the proposed IP addresses or enter a different one.

f. [Optional] Select the **Internet access** check box.
   If you select this setting, the recovery server will have access to the Internet during a production or test failover. By default, the TCP port 25 is open for outbound connections to public IP addresses.

g. [Optional] Select the **Use public IP address** check box.
   With a public IP address, the recovery server becomes accessible from the Internet during a failover or test failover. If you do not select this option, the server will be available only in your production network.
   The **Use public IP address** option requires the **Internet access** option to be selected.
   The public IP address will be shown after you complete the configuration. By default, TCP port 443 is open for inbound connections to public IP addresses.

   Note
   If you clear the **Use Public IP address** check box or delete the recovery server, its public IP address will not be reserved.

5. On the **Settings** tab, select **Set RPO threshold**, and then set the value.

The RPO threshold defines the maximum time interval between the last recovery point that is suitable for a failover and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

6. [Optional] [If the backups for the selected machine are encrypted by using encryption as a machine property], specify the password that will be automatically used when creating a virtual machine for the recovery server from the encrypted backup.
   a. Click **Enter password**, and then enter the password for the encrypted backup and define a name for the credentials.
      By default, you will see the most recent backup in the list.
   b. To view all the backups, select **Show all backups**.
   c. Click **Save**.

Note
Although the password that you specify will be stored in a secure credentials store, saving passwords might be against your compliance obligations.

7. On the **Cloud firewall rules** tab, edit the default firewall rules. For more information, see "Setting firewall rules for cloud servers" (p. 64).
8. Click **Create**.

The recovery server appears in the **Disaster Recovery** > **Servers** > **Recovery servers** tab of the Cyber Protect console.

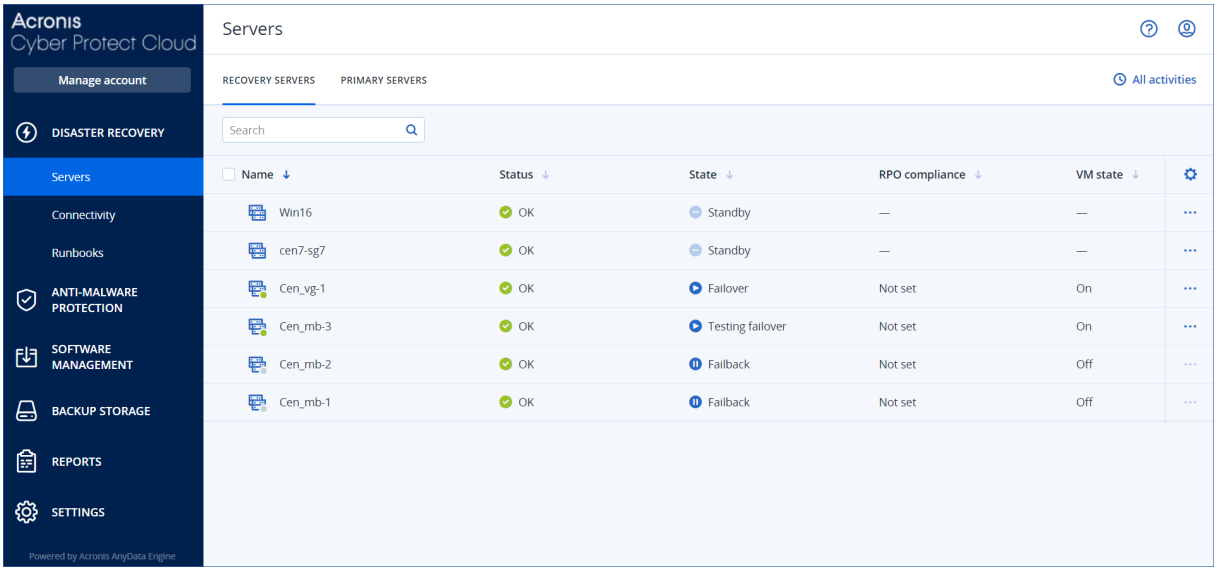

Operations with recovery servers

In the Cyber Protect console, primary servers are shown on the **Disaster Recovery** > **Servers** > **Recovery servers** tab.

Power on

To power on a recovery server

1. On the **Recovery servers** tab, click the recovery server.
2. Click **Power on**.

Power off

To power off a recovery server

1. On the **Recovery servers** tab, click the recovery server.
2. Click **Power off**.
3. In the **Power off server** screen, click **Power off**.

Force power off

To force power off a recovery server

1. On the **Recovery servers** tab, click the recovery server.
2. Click **Power off**.
3. In the **Power off server** screen, select **Force power off the server**, and then click **Power off**.

Stop

To stop a recovery server

1. On the **Recovery servers** tab, click the recovery server.
2. Click **Stop**.

Edit settings

To edit the settings of a recovery server

1. On the **Recovery servers** tab, click the recovery server.
2. Click **Stop**.
3. Click **Edit**, and then edit the settings.

Apply protection plan

To apply a plan to a primary server

1. On the **Primary servers** tab, click the primary server.
2. On the **Plan** tab, click **Create**.
   You will see a predefined protection plan where you can change only the schedule and retention rules. For more information, see "Backing up the cloud servers".

Configuring primary servers

A **primary server** is a virtual machine that does not have a linked machine on the local site if compared to a recovery server. Primary servers are used for protecting an application by replication, or running various auxiliary services (such as a web server).

Typically, a primary server is used for real-time data replication across servers running crucial applications. You set up the replication by yourself, using the application's native tools. For example,

Active Directory replication, or SQL replication, can be configured among the local servers and the primary server.

Alternatively, a primary server can be included in an AlwaysOn Availability Group (AAG) or Database Availability Group (DAG).

Both methods require a deep knowledge of the application and the administrator rights. A primary server constantly consumes computing resources and space on the fast disaster recovery storage. It needs maintenance on your side: monitoring the replication, installing software updates, and backing up. The benefits are the minimal RPO and RTO with a minimal load on the production environment (as compared to backing up entire servers to the cloud).

Primary servers are always launched only in the production network and have the following network parameters:

Parameter	Description
Cloud network	(required) The cloud network to which a primary server will be connected.
IP address in production network	(required) The IP address that the primary server will have in the production network. By default, the first free IP address from your production network is set.
Public IP address	(optional) The IP address for accessing a primary server from the Internet. If a server has no public IP address, it can be reached only from the local network, not through the Internet.
Internet access	(optional) Enables a primary server to access the Internet.

Creating a primary server

Prerequisites

- One of the connectivity types to the cloud site must be set.

To create a primary server

1. Go to **Disaster Recovery** > **Servers** > **Primary servers** tab.
2. Click **Create**.
3. In the **Create primary server** wizard, on the **Server configuration** tab, do the following:
   a. Select a template for the new virtual machine.
   b. Select the flavor of the configuration (number of virtual cores and the size of RAM).
   The following table shows the maximum total amount of disk space (GB) for each flavor.

Type	vCPU	RAM (GB)	Maximum total amount of disk space (GB)
F1	1	2	500
F2	1	4	1,000

Type	vCPU	RAM (GB)	Maximum total amount of disk space (GB)
F3	2	8	2,000
F4	4	16	4,000
F5	8	32	8,000
F6	16	64	16,000
F7	16	128	32,000
F8	16	256	64,000

c. [Optional] Change the virtual disk size. If you need more than one hard disk, click **Add disk**, and then specify the new disk size. You can add up to 10 disks for a primary server.

d. Change the default name of the recovery server.

e. Add a description.

4. On the **Network** tab, do the following:

a. Specify the cloud network in which the primary server will be included.

b. Select the **DHCP** option.

DHCP option	Description
Provided by cloud site	This is the default setting. The IP address of the server will be provided by an automatically configured DHCP server in the cloud.
Custom	The IP address of the server will be provided by your own DHCP server in the cloud.

c. Specify the **MAC address**.

The MAC address is a unique identifier that is assigned to the network adapter of the server. If you use custom DHCP, you can configure it to always assign a specific IP adresses to a specific MAC address. This ensures that the primary server always gets the same IP address. You can run applications that have licenses that are registered with the MAC address.

d. Specify the IP address that the server will have in the production network.

By default, the first free IP address from your production network is set.

Note

If you use a DHCP server, add this IP address to the server exclusion list in order to avoid IP address conflicts.

If you use a custom DHCP server, you must specify the same IP address in **IP address in production network** as the one configured in the DHCP server. Otherwise, test failover will not work properly, and the server will not be reachable via a public IP address.

e. [Optional] Select the **Internet access** check box.

If you select this option, the primary server will have access to the Internet. By default, TCP port 25 is open for outbound connections to public IP addresses.

f. [Optional] Select the **Use public IP address** check box.

With a public IP address, the primary server becomes accessible from the Internet. If you do not select this setting, the server will be available only in your production network.

The public IP address will be shown after you complete the configuration. By default, TCP port 443 is open for inbound connections to public IP addresses.

Note

If you clear the **Use Public IP address** check box or delete the recovery server, its public IP address will not be reserved.

5. [Optional] On the **Settings** tab, select **Set RPO threshold**, and then set the value.

   RPO threshold defines the maximum time interval between the last recovery point and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.
6. [Optional] On the **Cloud firewall rules** tab, edit the default firewall rules. For more information, see "Setting firewall rules for cloud servers" (p. 64).
7. Click **Create**.

The primary server becomes available in the production network. You can manage the server by using its console, RDP, SSH, or TeamViewer.

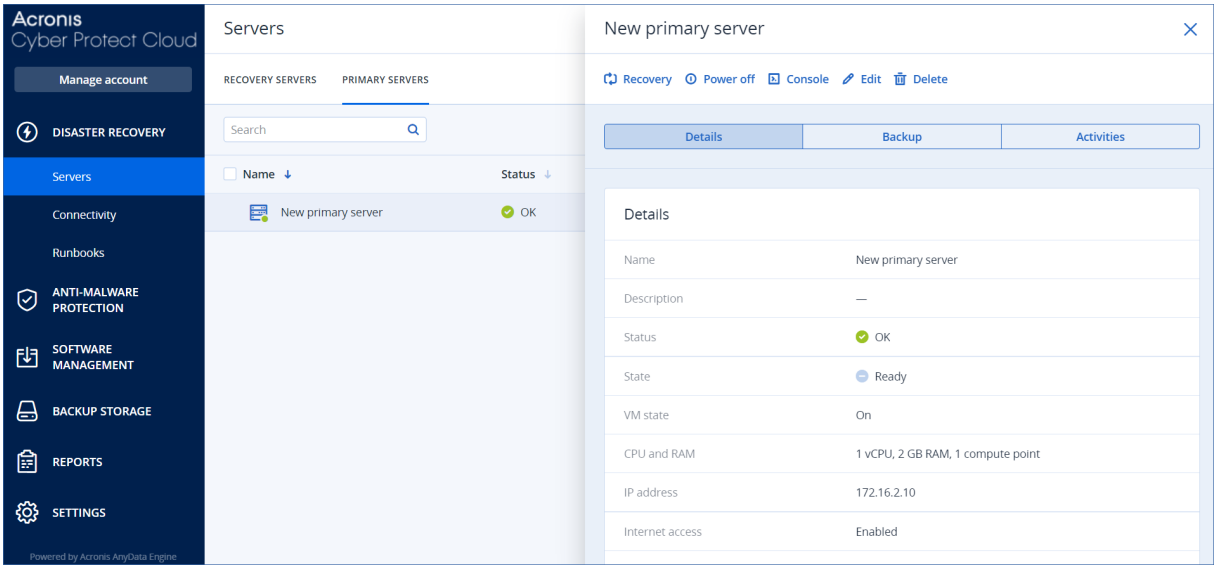

Operations with primary servers

In the Cyber Protect console, primary servers are shown on the **Disaster Recovery** > **Servers** > **Primary servers** tab.

Power on

To power on a primary server

1. On the **Primary servers** tab, click the primary server.
2. Click **Power on**.

Power off

To power off a primary server

1. On the **Primary servers** tab, click the primary server.
2. Click **Power off**.
3. In the **Power off server** screen, click **Power off**.

Force power off

To force power off a primary server

1. On the **Primary servers** tab, click the primary server.
2. Click **Power off**.
3. In the **Power off server** screen, select **Force power off the server**, and then click **Power off**.

Stop

To stop a primary server

1. On the **Primary servers** tab, click the primary server.
2. Click **Stop**.

Edit settings

To edit the settings of a primary server

1. On the **Primary servers** tab, click the primary server.
2. Click **Stop**.
3. Click **Edit**, and then edit the settings.

Apply protection plan

To apply a plan to a primary server

1. On the **Primary servers** tab, click the primary server.
2. On the **Plan** tab, click **Create**.
   You will see a predefined protection plan where you can change only the schedule and retention rules. For more information, see "Backing up the cloud servers".

Viewing details about cloud servers

To view the details of the cloud servers, go to **Disaster Recovery** > **Servers**. There are two tabs there: **Recovery servers** and **Primary servers**. To show all optional columns in the table, click the gear icon.

You can find the following information about each cloud server by selecting it.

Column	Description

name	
Name	A cloud server name defined by you
Status	The status reflecting the most severe issue with a cloud server (based on the active alerts)
State	A cloud server state
VM state	The power state of a virtual machine associated with a cloud server
Active location	The location where a cloud server is hosted. For example, **Cloud**.
RPO threshold	The maximum time interval allowed between the last suitable recovery point for failover and the current time. The value can be set within 15-60 minutes, 1-24 hours, 1-14 days.
RPO compliance	The RPO compliance is the ratio between the actual RPO and RPO threshold. The RPO compliance is shown if the RPO threshold is defined. It is calculated as follows: **RPO compliance = Actual RPO / RPO threshold** where **Actual RPO = current time – last recovery point time** **RPO compliance statuses** Depending on the value of the ratio between the actual RPO and RPO threshold, the following statuses are used: • **Compliant**. The RPO compliance < 1x. A server meets the RPO threshold. • **Exceeded**. The RPO compliance <= 2x. A server violates the RPO threshold. • **Severely exceeded**. The RPO compliance <= 4x. A server violates the RPO threshold more than 2x times. • **Critically exceeded**. The RPO compliance > 4x. A server violates the RPO threshold more than 4x times. • **Pending (no backups)**. The server is protected with the protection plan but the backup is being created and not completed yet.
Actual RPO	The time passed since the last recovery point creation
Last recovery point	The date and time when the last recovery point was created

Backups of cloud servers

Primary and recovery servers are backed up agentless on the cloud site. These backups have the following restrictions.

- The only possible backup location is the cloud storage. Primary servers are backed up to the **Primary servers backup** storage.

  Note
  Microsoft Azure backup locations are not supported.

- A backup plan cannot be applied to multiple servers. Each server must have its own backup plan, even if all of the backup plans have the same settings.
- Only one backup plan can be applied to a server.
- Application-aware backup is not supported.
- Encryption is not available.
- Backup options are not available.

When you delete a primary server, its backups are also deleted.

A recovery server is backed up only in the failover state. Its backups continue the backup sequence of the original server. When a failback is performed, the original server can continue this backup sequence. So, the backups of the recovery server can only be deleted manually or as a result of applying the retention rules. When a recovery server is deleted, its backups are always kept.

Note
The backup plans for cloud servers are performed according to UTC time.

Firewall rules for cloud servers

You can configure firewall rules to control the inbound and outbound traffic of the primary and recovery servers on your cloud site.

You can configure inbound rules after you provision a public IP address for the cloud server. By default, TCP port 443 is allowed, and all other inbound connections are denied. You can change the default firewall rules, and add or remove Inbound exceptions. If a public IP is not provisioned, you can only view the inbound rules, but cannot configure them.

You can configure outbound rules after when you provision Internet access for the cloud server. By default, TCP port 25 is denied, and all other outbound connections are allowed. You can change the default firewall rules, and add or remove outbound exceptions. If Internet access is not provisioned, you can only view the outbound rules, but cannot configure them.

Note
For security reasons, there are predefined firewall rules that you cannot change.

For inbound and outbound connections:

- Permit ping: ICMP echo-request (type 8, code 0) and ICMP echo-reply (type 0, code 0)
- Permit ICMP need-to-frag (type 3, code 4)
- Permit TTL exceeded (type 11, code 0)

For inbound connections only:

- Non-configurable part: Deny all

For outbound connections only:

- Non-configurable part: Reject all

Setting firewall rules for cloud servers

You can edit the default firewall rules for the primary and recovery servers in the cloud.

To edit the firewall rules of a server on your cloud site

1. In the Cyber Protect console, go to **Disaster Recovery**> **Servers**.
2. If you want to edit the firewall rules of a recovery server, click the **Recovery servers** tab. Alternatively, if you want to edit the firewall rules of a primary server, click the **Primary servers** tab.
3. Click the server, and then click **Edit**.
4. Click the **Cloud firewall rules** tab.
5. If you want to change the default action for the inbound connections:
   a. In the **Inbound** drop-down field, select the default action.

Action	Description
Deny all	Denies any inbound traffic. You can add exceptions and allow traffic from specific IP addresses, protocols, and ports.
Allow all	Allows all inbound TCP and UDP traffic. You can add exceptions and deny traffic from specific IP addresses, protocols, and ports.

Note
Changing the default action invalidates and removes the configuration of existing inbound rules.

   b. [Optional] If you want to save the existing exceptions, in the confirmation window, select

Save filled-in exceptions.

c. Click **Confirm**.

6. If you want to add an exception:
   a. Click **Add exception**.
   b. Specify the firewall parameters.

Firewall parameter	Description
Protocol	Select the protocol for the connection. The following options are supported: • **TCP** • **UDP** • **TCP+UDP**
Server port	Select the ports to which the rule applies. You can specify the following: • a specific port number (for example, `2298`) • a range of port numbers (for example, `6000-6700`) • any port number. Use `*` if you want the rule to apply to any port number.
Client IP address	Select the IP addresses to which the rule applies. You can specify the following: • a specific IP address (for example, `192.168.0.0`) • a range of IP addresses using the CIDR notation (for example, `192.168.0.0/24`) • any IP address. Use `*` if you want the rule to apply to any IP address.

7. If you want to remove an existing inbound exception, click the bin icon next to it.
8. If you want to change the default action for the outbound connections:
   a. In the **Outbound** drop-down field, select the default action.

Action	Description
Deny all	Denies any outbound traffic. You can add exceptions and allow traffic to specific IP addresses, protocols, and ports.
Allow all	Allows all outbound traffic. You can add exceptions and deny traffic from specific IP addresses, protocols, and ports.

Note
Changing the default action invalidates and removes the configuration of existing outbound rules.

b. [Optional] If you want to save the existing exceptions, in the confirmation window, select **Save filled-in exceptions**.

c. Click **Confirm**.

9. If you want to add an exception:

a. Click **Add exception**.

b. Specify the firewall parameters.

Firewall parameter	Description
Protocol	Select the protocol for the connection. The following options are supported: • **TCP** • **UDP** • **TCP+UDP**
Server port	Select the ports to which the rule applies. You can specify the following: • a specific port number (for example, `2298`) • a range of port numbers (for example, `6000-6700`) • any port number. Use `*` if you want the rule to apply to any port number.
Client IP address	Select the IP addresses to which the rule applies. You can specify the following: • a specific IP address (for example, `192.168.0.0`) • a range of IP addresses using the CIDR notation (for example, `192.168.0.0/24`) • any IP address. Use `*` if you want the rule to apply to any IP address.

10. If you want to remove an existing outbound exception, click the bin icon next to it.

11. Click **Save**.

Checking the cloud firewall activities

After an update of the configuration of the firewall rules of a cloud server, a log of the update activity becomes available in the Cyber Protect console. You can view the log and check the following information:

- user name of the user who updated the configuration
- date and time of the update
- firewall settings for inbound and outbound connections

- the default actions for inbound and outbound connections
- the protocols, ports and IP addresses of the exceptions for inbound and outbound connections

To view the details about a cloud firewall rules configuration change

1. In the Cyber Protect console, click **Monitoring** > **Activities**.
2. Click the corresponding activity, and click **All Properties**.
   The description of the activity should be **Updating cloud server configuration**.
3. In the **context** field, inspect the information that you are interested in.

Compute points

In Disaster Recovery, compute points are used for primary servers and recovery servers during test failover and production failover. Compute points reflect the compute resources used for running the servers (virtual machines) in the cloud.

The consumption of compute points during disaster recovery depends on the server's parameters, and the duration of the time period in which the server is in failover state. The more powerful the server and the longer the time period, the more compute points will be consumed. And the more compute points are consumed, the higher the price that you will be charged.

All servers that are running in the Acronis Cloud will be charged for compute points, depending on their configured flavor, and regardless of their state (powered on or powered off).

Recovery servers in Standby state do not consume compute points and will not be charged for compute points.

In the table below, you can see an example for eight servers in the cloud with different flavors, and the corresponding compute points that they will consume per hour. You can change the flavors of the servers in the **Details** tab.

Type	CPU	RAM	Compute points
F1	1 vCPU	2 GB	1
F2	1 vCPU	4 GB	2
F3	2 vCPU	8 GB	4
F4	4 vCPU	16 GB	8
F5	8 vCPU	32 GB	16
F6	16 vCPU	64 GB	32
F7	16 vCPU	128 GB	64
F8	16 vCPU	256 GB	128

Using the information in the table, you can easily estimate how many compute points a server (virtual machine) will consume.

For example, if you want to protect with Disaster Recovery one virtual machine with 4 vCPU* of 16 GB RAM, and one virtual machine with 2 vCPU with 8 GB of RAM, the first virtual machine will consume 8 compute points per hour, and the second virtual machine – 4 compute points per hour. If both virtual machines are in failover, the total consumption will be 12 compute points per hour, or 288 compute points for the whole day (12 compute points x 24 hours = 288 compute points).

* vCPU refers to a physical central processing unit (CPU) that is assigned to a virtual machine, and is a time-dependent entity.

Note
If the overage for the **Compute points** quota is reached, all primary and recovery servers will be shut down. It will not be possible to use these servers until the beginning of the next billing period, or until you increase the quota. The default billing period is a full calendar month.

Test failover

Performing a test failover means starting a recovery server in a test VLAN that is isolated from your production network. You can test several recovery servers at a time and check their interaction. In the test network, the servers communicate using their production IP addresses, but they cannot initiate TCP or UDP connections to the workloads in your local network.

During test failover, the virtual machine (recovery server) is not finalized. The agent reads the content of the virtual disks directly from the backup and randomly accesses different parts of the backup. This might make the performance of the recovery server in the test failover state slower than its normal performance.

Performing a test failover

Though performing a test failover is optional, we recommend that you make it a regular process with a frequency that you find adequate in terms of cost and safety. A good practice is creating a runbook – a set of instructions describing how to spin up the production environment in the cloud.

Important
You must create a recovery server in advance to protect your devices from a disaster.

You can perform failover only from recovery points (backups) that were created after the recovery server of the device was created.

At least one recovery point must be created before failing over to a recovery server. The maximum number of recovery points that is supported is 100.

To perform a test failover

1. Select the original machine or select the recovery server that you want to test.
2. Click **Disaster Recovery**.
   The description of the recovery server opens.
3. Click **Failover**.

4. Select the failover type **Test failover**.
5. Select the recovery point (backup), and then click **Start**.
6. If the backup that you selected is encrypted by using encryption as a machine property:
   a. Enter the encryption password for the backup set.

   Note
   The password will only be saved temporarily and will be used only for the current test failover operation. The password is automatically deleted from the credentials store if the test failover is stopped, or after the test failover completes.

   b. [Optional] To save the password for the backup set and use it in subsequent failover operations, select the **Store the password in a secure credentials store...** check box and then, in the **Credentials name** field, enter a name for the credentials.

   Important
   The password will be stored in a secured credentials store and will be applied automatically in subsequent failover operations. However, saving passwords might conflict with your compliance obligations.

   c. Click **Done**.

   When the recovery server starts, its state changes to **Testing failover**.

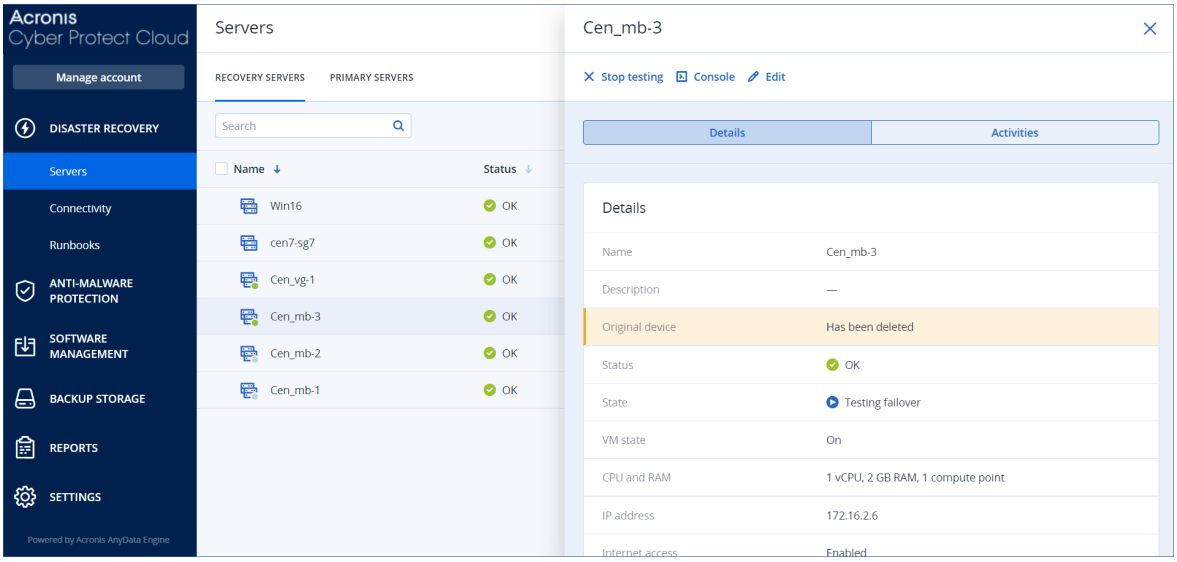

7. Test the recovery server by using any of the following methods:
   - In **Disaster Recovery** > **Servers**, select the recovery server, and then click **Console**.
   - Connect to the recovery server by using RDP or SSH, and the test IP address that you specified when creating the recovery server. Try the connection from both inside and outside the production network (as described in "Point-to-site connection").
   - Run a script within the recovery server.

     The script may check the login screen, whether applications are started, the Internet connection, and the ability of other machines to connect to the recovery server.

- If the recovery server has access to the Internet and a public IP address, you may want to use TeamViewer.

8. When the test is complete, click **Stop testing**.

   The recovery server is stopped. All changes made to the recovery server during the test failover are not preserved.

Automated test failover

With automated test failover, the recovery server is tested automatically once a month without any manual interaction.

The automated test failover process consists of the following parts:

1. creating a virtual machine from the last recovery point
2. taking a screenshot of the virtual machine
3. analyzing if the operating system of the virtual machine starts successfully
4. notifying you about the test failover status

Note
Automated test failover consumes compute points.

You can configure the automated test failover in the recovery server's settings. For more information, see "Configuring automated test failover" (p. 70).

Note that, in very rare cases, automated test failover might be skipped and might not be performed at the scheduled time. This is because production failover has higher priority than automated test failover, so the hardware resources (CPU and RAM) allocated for automated test failover might be temporarily limited to ensure that there are enough resources for a concurrent production failover.

If automated test failover is skipped for some reason, an alert will be raised.

Note
Automated test failover will fail if the backups of the original machine are encrypted by using encryption as a machine property, and the encryption password is not specified when creating the recovery server. For more information about specifying the encryption password, see "Creating a recovery server" (p. 54).

Configuring automated test failover

By configuring automated test failover, you can test your recovery server every month without performing any manual actions.

To configure automated test failover

1. In the console, go to **Disaster recovery** > **Servers** > **Recovery servers**, and then select the recovery server.
2. Click **Edit**.

3. On the **Automated test failover** tab, in the **Schedule** field, select **Monthly**.
4. [Optional] In **Screenshot timeout**, change the default value of the maximum time period (in minutes) for the system to try performing automated test failover.
5. [Optional] If you want to save the **Screenshot timeout** value as the default and have it populated automatically when you enable automated test failover for the other recovery servers, select **Set as default timeout**.
6. Click **Save**.

Viewing the automated test failover status

You can view the details of a completed automated test failover, such as status, start time, end time, duration, and the screenshot of the virtual machine.

Note
The screenshot of the virtual machine is kept until automated test failover runs again and generates a new screenshot.

To view the automated test failover status of a recovery server

1. In the console, go to **Disaster recovery** > **Servers** > **Recovery servers**, and then select the recovery server.
2. In the **Automated test failover** section, check the details of the last automated test failover.
3. [Optional] To view the screenshot of the virtual machine, click **Show screenshot**.

Disabling automated test failover

You can disable automated test failover if you want to save resources or you do not need automated test failover to be performed for a certain recovery server.

To disable automated test failover

1. In the console, go to **Disaster recovery** > **Servers** > **Recovery servers**, and then select the recovery server.
2. Click **Edit**.
3. In the wizard, click the **Automated test failover** tab.
4. Turn off the **Automated test failover** switch.
5. Click **Save**.

Production failover

Note
The availability of this feature depends on the service quotas that are enabled for your account.

When a recovery server is created, it stays in the **Standby** state. The corresponding virtual machine does not exist until you start a failover. Before starting a failover process, you must create at least one disk image backup (with bootable volume) of the original machine.

When starting the failover process, you select the recovery point (backup) of the original machine from which a virtual machine with the predefined parameters will be created. The failover operation uses the "run VM from a backup" functionality. The recovery server gets the transition state **Finalization**. This process implies transferring the server's virtual disks from the backup storage ('cold' storage) to the disaster recovery storage ('hot' storage).

Note
During the **Finalization**, the server is accessible and operable, although the performance is lower than normal. You can open the server console by clicking the **Console is ready** link. The link is available in the **VM State** column on the **Disaster Recovery** > **Servers** screen and in the server's **Details** view.

When the **Finalization** is completed, the server performance reaches its normal value. The server state changes to **Failover**. The workload is now switched from the original machine to the recovery server in the cloud site.

If the recovery server has a protection agent inside, the agent service is stopped in order to avoid interference (such as starting a backup or reporting outdated statuses to the backup component).

On the diagram below, you can see both the failover and failback processes.

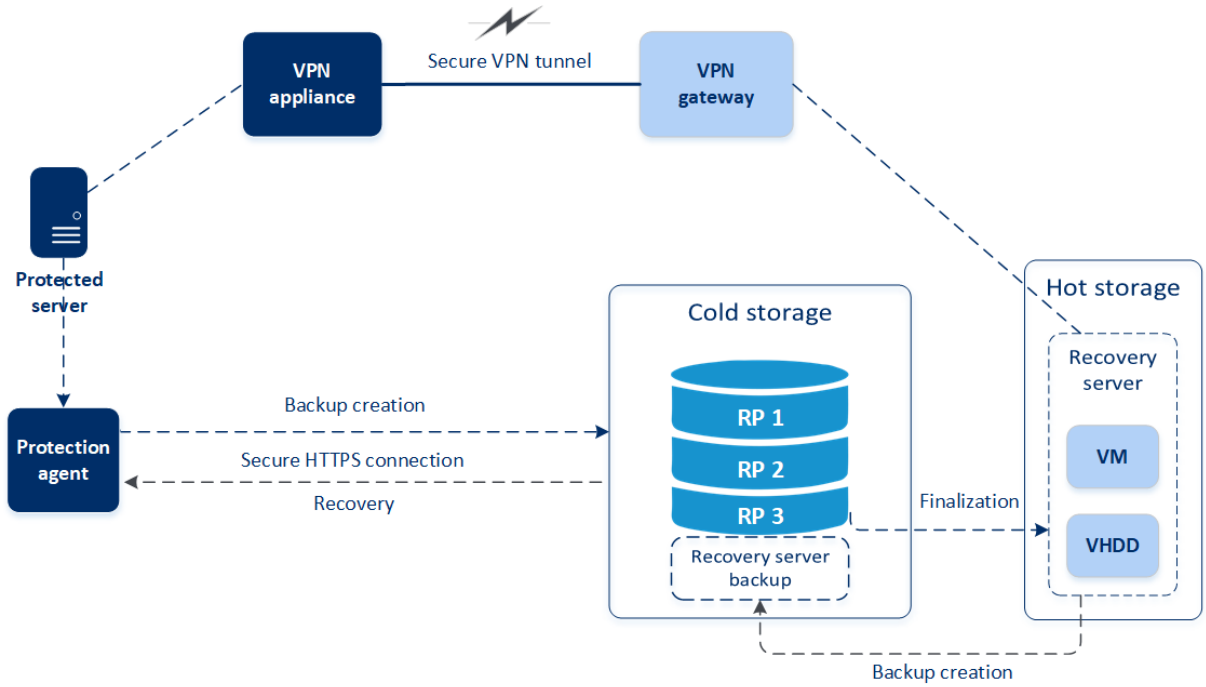

Performing a failover

Note
The availability of this feature depends on the service quotas that are enabled for your account.

A failover is a process of moving a workload from your premises to the cloud, and also the state when the workload remains in the cloud.

When you start a failover, the recovery server starts in the production network. To avoid interference and unwanted issues, ensure that the original workload is not online and cannot be accessed via VPN.

To avoid a backup interference into the same cloud archive, manually revoke the protection plan from the workload that is currently in **Failover** state. For more information about revoking plans, see Revoking a protection plan.

Important
You must create a recovery server in advance to protect your devices from a disaster.

You can perform failover only from recovery points (backups) that were created after the recovery server of the device was created.

At least one recovery point must be created before failing over to a recovery server. The maximum number of recovery points that is supported is 100.

You can follow the procedure below or watch the video tutorial.

To perform a failover

1. Ensure that the original machine is not available on the network.
2. In the Cyber Protect console, go to **Disaster recovery** > **Servers**> **Recovery servers**, and then select the recovery server.
3. Click **Failover**.
4. Select **Production failover**.
5. Select the recovery point (backup), and then click **Start**.
6. [If the backup that you selected is encrypted by using encryption as a machine property]
   a. Enter the encryption password for the backup set.

      Note
      The password will only be saved temporarily and will be used only for the current failover operation. The password is automatically deleted from the credentials store after the failover operation completes and the server returns to the **Standby** state.

   b. [Optional] To save the password for the backup set and use it in subsequent failover operations, select the **Store the password in a secure credentials store...** check box and

then, in the **Credentials name** field, enter a name for the credentials.

Important
The password will be stored in a secured credentials store and will be applied automatically in subsequent failover operations. However, saving passwords might conflict with your compliance obligations.

c. Click **Done**.

When the recovery server starts, its state changes to **Finalization**, and after some time to **Failover.**

Important
It is critical to understand that the server is available in both the **Finalization** and **Failover** states. During the **Finalization** state, you can access the server console by clicking the **Console is ready** link. The link is available in the **VM State** column on the **Disaster Recovery** > **Servers** screen and in the server's **Details** view.

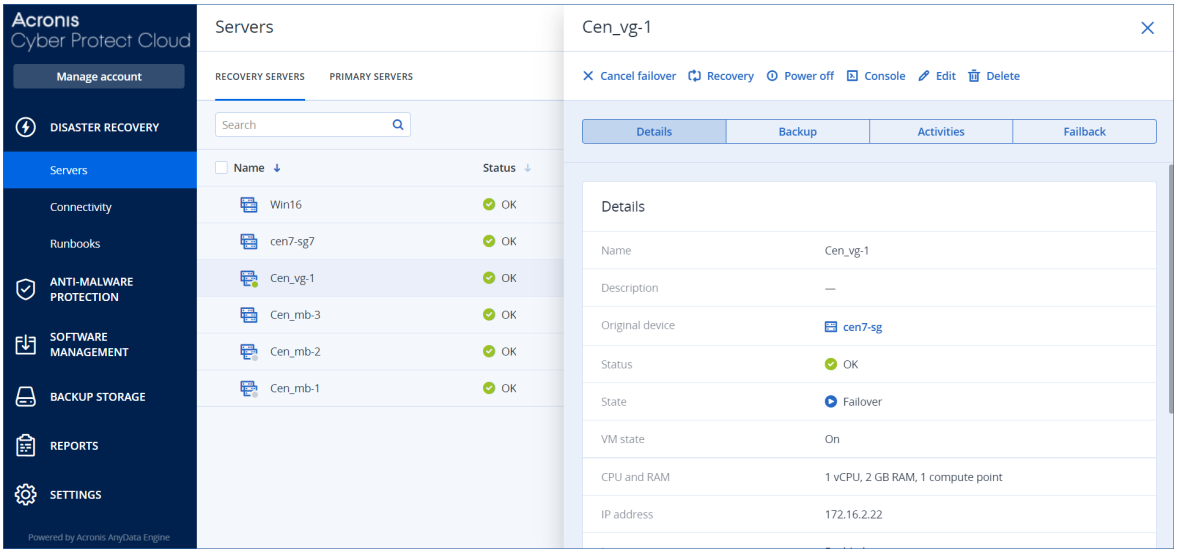

7. Ensure that the recovery server is started by viewing its console. Click **Disaster Recovery** > **Servers**, select the recovery server, and then click **Console**.
8. Ensure that the recovery server can be accessed using the production IP address that you specified when creating the recovery server.

Once the recovery server is finalized, a new protection plan is automatically created and applied to it. This protection plan is based on the protection plan that was used for creating the recovery server, with certain limitations. In this plan, you can change only the schedule and retention rules. For more information, refer to "Backing up the cloud servers".

How to perform failover of servers using local DNS

If your local site uses DNS servers to resolve machine names, recovery servers might fail to communicate after a failover. This happens because the DNS servers in the cloud are different from

those on the local site. By default, newly created cloud servers use the DNS servers of the cloud site, but you can configure custom DNS settings. For more information, see "Configuring custom DNS servers" (p. 49).

How to perform failover of a DHCP server

Your local infrastructure may have the DHCP server located on a Windows or Linux host. When such a host is failed over to the cloud site, the DHCP server duplication issue occurs because the VPN gateway in the cloud also performs the DHCP role. To resolve this issue, do one of the following:

- If only the DHCP host was failed over to the cloud, while the rest local servers are still on the local site, then you must log in to the DHCP host in the cloud and turn off the DHCP server on it. Thus, there will be no conflicts and only the VPN gateway will work as the DHCP server.
- If your cloud servers already got the IP addresses from the DHCP host, then you must log in to the DHCP host in the cloud and turn off the DHCP server on it. You must also log in to the cloud servers and renew the DHCP lease to assign new IP addresses allocated from the correct DHCP server (hosted on the VPN gateway).

Note
The instructions are not valid when your cloud DHCP server is configured with the **Custom DHCP** option, and some of the recovery or primary servers get their IP address from this DHCP server.

Stopping a failover

You can stop a production failover at any moment, during each phase of the process.

Note
Stopping a failover reverts all changes that were made from the moment the failover was started, except the recovery server backups.

To stop a production failover

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers** > **Recovery servers**.
2. Select the recovery server that is in status **Failover**.
3. Click the recovery server.
4. Click **Stop failover**.
5. In the confirmation window that appears, select the check box, and then click **Stop failover**.
   The failover is stopped. The recovery server returns to the **Standby** state.

Failback

Note
The availability of this feature depends on the service quotas that are enabled for your account.

A failback is a process of moving the workload from the cloud back to a physical or virtual machine on your local site. You can perform a failback on a recovery server in **Failover** state, and continue using the server on your local site.

You can perform automated failover to a virtual or physical target machine on your local site. During the failback, you can transfer the backup data to your local site while the virtual machine in the cloud continues to run. This technology helps you to achieve a very short downtime period, which is estimated and displayed in the Cyber Protect console. You can view it and use this information to plan your activities and, if necessary, warn your clients about an upcoming downtime period. If you perform agent-based failback via bootable media, the downtime is even shorter, as only the delta changes will be transferred to the local site.

For a failback to a target physical machine, you can use the agent-based failback via bootable media. For more information, see "Performing agent-based failback via bootable media" (p. 77).

For a failback to a target virtual machine, you can use the agent-based failback via bootable media or the agentless failback via hypervisor agent. For more information, see "Performing agent-based failback via bootable media" (p. 77) and "Performing agentless failback via a hypervisor agent" (p. 82).

In specific cases when you cannot use the automated failback procedure, you can perform a manual failback. For more information, see "Manual failback" (p. 85).

Note
Runbook operations support the failback in manual mode only. This means that if you start the failback process by executing a runbook that includes a **Failback server** step, the procedure will require a manual interaction: you must manually recover the machine, and confirm or cancel the failback process from the **Disaster Recovery** > **Servers** tab.

Agent-based failback via bootable media

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The agent-based failback via bootable media process is optimized for performing a failback to the original physical or virtual machine. During this process, only the delta changes are transferred to the local site.

The agent-based failback process via bootable media to a target physical or virtual machine consists of the following phases:

1. **Planning**. During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.
2. **Data transfer**. During this phase, the data is transferred from the cloud site to the local site while the virtual machine in the cloud continues to run. You can start the next phase - switchover

- at any time during the data transfer phase, but you should consider the following relations.

The longer you remain in the data transfer phase,

- the longer the virtual machine in the cloud continues to run.
- the more data will be transferred to your local site.
- the higher the cost you will pay (you spend more compute points).
- the shorter the downtime period that you will experience during the switchover phase.

If you want to minimize the downtime, start the switchover phase after more than 90% of the data is transferred to the local site.

If you can afford to experience a longer downtime period, and do not want to spend more compute points for running the virtual machine in the cloud, you can start the switchover phase earlier.

Note

The data transfer process uses a flashback technology. This technology compares the data that is available on the target machine to the data of the virtual machine in the cloud. If part of the data is already available on the target machine, it will not be transferred again. This technology makes the data transfer phase faster.

For this reason, we recommend that you restore the server to the original machine on your local site.

3. **Switchover**. During this phase, the virtual machine in the cloud is turned off, and the remaining data - including the last backup increment - is transferred to the local site. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process.
4. **Validation**. During this phase, the machine on the local site is ready, and you can reboot it using a Linux-based bootable media. You can verify if the virtual machine is working correctly, and:
   - If everything is working as expected, confirm the failback. After the failback confirmation, the virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state. This is the end of the failback process.
   - If something is wrong, you can cancel the failover and return to the planning phase.

Note

After the bootable media has been rebooted, you will not be able to use it again. If, at the validation phase, you discover something wrong, you must register a new bootable media and start the failback process again.

However, as flashback technology will be used, the data that is already on the local site will not be transferred again, and the failback process will be much faster.

Performing agent-based failback via bootable media

Note

The availability of this feature depends on the service quotas that are enabled for your account.

You can perform agent-based failback via bootable media to a target physical or virtual machine on your local site.

Note
The data transfer process uses a flashback technology. This technology compares the data that is available on the target machine to the data of the virtual machine in the cloud. If part of the data is already available on the target machine, it will not be transferred again. This technology makes the data transfer phase faster.

For this reason, we recommend that you restore the server to the original machine on your local site.

Prerequisites

- The agent that you will use to perform failback is online and is not currently used for another failback operation.
- Your Internet connection is stable.
- A registered bootable media is available. For more information, see "Creating bootable media to recover operating systems" in the Cyber Protection User Guide.
- The target machine is the original machine on your local site, or has the same firmware as the original machine.
- There is at least one full backup of the virtual machine in the cloud.

To perform a failback to a physical machine

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.
4. In the **Failback type** field, select **Agent-based via bootable media**.
5. In the **Target bootable media** field, click **Specify**, select the bootable media, and then click **Done**.

   Note
   We recommend that you use ready-made bootable media as it is already configured. For more information, see "Creating bootable media to recover operating systems" in the Cyber Protection User Guide.

6. [Optional] To change the default disk mapping, in the **Disk mapping** field, click **Specify**, map the disks of the backup to the disks of the target machine, and then click **Done**.
7. Click **Start data transfer** and then, in the confirmation window, click **Start**.

   Note
   If there is no backup of the virtual machine in the cloud, the system will perform a backup automatically before the data transfer phase.

   The data transfer phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows how much data is already transferred to the local site, and the total amount of data that must be transferred. The total amount of data includes the data from the last backup before the data transfer phase was started, and the backups of the newly generated data (backup increments), as the virtual machine continues to run during the data transfer phase. For this reason, the **Progress** values increase with time. As the system uses a flashback technology during the data transfer and does not transfer the data that is already available on the target machine, the progress might be faster than what is initially calculated by the console.
Downtime estimation	This parameter shows how much time the virtual machine in the cloud will be unavailable if you start the switchover phase now. The value is calculated based on the values of the **Progress** parameter, and decreases with time. As the system uses a flashback technology during the data transfer and does not transfer the data that is already available on the target machine, the downtime might be much shorter than the value that is initially displayed in the console.

8. Click **Switchover** and then, in the confirmation window, click **Switchover** again.
   The switchover phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows the progress of restoring the machine on the local site.
Estimated time to finish	This parameter shows the approximate time when the switchover phase will be completed and you will be able to start the machine on the local site.

Note
If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

9. After the **Switchover** phase completes, reboot the bootable media, and then verify that the physical machine on your local site is working as expected.
   For more information, see "Recovering disks using bootable media" in the Cyber Protection User Guide.
10. Click **Confirm failback** and then, in the confirmation window, click **Confirm** to finalize the process.
    The virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state.

Note

Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Agentless failback via a hypervisor agent

Note

The availability of this feature depends on the service quotas that are enabled for your account.

The agentless failback via a hypervisor agent process is optimized for performing a failback to a new virtual machine. If you want to perform a failback to the original virtual machine, follow the procedure for agent-based failback via bootable media.

The agentless failback via a hypervisor agent consists of four phases.

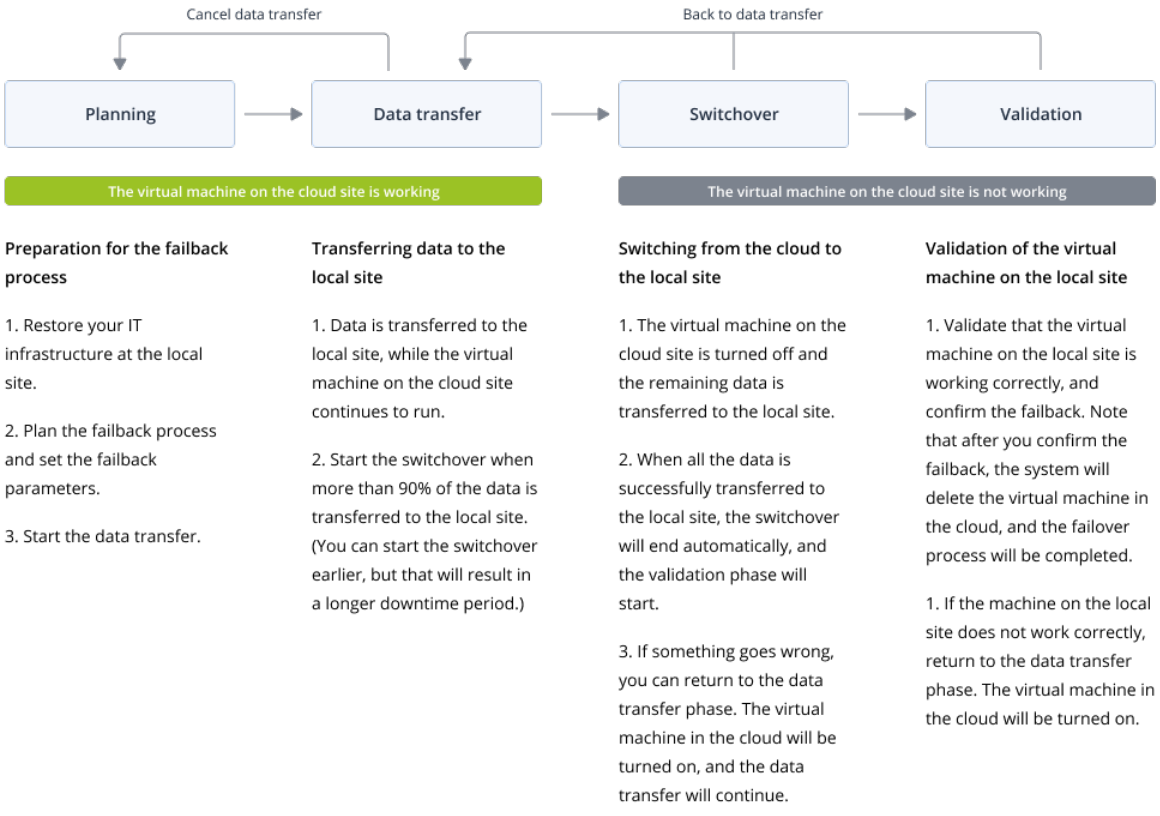

1. **Planning**. During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.

Note
To minimize the total time for the failback process, we recommend that you start the data transfer phase immediately after you set up your local servers, and then continue with the configuration of the network and the rest of the local infrastructure during the data transfer phase.

2. **Data transfer**. During this phase, the data is transferred from the cloud site to the local site while the virtual machine in the cloud continues to run. You can start the next phase - switchover - at any time during the data transfer phase, but you should consider the following relations.

   The longer you remain in the data transfer phase,
   - the longer the virtual machine in the cloud continues to run.
   - the more data will be transferred to your local site.
   - the higher the cost you will pay (you spend more compute points).
   - the shorter the downtime period that you will experience during the switchover phase.

   If you want to minimize the downtime, start the switchover phase after more than 90% of the data is transferred to the local site.

   If you can afford to experience a longer downtime period, and do not want to spend more compute points for running the virtual machine in the cloud, you can start the switchover phase earlier.

   If you cancel the failback process during the data transfer phase, the transferred data will not be deleted from the local site. To avoid potential issues, manually delete the transferred data before you start a new failback process. The following data transfer process will start from the beginning.
3. **Switchover**. During this phase, the virtual machine in the cloud is turned off, and the remaining data - including the last backup increment - is transferred to the local site. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process.

   You can view the estimated time to finish (downtime period) of this phase in the Cyber Protect console. When all the data is transferred to the local site (there is no data loss, and the virtual machine on the local site is an exact copy of the virtual machine in the cloud), the switchover phase completes. The virtual machine on the local site is recovered, and the validation phase starts automatically.
4. **Validation**. During this phase, the virtual machine on the local site is ready and automatically started. You can verify if the virtual machine is working correctly, and:
   - If everything is working as expected, confirm the failback. After the failback confirmation, the virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state. This is the end of the failback process.
   - If something is wrong, you can cancel the switchover and return to the data transfer phase.

Performing agentless failback via a hypervisor agent

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can perform an agentless failback to a target virtual machine on your local site via a hypervisor agent.

Prerequisites

- The agent that you will use to perform failback is online and is not currently used for another failback operation.
- Your internet connection is stable.
- There is at least one full backup of the virtual machine in the cloud.

To perform an agentless failback to a virtual machine via a hypervisor agent

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.
4. In the **Failback parameters** section, in the **Failback type** field, select **Agentless via hypervisor agent**, and then configure the other parameters.

   Note that by default, some of the **Failback parameters** are populated automatically with suggested values, but you can change them.

   The following table provides more information about the **Failback parameters**.

Parameter	Description
Backup size	Amount of data that will be transferred to your local site during the failback process. After you start the failback process to a target virtual machine, the **Backup size** will be increasing during the data transfer phase, because the virtual machine in the cloud will continue to run and generate new data. To calculate the estimated downtime period during the failback process to a target virtual machine, take 10% of the **Backup size** value (as we recommend that you start the switchover phase after 90% of the data is transferred to your local site), and divide it by the value of your Internet speed. **Note** The value of the Internet speed will decrease when you perform several failback processes at the same time.
Target	Failback location: a VMware ESXi host or a Microsoft Hyper-V host.

Parameter	Description
machine location	You can select from all the hosts that have an agent which is registered with the Cyber Protection service.
Agent	Agent which will perform the failback operation. You can use one agent to perform one failback operation at the same time. You can select an agent that is online and is not currently used for another failback process, has a version which supports the failback functionality, and has rights to access the backup. Note that you can install several agents on VMware ESXi hosts, and start a separate failback process using each of them. These failback processes can be performed at the same time.
Target machine settings	Virtual machine settings: • **Virtual processors**. Select the number of virtual processors. • **Memory**. Select how much memory the virtual machine will have. • **Units**. Select the units for the memory. • [Optional] **Network adapters**. To add a network adapter, click **Add**, and select a network in the **Network** field. When you are ready with the changes, click **Done**.
Path	(For Microsoft Hyper-V hosts) Folder on the host where your machine will be stored. Ensure that there is enough free memory space on the host for the machine.
Datastore	(For VMware ESXi hosts) Datastore on the host where your machine will be stored. Ensure that there is enough free memory space on the host for the machine.
Provisioning mode	Method of allocation of the virtual disk. For Microsoft Hyper-V hosts: • **Dynamically expanding** (default value). • **Fixed size**. For Microsoft Hyper-V hosts: • **Thin** (default value). • **Thick**.
Target machine name	Name of the target machine. By default, the target machine name is the same as the recovery server name. The target machine name must be unique on the selected **Target machine location**.

5. Click **Start data transfer**, and then in the confirmation window, click **Start**.

Note
If there is no backup of the virtual machine in the cloud, the system will perform a backup automatically before the data transfer phase.

The **Data transfer** phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows how much data is already transferred to the local site, and the total amount of data that must be transferred. The total amount of data includes the data from the last backup before the data transfer phase was started, and the backups of the newly generated data (backup increments), as the virtual machine continues to run during the data transfer phase. For this reason, both values of the **Progress** parameter increase with time.
Downtime estimation	This parameter shows how much time the virtual machine in the cloud will be unavailable if you start the switchover phase now. The value is calculated based on the values of the **Progress** parameter, and decreases with time.

6. Click **Switchover** and then, in the confirmation window, click **Switchover** again.
   The switchover phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows the progress of restoring the machine on the local site.
Estimated time to finish	This parameter shows the approximate time when the switchover phase will be completed and you will be able to start the machine on the local site.

Note
If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

7. After the **Switchover** phase completes and the virtual machine at your local site is started automatically, validate that it is working as expected.
8. Click **Confirm failback**, and then in the confirmation window, click **Confirm** to finalize the process.
   The virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state.

Note
Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Manual failback

Note
We recommend that you use the failback process in a manual mode only when you are advised to do so by the Support team.

You can also start a failback process in a manual mode. In this case, the data transfer from the backup in the cloud to the local site will not be done automatically. It must be done manually after the virtual machine in the cloud is powered off. This makes the failback process in a manual mode much slower, and you should expect a longer downtime period.

The failback process in a manual mode consists of the following phases:

1. **Planning**. During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.
2. **Switchover**. During this phase, the virtual machine in the cloud is turned off, and the newly generated data is backed up. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process. When the backup is complete, you recover the machine to the local site manually. You can either recover the disk by using bootable media, or recover the entire machine from the cloud backup storage.
3. **Validation**. During this phase, you verify that the physical or virtual machine at the local site is working correctly, and confirm the failback. After the confirmation, the virtual machine on the cloud site is deleted, and the recovery server returns to the **Standby** state.

Performing manual failback

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can perform a manual failback to a target physical or virtual machine on your local site.

To perform a manual failback

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.

4. In the **Target** field, select **Physical machine**.
5. Click the gear icon, and then enable the **Use manual mode** switch.
6. [Optional] Calculate the estimated downtime period during the failback process, by dividing the **Backup size** value by the value of your Internet speed.

   Note
   The value of the Internet speed will decrease when you perform several failback processes at the same time.

7. Click **Switchover**, and then in the confirmation window, click **Switchover** again.
   The virtual machine on the cloud site is turned off.

   Note
   If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

8. Recover the server from the cloud backup to the physical or virtual machine on your local site. For more information, see "Recovering a machine" in the Cyber Protection User Guide.
9. Ensure that the recovery is completed and the recovered machine works properly, and click **Machine is restored**.
10. If everything is working as expected, click **Confirm failback**, and then in the confirmation window, click **Confirm** again.
    The recovery server and recovery points become ready for the next failover. To create new recovery points, apply a protection plan to the new local server.

    Note
    Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Failback from Cyber Protect Cloud to an Azure virtual machine

You can perform a failback from Cyber Protect Cloud to the original Azure virtual machine by following the procedure for "Performing manual failback" (p. 85) and using one of the following recovery options in step 8:

Recovery options

- Agentless recovery
  Supports recovery only to a new Azure VM that is created automatically.
  Configure the Azure connection in the Cyber Protect console (**Devices** > **Add** > **Microsoft Azure virtual machine**).
  A backup appliance VM is deployed in the Azure subscription to manage the recovery.
  Recovery flow:

1. On the **Backup storage** screen, select a backup.
2. Recover as an Azure virtual machine.

You can use the same flow for physical, virtual, or agent-based backups.

To reduce incurred costs, the appliance virtual machine, when used for recovery only, can be powered on only during recovery, and then turned off manually.

- Agent-based recovery

  Supports recovery to the same Azure virtual machine (if the original virtual machine with the agent is available) or a new Azure virtual machine that has a new agent installed.

  The process consists of the following steps:

  1. Manually create a clean Windows or Linux virtual machine in Azure.
  2. Install the protection agent.
  3. Use the agent to browse and recover backups from Acronis Cloud Storage.

  For more information, see Recovering Microsoft Azure and Amazon EC2 machines.

Orchestration (runbooks)

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

A runbook is a set of instructions describing how to launch the production environment in the cloud. You can create runbooks in the Cyber Protect console.

With runbooks, you can:

- Automate the failover of one or multiple servers.
- Automatically check the failover result by pinging the server IP address and checking the connection to the port you specify.
- Set the sequence of operations for servers running distributed applications.
- Include manual operations in the workflow.
- Verify the integrity of your disaster recovery solution, by executing runbooks in the test mode.

To access the **Runbooks** screen, select **Disaster recovery** > **Runbooks**.

Creating a runbook

A runbook consists of steps that are executed consecutively. A step consists of actions that start simultaneously.

To create a runbook, follow the instruction from the following procedure or from the video tutorial.

To create a runbook

1. In the Cyber Protection console, go to **Disaster recovery** > **Runbooks**.
2. Click **Create runbook**.

3. Click **Add step**.
4. Click **Add action**, and then select the action that you want to add to the step.

Action	Description
Failover server	Performs a failover of a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these parameters, see "Runbook parameters" (p. 90). **Note** If the backup of the server that you select is encrypted by using encryption as a machine property, the **Failover server** action will be paused and will be changed automatically to **Interaction required**. To proceed with the execution of the runbook, you will have to provide the password for the encrypted backup.
Failback server	Performs a failback of a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see "Runbook parameters" (p. 90). **Note** Runbook operations support the failback in manual mode only. This means that if you start the failback process by executing a runbook that includes a **Failback server** step, the procedure will require a manual interaction: you must manually recover the machine, and confirm or cancel the failback process from the **Disaster Recovery** > **Servers** tab.
Start server	Starts a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see "Runbook parameters" (p. 90). **Note** The **Start server** action is not applicable for test failover operations in runbooks. If you try executing such an action, it will fail with the following error message: Failed: The action is not applicable to the current server state.
Stop server	Stops a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see "Runbook parameters" (p. 90). **Note** The **Stop server** action is not applicable for test failover operations in runbooks. If you try executing such an action, it will fail with the following error message: Failed: The action is not applicable to the current server state.
Manual operation	A manual operation requires an interaction from a user. To define this action, you must enter a description.

Action	Description
	When a runbook sequence reaches a manual operation, the runbook will be paused and will not proceed until a user performs the required manual operation, such as clicking the confirmation button.
Execute runbook	Executes another runbook. To define this action, you must choose a runbook. A runbook can include only one execution of a given runbook. For example, if you added the action "execute Runbook A", you can add the action "execute Runbook B", but cannot add another action "execute Runbook A".

5. Define the runbook parameters for the action. For more information about these parameters, see "Runbook parameters" (p. 90).
6. [Optional] To add a description of the step:
   a. Click the ellipsis icon, and then click **Description**.
   b. Enter a description of the step.
   c. Click **Done**.
7. Repeat steps 3-6 until you create the desired sequence of steps and actions.
8. [Optional] To change the default name of the runbook:
   a. Click the ellipsis icon.
   b. Enter the name of the runbook.
   c. Enter a description of the runbook.
   d. Click **Done**.
9. Click **Save**.
10. Click **Close**.

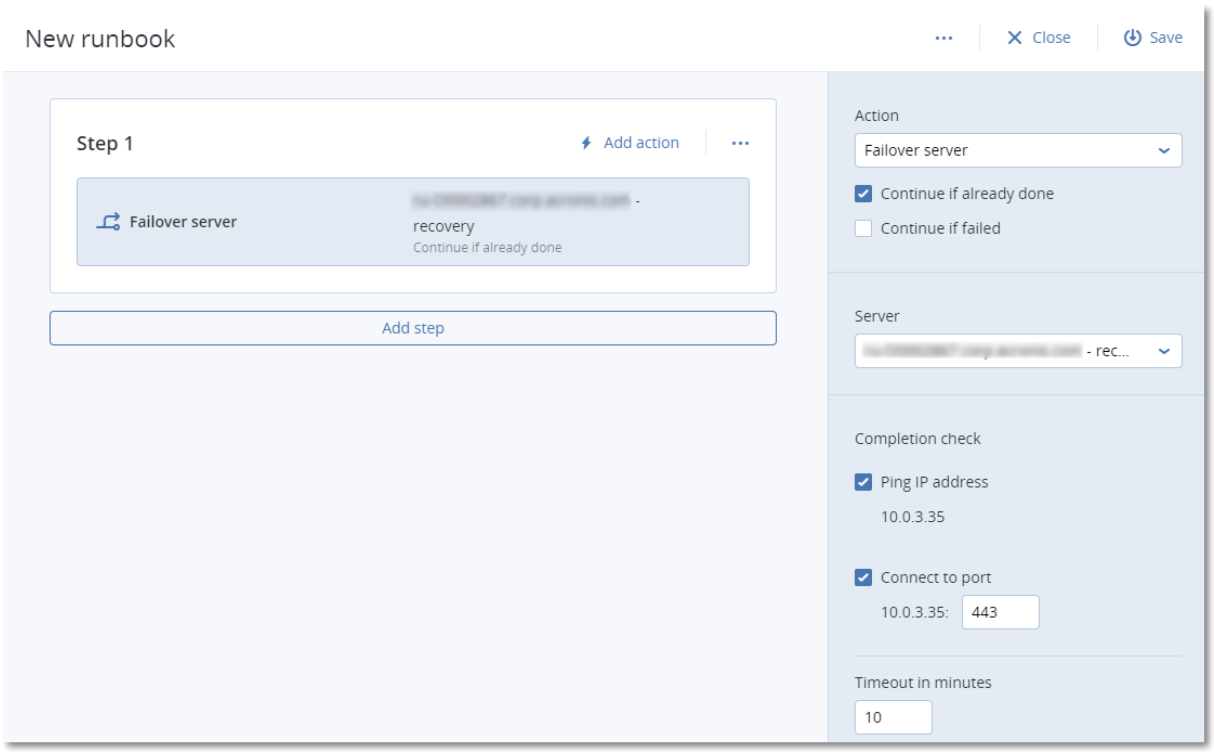

Runbook parameters

Runbook parameters are specific settings that you must configure to define a runbook action. There are two categories of runbook parameters - action parameters and completion check parameters.

Action parameters define the runbook behavior depending on the action initial state or result.

Completion check parameters ensure that the server is available and provides the necessary services. If a completion check fails, the action is considered failed.

The following table describes the configurable runbook parameters for each action.

Runbook parameter	Category	Available for action	Description
Continue if already done	Action parameter	• **Failover server** • **Start server** • **Stop server** • **Failback server**	This parameter defines the runbook behavior when the required action is already done (for example, a failover has already been performed or a server is already running). When enabled, the runbook issues a warning and proceeds. When disabled, the action fails, and then the runbook fails too. By default, this parameter is enabled.
Continue if failed	Action parameter	• **Failover server** • **Start server** • **Stop server** • **Failback server**	This parameter defines the runbook behavior when the required action fails. When enabled, the runbook issues a warning and proceeds. When disabled, the action fails, and then the runbook fails too. By default, this parameter is disabled.
Ping IP address	Completion check	• **Start server**	The software will ping the production IP address of the cloud server until the server replies or the timeout expires, whichever comes first.
Connect to port (443 by default)	Completion check	• **Failover server** • **Start server**	The software will try to connect to the cloud server by using its production IP address and the port you specify, until the connection is established or the timeout expires, whichever comes first. This way, you can check if the application that listens on the specified port is running.
Timeout in minutes	Completion check	• **Failover server** • **Start server**	The default timeout is 10 minutes.

Operations with runbooks

Note
The availability of this feature depends on the service quotas that are enabled for your account.

To access the list of operations, hover on a runbook and click the ellipsis icon. When a runbook is not running, the following operations are available:

- **Execute**
- **Edit**
- **Clone**
- **Delete**

Executing a runbook

Every time you click **Execute**, you are prompted for the execution parameters. These parameters apply to all failover and failback operations included in the runbook. The runbooks specified in the **Execute runbook** operations inherit these parameters from the main runbook.

- **Failover and failback mode**
  Choose whether you want to run a test failover (by default) or a real (production) failover. The failback mode will correspond to the chosen failover mode.
- **Failover recovery point**
  Choose the most recent recovery point (by default) or select a point in time in the past. If the latter is the case, the recovery points closest before the specified date and time will be selected for each server.

Stopping a runbook execution

During a runbook execution, you can select **Stop** in the list of operations. The software will complete all of the already started actions except for those that require user interaction.

Viewing the execution history

When you select a runbook on the **Runbooks** tab, the software displays the runbook details and execution history. Click the line corresponding to a specific execution to view the execution log.

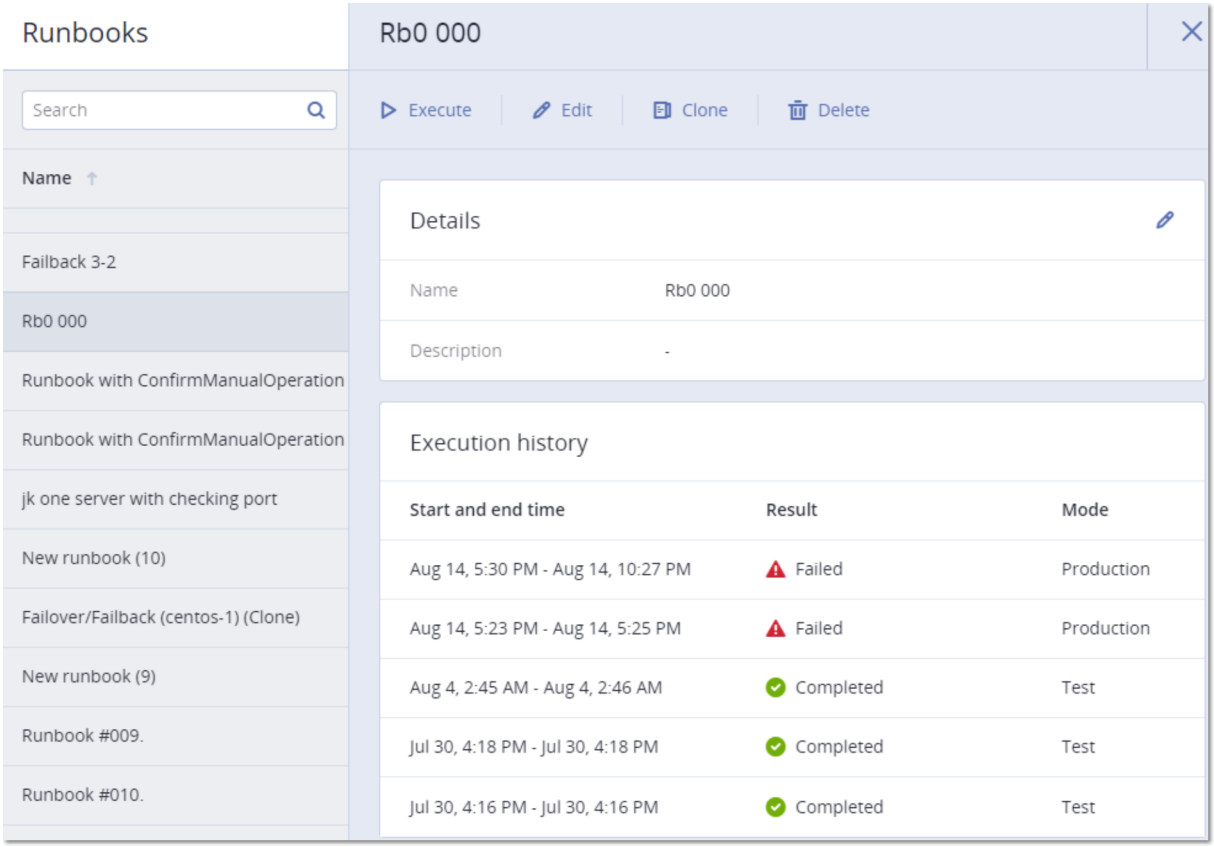

Removing the disaster recovery site

You can remove the disaster recovery site. This action will automatically delete the VPN gateway, VPN connections, and all runbooks that were configured on the site.

Prerequisites

No cloud servers are available on the disaster recovery site.

To remove the disaster recovery site

1. In the Cyber Protect console, go to **Disaster Recovery** > **Connectivity**.
2. Click **Show properties**.
3. Click **Remove disaster recovery site**.
4. In the confirmation window, click **Remove**.

Disaster Recovery to Microsoft Azure

Disaster Recovery to Microsoft Azure is a cost-efficient disaster recovery solution integrated with the Acronis Cyber Protection service. It leverages the power and flexibility of the Microsoft Azure enterprise-grade platform as a target DR site. This solution uses cold backups (recovery points), which can be stored in Microsoft Azure, Acronis Cyber Protect Cloud, or a partner-hosted cloud storage. The entire DR site configuration and orchestration are managed centrally from the Cyber Protect console. Your own Microsoft Azure subscription is used as the target DR site.

Automation and orchestration include the following capabilities:

- Initial DR site configuration
- Automated test failover with AI-powered screenshot verification
- Orchestrated failover
- Automated failback to both physical and virtual machines, with near-zero downtime
- Runbooks to automate key disaster recovery scenarios
- On-demand workers (temporary agents) that eliminate the need for permanent virtual appliances

- Cold DR mode – cost-effective recovery with minimal Azure infrastructure usage
- Warm DR mode – with near-zero RTO, coming in future releases

You retain full control over the underlying infrastructure and Microsoft Azure capabilities, with the flexibility to use native Azure services or integrate custom solutions - such as third-party firewalls and SD-WAN appliances - by connecting them to the selected recovery networks at the DR site. This solution also supports failover of Windows desktops running Windows 10 or Windows 11.

Note
To use Disaster Recovery to Microsoft Azure, you must have an active subscription to Microsoft Azure.

Software requirements for Disaster Recovery to Microsoft Azure

Supported operating systems

Protection with a recovery server in Microsoft Azure has been tested for the following operating systems:

- Ubuntu 20.x, 21.x, 22.10, 23.04
- Debian 10.x, 11.x
- Red Hat Enterprise Linux 8.x, 9.x
- Windows Server 2008 R2
- Windows Server 2012/2012 R2

- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server
- Windows Server 2022 – all installation options, except for Nano Server
- Windows Server 2025 – all installation options, except for Nano Server
- Windows 10
- Windows 11

The software may work with other Windows operating systems and Linux distributions, but this is not guaranteed.

Supported virtualization platforms

Protection of virtual machines with a recovery server has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7, 7.0
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2022 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016
- Kernel-based Virtual Machines (KVM) — fully virtualized guests (HVM) only. Paravirtualized guests (PV) are not supported.

Linux workloads that have agentless backups from a guest OS and have volumes with the Logical Volume Manager (LVM) configurations are supported.

Windows workloads that have agentless backups from a guest OS and have dynamic disks (LDM) configurations are supported.

The software might work with other virtualization platforms and versions, but this is not guaranteed.

Limitations

The following platforms and configurations are not supported in Disaster Recovery to Microsoft Azure:

1. Unsupported platforms:
   - Agents for Virtuozzo
   - macOS
2. Unsupported configurations:
   Microsoft Windows

- Active Directory service with FRS replication is not supported.
- Removable media without either GPT or MBR formatting (so-called "superfloppy") are not supported.

Linux

- File systems without a partition table.
- Linux workloads that are backed up with an agent from a guest OS and have volumes with the following advanced Logical Volume Manager (LVM) configurations: Striped volumes, Mirrored volumes, RAID 0, RAID 4, RAID 5, RAID 6, or RAID 10 volumes.

Note
Workloads with multiple operating systems installed are not supported.

3. Unsupported tenant modes:
   - Disaster recovery is not available when Compliance mode is enabled for the tenant.
4. Unsupported backup types:
   - Continuous data protection (CDP) recovery points are incompatible.

     Important
     If you create a recovery server from a backup having a CDP recovery point, then during the failback or creating backup of a recovery server, you will loose the data contained in the CDP recovery point.

   - Forensic backups cannot be used for creating recovery servers.

Cloud servers are not encrypted.

Licensing for Disaster Recovery to Microsoft Azure

To enable Disaster Recovery to Microsoft Azure, the **DR and direct backup to Azure** offering item must be enabled for your tenant. This offering item enables:

- Disaster Recovery (DR) to Azure that uses the customer's own Azure subscription.
- Direct backup to Azure that does not require an Advanced backup license.

One quota from the offering item is consumed when a recovery server is created or direct backup to Azure is enabled. Only one quota per workload is used, even if both DR and direct backup are active.

The quota for this offering item that is assigned to your tenant represents the maximum number of workloads that can be protected. A device that uses both Disaster Recovery to Azure and direct backup to Azure protection consumes one quota.

Hard quota overage

When the hard quota for the offering item is decreased, existing recovery servers might become unlicensed. The number of unlicensed recovery servers depends on the overage. Recovery servers that were in test or production failover remain functional but become unlicensed. Unlicensed

recovery servers in the Standby state have limited DR operations. You cannot power on unlicensed servers until they are licensed again.

Increasing the offering item quota automatically assigns licenses to unlicensed devices and clears related alerts.

Generated alerts

The following alerts are generated when there are issues that are caused by missing licenses.

- **Recovery server is unlicensed** - This alert is generated when a server becomes unlicensed.
- **Disaster Recovery protection was disabled for a workload** - This alert is generated when there are no available licenses.
- **Automated test failover failed** - This alert is generated when a failover is blocked due to a missing license.

Working with Disaster Recovery to Microsoft Azure

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

The basic workflow for using disaster recovery is the following:

1. Configure your DR site in Microsoft Azure. For more information, see "Creating a disaster recovery site in Microsoft Azure" (p. 102).
2. Create a recovery server of the workload that you want to protect. For more information, see "Creating recovery servers in Microsoft Azure" (p. 110).
3. [Optional] Configure the connectivity from your local site to the cloud site in Microsoft Azure by using Azure native services, such as Azure site-to-site VPN or Azure ExpressRoute. For more information, see "Connectivity and networks in Microsoft Azure" (p. 104).
4. [Optional] Configure runbooks. For more information, see "Creating a runbook in Microsoft Azure" (p. 133).
5. Configure automated test failover or perform a test failover. For more information, see "Configuring automated test failover in Microsoft Azure" (p. 119) and "Performing a test failover in Microsoft Azure" (p. 117).
6. [When a disaster occurs] Perform a production failover. For more information, see "Performing a production failover in Microsoft Azure" (p. 115).
7. [After the disaster] Perform a failback to the local site. For more information, see "Failback in Microsoft Azure" (p. 121).

Managing access to your Microsoft Azure subscription

Disaster Recovery to Microsoft Azure requires that you connect to a Microsoft Azure subscription in the Cyber Protect console.

You can configure Microsoft Azure subscriptions in the **Infrastructure** > **Public clouds** screen. There, you can also manage your subscriptions, by performing the following tasks: renewing access to the subscription, viewing subscription properties and activities, or removing the subscription.

The following table provides the links to the corresponding procedure for the task that you want to accomplish.

Task	Link
Add access to a Microsoft Azure subscription	"Adding access to a Microsoft Azure subscription" (p. 97)
Renew access to a Microsoft Azure subscription	"Renewing access to a Microsoft Azure subscription" (p. 98)
Remove access to a Microsoft Azure subscription	"Removing access to a Microsoft Azure subscription" (p. 99)

Adding access to a Microsoft Azure subscription

By adding a Microsoft Azure subscription in the Cyber Protect console, Acronis can securely access your subscription and directly back up the relevant workloads to Microsoft Azure.

To add access to a Microsoft Azure subscription

1. In the Cyber Protect console, go to **Infrastructure > Public clouds**.
2. Click **Add**, and from the displayed list of options, select **Microsoft Azure**.
3. In the displayed dialog, click **Sign in**. You are redirected to the Microsoft login page.

   Note
   You must be assigned with one of the following roles in Microsoft Azure AD in order to complete the connection to the subscription: Cloud Application Administrator, Application Administrator, or Global Administrator. You must also be assigned the Owner role for each selected subscription.

4. In the Microsoft login screen, enter your login credentials and accept the requested permissions. The connection process starts, and may take several minutes.
   For more information about securely accessing your Microsoft Azure and subscription, refer to article Microsoft Azure connection security and audit (72684).
5. When the connection is complete, do the following:
   a. In the **Microsoft Azure subscription** field, select the relevant subscription from the list.
   b. [Optional] In the **Azure region** field, select the region in which to deploy the system resources.

Note
The system preselects the region in which most of the resource groups are located but you can change it depending on your preference.

6. Click **Add subscription**.
   The subscription is added to the list of public clouds.
   To renew the annual access certificate for the subscription, see "Renewing access to a Microsoft Azure subscription" (p. 98).
   To remove access to the subscription, see "Removing access to a Microsoft Azure subscription" (p. 99).

Note
If the Microsoft Azure account you are logged into includes access to multiple Microsoft Azure ADs, including ADs in which you were invited as a guest user, only the default user directory is selected. If you want to use a directory in which you are a guest user, you need to create a new user in that specific Microsoft Azure AD. You can then log in to that account and connect to the relevant subscription.

Renewing access to a Microsoft Azure subscription

Once registered in the Cyber Protect console, access to a Microsoft Azure subscription is automatically set for one year by Acronis using a free and unique access certificate. When the certificate nears its expiry date, you can quickly and easily renew it.

To renew the access certificate for your Microsoft Azure subscription

1. In the Cyber Protect console, go to **Infrastructure > Public clouds**.
2. Select the relevant subscription from the displayed list.

Note
The **Access status** column indicates the current status of the access certificate for each subscription and shows one of two statuses: **OK** or **Expired**.

3. In the right pane, click **Renew access**.
   Alternatively, click the **Subscription** tab, and then click **Renew** in the **Access expiration date** field.

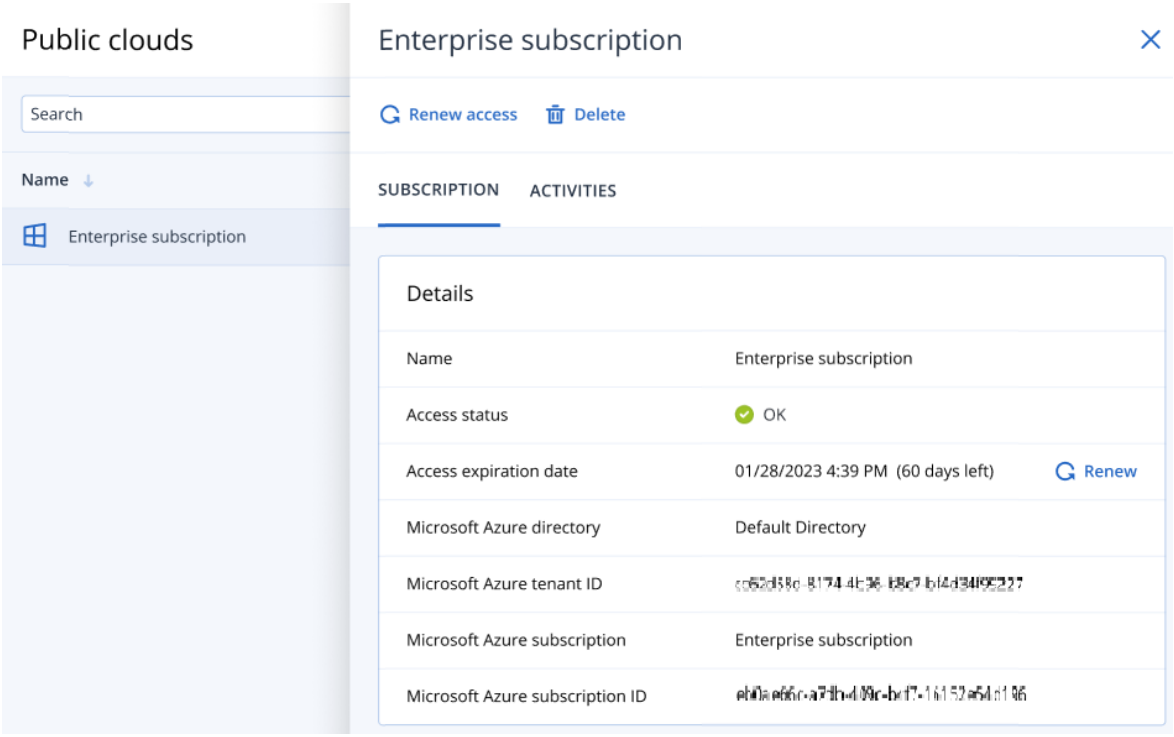

4. In the Microsoft login screen, enter your login credentials and accept the requested permissions.
   The connection process starts, and may take several minutes.
   When the authentication is successful, access is automatically renewed for one year.
   For more information about the required permissions, refer to article Microsoft Azure connection security and audit (72684).

Removing access to a Microsoft Azure subscription

You should remove access to the Microsoft Azure subscription if you are not backing up workloads to Microsoft Azure.

To remove access to a Microsoft Azure subscription

Important
You cannot remove a subscription if it is currently being used to backup to Microsoft Azure.

1. In the Cyber Protect console, go to **Infrastructure > Public clouds**.
2. Select the relevant subscription from the displayed list.
3. In the right pane, click **Delete**.

   Note
   You can only remove a subscription you added. You can also remove a subscription if you are a Company administrator or Unit administrator, or were assigned the role of Cyber administrator or Administrator in the Cyber Protection service.

4. In the displayed confirmation message, click **Remove**.

Cross-subscription configuration issues in Microsoft Azure

If you are using two different subscriptions for Microsoft Azure - for example, one for direct backup to Microsoft Azure ("Subscription 1"), and another one for the configuration of the DR site in Microsoft Azure ("Subscription 2") - but you remove the access to "Subscription 1", the following issues will occur:

- Failover will fail

  You cannot start a failover, as the backup data is stored in "Subscription 1".
- Failback will fail

  If the access to "Subscription 1" is removed after a failover was performed, it will not be possible to access all backup data that is stored in the subscription. Therefore, it will not be possible to perform backup operations of VM in failover and failback. While backup data remains in "Subscription 1", failover or recovery operations rely on valid access and permissions to the original storage location. Removing access breaks this dependency, even if DR operations are started in a different subscription.

Important
Do not remove or revoke access to the original Azure subscription if backups that are stored there are still used or needed.

Creating a disaster recovery protection plan with Microsoft Azure

The disaster recovery protection plan is a protection plan in which the **Disaster Recovery** module is enabled.

For Microsoft Azure, the DR site location must be configured. It is not possible to apply a protection plan with an Azure DR location if it was not configured.

Note

- Applying a disaster recovery protection plan creates cloud recovery servers. Existing cloud networks are not changed or recreated.
- After you configure disaster recovery, you will be able to perform a test or production failover from any of the recovery points (backups) that were generated after the recovery server for the device was created. You cannot use recovery points that were generated before the device was protected with disaster recovery (before the recovery server was created).
- A disaster recovery protection plan cannot be enabled if the IP address of a device cannot be detected. For example, when virtual machines are backed up agentless and are not assigned an IP address. In this case, we recommend that you create a recovery server manually.
- When you apply a protection plan, recovery servers are configured in the subnet that was configured in the mapping rules during the configuration of the DR site location, based on the IP address of the original device. If the IP address matches any of the specified source local networks, the recovery server will be created in the corresponding Azure recovery network and subnet. The last octet of the private IP will be taken from the original machine's IP address.

To create a disaster recovery protection plan

1. In the Cyber Protect console, go to **Devices** > **All devices**.
2. Select the machines that you want to protect.
3. Click **Protect**, and then click **Create plan**.
   The protection plan default settings open.
4. Configure the backup options.
   To use the disaster recovery functionality, the plan must back up the entire machine, or only the disks, required for booting up and providing the necessary services, to the Cloud storage or Microsoft Azure storage.
5. Enable the **Disaster recovery** module by turning on the switch next to the module name.
6. In the **Location** field, select **Microsoft Azure**.
7. Click **Create**.
   The plan is created and applied to the selected machines. The recovery servers with default parameters are created.

Managing the disaster recovery site in Microsoft Azure

You can create and configure your disaster recovery (DR) site in Microsoft Azure either as part of the disaster recovery protection plan creation or as a separate procedure, from the **Disaster recovery** screen.

When a disaster recovery protection plan is applied, the recovery server cloud network infrastructure is created only if it does not already exist. Existing cloud servers and networks are not changed or recreated.

A disaster recovery protection plan cannot be enabled if the IP address of a device cannot be detected. For example, when virtual machines are backed up agentless and are not assigned an IP address. In this case, we recommend that you create a recovery server manually.

When you apply a protection plan, recovery servers are configured in the subnet that is configured in the mapping rules during the configuration of the DR site location, and are based on the IP address of the original device. If the IP address matches any of the specified source local networks, the recovery server will be created in the corresponding Azure recovery network and subnet. The last octet of the private IP will be taken from the original machine's IP address.

Creating a disaster recovery site in Microsoft Azure

Prerequisites

- You have a subscription to Microsoft Azure.
- Your Microsoft account has one of the following Entra ID roles: Cloud Application Administrator, Application Administrator, or Global Administrator.
- Your Microsoft account has the Owner role for the Azure subscription.
- The DR and direct backup to Azure offering item is enabled for your tenant.

From All Devices

To create a DR site in Microsoft Azure

1. In the Cyber Protect console, go to **Devices** > **All devices**.
2. Select the workloads that you want to protect.
3. In the **Actions** menu, click **Protect**.
4. Create a protection plan and configure the following settings:
   a. In the **Backup** module, in the **Where to back up** field, select the storage location for the backups of your workloads.
   b. Enable the **Disaster recovery** module, and then click the **Location** field.

      Note
      If only **DR and direct backup to Azure** offering item is enabled for your tenant, the location will not be preselected, and you will see a link **Configure**.
      If both **DR to Acronis or hybrid cloud** and **DR and direct backup to Azure** offering items are enabled for your tenant, Cyber Protect Cloud will be preselected as a location.

   c. In the **Disaster recovery site configuration** wizard, on the **Site location** tab, select **Microsoft Azure Cloud**, and then click **Next**.
   d. On the **Azure subscription** tab, do the following:
      - If your Microsoft Azure subscription was already added to the Cyber Protect console, click it, and then click **Next**.
      - If you want to add a new Microsoft Azure subscription, click **Add subscription**, add the subscription, click it, and then click **Next**.
   e. On the **Target region** tab, in the **Azure region** field, select the Azure region for the DR site.

f. On the **Recovery network** tab, configure the recovery networks for production and test failover, and then click **Next**.

Option	Description
Default configuration	Use this option if you want the production and test networks in Azure to be created automatically, with one network for each environment.
Advanced configuration	Use this option if you want to select your existing Azure networks as the production and test networks and configure network mapping.

g. On the **Summary** tab, review the parameters of your DR site, and then click **Configure**.

The protection plan is successfully applied to the selected workloads. The DR site is created in Azure, located in the selected Azure region. On the **Disaster recovery** > **Connectivity** screen, you can view the production and test networks, and the recovery servers that were created, in **Standby** mode.

From Connectivity

To create a DR site in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery**.
2. Click **Configure**.
3. In the **Disaster recovery site configuration** wizard, on the **Site location** tab, select **Microsoft Azure Cloud**, and then click **Next**.
4. On the **Azure subscription** tab, do the following:
   - If your Microsoft Azure subscription was already added to the Cyber Protect console, click it, and then click **Next**.
   - If you want to add a new Microsoft Azure subscription, click **Add subscription**, add the subscription, click it, and then click **Next**.
5. On the **Target region** tab, in the **Azure region** field, select the Azure region for the DR site.
6. On the **Recovery network** tab, configure the recovery networks for production and test failover, and then click **Next**.

Option	Description
Default configuration	Use this option if you want the production and test networks in Azure to be created automatically, with one network for each environment.
Advanced configuration	Use this option if you want to select your existing Azure networks as the production and test networks and configure network mapping.

7. On the **Summary** tab, review the parameters of your DR site, and then click **Configure**.

   The DR site is created in Azure, located in the selected Azure region. On the **Disaster recovery** > **Connectivity** screen, you can view the production and test networks.

Removing the DR site from Microsoft Azure

You can remove the DR site from Microsoft Azure if you no longer need it or if you want to change the location of the DR site.

Prerequisites

There are no recovery servers on the DR site.

To remove the DR site from Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Connectivity**.
2. Click **Remove disaster recovery site**.
3. In the **Remove disaster recovery site** window, enter your login, and the click **Remove**.

Connectivity and networks in Microsoft Azure

Disaster Recovery to Microsoft Azure enables the orchestration of testing, failover, and failback, and the selection of recovery networks (Azure VNets and subnets) for both production and test failover scenarios. You can configure these networks during the initial DR site configuration or later: from the **Connectivity** screen, or individually for each recovery server.

You retain full control over Azure networking and connectivity, and have the flexibility to leverage native Azure platform capabilities or bring custom solutions, such as third-party firewalls or SD-WAN appliances into the DR site, by connecting them to the selected recovery networks.

Below is an overview of key Azure networking services, their relevance for disaster recovery use cases, and links for more information.

Azure Firewall

Azure Firewall provides centralized, stateful network traffic filtering across multiple VNets and subnets. It helps enforce security rules between workloads after a failover and supports scenarios where DR environments must meet corporate compliance or segmentation policies.

You can use Azure Firewall to control outbound and inbound traffic to and from failovered VMs (for example, to limit Internet access and allow only allowlisted sources). For managed exposure, position Azure Firewall between the on-premises network (via VPN) and the Azure DR VNets.

For more information, see the Microsoft Azure documentation.

Network security groups (NSGs)

NSGs allow defining granular rules to allow or deny network traffic to VMs or subnets. NSGs operate at the VM NIC and subnet level.

You can apply NSGs to control access to recovery VMs, ensure isolation of test failovers, or expose only necessary ports (for example, RDP and HTTP). NSGs are essential for securely enabling connectivity based on whether the recovery VM is in a test or production mode.

For more information, see the Microsoft Azure documentation.

DNS servers

Azure VNets support custom DNS server configurations or integration with Azure DNS. This controls how name resolution works in the DR environment.

Ensure that DR VMs resolve names correctly, either to other recovered services in Azure or to external systems. Custom DNS is essential when restoring AD-integrated environments or for DNS forwarding to on-prem systems.

For more information, see the Microsoft Azure documentation.

Subnet routing (User-defined routes)

Azure allows the creation of custom routes (UDRs) to control traffic flow between subnets and to on-premises networks, overriding default system routing.

Create routes to direct traffic from recovery VMs through firewalls, to VPN gateways, or inspection points. This helps to enforce routing policies and ensure connectivity back to on-premises via VPN or ExpressRoute.

For more information, see the Microsoft Azure documentation.

Public IP addresses

Azure VMs can be assigned public IP addresses (static or dynamic) to allow direct Internet access when needed. This is useful for exposing services or remote access.

Assign public IPs to failover workloads that require external access (for example, web servers and remote management). Avoid for sensitive workloads. Use Bastion or VPN instead where possible.

For more information, see the Microsoft Azure documentation.

Azure Bastion

Azure Bastion enables secure, browser-based RDP/SSH access to VMs without exposing public IPs. Operates via the Azure portal and uses TLS encryption.

Access recovered VMs securely for diagnostics or manual reconfiguration after failover. This is especially useful in test failover where exposing to the internet is not desired.

For more information, see the Microsoft Azure documentation.

Azure Site-to-Site VPN

Site-to-Site VPN provides an encrypted IPsec connection between your on-premises network and Azure VNets, enabling hybrid connectivity.

Ensure seamless access to recovered workloads from on-premises networks. Critical for accessing internal apps or restoring cross-dependencies with systems still running on-premises.

For more information, see the Microsoft Azure documentation.

Azure ExpressRoute

ExpressRoute offers dedicated private connectivity between your data center and Azure, bypassing the Internet for improved speed, reliability, and security.

Recommended for enterprise-grade DR with large datasets or low-latency requirements. Use to connect on-premises systems with DR VMs in Azure without relying on VPN.

For more information, see Microsoft Azure documentation.

Network management in Microsoft Azure

Recovery networks are Azure Virtual Networks (VNets) and subnets where your backup systems (recovery servers) will run if your main systems fail.

There are two kinds of recovery networks: production recovery networks and test recovery networks.

Production recovery network

Production recovery networks are used during a real disaster when you need to move your services to Azure. Failovered workloads (VMs) are connected to these recovery networks so that business operations are resumed.

This is an example of a production network:

VNet: `dr-prod-vnet-91a4e5bf`

Subnet: `subnet-one (10.0.10.0/24)`

Test recovery network

Test recovery networks are used during manual or automated test failovers to verify that disaster recovery works without impacting your production systems.

To prevent IP conflicts or data leakage, test recovery networks must be isolated from the production environment. To avoid overlapping with production networks, it is best practice to use a dedicated VNet.

For more information about VNets planning and design, see the Microsoft Azure documentation.

Best practices for Disaster Recovery network configuration

When you configure the recovery networks in Disaster Recovery to Microsoft Azure, we recommend that you follow the following best practices:

- Separate production and test VNets and subnets
  Keep test and production environments isolated.
- Use Azure tags and naming conventions
  Clearly label networks (for example, `dr-prod` and `dr-test`) for easy identification and automation.
- Plan IP ranges carefully
  Ensure that the VNets and subnets do not conflict with on-premises networks if there is a connectivity between the DR site and on-premises, for example via IPsec VPN or ExpressRoute.
  Maintain consistent address schemes for easier routing and identity integration.
- Preconfigure required network resources
  NSGs, route tables, custom DNS, and public IPs should be configured in advance for both test and production recovery networks.

Recommendations for the Active Directory Domain Services availability

If your protected workloads need to authenticate in a domain controller, we recommend that you have an Active Directory Domain Controller (AD DC) instance at the DR site in Microsoft Azure.

Recommendations for AD DS availability in DR site in Azure

The recommendations for a dedicated AD DC instance on the DR site are the following:

- Turn off Windows Firewall.
- Join the AD DC to the Active Directory service.
- Ensure that the Azure VM has Internet access.
- Add the Active Directory feature.
- Deploy at least one domain controller in the DR site in Azure
  Recovered workloads need to authenticate, apply Group Policies, and resolve names. Deploy an additional domain controller VM in Azure in advance, before the failover.
- Use a replicated domain controller (Not read-only)
  Read-only domain controllers (RODCs) may not support all authentication scenarios after a failover. Deploy a writeable domain controller and replicate it with your on-premises AD forest.
- Ensure proper DNS configuration
  Recovered VMs must resolve domain names and locate domain controllers. Configure the recovery VNet to use the IP address of the Azure-based domain controller(s) as the custom DNS server.
- Replicate SYSVOL and ensure time syncronization

Group Policies and domain operations rely on SYSVOL replication and correct time. Ensure that SYSVOL is synchronized and configure NTP or time synchronization settings for consistency between Azure and on-premises environments.

Adding a production recovery network from Microsoft Azure

After the DR site is configured, you can add additional production recovery networks from the ones that exist in Microsoft Azure.

To add a production recovery network from Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Connectivity**.
2. In the **Production networks** pane, click **Add network**.
3. In the **Add production network** window, in the **Virtual network** field, select a virtual network.
4. In the **Subnet** field, select the subnet.
5. [Optional] To make this subnet default, select **Set this subnet as 'Default' for mapping**.
   When applying the disaster recovery protection plan to the original devices, recovery servers are configured in the default subnet.
6. In the **Mapping to local network** field, click **Add**.
   You can define mapping rules by entering one or more source local networks in the CIDR format. The service will then compare the IP address of the original device with the mapping rules. If the IP address matches any of the specified source local networks, the recovery server will be created in the corresponding Azure recovery network and subnet.
7. Enter one or more source local networks in the CIDR format.
8. To add the network, click **Done**.

Adding a test recovery network from Microsoft Azure

After the DR site is configured, you can add additional test recovery networks from the ones that exist in Microsoft Azure.

To add a test recovery network from Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Connectivity**.
2. In the **Test networks** pane, click **Add network**.
3. In the **Add test network** window, in the **Virtual network** field, select a virtual network.
4. In the **Subnet** field, select the subnet.
5. [Optional] To make this subnet default, select **Set this subnet as 'Default' for mapping**.
   When applying the disaster recovery protection plan to the original devices, recovery servers are configured in the default subnet.
6. In the **Mapping to local network** field, click **Add**.
   You can define mapping rules by entering one or more source local networks in the CIDR format. The service will then compare the IP address of the original device with the mapping rules. If the IP address matches any of the specified source local networks, the recovery server will be created in the corresponding Azure recovery network and subnet.

7. Enter one or more source local networks in the CIDR format.
8. To add the network, click **Done**.

Editing recovery networks from Microsoft Azure

You can change the test or production recovery networks from Microsoft Azure.

To edit the settings of a recovery network in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Connectivity**.
2. In the **Production networks** or **Test network** pane, click the network that you want to change, and then click **Edit**.
3. Select another network from the ones that are available Microsoft Azure.
4. Click **Done**.

Recovery servers in Micrsoft Azure

A recovery server is a replica of the original machine that is created from the protected server's backup (recovery point) that is stored in the cloud - Cyber Protect Cloud or Microsoft Azure. In case of a disaster, the workload is switched from the original server to the recovery server in Microsoft Azure.

Recovery servers are either created manually, or automatically - when you apply a disaster recovery protection plan to a workload.

No compute points are charged for running your recovery servers in Microsoft Azure. All compute usage is billed directly to your Microsoft Azure subscription.

MAC address configuration for recovery servers is not available in Disaster Recovery to Microsoft Azure.

Managing recovery servers

The following operations with recovery servers are available in Disaster Recovery to Microsoft Azure:

Operation	Description
Power on	(For servers in the Failover state) Power on the Azure VM (recovery server).
Power off	[For servers in the Failover state] Power off the Azure VM (recovery server). The power off operation stops the Azure VM but does not deallocate resources. The VM will be in the Stopped (Allocated) state. In this state, an Azure VM still reserves CPU and memory, incurring compute charges as if it is running. This state preserves the IP address and server placement of the VM.
Force power	[For servers in the Failover state] Forcefully shut down the recovery server.

Operation	Description
off	
Edit settings	Modify the recovery server settings, such as network configurations or RPO thresholds from the Cyber Protect console.
Production failover	Switch workloads to the recovery server in the production network.
Test failover	Test the recovery server in the isolated test network without impacting production. To avoid conflicts during failover, ensure that production and test networks are configured properly. Regularly test failover operations to validate the recovery server functionality.
Connect (to console)	[For servers in the Failover state] After clicking **Connect** and being redirected to Azure, you can connect to the Azure virtual machine by using native Azure options, such as: • Assigning a public IP address and connecting via Remote Desktop Protocol (RDP) or SSH. • Using Azure Bastion, a secure service for connecting to the virtual machine without a public IP.

Creating recovery servers in Microsoft Azure

Recovery servers are automatically created when you apply a disaster recovery protection plan to a workload. If a disaster recovery protection plan is not applied to the workload, you can create a recovery server manually.

Prerequisites

- A backup plan is applied to the workload.
- The DR site in Microsoft Azure is configured.

To create a recovery server in Microsoft Azure

1. In the Cyber Protect console, go to **Devices** > **All devices**.
2. Click the workload that you want to protect with Disaster Recovery, and then, in the **Actions** menu, click **Disaster recovery**.
3. Click **Create recovery server**.
4. In the **Create recovery server** wizard, on the **Server configuration** tab, configure the settings, and then click **Next**.

Setting	Description
CPU and	Size of the Azure VM. Compute usage is charged directly to your Microsoft

Setting	Description
RAM	Azure subscription by Microsoft or your partner. If some Azure VM sizes are not available, check your Azure subscription limitations. The following Azure VM types are excluded from selection: • A-series (deprecated in Microsoft Azure) • VM types that are based on ARM CPU architecture The default settings are automatically determined based on the original device CPU and RAM configuration. The RAM is matched by rounding up to the nearest B-family VM size that meets or exceeds the original RAM value, and selecting the lowest available CPU core number that satisfies the RAM requirement. If RAM data from the original machine is unavailable (for example, for Azure VMs that use agentless backups), a minimal B-family VM size is selected by default. If the selected VM series or size is not available in the target region or subscription, the system automatically selects the closest available size from any B-family series within that region.
Disk type	Disk type of the Azure VM. The disk type that you select will be applied to all recovery server disks. Only locally redundant storage (LRS) disk types are available for selection. Premium SSD v2 and Ultra SSD are not available for selection. Premium SSD v2 is automatically assigned during a failover, if 4K sector disks are detected in the backup of the original workload.
Name	Name for the recovery server that is visible in the Cyber Protect console. This name is not used for the Azure VM.
Description	Description of the recovery server

5. On the **Network** tab, configure the settings for the production and test network, and then click **Next**.

Setting	Description
Network	(For Production network) The Azure Virtual Network (VNet) and subnet for production failover. During a production failover, the server will be connected to this Azure network.
IP address in production network	(For Production network) By default, the last octet of the original machine's IP address is derived within the production network range, but you can change the IP address at any time before the failover. When the server is in the Failover state, you can modify the IP address directly in Azure.
Network	(For Test network) The Azure VNet and subnet for test failover. We recommend that the test

Setting	Description
	network is isolated within a separate VNet. During a test failover, the server will be connected to this Azure network.
IP address in test network	(For Test network) By default, the last octet is derived from the original machine, within the test network range, but you can change the IP address at any time before the failover. When the server is in the Testing failover state, you can modify the IP address directly in Azure.

Note

- By default, no public IP is assigned to the recovery server, for security reasons. Without a public IP, the recovery server is only accessible from the local network. If necessary, assign a public IP directly in the Azure portal.
- By default, Internet access is enabled for resources in Azure. No additional configuration is required to allow outbound internet traffic. If you need to restrict or isolate outgoing Internet access in a test network, you must configure appropriate security controls, such as Network Security Group (NSG) rules, User Defined Routes (UDRs), or Azure Firewall policies, depending on your requirements.

6. On the **Automated test failover** tab, do the following:
   a. [Optional] Turn on the **Automated test failover** switch.
   b. [Optional] Configure the settings.

Setting	Description
Schedule	Automated test failover runs once per month.
VM startup timeout / Minutes	The maximum period during which the system tries to start a virtual machine in Azure and take a screenshot to verify if the operating system loaded successfully. This period does not include the time for restoring the data from a cold backup archive, as this duration depends on the size of the archive. Additionally, Azure VM compute hours are not counted during the data restoration time.
Use as default timeout	Select this checkbox if you want to save the **VM startup timeout / Minutes** value as default. In this case, the value will be populated automatically when you enable automated test failover for other recovery servers.

   c. Click **Next**.
7. On the **Settings** tab, do the following:
   a. [Optional] RPO threshold defines the maximum allowable time interval between the last recovery point and the current time. You can set a value within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

b. [Optional] [If the backups for the selected machine are encrypted by using encryption as a machine property], specify the password that will be automatically used when creating a virtual machine for the recovery server from the encrypted backup.
   a. Click **Enter password**, and then enter the password for the encrypted backup and define a name for the credentials.
      By default, you will see the most recent backup in the list.
   b. To view all the backups, select **Show all backups**.
   c. Click **Save**.

Note
Although the password that you specify will be stored in a secure credentials store, saving passwords might be against your compliance obligations.

8. Click **Create**.

The recovery server is created and is in the Standby state. No compute points are charged. All compute usage is billed directly to your Azure subscription.

Note
You can configure firewall rules for the VM only in the in the Azure portal. By default, for VMs in test and production failover, all inbound connections are prohibited, and all outbound connections to Internet are allowed within the production and test VNet.

Editing the recovery server settings

When you create and apply a disaster recovery protection plan, a recovery server is created with default settings. You can edit these default settings when necessary.

To edit the settings of a recovery server in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Recovery servers**.
2. Click the server whose settings you want to edit, and then click **Edit**.
3. Edit the recovery server settings.
4. Click **Save**.

Deleting a recovery server

You can delete the recovery servers that you created in Microsoft Azure.

To delete a recovery server in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Recovery servers**.
2. Click the server that you want to delete, and then click **Delete**.
3. In the confirmation window, click **Delete**.

Failover in Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

Production failover

Failover is the process switching the workload from the original server at your local site to the recovery server.

When a recovery server is created, it stays in the **Standby** state. The corresponding virtual machine does not exist in Microsoft Azure until you start a failover. Before you start a failover, you must create at least one disk image backup (with bootable volume) of the original machine.

When you start a failover, you select the recovery point (backup) of the original machine from which a virtual machine with the predefined parameters will be created in Microsoft Azure.

When the failover completes, the recovery server state changes to **Failover**. The workload is now switched from the original machine to the recovery server in Microsoft Azure.

If the recovery server has a protection agent, to avoid interference (such as starting a backup or reporting outdated statuses to the backup component), the agent service is stopped.

Test failover

Test failover is a process of createing a temporary VM in an isolated Azure virtual network (VNet) and test recovery procedures, configurations, and applications functionality. For more information, see "Test failover in Microsoft Azure" (p. 115).

Automated test failover

Automated test failover in Microsoft Azure validates backup integrity by booting a recovery server VM from the latest backup and capturing a screenshot to confirm that the operating system started successfully. If enabled, automated test failover is started once a month. For more information, see "Automated test failover in Microsoft Azure" (p. 118).

IP Address conflict handling in failover

If the configured IP address in the production network is already in use at the time of the production failover, another available IP address from the same network will be assigned automatically.

If the configured test network IP address is already in use at the time of the test failover, another available IP address from the test network will be assigned.

Recovery of recovery server in failover to a previous point in time

To recover a server in failover, initiate a new failover from a different recovery point to restore operations.

Failover widgets

During production and test failover, you can see information about the failover performance (recovery speed) and bottlenecks on the **Activities** tab of the recovery server details. To view the **Bottleneck** widget, expand the **Creating virtual machine** activity, and then in the **Copying data from the backup to the virtual machine disks** subactivity, expand **Bottleneck**.

Performing a production failover in Microsoft Azure

When you perform a production failover, the workload is switched from the original machine to the recovery server in Microsoft Azure.

To perform a production failover in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Recovery servers**.
2. Click the server that you want to fail over, and then click **Failover**.
3. In the **Server failover** window, select **Production failover**, and then select the recovery point (backup) from which the cloud server will be started.
4. Click **Start**.

When failover completes, the service starts running on the virtual machine in Azure. Clicking **Connect** will redirect you to the virtual machine in Azure.

Test failover in Microsoft Azure

A test failover is a vital part of the Disaster Recovery as a Service (DRaaS) strategy in Microsoft Azure. It enables organizations to validate recovery processes without affecting production environments. This is done by simulating the recovery of a virtual machine (VM) from a selected recovery point (backup). The process creates a temporary VM in an isolated Azure virtual network (VNet) to test recovery procedures, configurations, and applications functionality. Although they are optional, regular test failovers are highly recommended to ensure reliable and up-to-date recovery processes. You can establish a testing schedule that is based on your organization's cost and safety requirements.

You can test several recovery servers at a time and check their interaction. In the test network, the servers communicate using their production IP addresses, but they cannot initiate TCP or UDP connections to the workloads in your local network.

Test recovery network

To ensure that a test failover does not interfere with production operations, configure an isolated virtual network (VNet) in Azure for testing purposes. Confirm that the test VNet does not have

routing or peering connections to the production VNet. Test the connectivity from a virtual machine in the test VNet to ensure that it cannot access production resources.

Test failover process

The test failover process consists of the following phases.

Initiation

During this phase, you select a recovery point and start the test failover.

Worker (temporary agent) deployment

During this phase, a worker that is used for the test failover operation is automatically deployed. The initial deployment of the worker for a test or production failover may take several minutes. Starting workers for subsequent failovers should be faster.

Data restoration

During this phase, data is copied from the backup storage to a temporary Azure Blob Storage container. The time that it takes to copy or restore data depends on the workload size. After the data is copied, the Azure Blob Storage content is converted into a managed disk, which is then used to start the temporary VM.

Recovery VM creation

The recovery VM is connected to the preconfigured isolated Azure VNet and subnet. By default, the VM is assigned an IP address where the last octet matches the original machine's IP address. You can modify the IP address before test failover in the recovery server's setting. During the test failover, you can do it directly in the Azure portal.

Ensure the the VNet is isolated from the production network to prevent unintended interactions.

Verification

After the VM is created, clicking the **Connect** button will redirect you to the specific VM in the Azure portal.

Perform all necessary tests to verify application behavior, connectivity, and recovery objectives in the isolated environment.

The test failover does not overwrite existing recovery points, thus ensuring that your backups remain intact.

Stopping the test failover

Stopping the test failover deletes the temporary VM and associated resources and the worker. If an encryption password was used, it is automatically deleted from the credential store upon stopping or completing the test failover.

You can stop the test failover at any time from the Azure portal.

Performing a test failover in Microsoft Azure

Though performing a test failover is optional, we recommend that you make it a regular process with a frequency that you find adequate in terms of cost and safety.

Important
You can perform failover only from recovery points (backups) that were created after the recovery server of the device was created.

At least one recovery point must be created before failing over to a recovery server. The maximum number of recovery points that is supported is 100.

Prerequisites

- The recovery server is configured in Azure location and has at least one recovery point created after the recovery server was created.
- An isolated Azure VNet and subnet for the test failover to ensure no interference with production environments.
- Network security group (NSG) rules are configured to meet your requirements.

To perform a test failover in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Recovery servers**.
2. Click the server that you want to fail over, and then click **Failover**.
3. In the **Server failover** window, select **Test failover**, and then select the recovery point (backup) from which the cloud server will be started.
4. Click **Start**.
5. If the backup that you selected is encrypted by using encryption as a machine property:
   a. Enter the encryption password for the backup set.

      Note
      The password will only be saved temporarily and will be used only for the current test failover operation. The password is automatically deleted from the credentials store if the test failover is stopped, or after the test failover completes.

   b. [Optional] To save the password for the backup set and use it in subsequent failover operations, select the **Store the password in a secure credentials store...** check box and then, in the **Credentials name** field, enter a name for the credentials.

      Important
      The password will be stored in a secured credentials store and will be applied automatically in subsequent failover operations. However, saving passwords might conflict with your compliance obligations.

   c. Click **Done**.

When the recovery server starts, its state changes to **Testing failover**.

6. Test the recovery server by using any of the following methods:
   - In **Disaster Recovery** > **Recovery servers**, select the recovery server, and then click **Connect**.
   - Connect to the recovery server by using RDP or SSH and the test IP address that you specified when creating the recovery server. Try the connection from both inside and outside the production network.
   - Run a script within the recovery server.

     The script may check the login screen, whether applications are started, the Internet connection, and the ability of other machines to connect to the recovery server.
7. When the test is complete, click **Stop testing**.

   The recovery server is stopped. Changes that are made to the recovery server during the test failover are not saved.

Automated test failover in Microsoft Azure

Automated test failover in Microsoft Azure validates backup integrity by booting a recovery server VM from the latest backup and capturing a screenshot to confirm that the operating system started successfully.

With automated test failover, the recovery server is tested automatically once a month without any manual interaction.

The automated test failover process consists of the following parts:

1. Creating a virtual machine in Microsoft Azure from the latest recovery point.
2. Taking a screenshot of the virtual machine.
3. Analyzing if the operating system of the virtual machine starts successfully.
4. Notifying you about the test failover status.

Note
Automated test failover consumes Azure VM compute hours.

You can configure the automated test failover in the recovery server's settings. For more information, see "Configuring automated test failover in Microsoft Azure" (p. 119).

If automated test failover is skipped for some reason, an alert will be raised.

Note
Automated test failover will fail if the backups of the original machine are encrypted by using encryption as a machine property, and the encryption password is not specified when creating the recovery server. For more information about specifying the encryption password, see "Creating recovery servers in Microsoft Azure" (p. 110).

Configuring automated test failover in Microsoft Azure

By configuring automated test failover, you can test your recovery server every month without performing any manual actions.

To configure automated test failover of a recovery server in Microsoft Azure

1. In the Cyber Protect console, go to **Disaster recovery** > **Dashboard**.
2. In the **Automated test failover** widget, click **Configure**.
3. In the **Automated test failover configuration** dialog, select one or more recovery servers for which you want to configure automated test failover, and then click **Next**.
4. Turn on the **Automated test failover** switch.
5. In the **Schedule** field, select **Monthly**.
6. [Optional] In **VM startup timeout / Minutes**, change the default value of the maximum period during which the system tries to start the virtual machine in Azure and take a screenshot to verify if the operating system loaded successfully.

   Note
   - This timeout does not include the time required to restore data from a cold backup archive. The timer starts when the virtual machine starts booting.
   - During cold data restore, Azure temporarily creates a standard Blob storage, which is later converted into a managed disk for the VM.
   - Azure VM compute hours are not counted during the data restoration process.

7. [Optional] To save the **VM startup timeout / Minutes** value as the default and have it populated automatically when you enable automated test failover for the other recovery servers, select **Set as default timeout**.
8. Click **Configure**.

On the **Recovery servers** tab, you can see that the automated test failover started. After it completes, the virtual machine will be deleted and you can find the link to the screenshot of the validated operating system in the recovery server details.

Viewing the automated test failover status

You can view the details of a completed automated test failover of a recovery server in Microsoft Azure, such as status, start time, end time, duration, and the screenshot of the virtual machine.

Note
The screenshot of the virtual machine is kept until automated test failover runs again and generates a new screenshot.

To view the automated test failover status of a recovery server in Microsoft Azure

1. In the console, go to **Disaster recovery** > **Recovery servers**, and then select the recovery server.
2. In the **Automated test failover** section, check the details of the last automated test failover.
3. [Optional] To view the screenshot of the virtual machine, click **Show screenshot**.

Disabling automated test failover

You can disable automated test failover if you want to save resources or you do not need automated test failover to be performed for a certain recovery server in Microsoft Azure.

To disable automated test failover for a recovery server in Microsoft Azure

1. In the console, go to **Disaster recovery** > **Recovery servers**, and then select the recovery server.
2. Click **Edit**.
3. In the wizard, click the **Automated test failover** tab.
4. Turn off the **Automated test failover** switch.
5. Click **Save**.

Requirements and limitations for failover of Linux VMs to Microsoft Azure

This section outlines key requirements and limitations for failing over Linux workloads to Microsoft Azure, especially in environments without Internet access.

Requirements

Azure VM Agent installation with Internet access

The Azure VM Agent is automatically installed during test and production failover.

The required Linux tools (for example, **cloud-init**, network drivers) are fetched from public repositories (for example, `archive.ubuntu.com` or `mirror.centos.org`).

For security, you can allow outbound https only to specific repositories.

Azure VM Agent installation without Internet access

You must manually install the Azure VM Agent on the source Linux machine before failover.

Without the agent, the failover may fail or result in limited VM functionality.

Limitations

Internet dependency

No Internet access on the target Azure VM will require an Azure agent to be preinstalled on the original workload, which adds configuration overhead.

No access to Internet might prevent post-failover updates (for example, networking or backup tools).

Recommendations

- Allow Internet access for production and test VNet before a failover. This is a default setup.
- Restrict outbound access to only specific Linux repositories if you know them before a failover.
- Restrict access to specific production services during test failovers.
- You can use internal mirrors or VPN to minimize Internet reliance.
- Run regular test failovers to ensure environment readiness.

Failback in Microsoft Azure

Failback is a process of moving the workload from the cloud back to a physical or virtual machine on your local site. You can perform a failback when a recovery server is in the **Failover** state, and then continue using the server on your local site.

You can perform automated failover to a virtual or physical target machine on your local site. During the failback, you can transfer the backup data to your local site while the virtual machine in the cloud continues to run. This technology helps you achieve a very short downtime period, which is estimated and displayed in the Cyber Protect console. You can view it and use this information to plan your activities and, if necessary, warn your clients about an upcoming downtime period. If you perform agent-based failback via bootable media, the downtime is even shorter, as only the delta changes will be transferred to the local site.

For a failback to a target physical machine, you can use the agent-based failback via bootable media. For more information, see "Agent-based failback via bootable media from Microsoft Azure" (p. 121).

For a failback to a target virtual machine, you can use the agent-based failback via bootable media or the agentless failback via hypervisor agent. For more information, see "Performing agent-based failback via bootable media from Microsoft Azure" (p. 123) and "Performing agentless failback via a hypervisor agent from Microsoft Azure" (p. 127).

When you cannot use the automated failback procedure, you can perform a manual failback. For more information, see "Manual failback from Microsoft Azure" (p. 130).

Note
Runbook operations support the failback in manual mode only. This means that if you start the failback process by executing a runbook that includes a **Failback server** step, the procedure will require a manual interaction: you must manually recover the machine, and confirm or cancel the failback process from the **Disaster Recovery** > **Servers** tab

Agent-based failback via bootable media from Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The agent-based failback via bootable media process is optimized for performing a failback to the original physical or virtual machine. During this process, only the delta changes are transferred to the local site.

The agent-based failback process via bootable media to a target physical or virtual machine consists of the following phases:

1. **Planning**. During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.
2. **Data transfer**. During this phase, the data is transferred from the cloud site to the local site while the virtual machine in the cloud continues to run. You can start the next phase - switchover - at any time during the data transfer phase, but you should consider the following relations.

   The longer you remain in the data transfer phase,
   - the longer the virtual machine in the cloud continues to run.
   - the more data will be transferred to your local site.
   - the higher the cost you will pay (you spend more compute points).
   - the shorter the downtime period that you will experience during the switchover phase.

   If you want to minimize the downtime, start the switchover phase after more than 90% of the data is transferred to the local site.

   If you can afford to experience a longer downtime period, and do not want to spend more compute points for running the virtual machine in the cloud, you can start the switchover phase earlier.

   Note

   The data transfer process uses a flashback technology. This technology compares the data that is available on the target machine to the data of the virtual machine in the cloud. If part of the data is already available on the target machine, it will not be transferred again. This technology makes the data transfer phase faster.

   For this reason, we recommend that you restore the server to the original machine on your local site.

3. **Switchover**. During this phase, the virtual machine in the cloud is turned off, and the remaining data - including the last backup increment - is transferred to the local site. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process.
4. **Validation**. During this phase, the machine on the local site is ready, and you can reboot it using a Linux-based bootable media. You can verify if the virtual machine is working correctly, and:
   - If everything is working as expected, confirm the failback. After the failback confirmation, the virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state. This is the end of the failback process.
   - If something is wrong, you can cancel the failover and return to the planning phase.

Note
After the bootable media has been rebooted, you will not be able to use it again. If, at the validation phase, you discover something wrong, you must register a new bootable media and start the failback process again.
However, as flashback technology will be used, the data that is already on the local site will not be transferred again, and the failback process will be much faster.

Performing agent-based failback via bootable media from Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can perform agent-based failback via bootable media from Microsoft Azure to a target physical or virtual machine on your local site.

Note
The data transfer process uses a flashback technology. This technology compares the data that is available on the target machine to the data of the virtual machine in the cloud. If part of the data is already available on the target machine, it will not be transferred again. This technology makes the data transfer phase faster.

For this reason, we recommend that you restore the server to the original machine on your local site.

Prerequisites

- The agent that you will use to perform failback is online and is not currently used for another failback operation.
- Your Internet connection is stable.
- A registered bootable media is available.
  For more information, see Creating bootable media to recover operating systems.
- The target machine is the original machine on your local site, or has the same firmware as the original machine.
- There is at least one full backup of the virtual machine in the cloud.

To perform a failback to a physical machine

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.
4. In the **Failback type** field, select **Agent-based via bootable media**.
5. In the **Target bootable media** field, click **Specify**, select the bootable media, and then click **Done**.

Note
We recommend that you use ready-made bootable media as it is already configured. For more information, see Creating bootable media to recover operating systems.

6. [Optional] To change the default disk mapping, in the **Disk mapping** field, click **Specify**, map the disks of the backup to the disks of the target machine, and then click **Done**.
7. Click **Start data transfer** and then, in the confirmation window, click **Start**.

Note
If there is no backup of the virtual machine in the cloud, the system will perform a backup automatically before the data transfer phase.

The data transfer phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows how much data is already transferred to the local site, and the total amount of data that must be transferred. The total amount of data includes the data from the last backup before the data transfer phase was started, and the backups of the newly generated data (backup increments), as the virtual machine continues to run during the data transfer phase. For this reason, the **Progress** values increase with time. As the system uses a flashback technology during the data transfer and does not transfer the data that is already available on the target machine, the progress might be faster than what is initially calculated by the console.
Downtime estimation	This parameter shows how much time the virtual machine in the cloud will be unavailable if you start the switchover phase now. The value is calculated based on the values of the **Progress** parameter, and decreases with time. As the system uses a flashback technology during the data transfer and does not transfer the data that is already available on the target machine, the downtime might be much shorter than the value that is initially displayed in the console.

8. Click **Switchover** and then, in the confirmation window, click **Switchover** again.
   The switchover phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows the progress of restoring the machine on the local site.
Estimated time to finish	This parameter shows the approximate time when the switchover phase will be completed and you will be able to start the machine on the local site.

Note
If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

9. After the **Switchover** phase completes, reboot the bootable media, and then verify that the physical machine on your local site is working as expected.
   For more information, see Recovering disks by using bootable media.
10. Click **Confirm failback** and then, in the confirmation window, click **Confirm** to finalize the process.
    The virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state.

Note
Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Agentless failback via a hypervisor agent from Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

The agentless failback via a hypervisor agent process is optimized for performing a failback from Microsoft Azure to a new virtual machine. If you want to perform a failback to the original virtual machine, follow the procedure for agent-based failback via bootable media.

The agentless failback via a hypervisor agent consists of four phases.

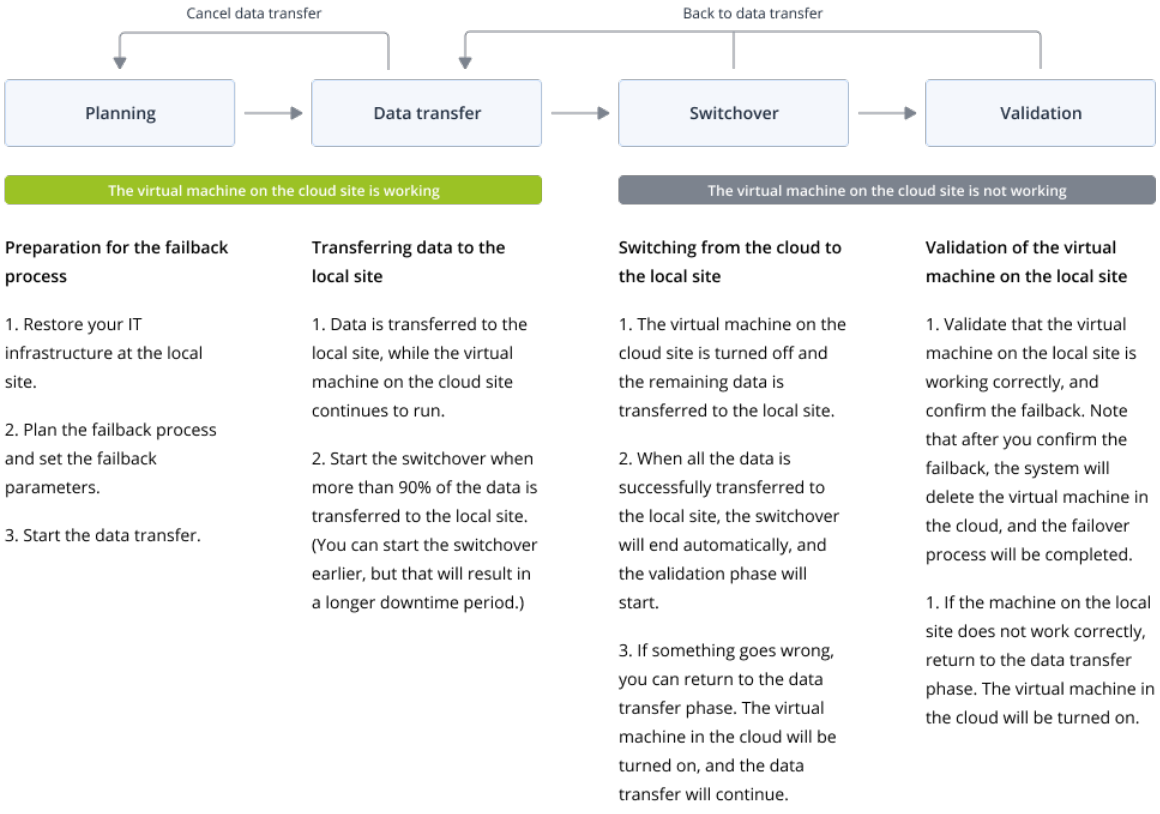

1. **Planning**. During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.

   Note
   To minimize the total time for the failback process, we recommend that you start the data transfer phase immediately after you set up your local servers, and then continue with the configuration of the network and the rest of the local infrastructure during the data transfer phase.

2. **Data transfer**. During this phase, the data is transferred from the cloud site to the local site while the virtual machine in the cloud continues to run. You can start the next phase - switchover - at any time during the data transfer phase, but you should consider the following relations.

   The longer you remain in the data transfer phase,
   - the longer the virtual machine in the cloud continues to run.
   - the more data will be transferred to your local site.
   - the higher the cost you will pay (you spend more compute points).
   - the shorter the downtime period that you will experience during the switchover phase.

   If you want to minimize the downtime, start the switchover phase after more than 90% of the data is transferred to the local site.

If you can afford to experience a longer downtime period, and do not want to spend more compute points for running the virtual machine in the cloud, you can start the switchover phase earlier.

If you cancel the failback process during the data transfer phase, the transferred data will not be deleted from the local site. To avoid potential issues, manually delete the transferred data before you start a new failback process. The following data transfer process will start from the beginning.

3. **Switchover**. During this phase, the virtual machine in the cloud is turned off, and the remaining data - including the last backup increment - is transferred to the local site. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process.

   You can view the estimated time to finish (downtime period) of this phase in the Cyber Protect console. When all the data is transferred to the local site (there is no data loss, and the virtual machine on the local site is an exact copy of the virtual machine in the cloud), the switchover phase completes. The virtual machine on the local site is recovered, and the validation phase starts automatically.
4. **Validation**. During this phase, the virtual machine on the local site is ready and automatically started. You can verify if the virtual machine is working correctly, and:
   - If everything is working as expected, confirm the failback. After the failback confirmation, the virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state. This is the end of the failback process.
   - If something is wrong, you can cancel the switchover and return to the data transfer phase.

Performing agentless failback via a hypervisor agent from Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can perform an agentless failback from Microsoft Azure to a target virtual machine on your local site via a hypervisor agent.

Prerequisites

- The agent that you will use to perform failback is online and is not currently used for another failback operation.
- Your internet connection is stable.
- There is at least one full backup of the virtual machine in the cloud.

To perform an agentless failback to a virtual machine via a hypervisor agent

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.

4. In the **Failback parameters** section, in the **Failback type** field, select **Agentless via hypervisor agent**, and then configure the other parameters.

   Some of the **Failback parameters** are populated automatically with suggested values, but you can change them.

   The following table provides more information about the **Failback parameters**.

Parameter	Description
Backup size	Amount of data that will be transferred to your local site during the failback process. After you start the failback process to a target virtual machine, the **Backup size** will be increasing during the data transfer phase, because the virtual machine in the cloud will continue to run and generate new data. To calculate the estimated downtime period during the failback process to a target virtual machine, take 10% of the **Backup size** value (as we recommend that you start the switchover phase after 90% of the data is transferred to your local site), and divide it by the value of your Internet speed. **Note** The value of the Internet speed will decrease when you perform several failback processes at the same time.
Target machine location	Failback location: a VMware ESXi host or a Microsoft Hyper-V host. You can select from all the hosts that have an agent which is registered with the Cyber Protection service.
Agent	Agent which will perform the failback operation. You can use one agent to perform one failback operation at the same time. You can select an agent that is online and is not currently used for another failback process, has a version which supports the failback functionality, and has rights to access the backup. You can install several agents on VMware ESXi hosts, and start a separate failback process using each of them. These failback processes can be performed at the same time.
Target machine settings	Virtual machine settings: • **Virtual processors**. Select the number of virtual processors. • **Memory**. Select how much memory the virtual machine will have. • **Units**. Select the units for the memory. • [Optional] **Network adapters**. To add a network adapter, click **Add**, and select a network in the **Network** field. When you are ready with the changes, click **Done**.
Path	(For Microsoft Hyper-V hosts) Folder on the host where your machine will

Parameter	Description
	be stored. Ensure that there is enough free memory space on the host for the machine.
Datastore	[For VMware ESXi hosts] Datastore on the host where your machine will be stored. Ensure that there is enough free memory space on the host for the machine.
Provisioning mode	Method of allocation of the virtual disk. For Microsoft Hyper-V hosts: • **Dynamically expanding** (default value). • **Fixed size**. For Microsoft Hyper-V hosts: • **Thin** (default value). • **Thick**.
Target machine name	Name of the target machine. By default, the target machine name is the same as the recovery server name. The target machine name must be unique on the selected **Target machine location**.

5. Click **Start data transfer**, and then in the confirmation window, click **Start**.

Note
If there is no backup of the virtual machine in the cloud, the system will perform a backup automatically before the data transfer phase.

The **Data transfer** phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows how much data is already transferred to the local site, and the total amount of data that must be transferred. The total amount of data includes the data from the last backup before the data transfer phase was started, and the backups of the newly generated data (backup increments), as the virtual machine continues to run during the data transfer phase. For this reason, both values of the **Progress** parameter increase with time.
Downtime estimation	This parameter shows how much time the virtual machine in the cloud will be unavailable if you start the switchover phase now. The value is calculated based on the values of the **Progress** parameter, and decreases with time.

6. Click **Switchover** and then, in the confirmation window, click **Switchover** again.
   The switchover phase starts. The console displays the following information:

Field	Description
Progress	This parameter shows the progress of restoring the machine on the local site.
Estimated time to finish	This parameter shows the approximate time when the switchover phase will be completed and you will be able to start the machine on the local site.

Note
If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

7. After the **Switchover** phase completes and the virtual machine at your local site is started automatically, validate that it is working as expected.
8. Click **Confirm failback**, and then in the confirmation window, click **Confirm** to finalize the process.
   The virtual machine in the cloud is deleted, and the recovery server returns to the **Standby** state.

Note
Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Manual failback from Microsoft Azure

Note
We recommend that you use the failback process in a manual mode only when you are advised to do so by the Support team.

You can also start a failback process in a manual mode. In this case, the data transfer from the backup in the cloud to the local site will not be done automatically. It must be done manually after the virtual machine in the cloud is powered off. This makes the failback process in a manual mode much slower, and you should expect a longer downtime period.

The failback process in a manual mode consists of the following phases:

1. **Planning**. During this phase, you restore the IT infrastructure at your local site (such as the hosts and the network configurations), configure the failback parameters, and plan when to start the data transfer.
2. **Switchover**. During this phase, the virtual machine in the cloud is turned off, and the newly generated data is backed up. If no backup plan is applied on the recovery server, a backup will be performed automatically during the switchover phase, which will slow down the process. When the backup is complete, you recover the machine to the local site manually. You can either

recover the disk by using bootable media, or recover the entire machine from the cloud backup storage.

3. **Validation**. During this phase, you verify that the physical or virtual machine at the local site is working correctly, and confirm the failback. After the confirmation, the virtual machine on the cloud site is deleted, and the recovery server returns to the **Standby** state.

Performing a manual failback from Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

You can perform a manual failback from Microsoft Azure to a target physical or virtual machine on your local site.

To perform a manual failback

1. In the Cyber Protect console, go to **Disaster recovery** > **Servers**.
2. Select the recovery server that is in the **Failover** state.
3. Click the **Failback** tab.
4. In the **Target** field, select **Physical machine**.
5. Click the gear icon, and then enable the **Use manual mode** switch.
6. [Optional] Calculate the estimated downtime period during the failback process, by dividing the **Backup size** value by the value of your Internet speed.

   Note
   The value of the Internet speed will decrease when you perform several failback processes at the same time.

7. Click **Switchover**, and then in the confirmation window, click **Switchover** again.
   The virtual machine on the cloud site is turned off.

   Note
   If no backup plan is applied to the virtual machine in the cloud, a backup will be performed automatically during the switchover phase, which will cause a longer downtime.

8. Recover the server from the cloud backup to the physical or virtual machine on your local site.
   For more information, see Recovery of virtual machines and Recovery of physical machines.

   Note
   If you are performing a failback from a Microsoft Azure virtual machine to the original Azure virtual machine, use the recovery options that are described in "Failback from Microsoft Azure to an Azure virtual machine" (p. 132).

9. Ensure that the recovery is completed and the recovered machine works properly, and click **Machine is restored**.

10. If everything is working as expected, click **Confirm failback**, and then in the confirmation window, click **Confirm** again.

    The recovery server and recovery points become ready for the next failover. To create new recovery points, apply a protection plan to the new local server.

    Note
    Applying a protection plan on the recovered server is not part of the failback process. After the failback process completes, apply a protection plan on the recovered server to ensure that it is protected again. You may apply the same protection plan that was applied on the original server, or a new protection plan that has the **Disaster Recovery** module enabled.

Failback from Microsoft Azure to an Azure virtual machine

You can perform a failback from a Microsoft Azure virtual machine to the original Azure virtual machine by following the procedure for "Performing a manual failback from Microsoft Azure" (p. 131) and using one of the following recovery options in step 8:

Recovery options

- Agentless recovery

  Supports recovery only to a new Azure VM that is created automatically.

  Process: Configure the Azure connection in the Cyber Protect console (**Devices** > **Add** > **Microsoft Azure virtual machine**).

  A backup appliance VM is deployed in the Azure subscription to manage the recovery.

  Recovery flow:

  On the **Backup storage** screen, select a backup. Recover as an Azure VM.

  You can use the same flow for physical, virtual, or agent-based backups.

  Cost consideration: The appliance VM, when used for recovery only, can be powered on only during recovery, then turned off manually.
- Agent-based recovery

  Supports recovery to the same Azure VM (if the original VM with the agent is available) or a new Azure VM that has a new agent installed.

  Process: Manually create a clean Windows/Linux VM in Azure. Install the protection agent. Use the agent to browse and recover backups from Acronis Cloud Storage.

  For more information, see Recovering Microsoft Azure and Amazon EC2 machines.

Runbooks in Micrsoft Azure

A runbook is a set of instructions describing how to launch the production environment in the cloud.

With runbooks, you can:

- Automate the failover of one or multiple servers.
- Automatically check the failover result by pinging the server IP address and checking the connection to the port you specify.
- Set the sequence of operations for servers running distributed applications.
- Include manual operations in the workflow.
- Verify the integrity of your disaster recovery solution, by executing runbooks in the test mode.

Runbooks to automate failover operation and ensure that your systems are recovered in the correct order to address dependencies between applications.

Runbooks let you automate a failover of one or multiple servers. You can set the correct sequence of failover operations for servers running distributed applications. You can execute runbooks in either test or production mode, to check the integrity of your disaster recovery solution.

Creating a runbook in Microsoft Azure

A runbook consists of steps that are executed consecutively. A step consists of actions that start simultaneously.

To create a runbook in Microsoft Azure

1. In the Cyber Protection console, go to **Disaster recovery** > **Runbooks**.
2. Click **Create runbook**.
3. Click **Add step**.
4. Click **Add action**, and then select the action that you want to add to the step.

Action	Description
Failover server	Performs a failover of a recovery server. To define this action, you must select a recovery server and configure the runbook parameters that are available for this action. For more information about these parameters, see "Runbook parameters in Microsoft Azure" (p. 135). **Note** If the backup of the server that you select is encrypted by using encryption as a machine property, the **Failover server** action will be paused and will be changed automatically to **Interaction required**. To proceed with the execution of the runbook, you will have to provide the password for the encrypted backup.
Failback server	Performs a failback of a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these parameters, see "Runbook parameters in Microsoft Azure" (p. 135).

Action	Description
	Note Runbook operations support the failback in manual mode only. This means that if you start the failback process by executing a runbook that includes a **Failback server** step, the procedure will require a manual interaction: you must manually recover the machine, and confirm or cancel the failback process from the **Disaster Recovery** > **Servers** tab.
Start server	Starts a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these settings, see For more information about these parameters, see "Runbook parameters in Microsoft Azure" (p. 135). **Note** The **Start server** action is not applicable for test failover operations in runbooks. If you try executing such an action, it will fail with the following error message: Failed: The action is not applicable to the current server state.
Stop server	Stops a cloud server. To define this action, you must select a cloud server and configure the runbook parameters that are available for this action. For more information about these parameters, see "Runbook parameters in Microsoft Azure" (p. 135). **Note** The **Stop server** action is not applicable for test failover operations in runbooks. If you try executing such an action, it will fail with the following error message: Failed: The action is not applicable to the current server state.
Manual operation	A manual operation requires an interaction from a user. To define this action, you must enter a description. When a runbook sequence reaches a manual operation, the runbook will be paused and will not proceed until a user performs the required manual operation, such as clicking the confirmation button.
Execute runbook	Executes another runbook. To define this action, you must choose a runbook. A runbook can include only one execution of a given runbook. For example, if you added the action "execute Runbook A", you can add the action "execute Runbook B", but cannot add another action "execute Runbook A".

5. Define the runbook parameters for the action. For more information about these parameters, see "Runbook parameters" (p. 90).
6. [Optional] To add a description of the step:
   a. Click the ellipsis icon, and then click **Description**.
   b. Enter a description of the step.
   c. Click **Done**.
7. Repeat steps 3-6 until you create the desired sequence of steps and actions.

8. [Optional] To change the default name of the runbook:
   a. Click the ellipsis icon.
   b. Enter the name of the runbook.
   c. Enter a description of the runbook.
   d. Click **Done**.
9. Click **Save**.
10. Click **Close**.

Runbook parameters in Microsoft Azure

Runbook parameters are specific settings that you must configure to define a runbook action. They define the runbook behavior depending on the action initial state or result.

The following table describes the configurable runbook parameters for each action.

Runbook parameter	Available for action	Description
Continue if already done	• **Failover server** • **Start server** • **Stop server** • **Failback server**	This parameter defines the runbook behavior when the required action is already done (for example, a failover has already been performed or a server is already running). When enabled, the runbook issues a warning and proceeds. When disabled, the action fails, and then the runbook fails too. By default, this parameter is enabled.
Continue if failed	• **Failover server** • **Start server** • **Stop server** • **Failback server**	This parameter defines the runbook behavior when the required action fails. When enabled, the runbook issues a warning and proceeds. When disabled, the action fails, and then the runbook fails too. By default, this parameter is disabled.

Operations with runbooks in Microsoft Azure

Note
The availability of this feature depends on the service quotas that are enabled for your account.

When a runbook is not running, the following operations are available: execute, edit, clone, vie details, and delete.

Execute

Every time you click **Execute**, you are prompted for the execution parameters. These parameters apply to all failover and failback operations that are included in the runbook. If there are runbooks that are specified in the **Execute runbook** operations, they will inherit these parameters from the main runbook.

To execute a runbook

1. In the Cyber Protect console, go to **Disaster Recovery** > **Runbooks**.
2. Click the runbook that you want to execute, and then click **Execute**.
3. In the **Execution parameters** window, configure the parameters.

Parameter	Description
Failover and failback mode	Select whether you want to run a test failover (by default) or a real (production) failover. The failback mode will correspond to the selected failover mode.
Failover recovery point	Select the most recent recovery point (by default) or select a point in time in the past. If you select a point in time in the past, the recovery points that are closest before the specified date and time will be selected for each server.

4. Click **Start**.

The runbook execution starts. You can stop the runbook execution. The software will complete all of the already started actions except the ones that require user interaction.

Edit

To edit a runbook

1. In the Cyber Protect console, go to **Disaster Recovery** > **Runbooks**.
2. Click the runbook that you want to execute, and then click **Edit**.
3. Edit the runbook.
4. Click **Save**.

Clone

When you clone a runbook the history of the original runbook is not cloned.

To clone a runbook

1. In the Cyber Protect console, go to **Disaster Recovery** > **Runbooks**.
2. Click the runbook that you want to execute, and then click **Clone**.
3. [Optional] In the clone window, edit the default name and enter a description.
4. Click **Clone**.

View details

To view details about a runbook

1. In the Cyber Protect console, go to **Disaster Recovery** > **Runbooks**.
2. Click the runbook whose details you want to view.
   The runbook details and execution history are shown.
3. [Optional] To view the execution log of a specific execution, click the line that corresponds to it.

Delete

To delete a runbook

1. In the Cyber Protect console, go to **Disaster Recovery** > **Runbooks**.
2. Click the runbook that you want to execute, and then click **Delete**.
3. In the confirmation window, click **Delete**.

Workers in Microsoft Azure

Workers are temporary, on-demand agents that run in your Azure subscription in the system resource group (with the prefix: "cyber-protect-rg*") during disaster recovery operations, such as failover, back up after failover, and failback. One worker is deployed for each operation and is deleted after operation completes. This on-demand model helps reduce costs by automatically deploying workers only during active DR operations and removing them afterward.

The initial deployment of the worker for a test or production failover might take several minutes. Starting workers for the subsequent failovers should be faster.

To monitor the status of active workers, go to **Public clouds**, click the connection to your Microsoft Azure subscription, and then click the **Workers** tab.

If a disaster recovery operation fails because of an error with the worker, you can generate a report and view more information about the error. To do this, go to the **Workers** tab, turn on the **Troubleshooting** switch, click the **Instructions** link, and then follow the instructions.

Azure resources that are created during the DR site configuration and failover

When you add a connection to your Azure subscription, the system creates a resource group with the prefix `cyber-protect-rg*` in the Azure region that you selected during the configuration of the Azure connection. This resource group has the following tag: **Application:CyberProtect**. For more information, see Microsoft Azure connection security and audit (72684).

When the (DR) site configuration is completed, a resource group with the prefix `dr-rg*` is created to aggregate the resources for failover. This resource group has the following tag: **Application:DisasterRecovery**.

During disaster recovery operations, temporary workers (agents) are deployed to orchestrate failover, failback, and post-failover backup operations. These are temporary Azure VMs that are running in the `cyber-protect-rg*` resource group. All Azure resources that are required for a disaster recovery failover, such as storage accounts, and failed-over VMs are created in the `dr-rg*` resource group. You can manage the Azure region for the Azure connection resource group, the DR resource group, and their associated resources during the DR site configuration.

Soft deletion of tenants that have a disaster recovery site in Microsoft Azure

The soft deletion of a tenant feature (Recycle bin) is supported for the Microsoft Azure DR site configuration.

To ensure business continuity during the soft and hard deletion, Azure VMs in failover will not be stopped or deleted.

Corner case

If you disable the **DR and direct backup to Azure** offering item (either intentionally or accidentally), a soft delete of the Azure DR configuration is triggered. This means that the configuration is removed but remains in the Recycle bin for 30 days. If during this period you configure a new DR site in a different location (not Microsoft Azure), to recover the initial DR configuration to Microsoft Azure, you must enable the **DR and direct backup to Azure** offering item, remove the new configuration, and then recover the initial one.

Disaster recovery dashboard

The Disaster recovery **Dashboard** page contains widgets that provide a real-time overview and actionable insights for your disaster recovery site throughout its lifecycle. This helps you to:

- Quickly detect and resolve problems by monitoring the status of recovery and primary servers.
- Avoid overuse of compute points by monitoring the number of servers that consume them.

Disaster recovery - eligible devices

The widget shows the total number of devices that are protected by Disaster recovery (have a recovery server) and the total number of devices that are eligible for protection by Disaster recovery.

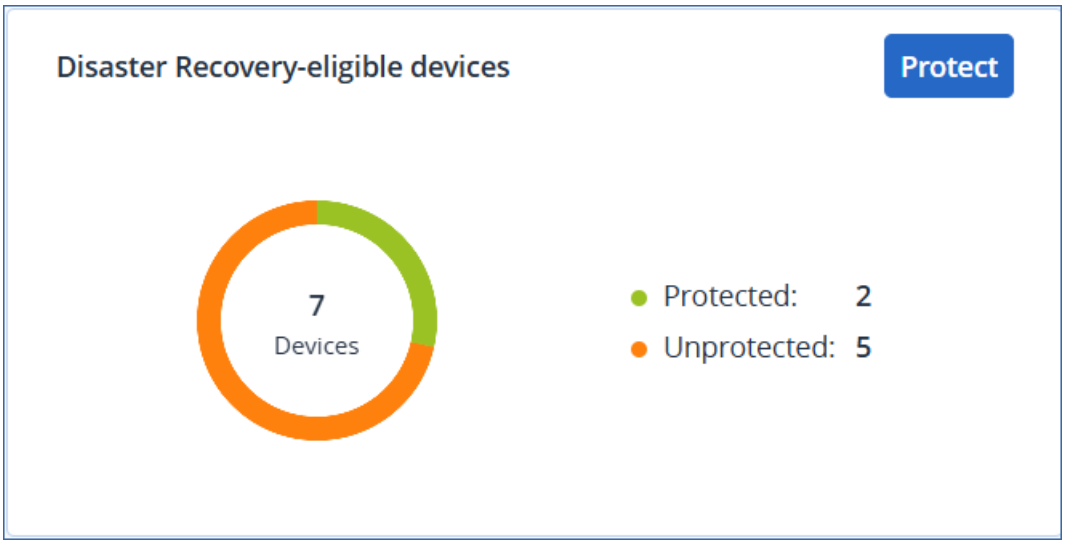

To go to the **All devices** page where you can configure Disaster recovery for the eligible devices, click **Protect**.

Health check

The widget shows information about the health of your disaster recovery infrastructure. You can check the state of the site configuration, network availability (available for Disaster Recovery to Cyber Protect Cloud only), and if there are missing service quotas.

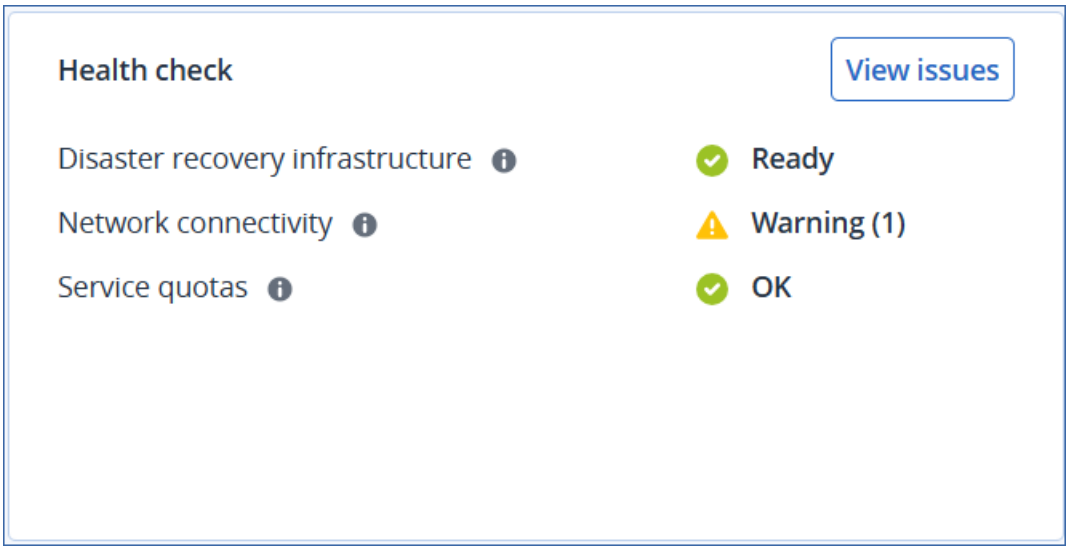

To view more information about detected issues (warnings or errors), click **View issues**.

Automated test failover

The widget shows information about the automated test failover operations of your recovery servers.

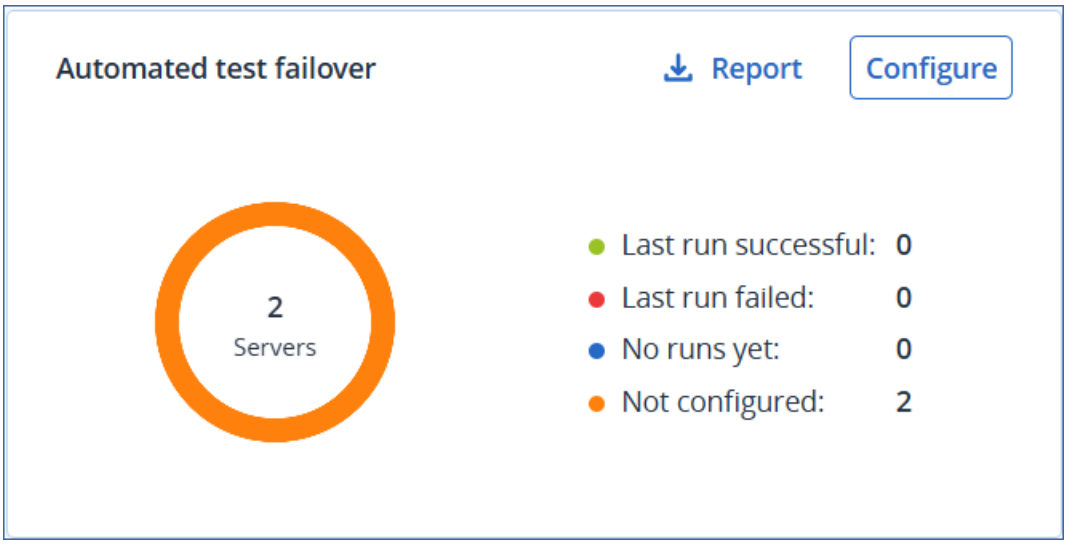

To start configuring automated test failover for your servers, click **Configure**.

To download the data from the widget, click **Report**.

Recovery servers in failover

The widget shows the number and status of the recovery servers that are in production or test failover.

If there are no recovery servers currently running in the cloud, you will see a zero. If there is an issue with any server, you will see a warning or an error status.

The displayed compute points usage per hour is a real-time snapshot, not a historical value. This means that if, for example, you have two recovery servers and each one of them uses eight points, you will see a total of 16 points displayed in the widget.

You can use this information to estimate the cost of running these servers, and also as a reminder to stop a failover of a server that you do not need anymore.

Primary servers

The widget is available for Disaster Recovery to Cyber Protect Cloud only. It shows the number and status of the primary servers in your environment, and their compute point usage per hour. You can use this information to quickly spot and resolve any issues.

The displayed compute points usage per hour is a real-time snapshot, not a historical value. This means that if, for example, you have two recovery servers and each one of them uses eight points, you will see a total of 16 points displayed in the widget.

You can use this information to estimate the cost of running these servers, and also as a reminder to stop a failover of a server that you do not need anymore.

Cloud server alerts

The widget shows the latest alerts by severity, so that you can see critical alerts at a glance. The **Alert type** and **Recovery server** values in the widget are links that open the alert details and recovery server details, respectively.

Disaster Recovery compatibility with encryption software

Disaster recovery is compatible with the following disk-level encryption software:

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

Note

- For workloads with disk-level encryption, we recommend that you install the protection agent in the guest operating system of the workload, and perform agent-based backups.
- Failover and failback will not be supported for agentless backups of encrypted workloads.

For more information about compatibility with encryption software, see the Cyber Protection User Guide.

Site-to-site Open VPN - Additional information

When you create a recovery server, you configure its **IP address in production network** and its **Test IP address**.

After you perform failover (run the virtual machine in the cloud), and log in to the virtual machine to check the IP address of the server, you see the **IP address in production network**.

When you perform test failover, you can reach the test server only by using the **Test IP address**, which is visible only in the configuration of the recovery server.

To reach a test server from your local site, you must use the **Test IP address**.

Note

The network configuration of the server always shows the **IP address in production network** (as the test server mirrors how the production server would look). This happens because the test IP address does not belong to the test server, but to the VPN gateway, and is translated to the production IP address using NAT.

The diagram below shows an example of the Site-to-site Open VPN configuration. Some of the servers in the local environment are recovered to the cloud using failover (while the network infrastructure is ok).

1. The customer enabled Disaster Recovery by:
   a. configuring the VPN appliance (14), and connected it to the dedicated cloud VPN server (15)
   b. protecting some of the local servers with Disaster Recovery (1, 2, 3, x8, and x10)

      Some servers on the local site (like 4) are connected to networks which are not connected to the VPN appliance. Such servers are not protected with Disaster Recovery.
2. Part of the servers (connected to different networks) work in the local site: (1, 2, 3, and 4)
3. The protected servers (1, 2, and 3) are being tested with test failover (11, 12, and 13)
4. Some servers in the local site are unavailable (x8, x10). After performing failover, they become available in the cloud (8, and 10)
5. Some primary servers (7, and 9), connected to different networks, are available in the cloud environment

6. (5) is a server in the Internet with a public IP address
7. (6) is a workstation connected to the cloud using a Point-to-site VPN connection (p2s)

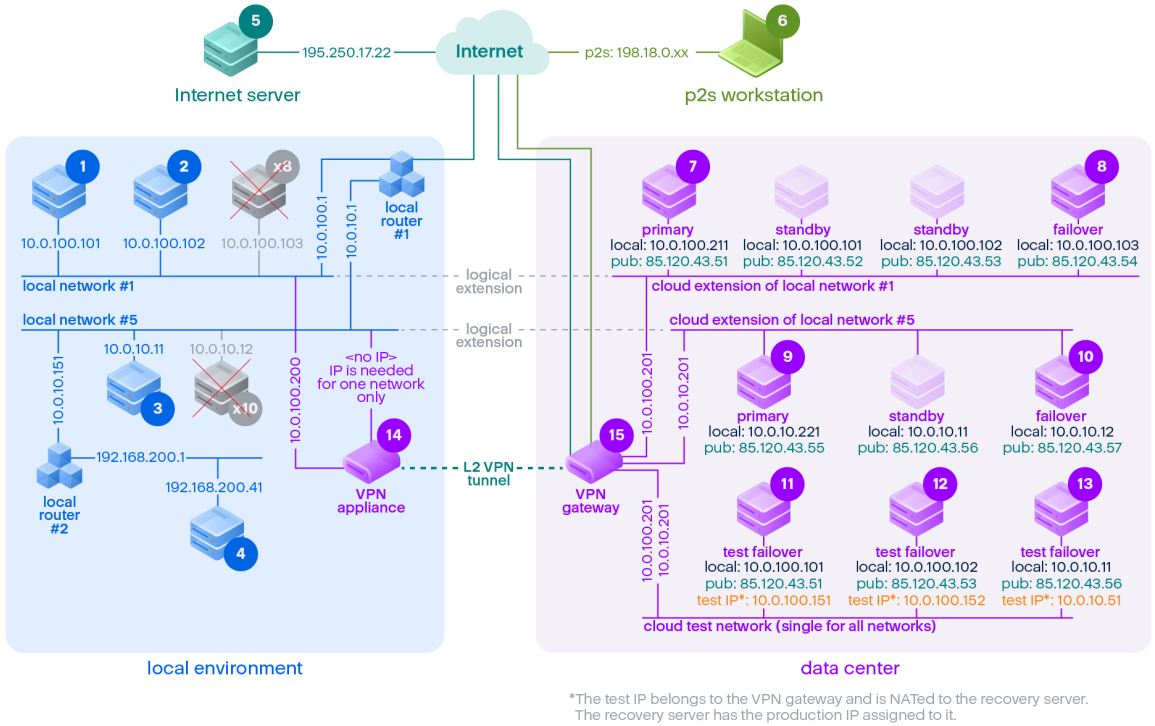

In this example, the following connection setup is available (for example, "ping") from a server in the **From:** row to a server in the **To:** column.

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
From:		local	local	local	local	internet	p2s	primary	failover	primary	failover	test failover	test failover	test failover	VPN appliance	VPN server

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	local		direct	via local router 1	via local router 2	via local router 1 and Internet	no	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: NAT (VPN server) via local router 1 and Internet: pub	via tunnel: NAT (VPN server) via local router 1 and Internet: pub	via local router 1 and tunnel: NAT (VPN server) via local router 1 and Internet: pub	direct	no
2	local	direct		via local router 1	via local router 2	via local router 1 and Internet	no	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: local via local router 1 and Internet: pub	via tunnel: NAT (VPN server) via local router 1 and Internet: pub	via tunnel: NAT (VPN server) via local router 1 and Internet: pub	via local router 1 and tunnel: NAT (VPN server) via local router 1 and Internet: pub	direct	no
3	local	via local	via local		via local	via local router 1	no	via tunnel:	via tunnel:	via tunnel:	via tunnel:	via tunnel:	via tunnel:	via local	via local router	no

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
		router 1	router 1		router 2	and Internet		local via local router 1 and Internet: pub	local via local router 1 and Internet: pub	local via local router 1 and Internet: pub	local via local router 1 and Internet: pub	NAT (VPN server) via local router 1 and Internet: pub	NAT (VPN server) via local router 1 and Internet: pub	router 1 and tunnel: NAT (VPN server) via local router 1 and Internet: pub		
4	local	via local router 2 and router 1	via local router 2 and router 1	via local router 2		via local router 2, and router 1, and Internet	no	via local router 2 and tunnel: local via local router 2, and local router 1, and Internet: pub	via local router 2 and tunnel: local via local router 2, and local router 1, and Internet: pub	via local router 2 and tunnel: local via local router 2, and local router 1, and Internet: pub	via local router 2 and tunnel: local via local router 2, and local router 1, and Internet: pub	via tunnel: NAT (VPN server) via local router 2, and router 1, and Internet: pub	via tunnel: NAT (VPN server) via local router 2, and router 1, and Internet: pub	via tunnel: NAT (VPN server) via local router 2, and router 1, and Internet: pub	via local router 2	no
5	internet	no	no	no	no		n/a	via Interne	via Interne	via Interne	via Interne	via Interne	via Interne	via Interne	no	no

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
								t: pub	t: pub	t: pub	t: pub	t: pub	t: pub	t: pub		
6	p2s	no	no	no	no	via Internet		via p2s VPN (VPN server): local via Internet: pub	via p2s VPN (VPN server): local via Internet: pub	via p2s VPN (VPN server): local via Internet: pub	via p2s VPN (VPN server): local via Internet: pub	via p2s VPN - NAT (VPN server) via Internet: pub	via p2s VPN - NAT (VPN server) via Internet: pub	via p2s VPN - NAT (VPN server) via Internet: pub	no	no
7	primary	via tunnel	via tunnel	via tunnel and local router 1	via tunnel and local router 1 and 2	via Internet (via VPN server)	no		direct in cloud: local	via tunnel and local router 1: local	via tunnel and local router 1: local	via VPN server: NAT	via VPN server: NAT	via tunnel and local router 1: NAT	no	DHCP and DNS protocols only
8	failover	via tunnel	via tunnel	via tunnel and local router 1	via tunnel and local router 1 and 2	via Internet (via VPN server)	no	direct in cloud: local		via tunnel and local router 1: local	via tunnel and local router 1: local	via VPN server: NAT	via VPN server: NAT	via tunnel and local router 1: NAT	no	DHCP and DNS protocols only
9	primary	via tunne	via tunne	via tunne	via tunne	via Internet	no	via tunnel	via tunnel		direct in	via tunnel	via tunnel	via VPN server:	no	DHCP and DNS

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
		l and local router 1	l and local router 1	l	l	(via VPN server)		and local router 1: local	and local router 1: local		cloud: local	and local router 1: NAT	and local router 1: NAT	NAT		protocol s only
10	failover	via tunne l and local router 1	via tunne l and local router 1	via tunne l	via tunne l	via Internet (via VPN server)	no	via tunnel and local router 1: local	via tunnel and local router 1: local	direct in cloud: local		via tunnel and local router 1: NAT	via tunnel and local router 1: NAT	via VPN server: NAT	no	DHCP and DNS protocol s only
11	test failover	no	no	no	no	via Internet (via VPN server)	no	no	no	no	no		direct in cloud: local	via VPN server: local (routin g)	no	DHCP and DNS protocol s only
12	test failover	no	no	no	no	via Internet (via VPN server)	no	no	no	no	no	direct in cloud: local		via VPN server: local (routin g)	no	DHCP and DNS protocol s only
13	test failover	no	no	no	no	via Internet (via VPN server)	no	no	no	no	no	via VPN server: local (routin g)	via VPN server: local (routin g)		no	DHCP and DNS protocol s only

	To:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
14	VPN appliance	direct	direct	via local router 1	via local router 2	via Internet (local router 1)	no	no	no	no	no	no	no	no		no
15	VPN server	no	no	no	no	no	no	no	no	no	no	no	no	no	no	

Glossary

C

Cloud server

General reference to a recovery or a primary server.

Cloud site (or DR site)

Remote site hosted in the cloud and used for running recovery infrastructure, in case of a disaster.

F

Failback

The process of restoring servers to the local site after they have been shifted to the cloud site during the failover.

Failover

Switching the workload or application to the cloud site in case of a natural or man-made disaster on the local site.

Finalization

The intermediate state for production failover or recovery process of the cloud server. This process implies transferring the server's virtual disks from the backup storage ("cold" storage) to the disaster recovery storage ("hot" storage). During the finalization, the server is accessible and operable although the performance is lower than normal.

L

Local site

The local infrastructure deployed on your company's premises.

P

Point-to-site (P2S) connection

A secure VPN connection from outside to the cloud and local sites by using your endpoint devices (such as a computer or laptop).

Primary server

A virtual machine that does not have a linked machine on the local site (such as a recovery server). Primary servers are used for protecting an application or running various auxiliary services (such as a web server).

Production network

The internal network extended by means of a VPN tunneling and covering both local and cloud sites. Local servers and cloud servers can communicate with each other in the production network.

Protected server

A physical or virtual machine owned by a customer and which is protected with the service.

Public IP address

An IP address that is needed to make cloud servers available from the Internet.

R

Recovery point objective (RPO)

Amount of data lost from outage, measured as the amount of time from a planned outage or disaster event. RPO threshold defines the maximum time interval allowed between the last suitable recovery point for a failover and the current time.

Recovery server

A VM replica of the original machine, based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers, in case of a disaster.

Runbook

Planned scenario consisting of configurable steps that automate disaster recovery actions.

S

Site-to-site (S2S) connection

Connection extending the local network to the cloud, via a secure VPN tunnel.

T

Test IP address

An IP address that is needed in case of a test failover, to prevent duplication of the production IP address.

Test network

Isolated virtual network that is used to test the failover process.

V

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.

VPN gateway (formerly, VPN server or connectivity gateway)

A special virtual machine providing a connection between the local site and the cloud site networks via a secure VPN tunnel. The VPN gateway is deployed on the cloud site.

Index

V

W