

vSphere Replication for Disaster Recovery to Cloud

vSphere Replication 6.1



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 About Disaster Recovery to Cloud 4**
 - Updated Information 5
- 2 Disaster Recovery to Cloud System Requirements and Compatibility 6**
 - Roles and Permissions that Disaster Recovery to Cloud Requires 7
- 3 Installing and Configuring vSphere Replication to Cloud 8**
 - Installing vSphere Replication for Disaster Recovery to Cloud 8
 - Upgrading vSphere Replication from Earlier Product Versions 9
 - Configure NTP Synchronisation in Your Environment 9
 - How vSphere Replication Connects to Cloud 10
 - Configuring the Connection to the Cloud 12
- 4 Replicating Virtual Machines to Cloud 18**
 - Configure a Replication to Cloud for a Single Virtual Machine 19
 - Configure a Cloud Replication Task for Multiple Virtual Machines 22
 - Using Replication Seeds for Replications to Cloud 25
- 5 Reconfiguring Replications to the Cloud 26**
 - Reconfigure a Replication to Cloud 26
- 6 Recovering Virtual Machines to Cloud 27**
 - Test Recovery to Cloud 27
 - Planned Migration to Cloud 29
- 7 Configuring Replications from Cloud 31**
 - Configure a Replication From Cloud 32
 - Configure a Reverse Replication from Cloud 34
- 8 Monitoring and Managing Replication Tasks 36**
 - Monitoring the Status of Replication Tasks 36
 - Pause or Resume a Replication Task 37
 - Stop a Replication To Cloud 38
 - Stop a Replication From Cloud 39
- 9 Troubleshooting vSphere Replication for Disaster Recovery to Cloud 41**
 - vSphere Replication UI is Missing After a vCenter Server Upgrade 41

About Disaster Recovery to Cloud

1

You can subscribe to a Disaster Recovery service to protect your vSphere workloads.

Disaster Recovery lets administrators of small sites to protect their vSphere virtual workloads from a wide class of disasters by replicating those workloads into the cloud. Disaster Recovery uses the host-based replication feature of vSphere Replication to copy the protected source virtual machines into the infrastructure of the cloud provider. If a disaster occurs, the Disaster Recovery servers can convert the replicated data into vApps and virtual machines in the cloud.

Updated Information

This *vSphere Replication for Disaster Recovery to Cloud* document is updated with each release of the product or when necessary.

This table provides the update history of the *vSphere Replication for Disaster Recovery to Cloud* document.

Revision	Description
EN-001621-02	<ul style="list-style-type: none">Updated the information in topic Chapter 2 Disaster Recovery to Cloud System Requirements and Compatibility.Updated the information in topic Upgrading vSphere Replication from Earlier Product Versions.Updated the information about the Interoperability of vSphere Replication in topics Configure a Replication to Cloud for a Single Virtual Machine, and Configure a Cloud Replication Task for Multiple Virtual Machines.Updated the information in topic Select Recovery Networks on the Target Virtual Data Center.Updated the information in topic Configure a Replication From Cloud.
EN-001621-01	<ul style="list-style-type: none">Corrected information on supported vCenter Server versions in Chapter 2 Disaster Recovery to Cloud System Requirements and Compatibility.Corrected product name in topic Select Recovery Networks on the Target Virtual Data Center.
EN-001621-00	Initial release.

Disaster Recovery to Cloud System Requirements and Compatibility

2

To enable replications to the cloud, your environment must meet certain requirements in terms of additional configuration and specific versions of the VMware products that you use.

System Requirements

Disaster Recovery to Cloud has the same requirements to the environment as vSphere Replication. In addition, Disaster Recovery to Cloud requires that ports 10000 to 10010 of ESXi hosts are open for outgoing traffic. The required ports are open automatically when you install a VIB on each supported ESXi host in the environment where the vSphere Replication appliance is deployed. See [How vSphere Replication Connects to Cloud](#).

Product Compatibility

Replications to the cloud require that you run certain versions of VMware products on the source site and on the target site. Your cloud provider ensures that the target environment is configured for replications to cloud. You must verify that you run a supported version of the following products on the source site.

Table 2-1. Compatible Product Versions on the Source Site for Replications to the Cloud

Product	Supported Version
vSphere Replication appliance	6.x
ESXi host	5.0, 5.1.x, 5.5.x, and 6.0
vCenter Server	6.0
vSphere Web Client	6.0

Roles and Permissions that Disaster Recovery to Cloud Requires

Replications to the cloud require certain users, roles, and permissions.

vSphere Web Client

On the source vSphere side, you need the same credentials as the ones required for vSphere Replication. See the topic vSphere Replication Roles Reference in the *VMware vSphere Replication Administration* document.

vCloud User Credentials

When you create a connection to the target virtual data center, you provide two pairs of credentials.

Connection Credentials Used for authenticating within the cloud organization, these credentials initiate a user session with your cloud provider. The privileges for your user account are managed by your cloud provider.

- **com.vmware.hcs.{com.vmware.hcs}:ManageRight**
- **com.vmware.hcs.{com.vmware.hcs}:ViewRight**
- **Organization.View Organization Networks**
- **Organization.View Organizations**
- **Organization VDC.View Organization VDCs**

Credentials to the cloud are required for each target site, once per user session, and not per operation in the vSphere Web Client. When the authenticated user session to a target site expires, users are prompted to input their credentials again.

System Monitoring Credentials

Used at runtime to let the source and the target site communicate. These credentials are stored in the vSphere Replication appliance on the source site. The user name that you provide must be assigned the vSphere Replication role, or the following rights in your cloud organization.

- **com.vmware.hcs.{com,vmware.hcs}:ManageRight**
- **com.vmware.hcs.{com,vmware.hcs}:ViewRight**
- **Organization.View Organization Networks**
- **Organization.View Organizations**
- **Organization VDC.View Organization VDCs**

Although you can use the same credentials for both connection and system monitoring, a good practice is to use different pairs of credentials.

Installing and Configuring vSphere Replication to Cloud

3

Before you configure replications to the cloud, you must deploy the vSphere Replication appliance on the source site and set up your environment to enable connections to the cloud.

This section includes the following topics:

- [Installing vSphere Replication for Disaster Recovery to Cloud](#)
- [Upgrading vSphere Replication from Earlier Product Versions](#)
- [Configure NTP Synchronisation in Your Environment](#)
- [How vSphere Replication Connects to Cloud](#)
- [Configuring the Connection to the Cloud](#)

Installing vSphere Replication for Disaster Recovery to Cloud

vSphere Replication is distributed as an OVF virtual appliance.

You deploy vSphere Replication by using the vSphere OVF deployment wizard.

Depending on the version of the vCenter Server on which you install vSphere Replication, the deployment procedure might vary.

Table 3-1. vSphere Replication Deployment Procedures

vCenter Server Version	vSphere Replication Deployment Procedure
vCenter Server 5.5.x	See the topic Deploy the vSphere Replication Virtual Appliance in the <i>vSphere Replication 5.5 Administration</i> document.
vCenter Server 6.0	See the topic Deploy the vSphere Replication Virtual Appliance in the <i>vSphere Replication 6.1 Administration</i> document.

Important In these procedures, the steps for installing vSphere Replication on the target site apply to vCenter Server to vCenter Server replications. If you intend to use vSphere Replication only for replications to cloud, do not attempt to install vSphere Replication on the target site. Your cloud provider ensures that the target site is configured for replications to cloud.

After installing the vSphere Replication appliance, you must configure it to synchronize with an external NTP server. See [Configure NTP Synchronisation in Your Environment](#).

Upgrading vSphere Replication from Earlier Product Versions

You can upgrade vSphere Replication 5.5.x, 5.8 and 6.0.0.x to vSphere Replication 6.x.

To upgrade a previously installed version of vSphere Replication to vSphere Replication for Disaster Recovery to Cloud, you must mount the vSphere Replication ISO file on a system in your environment that is accessible from the vSphere Replication appliance, and apply the update through the virtual appliance administration interface (VAMI) on port 5480. See [Upgrade vSphere Replication by Using the Downloadable ISO Image](#).

After upgrading the vSphere Replication appliance, you must configure it to synchronize with an external NTP server. See [Configure NTP Synchronisation in Your Environment](#).

Configure NTP Synchronisation in Your Environment

You must synchronize the time on the vSphere Replication appliance in your environment with an NTP server.

By default, the vSphere Replication appliance is synchronized with the ESXi host on which it resides. You must disable the NTP synchronization with the host and configure the vSphere Replication appliance and the vCenter Server to synchronize with an external NTP server.

Procedure

- 1 Configure NTP synchronization on the vSphere Replication appliance.
 - a In the vSphere inventory tree, locate the vSphere Replication appliance, right-click and select **Edit Settings**.
 - b On the **VM Options** tab, click **VMware Tools**.
 - c Deselect the **Synchronize guest time with host** check box.
 - d In the virtual appliance console, run the command `chkconfig ntp on` to run NTP synchronization every time the vSphere Replication appliance starts up.
 - e To configure the vSphere Replication appliance to synchronize with an NTP server, edit the `/etc/ntp.conf` file to enter the address of an NTP server.

Add the following line in the `ntp.conf` file:

```
server <your_ntp_server_address>
```

- f Run the `service ntp start` command.
- 2 Configure the vCenter Server on the source site to synchronize with the NTP server that you configured in the vSphere Replication appliance.

How vSphere Replication Connects to Cloud

When you create a connection to the cloud, the vCloud Tunneling Agent in the vSphere Replication appliance creates a tunnel to secure the transfer of replication data to your cloud Organization.

When a tunnel is created, the vCloud Tunneling Agent opens a port on the vSphere Replication appliance. ESXi hosts connect to that port to send replication data to a cloud organization. The port is picked randomly from a configurable range. The default port range is 10000-10010 TCP.

By default, ports 10000-10010 are not open on ESXi hosts. When you power on the vSphere Replication appliance, a vSphere Installation Bundle (VIB) is installed on all supported ESXi hosts in the vCenter Server inventory where the appliance is deployed. The VIB creates a firewall rule, Replication-to-Cloud Traffic, that opens TCP ports 10000 to 10010 for outgoing traffic. The rule is enabled automatically and takes effect immediately when you power on the vSphere Replication appliance, or when a host is registered or connected in the vCenter Server. If an administrator removes the VIB from a host, for example by using the `esxccli` utility, the vSphere Replication appliance reinstalls the VIB the next time you restart the appliance or when a host is restarted or reconnected to the inventory. If you do not want ports 10000 to 10010 to be open on an ESXi host, and if you do not plan to use this host as a replication source, you can disable the Replication-to-Cloud Traffic rule. See [Allow or Deny Access to an ESXi Service or Management Agent with the vSphere Web Client](#).

To reduce the number of open ports or to change the ports that are used for communication between ESXi hosts and the vCloud Tunneling Agent, you can create a custom firewall rule and reconfigure the agent.

Change the Cloud Tunnel Ports on ESXi Hosts

When you power on the vSphere Replication appliance, it automatically configures all ESXi hosts in your environment to open TCP ports 10000-10010 for outgoing data transfers.

The vCloud Tunneling Agent in the vSphere Replication appliance uses ports 10000-10010 to receive data from ESXi instances that host replication sources.

If you do not want to have unused open ports on your ESXi hosts, if the number of open ports is insufficient, or if you want to change which ports are open, you can reconfigure your firewall settings.

To change the default ports that are used to transfer replication data from ESXi hosts to the vCloud Tunneling Agent, you must configure each ESXi instance that hosts a replication source virtual machine, and the vCloud Tunneling Agent.

Procedure

- 1 Disable the default **Replication-to-cloud Traffic** rule that is created by the vSphere Replication appliance.

For detailed procedure, see [Allow or Deny Access to an ESXi Service or Management Agent with the vSphere Web Client](#).

- 2 Create a custom firewall rule on each ESXi server that hosts replication source machines.
See [Creating custom firewall rules in VMware ESXi 5.0 \(KB 2008226\)](#).
- 3 Enable the custom firewall rule that you created on each ESXi host.
See [Allow or Deny Access to an ESXi Service or Management Agent with the vSphere Web Client](#).

What to do next

Configure the vCloud Tunneling Agent to use the ports that you configured on ESXi hosts.

Customize the Ports that vSphere Replication Uses for Tunneling

By default, the vCloud Tunneling Agent in the vSphere Replication appliance is configured to use TCP ports ranging between 10000 and 10010 to create tunnels to the cloud. All ESXi instances that might host replication source virtual machines must have their firewall configured to allow outgoing traffic on these ports.

For each tunnel to cloud, the vCloud Tunneling Agent allocates one unique port from the specified range. You can reconfigure ESXi hosts and the vCloud Tunneling Agent to reduce the number of open ports or to change the ports that are used to create tunnels to cloud.

After you reconfigure the ESXi hosts to use custom ports, you must configure the vCloud Tunneling Agent to use the same custom ports.

Prerequisites

- Verify that the ports you selected to use for cloud tunnels are open for outgoing traffic on all ESXi servers that host replication sources.
- Verify that you know the IP address of the vSphere Replication appliance in your environment. To check the IP address of the vSphere Replication appliance, select the vCenter Server in the inventory tree, navigate to the **Manage** tab, click **vSphere Replication**, and click **About**.
- Verify that you have root user credentials for the vSphere Replication appliance.
- Verify that TCP port 22 is open on the vSphere Replication appliance, and that SSH connections are enabled. See topic [Unable to Establish an SSH Connection to the vSphere Replication Appliance](#) in the *vSphere Replication Administration* document.

Procedure

- 1 Use an SSH client to connect to the vSphere Replication appliance and log in as the root user.
- 2 Run the following command to configure the ports for tunnel connections.

```
/opt/vmware/vcta/bin/cell-management-tool
    configure-vcta-server -prl LOW -prh HIGH
```

Where *LOW* and *HIGH* define the range of ports to be used for tunnel connections. To use only one port, type the port number as the value for *LOW* and *HIGH*.

For example, the following command configures the vCloud Tunneling Agent to use only port 10001.

```
/opt/vmware/vcta/bin/cell-management-tool
    configure-vcta-server -prl 10001 -prh 10001
```

Note You can designate any free TCP port in your environment for the communication between ESXi hosts and the vCloud Tunneling Agent, but you must verify that all ESXi hosts and the vCloud Tunneling Agent are configured to use the same ports.

- 3 Run the following command to restart the vCloud Tunneling Agent.

```
service vmware-vcd restart
```

Configuring the Connection to the Cloud

In addition to installing and configuring the vSphere Replication appliance, you must configure the connection to your cloud provider.

You can configure a connection to the cloud provider before you start the **Configure Replication** wizard or while you configure a replication task.

Connect to a Cloud Provider Site

Before you configure replication tasks to the cloud, you configure the connections between your vSphere environment and virtual data centers that belong to your cloud organizations.

You can connect a vCenter Server to multiple virtual data centers, and a virtual data center can be connected to multiple vCenter Server instances. However, you can have only one connection between a source vCenter Server and a target virtual data center.

Prerequisites

Verify that you have user credentials for a cloud organization in which vCloud Director is enabled. Your cloud provider enables the Disaster Recovery to Cloud service per your contract.

Procedure

- 1 Click the **Configure** tab, click **vSphere Replication > Target Sites**, and click the cloud connection icon .

The **Connect to a Cloud Provider** wizard opens.

- 2 On the Connection settings page, type the address of your cloud provider, the organization name, and credentials to authenticate with the cloud.

By default, vSphere Replication uses these credentials to establish a user session to the cloud and for system monitoring purposes. To enable system monitoring, these credentials will be stored in the vSphere Replication appliance, unless you select to use another user account for system monitoring.

- 3 (Optional) If you do not want to store the credentials that you used for authentication, select the **Use a different account for system monitoring** check box, and type the credentials to be used for system monitoring.

These credentials are encrypted and stored in the vSphere Replication database.

- 4 Click **Next**.

The **Connect to a Cloud Provider** wizard displays a list of virtual data centers to which you can connect. If a virtual data center is already connected to the vCenter Server, that data center does not appear in the list.

- 5 From the list of virtual data centers, select a target for the connection and click **Next**.

- 6 Review your settings and click **Finish**.

The connection to the cloud organization appears in the list of target sites. The status of the connection is *Missing network settings*.

What to do next

Select the networks on the target site that vSphere Replication must use for recovery operations. See [Select Recovery Networks on the Target Virtual Data Center](#)

Select Recovery Networks on the Target Virtual Data Center

To finalize the configuration of a connection to the target site, you must specify the networks that the Disaster Recovery to Cloud service should use for tests and recovery operations.

When you add a new connection to the cloud, at first it appears in *Missing mapping* status.

When you subscribe to the Disaster Recovery to Cloud service, VMware automatically creates two default networks for your service—an isolated network and an external routed network. The Edge Gateway for the routed network has a public IP address on an external interface so that it is accessible through the Internet. You can use these networks for your virtual machines protected by the Disaster Recovery to Cloud service, or create other networks in your cloud organization.

When you run a test recovery, vSphere Replication configures the replicated virtual machine on the target site to connect to the test network. This lets you access the target virtual machine and verify that it operates as expected and that data is replicated correctly per your replication settings.

The recovery network is used when you perform planned migrations and recovery operations. vSphere Replication configures the replicated virtual machine on the target site and connects it to the recovery network, so that you can have access.

Although you can use the same network for all recovery workflows, a good practice is to run test recoveries in a separate network.

Note You can configure only one pair of networks for a cloud virtual data center.

Prerequisites

Verify that you created a connection to a cloud virtual data center. See [Connect to a Cloud Provider Site](#).

Procedure

- 1 On the **vSphere Replication** tab under **Manage**, click the target network settings icon .

If your user session to the cloud has expired, the **Network Mappings** wizard prompts you to type your credentials.
- 2 From the drop-down menus, select a recovery network and a test network and click **Next**.

The drop-down menus display only the networks that are configured for vCloud Director based cloud.
- 3 On the Local recovery networks page, select the cloud networks from the left pane and select the local recovery networks from the right pane. Click **Add mappings** and click **Next**.

Note You can select a virtual data center (VDC) network or a vApp network. When you select a vApp network, the network mappings are configured only for the selected vApp. When you select a VDC network, the network mappings are configured for all VMs in that network.

- 4 On the Local test networks page, select the cloud networks from the left pane and select the local test networks from the right pane. Click **Add mappings** and click **Next**.
- 5 On the Ready to complete page, review your settings and click **Finish**.

What to do next

When you test a replication or perform a recovery operation, the vCloud Director based cloud automatically attaches the virtual machine to the test or recovery network respectively.

Disable the Automatic Export of MAC Addresses During Replication

By default, when you configure a virtual machine for replication to cloud, its NICs and MAC addresses are copied automatically to the target site as part of the provisioning of the placeholder virtual machine.

If the test network is not isolated from the production network and these networks have common routing, a test recovery of a replicated virtual machine might result in duplicate MAC addresses in your virtual data center.

To avoid duplicate MAC addresses in your data center, you can disable the automatic copying of network configurations from the source site to cloud sites.

Note Disabling the automatic copying of network configurations does not delete the configurations that are already replicated to the target site. See <http://kb.vmware.com/kb/2086292> .

Procedure

- 1 Use the vSphere Web Client on the source site to locate the HMS virtual machine and log in as the root user.

- 2 Navigate to folder `/opt/vmware/hms/conf/`.
- 3 Run the `vi hms-configuration.xml` command to open the `hms-configuration.xml` file for editing.
- 4 Locate the `<hms-dr2c-export-mac-address>` parameter, and modify the value to `false`:
`<hms-dr2c-export-mac-address>false</hms-dr2c-export-mac-address>`
- 5 Run the `:wq` command to save the change, and run the following command to restart the HMS service.

```
# service hms restart
```

The automatic copying of network configurations to target cloud sites is disabled for all newly configured replications.

Cloud Connection States Displayed in the vSphere Web Client

In the vSphere Web Client, on the **vSphere Replication** tab under **Manage**, you can check the status of connections between your vSphere environment and the virtual data centers on the remote site.

The following table lists the cloud connection states that you can observe, their meanings, and what you can do to change a state back to normal.

Table 3-2. Cloud Connection States

Icon	Status	Description	Remediation
	Connected	The connection between the source site and the target site is working properly.	Not needed.
	Not authenticated	The remote site is online, but your user session to the cloud has expired. In this state, you need to provide credentials to manage replication tasks. Already configured replications are running in the background.	Reconnect the sites. <ol style="list-style-type: none"> 1 Select the cloud organization that indicates the Not authenticated status. 2 Click the Reconnect icon  above the list of target sites. 3 Click Yes to confirm. 4 In the Reconnect Sites dialog box, type the credentials for the remote site and click OK.

Table 3-2. Cloud Connection States (Continued)

Icon	Status	Description	Remediation
	Missing network settings	<p>You have not selected the networks to use for recovery and test recovery operations on the target site.</p> <p>In this state, when you start the configure replication wizard, you are prompted to configure the networks on the target site first.</p>	<p>Configure the network settings.</p> <ol style="list-style-type: none"> 1 Select the cloud organization that indicates the Missing network settings status. 2 Click the network configuration icon  above the list of target sites. 3 Select both a recovery network and a test network and click Next. 4 On the Ready to complete page, verify that you selected the correct networks and click Finish.
	Connection issue	<ul style="list-style-type: none"> ■ The SSL certificate on the remote site has been changed. ■ The network connection between the source site and the target site is not functioning properly, or the remote site is offline. ■ The cloud user that is used for connection or system monitoring might be disabled or deleted. <p>In this state, configured replications might not be running.</p>	<ul style="list-style-type: none"> ■ Select the cloud organization that indicates the Connection issue status and click the Reconnect icon . <p>If the SSL certificate on the remote site has changed, the thumbprint of the new certificate is displayed for you to confirm.</p> <ul style="list-style-type: none"> ■ In the inventory tree, click the vCenter Server and navigate to the Events tab under Monitor to search for events related to vSphere Replication. ■ Contact your cloud provider to verify the status of the remote site.

Reconnect to a Cloud Provider Site

If the state of a connection to cloud is **Not authenticated**, your user session to the target virtual data center has expired.

Procedure

- 1 Select the cloud organization for which a **Not authenticated** state is displayed.
- 2 Click the **Reconnect** icon  above the list of target sites.
- 3 Click **Yes** to confirm.

4 In the Reconnect Sites dialog box, type the credentials for the remote site and click **OK**.

The connection state changes to Connected.

Replicating Virtual Machines to Cloud

4

You can configure replications from vSphere environments to cloud for a single virtual machine or for multiple virtual machines.

To replicate virtual machines to cloud, you must deploy the vSphere Replication 5.8 appliance at the source site, and your cloud provider must enable replications to the cloud in your cloud organization.

The source and target sites must be connected so that you can configure replications. Though you can create connections to the cloud while you configure replications, the good practice is to create cloud connections before you start the **Configure Replication** wizard. See [Connect to a Cloud Provider Site](#).

To avoid copying big volumes of data between the source site and the cloud over a network connection, you can create replication seeds on the target site and configure replication tasks to use them. See [Using Replication Seeds for Replications to Cloud](#).

For each replication task, you can set a recovery point objective (RPO) to a certain time interval depending on your data protection needs. vSphere Replication applies all changes made to replication source virtual machines to their replicas on the target site. This process reoccurs at the RPO interval that you set.

You can configure replications for powered-off virtual machines, but the data synchronization begins when the virtual machine is powered on. While the source virtual machine is powered off, the replication appears in Not active status.

You cannot use vSphere Replication to replicate virtual machine templates.

This section includes the following topics:

- [Configure a Replication to Cloud for a Single Virtual Machine](#)
- [Configure a Cloud Replication Task for Multiple Virtual Machines](#)
- [Using Replication Seeds for Replications to Cloud](#)

Configure a Replication to Cloud for a Single Virtual Machine

To start replicating virtual machines to your cloud organization, you configure replication from the source site by using the vSphere Web Client.

When you configure replication, you set a recovery point objective (RPO) to determine the maximum data loss that you can tolerate. For example, an RPO of 1 hour seeks to ensure that a virtual machine loses the data for no more than 1 hour during the recovery. For smaller RPO values, less data is lost in a recovery, but more network bandwidth is consumed keeping the replica up to date. The RPO value affects replication scheduling, but vSphere Replication does not adhere to a strict replication schedule. See the topic *How the Recovery Point Objective Affects Replication Scheduling* in the *vSphere Replication Administration* document.

Every time that a virtual machine reaches its RPO target, vSphere Replication records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To reduce the volume of data that is kept in the vCenter Server events database, limit the number of days that vCenter Server retains event data. See *Configure Database Retention Policy* in the *vCenter Server and Host Management Guide*. Alternatively, set a higher RPO value.

vSphere Replication guarantees crash consistency amongst all the disks that belong to a virtual machine. If you use quiescing, you might obtain a higher level of crash consistency amongst the disks that belong to a virtual machine. The available quiescing types are determined by the virtual machine's operating system. See [Interoperability Pages for vSphere Replication 6.1](#) for quiescing support for Windows and Linux virtual machines.

If you plan to use replication seeds, read and understand the information in topic [Using Replication Seeds for Replications to Cloud](#).

Note By default, when you configure a virtual machine for replication to cloud, its NICs and MAC addresses are copied automatically to the target site as part of the provisioning of the placeholder virtual machine. If the test network is not isolated from the production network and these networks have common routing, a test recovery of a replicated virtual machine might result in duplicate MAC addresses in your virtual data center. See [Disable the Automatic Export of MAC Addresses During Replication](#).

Prerequisites

- Verify that the vSphere Replication appliance is deployed in your environment.
- Verify that the Disaster Recovery to Cloud service is enabled in the target cloud organization.
- Configure a connection to the cloud organization to which you want to replicate data. See [Connect to a Cloud Provider Site](#).

Procedure

- 1 On the vSphere Web Client Home page, click **VMs and Templates**.

- 2 In the inventory tree, right-click the virtual machine that you want to replicate and select **All vSphere Replication Actions > Configure Replication**.

The **Configure Replication** wizard opens.

- 3 Select **Replicate to a cloud provider** and click **Next**.

- 4 Select the target site to which you want to replicate the virtual machine.

- If you have created a connection to the cloud provider, select the target virtual data center from the list and click **Next**.

If the status of the connection is `Not authenticated`, you must provide credentials to authenticate with the cloud organization. If you have not selected the networks on the target site to use for recovery operations, you are prompted to do so.

- If you have not created a connection to the cloud provider, click **New Provider VDC**, click **Next**, and follow the on-screen prompts to connect to the target cloud organization.

- 5 On the Target location page, select where to store replication data.

Option	Procedure
Use storage policy	From the drop-down menu, select the storage policy for replication placement and click Next .
Use replication seeds	<ol style="list-style-type: none"> a Click Next to navigate to the list of available seed vApps on the target site. b Select a seed vApp from the list and click Next. <p>Note If you remove a disk from a replication source virtual machine, the seed disk is not deleted from the datastore on the target site.</p>

- 6 (Optional) On the Replication options page, select the quiescing method for the guest operating system of the source virtual machine.

Note Quiescing options are available only for virtual machines that support quiescing.

- 7 (Optional) Select **Enable network compression for VR data**.

Compressing the replication data that is transferred through the network saves network bandwidth and might help reduce the amount of buffer memory used on the vSphere Replication server. However, compressing and decompressing data requires more CPU resources on both the source site and the server that manages the target datastore.

- 8 On the Recovery settings page, use the RPO slider or the time spinners to set the acceptable period for which data can be lost in the case of a site failure.

The available RPO range is from 15 minutes to 24 hours.

- 9 (Optional) To save multiple replication instances that can be converted to snapshots of the source virtual machine during recovery, select **Enable** in the Point in time instances pane, and adjust the number of instances to keep.

Note You can keep up to 24 instances for a virtual machine. This means that if you configure vSphere Replication to keep 6 replication instances per day, the maximum number of days you can set is 4 days.

The number of replication instances that vSphere Replication keeps depends on the configured retention policy, and requires that the RPO period is short enough for these instances to be created. Because vSphere Replication does not check whether the RPO settings will create enough instances to keep, and does not display a warning message if the number of instances is not sufficient, you must ensure that you set vSphere Replication to create the instances that you want to keep. For example, if you set vSphere Replication to keep 6 replication instances per day, the RPO period should not exceed 4 hours, so that vSphere Replication can create 6 instances in 24 hours.

- 10 Click **Next**.

- 11 On the Ready to complete page, review the replication settings, and click **Finish**.

A virtual machine configuration task appears in the Recent Tasks list at the bottom of the vSphere Web Client. A progress bar indicates that the source virtual machine is being configured for replication.

If the configuration operation completes successfully, the replication task that you created appears in the list of outgoing replications on the **vSphere Replication** tab under **Monitor**.

Note If the replication source virtual machine is powered off, the replication remains in Not Active state until you power on the virtual machine.

What to do next

On the **vSphere Replication** tab under **Monitor**, you can check the state of each replication. See [Monitoring the Status of Replication Tasks](#).

You can click a replication task in the list and use the tabs at the bottom of the vSphere Web Client to view details about the replication, the recovery status, and the latest performed test, if test results are not cleared yet.

Configure a Cloud Replication Task for Multiple Virtual Machines

To configure batches of virtual machines for replication to the cloud, you can select multiple virtual machines and start the **Configure Replication** wizard.

When you configure replication, you set a recovery point objective (RPO) to determine the maximum data loss that you can tolerate. For example, an RPO of 1 hour seeks to ensure that a virtual machine loses the data for no more than 1 hour during the recovery. For smaller RPO values, less data is lost in a recovery, but more network bandwidth is consumed keeping the replica up to date. The RPO value affects replication scheduling, but vSphere Replication does not adhere to a strict replication schedule. See the topic *How the Recovery Point Objective Affects Replication Scheduling* in the *vSphere Replication Administration* document.

Every time that a virtual machine reaches its RPO target, vSphere Replication records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To reduce the volume of data that is kept in the vCenter Server events database, limit the number of days that vCenter Server retains event data. See *Configure Database Retention Policy* in the *vCenter Server and Host Management Guide*. Alternatively, set a higher RPO value.

vSphere Replication guarantees crash consistency amongst all the disks that belong to a virtual machine. If you use quiescing, you might obtain a higher level of crash consistency amongst the disks that belong to a virtual machine. The available quiescing types are determined by the virtual machine's operating system. See [Interoperability Pages for vSphere Replication 6.1](#) for quiescing support for Windows and Linux virtual machines.

If you plan to use replication seeds, read and understand the information in topic [Using Replication Seeds for Replications to Cloud](#).

Note By default, when you configure a virtual machine for replication to cloud, its NICs and MAC addresses are copied automatically to the target site as part of the provisioning of the placeholder virtual machine. If the test network is not isolated from the production network and these networks have common routing, a test recovery of a replicated virtual machine might result in duplicate MAC addresses in your virtual data center. See [Disable the Automatic Export of MAC Addresses During Replication](#).

Prerequisites

- Verify that the vSphere Replication appliance is deployed in your environment.
- Verify that the Disaster Recovery to Cloud service is enabled in the target cloud organization.
- Configure a connection to the cloud organization to which you want to replicate data. See [Connect to a Cloud Provider Site](#).

Procedure

- 1 On the vSphere Web Client Home page, click **VMs and Templates**.

- 2 Select a data center, navigate to the **Related Objects** tab, and click the **Virtual Machines** tab.
- 3 Select the virtual machines for which you want to configure replications.
- 4 Right-click the virtual machines and select **All vSphere Replication Actions > Configure Replication**.

The **Configure Replication** wizard opens and vSphere Replication validates the virtual machines that can be configured for replication.

- 5 Verify the validation results and click **Next**.
- 6 Select **Replicate to a cloud provider** and click **Next**.
- 7 Select the target site to which you want to replicate the virtual machine.
 - If you have created a connection to the cloud provider, select the target virtual data center from the list and click **Next**.

If the status of the connection is `Not authenticated`, you must provide credentials to authenticate with the cloud organization. If you have not selected the networks on the target site to use for recovery operations, you are prompted to do so.
 - If you have not created a connection to the cloud provider, click **New Provider VDC**, click **Next**, and follow the on-screen prompts to connect to the target cloud organization.
- 8 On the Target location page, select where to store replication data.

Option	Procedure
Use storage policy	From the drop-down menu, select the storage policy for replication placement and click Next .
Use replication seeds	<ol style="list-style-type: none"> a Select the storage policy to use for virtual machines without seeds. b Select the Use replication seeds check box and click Next. c On the Replication seed page, assign seed vApps to source virtual machines, and click Next. <p>For all source virtual machines that do not have a seed vApp assigned, vSphere Replication applies the storage policy that you selected from the drop-down menu on the Target location page.</p> <p>Note If you remove a disk from a replication source virtual machine, the seed disk is not deleted from the datastore on the target site.</p>

- 9 (Optional) On the Replication options page, select the quiescing method for the guest operating system of the source virtual machine.

Note Quiescing options are available only for virtual machines that support quiescing.

- 10 (Optional) Select **Enable network compression for VR data**.

Compressing the replication data that is transferred through the network saves network bandwidth and might help reduce the amount of buffer memory used on the vSphere Replication server. However, compressing and decompressing data requires more CPU resources on both the source site and the server that manages the target datastore.

- 11 On the Recovery settings page, use the RPO slider or the time spinners to set the acceptable period for which data can be lost in the case of a site failure.

The available RPO range is from 15 minutes to 24 hours.

- 12 (Optional) To save multiple replication instances that can be converted to snapshots of the source virtual machine during recovery, select **Enable** in the Point in time instances pane, and adjust the number of instances to keep.

Note You can keep up to 24 instances for a virtual machine. This means that if you configure vSphere Replication to keep 6 replication instances per day, the maximum number of days you can set is 4 days.

The number of replication instances that vSphere Replication keeps depends on the configured retention policy, and requires that the RPO period is short enough for these instances to be created. Because vSphere Replication does not check whether the RPO settings will create enough instances to keep, and does not display a warning message if the number of instances is not sufficient, you must ensure that you set vSphere Replication to create the instances that you want to keep. For example, if you set vSphere Replication to keep 6 replication instances per day, the RPO period should not exceed 4 hours, so that vSphere Replication can create 6 instances in 24 hours.

- 13 Click **Next**.

- 14 On the Ready to complete page, review the replication settings, and click **Finish**.

For each source virtual machine, a configuration task appears in the Recent Tasks list in the bottom of the vSphere Web Client. A progress bar indicates that the source virtual machine is being configured for replication.

For each source virtual machine that is configured successfully, a replication task appears on the **vSphere Replication** tab under **Monitor**.

For source virtual machines that are powered on, the initial synchronization starts after the configuration. For source virtual machine that are powered off, the initial synchronization starts when you power on the virtual machines.

Note If a replication source virtual machine is powered off, the replication remains in Not Active state until you power on the virtual machine.

What to do next

On the **vSphere Replication** tab under **Monitor**, you can check the state of each replication. See [Monitoring the Status of Replication Tasks](#).

You can click a replication task in the list and use the tabs at the bottom of the vSphere Web Client to view details about the replication, the recovery status, and the latest performed test, if test results are not cleared yet.

Using Replication Seeds for Replications to Cloud

For each new replication that you configure, an initial full synchronization operation is performed. During this operation, vSphere Replication copies the whole data from the source virtual machine to a placeholder vApp on the target site.

If the source virtual machine is too big, or the bandwidth of your network connection to the cloud is too low, the initial full sync might take a long time. Therefore, you might choose to copy the source virtual machine to the target site by using removable media, or other means of data transfer. Then you can configure a replication and use the virtual machine copy on the target site as a replication seed. When a replication is configured to use a seed vApp, vSphere Replication does not copy the whole source virtual machine to the target site. Instead, it copies to the seed vApp only the different blocks between the source virtual machine and the seed.

Note vSphere Replication stores the replication data in the seed vApp. No copies of the seed vApp are created. Therefore, a seed vApp can be used for only one replication.

Creating Seed vApps in the Cloud

Seed vApps on the target site can be created in the following ways.

- **Offline data transfer:** You can export a virtual machine as an OVF package and let a Cloud service administrator import the package in your cloud organization.
- **Clone a virtual machine:** A virtual machine in the org virtual data center can be cloned to create a seed vApp. vSphere Replication calculates checksum and exchanges the different blocks from the replication source to the seed vApp.
- **Copy over the network:** A source virtual machine can be copied to the cloud organization by using means other than vSphere Replication to copy the initial source data to the target site.

Note The size and number of disks, and their assignment to disk controllers and bus nodes must match between the replication source and the seed virtual machine. For example, if the replication source machine has two disks of 2 GB each, one of them assigned to SCSI controller 0 at bus number 0, and the second one assigned to SCSI controller 1 at bus number 2, the seed vApp that you use must have exactly the same hardware configuration - 2 disks of 2 GB each, at SCSI 0:0 and at SCSI 1:2.

Reconfiguring Replications to the Cloud

5

You can reconfigure cloud replications to change the quiescing method for the guest operating system, the RPO, the network compression, and the retention of point in time instances.

Reconfigure a Replication to Cloud

You reconfigure a replication to change the RPO settings, the number of replication instances to keep, or the quiescing method that is applied when synchronizing the replication source virtual machine to your cloud organization.

Cloud replications appear in the **Outgoing Replications** list on the **vSphere Replication** tab under **Monitor**.

Procedure

- 1 On the vSphere Replication Home page, click the **Monitor** tab, and click **Outgoing Replications**.
- 2 Select the cloud replication that you want to reconfigure and click the **Reconfigure replication** icon , or right-click the replication source virtual machine and select **All vSphere Replication Actions > Reconfigure**.

The reconfiguration wizard opens.

- 3 If the connection to the cloud organization has expired, type your user credentials and click **Next** to reconnect.
- 4 (Optional) To reconfigure the quiescing method, use the drop-down menu on the Replication options page and click **Next**.
- 5 (Optional) To reconfigure the RPO, click **Next** until you reach the Recovery settings page, and modify the RPO value.
- 6 Click **Next**.
- 7 Click **Finish** to save your changes.

Recovering Virtual Machines to Cloud

6

You can check if virtual machines are properly replicated in the cloud, and migrate replicated virtual machines to your cloud organization.

This section includes the following topics:

- [Test Recovery to Cloud](#)
- [Planned Migration to Cloud](#)

Test Recovery to Cloud

Test recoveries allow you to verify that source data is replicated correctly on the target site.

When you initiate a replication task to cloud, Disaster Recovery to Cloud creates a placeholder virtual machine on the target virtual data center. If the replication uses a seed, that seed is the placeholder virtual machine. The placeholder virtual machine is not visible on the network and is not accessible until you recover it or run a test recovery.

Note During test recovery, Disaster Recovery to Cloud does not create a copy of the recovered virtual machine. When you run a test recovery, the placeholder virtual machine is reconfigured and connected to the selected test network so that you can log in and verify the replication progress.

Run a Test Recovery to Cloud

You run a test recovery to verify that data is replicated correctly from the source virtual machine to the target cloud organization.

Note If the test network is not isolated from the production network and these networks have common routing, a test recovery of a replicated virtual machine might result in duplicate MAC addresses in your virtual data center. If you use the same network for test recoveries and for production, and if you have not disabled the automatic copying of network configurations, see <http://kb.vmware.com/kb/2086292>.

Prerequisites

- Configure at least one replication task.
- Verify that the status of the replication task allows running test recoveries.

Test recoveries are allowed for the following replication statuses: OK, OK (RPO Violation), Error, Error (RPO Violation), Full Sync, Full Sync (RPO Violation), Not Active, Not Active (RPO Violation), Paused, Sync, and Sync (RPO Violation).

Note By default, the Test Status column is not visible in the list of outgoing replications . To view the column, right-click the table header, select **Show/Hide Columns**, select the **Test Status** check box, and click **OK**.

- If you have run test recoveries for the replication that you want to test, verify that you have cleared the test results.

Note You cannot run a test recovery before you clean up your previous test results for a replication.

Procedure

- 1 In the vSphere Web Client, click the **Monitor** tab, and click **vSphere Replication**.
- 2 In the **Outgoing Replications** list, click the replication for which you want to run a test recovery.
- 3 Click the **Run test recovery** icon .

The **Test Recovery** wizard opens. If the user session to the target cloud organization is expired, the wizard prompts you to enter user credentials.

- 4 On the Test recovery options page, select a data synchronization option and click **Next**.

Option	Description
Synchronize recent changes	vSphere Replication runs a synchronization task before it configures the placeholder virtual machine on the target site for testing.
Use a point in time instance	vSphere Replication configures the placeholder virtual machine for testing and uses the data that is copied on the target site at the point in time that you select from the list .

- 5 (Optional) To power on the test virtual machine on the target site after test configuration completes, select **Power on the test virtual machine** on the Ready to complete page.
- 6 Verify that the test configuration is correct, and click **Finish**.

The test recovery status appears under the list of replications, on the **Test** tab.

Note You cannot stop a replication while a test recovery for the replication is in progress.

What to do next

After you verify that data appears as expected in the test virtual machine, clean up the test results.

Clean Up a Test Recovery

You can run a test recovery or a planned migration for a replication only after the results of its previous test recovery are cleaned up.

In the vSphere Web Client, you can clean up test recovery results for replication tasks that appear under **vSphere Replication** on the **Monitor** tab.

Procedure

- 1 In the list of replications, click a replication to check its test recovery status.
- 2 At Below the list, click the **Test** tab to view details of the test status.
- 3 If the test status is other than Test recovery has not been run or has been cleaned up from the target site, click the **Run test cleanup** icon .
- 4 Click **Yes**.

Planned Migration to Cloud

Planned migration is an action that is available for replications to cloud. Planned migrations allow you to move your workloads from vCenter Server to your cloud organization.

When you run a planned migration operation, the replication source virtual machine is powered off. The placeholder virtual machine that is created in the cloud during replication is configured to run as a fully functional virtual machine. When the recovered virtual machine is powered on in the target cloud site, the replication task on the source is no longer active.

Migrate a Virtual Machine to Cloud

You can run a planned migration to move your workload from the vCenter Server to your cloud organization.

You might want to migrate replicated virtual machines to cloud if you plan maintenance on the source site.

Prerequisites

- Verify that the source site and the target site are online.
- Verify that you have enough privileges to initiate migrations to cloud.
- If you have run test recoveries for the replication that you want to migrate, verify that you cleaned up the test results.

Procedure

- 1 In the list of replications, click the replication that you want to migrate.

- 2 Click the **Run planned migration** icon .

The **Planned Migration** wizard opens. If the user session to the target cloud organization is expired, the wizard prompts you to enter user credentials.

- 3 On the Planned migration options page, select a data synchronization option and click **Next**.
- 4 On the Source VM shutdown page, select how to stop the source virtual machine and click **Next**.

Option	Description
Guest shutdown	Shuts down the operating system of the virtual machine within the timeout period that you set in the time spinners. This option uses VMware Tools. Select the Guest shutdown option only if VMware Tools is installed in the guest operating system.
Power off	Immediately shuts down the guest operating system or powers off the virtual machine. The guest operating system might not shut down properly. Select the Power off option only if VMware Tools is not installed on the guest operating system.

- 5 (Optional) To power on the recovered virtual machine on the target site at the end of the migration process, select **Power on the recovered virtual machine** on the Ready to Complete page.
- 6 Review your settings, and click **Finish**.

The replication status changes to Recovered, and the source virtual machine is no longer being replicated to the target site.

What to do next

To continue replicating the source virtual machine to the target site, stop the replication task that is in Recovered state and configure a new replication.

Configuring Replications from Cloud



You can replicate a virtual machine from your cloud environment to a vCenter Server if the virtual machine was recovered in the cloud.

You select whether to configure a new replication from cloud or a reverse replication from cloud depending on the condition of your local environment.

Configuring Replications from Cloud

When the local site does not contain data about an outgoing or incoming cloud replication for the virtual machine that you want to replicate, you can configure a replication from cloud for that machine.

In addition to simply replicating virtual machines from cloud to your local site, you can use replications from cloud to restore your site by using the data that was previously replicated in the cloud. For example, a partial or complete breakdown has occurred at your local site, and the source virtual machines that were used for replications to cloud are missing. Additionally, the data for outgoing cloud replications is missing, too. In your cloud organization, you have recovered some of the replicated virtual machines. To restore them back on your local site, you can configure replications from cloud for the recovered virtual machines.

Configuring Reverse Replications

On the local site, for an outgoing cloud replication that is in the Recovered state, you can reverse that replication to start transferring data from the recovered virtual machine in the cloud to the local virtual machine that served as the replication source before the recovery operation.

You can configure a reverse replication to update a replicated virtual machine on your local site with the changes that occurred on its restored copy in the cloud. For example, you replicated a virtual machine from the local site to the cloud and recovered the virtual machine to the cloud to use it while your local site is being maintained. While the local site was offline, changes occurred in the recovered virtual machine in the cloud. When your local site is back online, you can copy the changes from the cloud to your local environment, or even migrate the virtual machine from the cloud back to the local environment.

When you reverse a replication, you can only use the original replication settings. You cannot change the datastore location, RPO, PIT policy, and so on.

This section includes the following topics:

- [Configure a Replication From Cloud](#)
- [Configure a Reverse Replication from Cloud](#)

Configure a Replication From Cloud

You can use vSphere Replication to configure a replication from cloud to your local site.

If your local site was recovered from a major breakdown and you need to restore it, or you cannot configure a reverse replication, you can configure a new replication from cloud to synchronize data from cloud to your local site.

Note You can configure a replication from cloud for only one virtual machine in a vApp.

Prerequisites

- Verify that the cloud site is available and connected to the local site. See [Connect to a Cloud Provider Site](#).
- Verify that the list of incoming replications does not contain a replication for the virtual machine that you want to configure for replication from cloud. See [Stop a Replication From Cloud](#).

Procedure

- 1 Use the vSphere Web Client to connect to your local site.
- 2 Navigate to the **vSphere Replication** tab under **Monitor**, and click **Incoming Replications**.
- 3 Above the list of incoming replications, click the **Configure replication from cloud provider** icon .

The **Configure Replication From Cloud Provider** wizard opens.

- 4 On the Source site page, select the cloud provider site where the virtual machine is located.
 - If you have created a connection to the cloud provider, select the source virtual data center from the list and click **Next**.

If the status of the connection is `Not authenticated`, you must provide credentials to authenticate with the cloud organization.
 - If you have not created a connection to the cloud provider, click **New Provider VDC**, click **Next**, and follow the on-screen prompts to connect to the target cloud organization.
- 5 On the Available VMs page, select the virtual machine that you want to replicate.

You can select only one virtual machine from a vApp.
- 6 Accept the automatic assignment of a vSphere Replication server or select a particular server on the local site and click **Next**.

- 7 On the Target location page, click **Edit** to select the datastore where replication data will be saved.

If you want to use existing disks as seeds for the replication, browse the datastore to locate the folder where the seed disks are located.

- 8 (Optional) To configure the replication of individual disks, click the name of the source virtual machine.

The list of disks on the source virtual machine expands.

For each disk, you can select the virtual format, storage policy, and a datastore where it is replicated. If the source virtual machine contains more than one disk, you can disable the replication of a disk by clicking **Disable** in its Replication Enabled row.

- 9 (Optional) On the Replication options page, select the quiescing method for the guest operating system of the source virtual machine.

Note Quiescing options are available only for virtual machines that support quiescing. vSphere Replication does not support VSS quiescing on Virtual Volumes.

- 10 (Optional) Select **Network Compression**.

Compressing the replication data that is transferred through the network saves network bandwidth and might help reduce the amount of buffer memory used on the vSphere Replication server. However, compressing and decompressing data requires more CPU resources on both the source site and the server that manages the target datastore.

- 11 (Optional) On the Failback recovery settings page, select the Configure recovery settings check box, and then select a VM folder and host or resource pool.

- 12 On the Recovery settings page, use the RPO slider or the time spinners to set the acceptable period for which data loss is acceptable in the case of a site failure.

The available RPO range is from 15 minutes to 24 hours.

- 13 (Optional) To save multiple replication instances that can be converted to snapshots of the source virtual machine during recovery, select **Enable** in the Point in time instances pane, and adjust the number of instances to keep.

Note You can keep up to 24 instances for a virtual machine. This means that if you configure vSphere Replication to keep 6 replication instances per day, the maximum number of days you can set is 4 days.

The number of replication instances that vSphere Replication keeps depends on the configured retention policy, and requires that the RPO period is short enough for these instances to be created. Because vSphere Replication does not check whether the RPO settings will create enough instances to keep, and does not display a warning message if the instances are not enough, you must ensure that you set vSphere Replication to create the instances that you want to keep. For example, if you set vSphere Replication to keep 6 replication instances per day, the RPO period should not exceed 4 hours, so that vSphere Replication can create 6 instances in 24 hours.

14 On the Ready to complete page, review the replication settings, and click **Finish**.

A virtual machine configuration task appears in the Recent Tasks list at the bottom of the vSphere Web Client. A progress bar indicates that the source virtual machine is being configured for replication.

If the configuration operation completes successfully, the replication task that you created appears in the list of incoming replications on the **vSphere Replication** tab under **Monitor**.

Note If the replication source virtual machine is powered off, the replication remains in Not Active state until you power on the virtual machine.

What to do next

On the **vSphere Replication** tab under **Monitor**, you can check the state of each replication. See [Monitoring the Status of Replication Tasks](#).

Note You can pause, resume, sync, test, recover, and stop replications from cloud, but you cannot reconfigure or move these replications between vSphere Replication servers.

Configure a Reverse Replication from Cloud

You can use vSphere Replication to reverse a recovered outgoing replication and start copying data from the cloud to your local site.

If you replicated a virtual machine from the local site to the cloud and recovered the virtual machine at the cloud site to use it while your local site is being maintained, when your local site is back online, you can synchronize the changes from the cloud to your local environment, or migrate the virtual machine from the cloud back to the local environment.

When you reverse a replication, you can only use the original replication settings. You cannot change the datastore location, RPO, PIT policy, and so on.

Note When you reverse a replication, the source virtual machine on the local site is unregistered from the inventory and its disks are overridden by the disks that are replicated from the cloud. When the source virtual machine is unregistered, you can no longer use it unless you recover the replication.

Prerequisites

- Verify that the cloud site is available and connected to the local site. See [Connect to a Cloud Provider Site](#).
- In the list of outgoing replications, verify that the status of the replication that you want to reverse is Recovered. See [Migrate a Virtual Machine to Cloud](#).

Procedure

- 1 Use the vSphere Web Client to connect to your local site.
- 2 Navigate to the **vSphere Replication** tab under **Monitor**, and click **Outgoing Replications**.

- 3 In the list of outgoing replications, select the replication that you want to reverse, and click the **Reverse replication** icon ().

Note The replication status must be Recovered.

vSphere Replication validates the source and target virtual machine, and the Reverse Replication dialog box opens.

- 4 Review the settings for the reverse replication and click **OK**.

Caution The source virtual machine on the local site is unregistered from the inventory and becomes inaccessible until you recover the replication.

vSphere Replication starts synchronizing data from the cloud to your local environment.

The reversed replication is removed from the list of outgoing replications and appears in the list of incoming replications.

What to do next

You can recover the replication to migrate your virtual machine from cloud to your local environment.

Note You can pause, resume, sync, test, recover, and stop replications from cloud, but you cannot reconfigure or move these replications between vSphere Replication servers.

If the reverse replication cannot be configured, try configuring a new replication from cloud. See [Configure a Replication From Cloud](#).

Monitoring and Managing Replication Tasks

8

Outgoing replications are listed on the **vSphere Replication** tab under **Monitor**. You can monitor the state of replications to the cloud, control their running state, or stop them if you no longer need them.

This section includes the following topics:

- [Monitoring the Status of Replication Tasks](#)
- [Pause or Resume a Replication Task](#)
- [Stop a Replication To Cloud](#)
- [Stop a Replication From Cloud](#)

Monitoring the Status of Replication Tasks

In the vSphere Web Client, you can check the status of replication tasks for a vCenter Server. The list of outgoing replications is located on the **vSphere Replication** tab under **Monitoring**.

Table 8-1. Replication Statuses

Status	Description	Possible Cause	Solution
Not Active	The replication is not running at the moment.	<ul style="list-style-type: none">■ The source virtual machine is powered off.■ A communication problem might have occurred between the source ESXi host and the target site.	<ul style="list-style-type: none">■ Power on the source virtual machine.■ If all replications for an ESXi host are in Not Active state, verify that the security rule Replication-to-Cloud Traffic is enabled on the host. This rule opens TCP ports 10000 to 10010 for outgoing communication.
Paused	The replication is not running at the moment.	A vSphere Replication user has paused the replication.	In the list of replications, right-click the paused replication and select Resume .

Table 8-1. Replication Statuses (Continued)

Status	Description	Possible Cause	Solution
Error	The replication is not running at the moment.	<ul style="list-style-type: none"> ■ A configuration error occurred. ■ A replication error occurred. For example, the target site infrastructure is not accessible. 	<ul style="list-style-type: none"> ■ Reconfigure the replication. ■ Navigate to the Issues tab to check whether some problem occurred on the virtual machine.
Status (RPO violation)	<p>For replication status OK, Sync, or Full Sync, the replication is running, but the RPO that is set for the replication is not met and is violated.</p> <p>For replication status Not Active or Error, the replication is not running, and the RPO that is set for the replication is violated.</p>	<ul style="list-style-type: none"> ■ The network connection between the source and the target site is dropping intermittently. ■ The bandwidth of the connection between the source and the target site is too low. ■ The replication is not running, so data cannot be replicated on the target site. 	<ul style="list-style-type: none"> ■ Improve the network connection between the source and target site. ■ Increase the RPO period. ■ For replication status Not Active or Error, address the cause for the status and wait for the next sync.

Pause or Resume a Replication Task

To control the network traffic between the source and the target site, you can pause and resume replications.

Prerequisites

Verify that you have enough privileges to manage replications in the vSphere Web Client. See [Roles and Permissions that Disaster Recovery to Cloud Requires](#).

Procedure

- 1 On the vSphere Replication Home page, click the **Monitor** tab, and click **Outgoing Replications**.
- 2 Right-click the replication task that you want to pause or resume and select the corresponding menu item.

You can pause and resume multiple replications simultaneously only if they are replicated to the same virtual data center.

- 3 Click **Yes** to confirm.
- 4 If your user session to the cloud provider has expired, enter your credentials and click **OK** to reconnect.

Stop a Replication To Cloud

If you no longer need to replicate a virtual machine to the cloud, you can stop the replication permanently.

When you stop a replication, data is removed from both the source and the target site. Therefore, stopping a replication requires that both the source and the target site are online and connected.

If the target site is offline, you can force stop the replication task from the source site. When you force stop a replication, you remove the replication task only from the source site. The data on the target site remains intact. When the target site becomes available, you must delete the replication artifacts from the target site manually or contact your cloud provider.

Note For stopped replications that use replication seeds, the seed vApps are not deleted from the target site.

Prerequisites

Verify that you have enough privileges to manage replications in the vSphere Web Client. See [Roles and Permissions that Disaster Recovery to Cloud Requires](#).

Procedure

- 1 On the vSphere Replication Home page, click the **Monitor** tab, and click **Outgoing Replications**.
- 2 Right-click a replication and select **Stop**.

You can stop multiple replication tasks simultaneously only if they are replicated to the same virtual data center.

- 3 (Optional) To delete the replication only from the source site, select **Force stop replication** in the Stop Replication dialog box.

Note All data that was stored to the cloud during the replication remains on the target datastore, and the replication remains visible on the target site. You must manually delete the replication artifacts from the target site or contact your cloud provider to clear them from the cloud site.

- 4 Click **Yes** to confirm.
- 5 If your user session to the cloud provider has expired, enter your credentials and click **OK** to reconnect.

If both sites are online, Disaster Recovery to Cloud applies the following changes.

- On the source site, removes the replication entry from the list of outgoing replications, and removes the replication related configurations from the source virtual machine.
- On the cloud site, removes the task from the list of incoming replications, and deletes the replication data from the storage.

If only the source site is online and you selected to perform a force stop operation, the replication task is deleted from the list of outgoing replications, and replication related configurations are removed from the source virtual machine.

Stop a Replication From Cloud

If you no longer need to replicate a virtual machine from the cloud, you can stop the replication permanently.

When you stop a replication, data is removed from both the source and the target site. Therefore, stopping a replication requires that both the cloud and the local site are online and connected.

If the cloud site is offline, you can force stop the replication task from the local site. When you force stop a replication, you remove the replication task only from the source site. The data on the cloud site remains intact. When the cloud site becomes available, you must delete the replication artifacts from the cloud site manually or contact your cloud provider.

Note For stopped replications that use replication seeds, the seed VMs are not deleted from the local site.

Prerequisites

Verify that you have enough privileges to manage replications in the vSphere Web Client. See [Roles and Permissions that Disaster Recovery to Cloud Requires](#).

Procedure

- 1 Use the vSphere Web Client to connect to your local site.
- 2 Navigate to the **vSphere Replication** tab under **Monitor**, and click **Incoming Replications**.
- 3 Right-click a replication and select **Stop**.

You can stop multiple replication tasks simultaneously only if they are replicated from the same virtual data center.

- 4 (Optional) To delete the replication only from the local site, select **Force stop replication** in the Stop Replication dialog box.

Note The replication remains active on the cloud site. Contact your provider to clear the replication from the cloud site.

- 5 Click **Yes** to confirm.
- 6 If your user session to the cloud provider has expired, enter your credentials and click **OK** to reconnect.

If both sites are online, Disaster Recovery to Cloud applies the following changes.

- On the cloud site, removes the replication entry, and removes the replication related configurations from the source virtual machine.

- On the local site, removes the task from the list of incoming replications, and deletes the replication data from the storage.

If only the local site is online and you selected to perform a force stop operation, the replication task is deleted from the list of incoming replications, and replication data is deleted from the storage.

Troubleshooting vSphere Replication for Disaster Recovery to Cloud

9

Known troubleshooting information can help you diagnose and correct problems that occur while using vSphere Replication for Disaster Recovery to Cloud.

vSphere Replication UI is Missing After a vCenter Server Upgrade

After you upgrade the vCenter Server that contains the vSphere Replication 5.8 virtual appliance, the vSphere Replication user interface is no longer visible in the vSphere Web Client.

Problem

If you upgrade a vSphere Replication appliance that runs in a vCenter Server 5.1.x, and later upgrade the vCenter Server to version 5.5, the user interface components that are related to vSphere Replication are no longer visible in the vSphere Web Client interface.

Cause

This problem occurs because after the upgrade of the vCenter Server instance, vSphere Replication needs to update its extension registration in vCenter Server.

Solution

- 1 Use a supported browser to log in to the virtual appliance management interface (VAMI) of the vSphere Replication appliance that is managed by the updated vCenter Server.

The VAMI URL is `https://vr_appliance_address:5480`. For a list of browsers that vSphere Replication VAMI supports, see https://www.vmware.com/support/developer/studio/studio25/release_notes.html.

- 2 On the **VR** tab, click **Configuration**.
- 3 Under Actions, click **Save and Restart**.
- 4 After the save and restart operations complete, log out of the VAMI.
- 5 Clear the browser cache, log out of the vSphere Web Client, and log in again.