

# Using VMware Cloud Disaster Recovery

17 July 2023

VMware Cloud Disaster Recovery

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

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

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2020, 2021, 2022, 2023 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

<b>1</b>	<b>VMware Cloud DR</b>	<b>9</b>
	Onboard to VMware Cloud DR	10
	Available AWS Regions	12
	Ransomware Recovery	13
	PCI DSS Compliance	15
	About VMware Cloud DR	15
<b>2</b>	<b>Set Up VMware Cloud DR</b>	<b>17</b>
	Add VMware Cloud DR to Your Organization	17
	Choose a Purchase Option	18
	Create a Subscription	20
	Activate Recovery Region	22
	Deactivate Recovery Region	23
	Reactivate a Recovery Region	24
	Invite Users to Your Organization	24
	User Roles	25
	VMware Cloud DR Service Roles	25
	Service Roles and Permitted Operations	27
	Manage Recovery Region	30
	Create an API Token	31
	Add the API Token	31
	Change the API Token	32
<b>3</b>	<b>Configuring Access to VMware Cloud DR</b>	<b>33</b>
<b>4</b>	<b>VMware Cloud DR Dashboard</b>	<b>36</b>
	Recovery Region Summary	37
	Quick Setup Guide	38
	The Sites List	38
	The Topology Map	39
	Running Tasks and Recent Alarms	40
<b>5</b>	<b>Deploy a Cloud File System</b>	<b>43</b>
	Video: What is the Scale-Out Cloud File System?	44
	Choose an Availability Zone for Recovery	44
	Cloud File System Information	47
<b>6</b>	<b>Set Up Protected Sites</b>	<b>50</b>

AWS Direct Connect	54
Configure Direct Connect (Private VIF)	59
Unset Direct Connect	60
Delete a Private VIF	61
Switch Between Public Internet and Direct Connect for Protected Sites	62
Set Up Protected SDDC with Recovery in Different Region (Inter-Region DR)	62
Set Up Protected SDDC with Recovery in Same Region (Intra-Region DR)	63
Create Protected Site for On-premises vSphere	64
Edit or Delete a Protected Site	66
Remove a vCenter Server from a Protected Site	66
Remove a DRaaS Connector from a Protected Site	67
Send Support Bundle	67

## 7 Deploy the DRaaS Connector 68

System and Network Requirements for the DRaaS Connector	69
DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC	71
Service Public IP Addresses	72
(Optional) Install SSL Certificate Before DRaaS Connector Deployment	73
Download the DRaaS Connector OVA from VMware Cloud DR DR UI	74
Deploy the DRaaS Connector Using vCenter Server UI	75
Configure the DRaaS Connector VM from VM Console	76
Register vCenter Server	78
Custom Script for Creating a Restricted vCenter Server User	79
Privileges for a Restricted vCenter Server User	95
Create a Restricted vCenter Server User with the DRaaS Connector CLI	97
DRaaS Connector CLI Commands for Creating a Restricted vCenter Server User	98
Create a Restricted vCenter Server User	99
Register vCenter Server Using Restricted vCenter Server User	99
Reregister vCenter Server	100
DRaaS Connector CLI Commands for Registering/Re-registering vCenter Server	100
Refresh vCenter Server User Credentials	101
Power the DRaaS Connector On and Off	102
DRaaS Connector Connectivity Check	102
DRaaS Connector Performance Check	105

## 8 Using Google Cloud VMware Engine as a Protected Site 110

Before You Add Google Cloud VMware Engine as a Protected Site	111
Set Up Protected Site for Google Cloud VMware Engine	112
Download the DRaaS Connector OVA from VMware Cloud DR DR UI	113
Deploy the DRaaS Connector Using vCenter Server UI	113
Configure the DRaaS Connector VM from VM Console	115



Register vCenter for Google Cloud VMware Engine Protected Site 117

## 9 Using Protection Groups 118

Protection Group Caveats 122

Video: Configuring Protection Group Policies 123

Create a Protection Group 123

VM Name Pattern 126

Protection Group Health Status 127

Snapshots 128

Standard-frequency Snapshots 130

Run a Quiesce Compatibility Check 131

Enabling Quiescing for Linux VMs 132

High-frequency Snapshots 133

Caveats for High-Frequency Snapshots 134

Run a Host Compatibility Check for High-Frequency Snapshots 136

Preview: 15 Minute RPO with High-frequency Snapshots 141

Deactivate High-frequency Snapshots from VMs 143

Take a Manual Snapshot 146

Restore an Individual VM 147

Guest File Recovery 148

Recover VM Guest Files 149

Snapshot Retention 152

Delete Snapshots 153

Edit Snapshot Name 153

Edit Snapshot Schedules 153

Deactivate Snapshot Schedule 154

Cancel a Snapshot Task 154

Snapshot Events and Logs 157

Edit a Protection Group 159

Delete a Protection Group 160

Throttle Replication 160

Export VM-to-Protection Group Mappings 161

Unprotect a VM 161

Replication Progress Statistics 162

## 10 Deploy a Recovery SDDC 164

Before You Start Using an SDDC for Recovery 165

Deploy a New recovery SDDC 168

Add an Existing SDDC for Recovery 171

Create Firewall Rule for PCI-Hardened Recovery SDDC 172

Delete a recovery SDDC 174

Daily Usage Email Reminder	174
Adding and Removing Hosts	175
Add a Host	175
Remove a Host	176
Adding, Attaching, Deleting Clusters	176
Add a Cluster	176
Attach a Cluster	177
Delete a Cluster	177
Add a Network to a recovery SDDC	178
Request Public IP Addresses	179
Add NAT Rules	179
Add a Firewall Rule to the recovery SDDC Network	180
Recovery SDDC Firewall Rules Naming Convention	180
Create a Firewall Rule for Public IP Addresses Accessing vCenter	181
Access Recovery SDDC on VMware Cloud on AWS	182

## **11 Set Up Recovery Plans in VMware Cloud DR** 183

'Just-in-Time' DR	184
VM Customization During Failover/Failback	184
Create a Recovery Plan	185
Configure Recovery Plans	186
Selecting Sites	186
Select Protection Groups	186
Configure Plan Mappings	187
vCenter Mapping	187
Compute Resources Mapping	188
Virtual Networks Mapping	188
IP Address Mapping	189
Configure Script VM	194
Configure Recovery Steps	195
Configure Plan for Ransomware Recovery	199
Alerts	201
Recovery Plan Compliance Checks	201
View Compliance Checks	207
Viewing Recovery Plans	208
Activating/Deactivating Recovery Plans	208

## **12 Running a Recovery Plan for Failover** 209

How a Failover Recovery Plan Runs	210
Configure VM Storage for Failover	211
Run a Failover Plan	214

Migrate a VM from the Cloud File System to the SDDC Datastore	217
System Behavior During Failover	217
Recovery Plan States	220
User Input During Failover	221
Failover Error Handling	222
Failover Completion	223
Running a Test Failover Recovery Plan	227
Test Failover Example	228

## **13 Run a Failback Recovery Plan 232**

Failback Process	233
Failback Example	234
Preparing for Failback	235
Set the Plan Datastore	235
Run a Failback Recovery Plan	236
Cleaning up Failback Source	236
Performing Repeated Failbacks with Same Snapshot Issue	236

## **14 Ransomware Recovery 239**

Ransomware Recovery and Disaster Recovery	241
The Isolated Recovery Environment (IRE)	242
Ransomware Recovery Workflow and Tasks	243
Video: Ransomware Recovery Product Demo	246
Activate Ransomware Recovery Services	246
Access to Carbon Black Cloud	247
Change Country for Ransomware Data Analysis	248
Network Isolation Levels	248
Video: Network Isolation for Ransomware Recovery	254
Configure Plan for Ransomware Recovery	254
Run a Ransomware Recovery Plan	257
Start VMs in Recovery SDDC	259
Search for VMs in a Plan	261
Try a Different Snapshot	262
Snapshot Timeline: Change Rate and Entropy Rate	263
Integrated Security and Vulnerability Analysis	266
Video: Next Generation Antivirus and Behavioral Analysis with Ransomware Recovery	269
Discard or Detach VMs	270
Open Security Console	271
Guest File Recovery	271
Recover VM Guest Files	273
Copy IP Address for VM Access	275

Badging Snapshots	276
User Notes	277
Restart Validation from the Staging Snapshots on recovery SDDC	278
Restart Validation from Protected Site Snapshots	279
Power Off and Stage VMs	279
Recover VMs in Protected Site	281
Recover VMs in Other Protected Site	282
End Ransomware Recovery	283
Ransomware Events	285
Manual Sensor Installation	285
Run Plan and Install Windows Sensor	286
Run Plan and Install Linux Launcher and Sensor	289
Uninstalling Sensors	291
Create a Custom Network Isolation Level	293
Create a Group	293
Create Firewall Rules and Apply to the SDDC Group	294

## **15 Reports 298**

## **16 Monitor Events, Tasks, and Alarms 300**

Forward Events to vRealize Log Insight Cloud	301
Update Event Forwarding API Key	303
Stop Event Forwarding	304
SLA Status	304
Viewing Events	307
Viewing Tasks	311
Alarms	311
Running Tasks and Recent Alarms	312
Replication Progress Statistics	314
Configure Email Alerts	315

## **17 Upgrade Process 317**

## **18 Deactivate VMware Cloud DR 318**

# VMware Cloud DR

# 1

VMware Cloud DR protects VMware vSphere VMs by replicating them to a cloud file system site and recovering them as needed to a VMware Cloud on AWS SDDC.

## Service Components

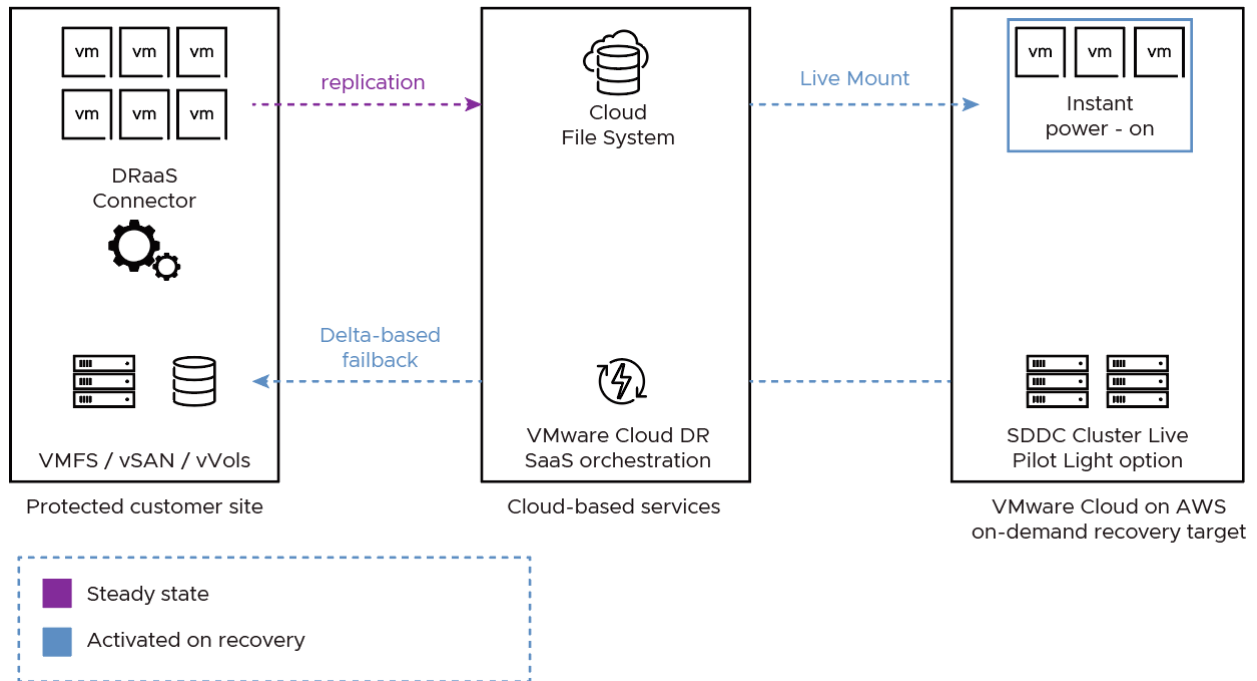
The VMware Cloud DR (sometimes called VCDR) service consists of the following components:

- Cloud file system. A cloud component that enables the efficient storage of snapshots of protected VMs in cloud storage and allows VMs to be recovered quickly, without requiring data rehydration.
- Orchestrator. A cloud component that presents a user interface (UI) to automate the disaster recovery process.
- DRaaS Connector. A virtual appliance installed in the VMware vSphere environment to protect VMs using snapshot replication from protection groups.
- Protection groups. A configuration component that allows you to create regularly scheduled snapshots of VMs which are replicated to the cloud file system.
- Recovery plan. An orchestration component that defines the steps required to recover VMs from snapshots from the cloud file system to a recovery SDDC, or to recover VMs from a ransomware attack.

VMware Cloud DR cloud components (cloud file system and orchestrator) are deployed and managed by a VMware Cloud on AWS account dedicated to each tenant.

## Service Architecture

The recovery SDDC is created immediately prior to performing a recovery and doesn't have to be provisioned to support replication in the steady state.



## End to End Security

Data transfers to and from protected sites to recovery SDDCs use secure replication, which ensures an SSL connection is established and used for all data transfers.

Read the following topics next:

- [Onboard to VMware Cloud DR](#)
- [Available AWS Regions](#)
- [Ransomware Recovery](#)
- [PCI DSS Compliance](#)
- [About VMware Cloud DR](#)

## Onboard to VMware Cloud DR

Onboarding to VMware Cloud DR begins with learning about and signing up to buy the service.

### 1 Learn, Try, and Buy VMware Cloud DR

Use tools and resources on the VMware Cloud Launchpad to learn about VMware Cloud DR, estimate costs, how to use it, and how to buy it:

<a href="#">Launchpad</a>	A one-stop-shop to learn, try, and buy VMware Cloud DR. Explore multiple resources and tools that help you to effectively plan and execute your data protection strategy.
<a href="#">TCO calculator</a>	Discover how much you can save with VMware Cloud DR.
<a href="#">VMware Cloud DR Planner</a>	Plan for DR by using this tool to size and estimate associated costs of your DR solution.
<a href="#">Hands On Lab</a>	Learn how to use VMware Cloud DR with the Hands On Lab.
<a href="#">Buy the service online</a>	It's easy to buy VMware Cloud DR online.

## 2 Set Up the VMware Cloud DR Service

Once you purchase VMware Cloud DR, perform these tasks to set up the service:

Task	
<a href="#">Fill out the pre-deployment checklist.</a>	Gather relevant information needed for deploying and setting up the service.
<a href="#">Add VMware Cloud DR to Your Organization</a>	Add the VMware Cloud DR service to your VMware Cloud Organization.
<a href="#">Create a Subscription</a>	Create a subscription for protected capacity based upon your DR requirements.
<a href="#">Activate Recovery Region</a>	After you add the service, you must activate an AWS region to use for recovery and data protection operations.
<a href="#">Manage Recovery Region</a>	Open VMware Cloud DR when you are ready to start configuring the main components.
<a href="#">Invite Users to Your Organization</a>	Invite users to your organization and grant them specific roles, based on their intended use of the service.
<a href="#">Create an API Token</a>	Create an API token from the VMware Cloud Services console to use with VMware Cloud DR.
<a href="#">Add the API Token</a>	When you log in to VMware Cloud DR, add an API token you created in the VMware Cloud Services console.

## 3 Configure VMware Cloud DR Main Components

To prepare your environment for disaster recovery operations, perform the following tasks:

Task	Description
<a href="#">Choose an Availability Zone for Recovery</a>	When you deploy a cloud file system for the first time, you must select an AWS Availability Zone (AZ) that is used for disaster recovery failover operations.
<a href="#">Chapter 5 Deploy a Cloud File System</a>	Deploy a cloud file system for snapshot replication and cloud backup.
<a href="#">Chapter 6 Set Up Protected Sites</a>	Set up a protected site, either a customer-managed vSphere or VMware Cloud on AWS SDDC.
<a href="#">Chapter 7 Deploy the DRaaS Connector</a>	Once you set up a protected site, deploy the DRaaS Connector on it so you can begin replicating VM snapshots to the cloud file system.

## 4 Start Using VMware Cloud DR

To start using the service:

<a href="#">Create a Protection Group</a>	Create a protection group to enable snapshot replication to a cloud file system. Snapshots are later used for failover operations to a recovery SDDC.
<a href="#">Chapter 10 Deploy a Recovery SDDC</a>	Deploy a new or add an existing SDDC for recovery operations.
<a href="#">Configure Recovery Plans</a>	Configure a recovery plan to orchestrate failover of a protected site to a recovery SDDC.
<a href="#">Chapter 12 Running a Recovery Plan for Failover</a>	Once you configure a recovery plan, you can run it as a failover, either for actual disaster recovery or for test operations.
<a href="#">Run a Ransomware Recovery Plan for ransomware recovery.</a>	You can choose to run a plan for ransomware recovery so you can analyse historical snapshots and perform forensics and recover composed VMs.
<a href="#">Chapter 13 Run a Failback Recovery Plan</a>	When you are ready to restore a protected site, run a failback recovery plan to restore the site to its last best configuration.
<a href="#">Run a Ransomware Recovery Plan</a>	You can test or run an actual ransomware recovery plan to analyze, inspect, and clean VMs attacked by ransomware and then restore them to a production site.

## Available AWS Regions

VMware Cloud DR is available in several AWS regions.

Region Name	Code
Africa (Cape Town)	af-south-1
Asia Pacific (Hong Kong)	ap-east-1



Region Name	Code
Asia Pacific (Melbourne)	ap-southeast-4
Asia Pacific (Mumbai)	ap-south-1
Asia Pacific (Osaka)	ap-northeast-3
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Tokyo)	ap-northeast-1
Canada (Central)	ca-central-1
Europe (Frankfurt)	eu-central-1
Europe (Ireland)	eu-west-1
Europe (London)	eu-west-2
Europe (Milan)	eu-south-1
Europe (Paris)	eu-west-3
Europe (Stockholm)	eu-north-1
Europe (Zurich)	eu-central-2
Middle East (Bahrain)	me-south-1
South America (São Paulo)	sa-east-1
US West (N. California)	us-west-1
US East (Ohio)	us-east-2
US West (Oregon)	us-west-2
US East (N. Virginia)	us-east-1

## Ransomware Recovery

VMware Ransomware Recovery provides an on-demand, cloud-based isolated recovery environment (IRE) with integrated security and behavior analysis tools that help you recover from a ransomware attack using cloud backups (snapshots).

## The Problem

Ransomware has emerged as a dominant threat to enterprise IT, with Gartner estimating that 75% of organizations will be affected by ransomware by 2025. Businesses affected by ransomware can often recover data from backups, although the cost of recovery in terms of time, loss of business, and partial data loss remains high.

Traditional backup and restore solutions are not designed to easily recover from a ransomware attack, and the process is costly and time consuming. After a ransomware attack, it is nearly impossible to be certain that backups are not infected without explicit validation. In many cases, the most recent backup data is likely compromised.

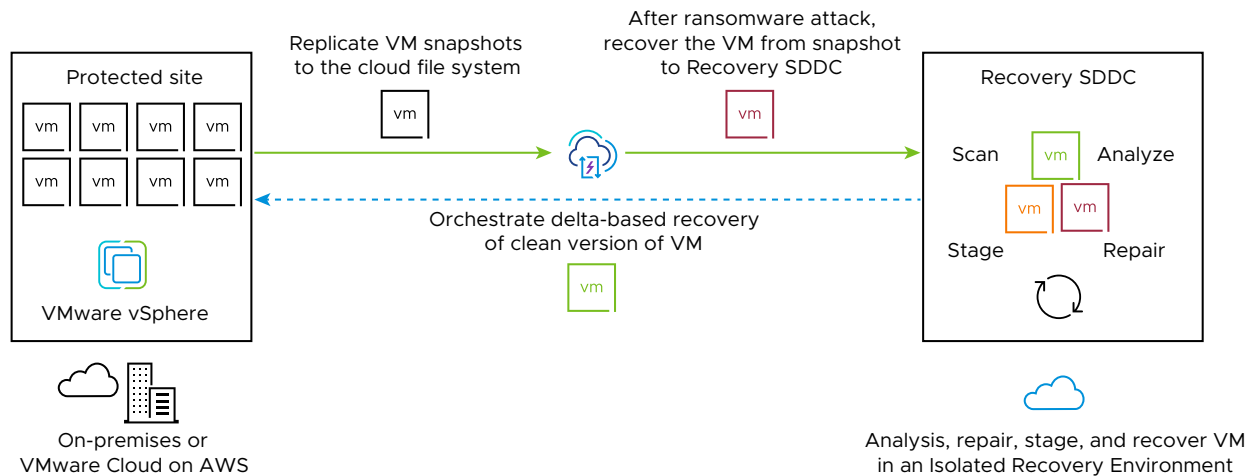
The ransomware recovery administrator must assume that malware is embedded in the backup data. Unlike disaster recovery, backups must be either validated or cleansed during the ransomware recovery process to avoid reintroducing ransomware into a production environment.

What the ransomware administrator needs is an air-gapped staging area for infected VMs that is isolated from other networks, so the forensic and remediation processes can occur without external ransomware triggers and without the risk of infecting other workloads in production.

## The Solution

To solve this problem, the VMware Cloud DR (sometimes called VCDR) provides these key capabilities with VMware Ransomware Recovery:

- **On-demand, cloud-based (IRE) with predefined network isolation levels.** The VMware Cloud DR recovery SDDC provides a network-restricted IRE on VMware Cloud on AWS as a cost effective solution which does not require building an environment from scratch and patching together different tools and hardware. After a ransomware attack, you can launch a Recovery Plan and select VMs from a deep snapshot history to be placed into an IRE for forensic analysis and validation.
- **Deep integrated security and vulnerability analysis.** VMware Ransomware Recovery provides access to continuous cloud scanning systems that analyze each VM in recovery for suspicious OS behaviors, malware file signatures, and known vulnerabilities.
- **Recovery orchestration.** You can configure recovery plans to automatically move protected VMs to the recovery SDDC for analysis and validation. When you have succeeded in finding clean VMs, you can orchestrate those VMs back to a protected production site.



For more information, see [Chapter 14 Ransomware Recovery](#).

## PCI DSS Compliance

VMware Cloud DR is now compatible with compliance hardening requirements for Payment Card Industry Data Security Standard (PCI DSS).

VMware Cloud DR compliance hardening uses a shared accountability model that distributes security and compliance responsibilities among AWS, VMware, and you the customer. Read the VMware Cloud DR [shared responsibility](#) document for supplemental guidance covering the responsibilities and ownership of compliance hardening functions in VMware Cloud DR.

As part of your shared responsibility, you should harden your recovery SDDC for PCI DSS Compliance. For more information, see [Configure SDDC Compliance Hardening](#).

You can also configure [Chapter 3 Configuring Access to VMware Cloud DR](#) in VMware Cloud DR to help you achieve PCI DSS compliance. Access lists allow you to control access from the DRaaS Connector to the cloud file system and the Orchestrator, and control which users can access the VMware Cloud DR UI.

---

**Note** If you need PCI hardening for your recovery SDDC, contact VMware Support for assistance.

---

## About VMware Cloud DR

You can view general information about VMware Cloud DR software from the settings page.

When you navigate to **Settings > About VMware Cloud DR**, you can view the following information about VMware Cloud DR software:

Item	Description
VMware Cloud DR ID	Support uses this ID to track customer cases and support calls.
Orchestrator FQDN	You can use the Orchestrator FQDN when configuring networking for a protected site. The Orchestrator is assigned an IP address if you are using AWS Direct Connect for your protected site.
Appliance OVA URL	The URL to the DRaaS Connector OVA, which downloads the connector to your protected site vCenter Server.
Recovery region	Recovery region where VMware Cloud DR is deployed.
Recovery AZ	The AWS AZ where VMware Cloud DR is deployed.

# Set Up VMware Cloud DR

## 2

After you have signed up for VMware Cloud DR, you are ready to set up the service.

When you set up the VMware Cloud DR service, you do the following:

- Add the service to your VMware Cloud organization.
- Choose a purchase option.
- Activate a recovery region.
- Create a subscription.
- Invite users to your organization.
- Open the VMware Cloud DR UI.
- Create an API token.

Read the following topics next:

- [Add VMware Cloud DR to Your Organization](#)
- [Choose a Purchase Option](#)
- [Create a Subscription](#)
- [Activate Recovery Region](#)
- [Invite Users to Your Organization](#)
- [Manage Recovery Region](#)
- [Create an API Token](#)

## Add VMware Cloud DR to Your Organization

When you add VMware Cloud DR service to your organization, it becomes accessible to your users from the VMware Cloud console.

## Procedure

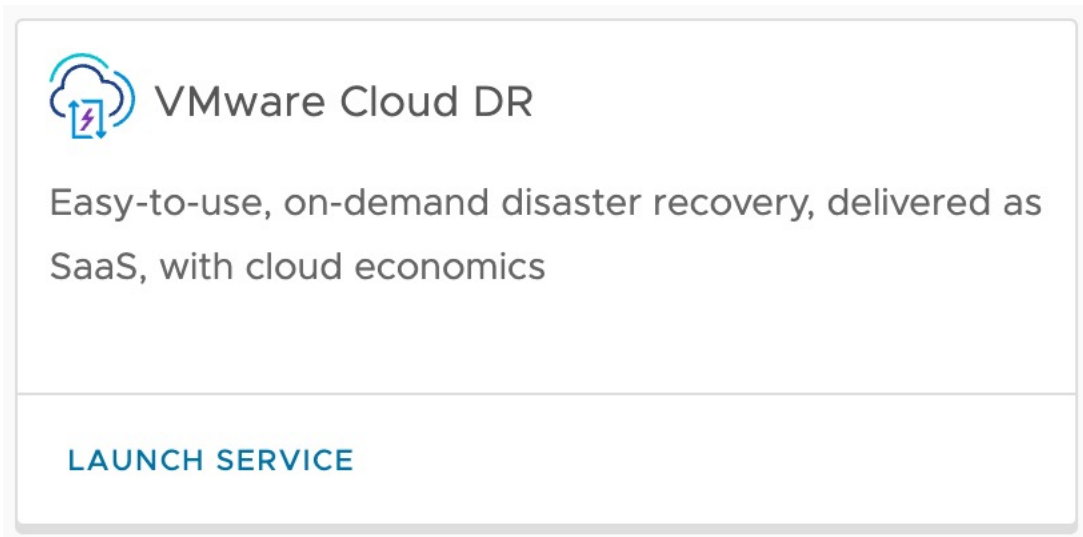
- 1 Log in to VMware Cloud Services using your VMware account at <https://console.cloud.vmware.com>.

If your account already has VMware Cloud on AWS activated, then VMware Cloud DR tile appears in the console under **Services > My Services**.

If you are a new customer or have not yet activated VMware Cloud on AWS, find VMware Cloud on AWS under Other Services. Once VMware Cloud on AWS is activated, then the VMware Cloud DR tile appears under **My Services**.

If you are purchasing VMware Cloud DR from an MSP, you can contact your MSP representative to request the service.

- 2 Once the service has been added to your organization, the VMware Cloud DR tile appears under **Services > My Services**. Click the VMware Cloud DR tile to [Activate Recovery Region](#).



## Choose a Purchase Option

You can choose several different VMware Cloud DR purchase options from the VMware website.

### Purchase Options

From the VMware Cloud DR [pricing page](#), you can view purchase options and prices:

- [Create a Subscription](#). You can pay up front or monthly for a 1-year or 3-year for a fixed amount of protected storage capacity.
- Protected VM term subscriptions. You can pay up front or monthly for 1-year or 3-year subscriptions to protect individual VMs.

- Term subscriptions for both protected capacity and protected VMs can be paid for all up front or paid monthly. Alternatively, you can pay for VMware Cloud DR completely on-demand. If you buy on-demand, protected capacity and protected VMs is metered hourly and billed monthly.

---

**Note** A minimum 10 TiB of protected capacity per-Orchestrator recovery region applies across a subscription region, irrespective of usage.

---

**Note** If you are not sure how much protected capacity you need, use the [VMware Cloud DR Planner](#) to estimate your needs.

---

## Payment Methods

To purchase VMware Cloud DR, you can use VMware funds, credit cards or link an unrestricted Pay by Invoice account. For more information, see [Managing Payment Methods](#).

## Sellers

You can choose VMware, Amazon Web Services (AWS), or any approved two-tier partners as sellers of VMware Cloud DR.

You can also choose different sellers for different sales regions, if there are no active subscriptions in the subscription region.

---

**Note** If the seller of record has already been selected for the initial subscription in a sales region, or if there is already an active term subscription in a region, then seller of record cannot be changed until all active subscriptions expire in that region.

---

If you purchase through VMware:

- VMware is the seller of record.
- Billing is done by VMware.
- You deploy the service from the VMware Cloud DR Global DR Console.
- You manage your protected sites, recovery sites, recovery plans, and protection groups using the VMware Cloud DR UI.

When you purchase through AWS:

- AWS is the seller of record.
- Billing is done by AWS.
- You deploy the service from the VMware Cloud DR Global DR Console.
- You manage your protected sites, recovery sites, recovery plans, and protection groups using the VMware Cloud DR UI.

You also have the option to buy VMware Cloud DR through a Managed Service Provider (MSP), the MSP handles billing, support, deployment, and management of your VMware Cloud on AWS infrastructure. Consult your MSP for more information.

## 2-Tier Support

With a 2-Tier option, instead of purchasing Subscription Purchasing Program (SPP) credits up front, distributors can provide volume discounts for a reseller/end customer combination, and then make payments monthly on their upfront commitment by signing a Commitment Based Contract (CBC) with VMware, committing to spend a certain amount of money on behalf of the reseller/end customer combination over a specific period. Distributors are charged monthly by VMware based on your consumption of VMware Cloud DR.

For more information, see the [VMware Cloud on AWS Release Notes](#).

## Create a Subscription

You can create a protected capacity subscription for the VMware Cloud DR service from the Global Console.

When you create a protected capacity or protected VM subscription, you can choose a one or three year term, and pay the full amount up front, monthly, or as you go with on-demand mode. Paying the full amount up front saves considerable costs.

Protected storage capacity is the sum of the logical storage size, measured in tebibytes (TiB), of your protected VMs and the incremental snapshots you replicate to a cloud file system.

---

**Note** 1 TiB = approximately 1.10 TB

---

A minimum charge of 10 TiB of data capacity per-Orchestrator recovery region applies across each subscription region. If you do not have enough subscription capacity to cover the minimum charge or actual usage, on-demand charges will apply to meet the minimum charge.

---

**Note** Recovery regions are associated with VMware Cloud subscription regions. Each subscription region contains several AWS regions that the subscription covers, and each subscription region can be associated with different sellers.

---

### Minimum Charges and Overages

If you are using more than one recovery region within a subscription region, you will be charged 10 TiB minimum per recovery region. If you activate multiple recovery regions within one subscription region and your total usage exceeds the total minimum protected capacity across the entire subscription region, your usage will be charged on-demand. If your recovery regions are in different subscription regions, each subscription region will adhere to its own 10 TiB minimum charge per recovery region.



For example, in scenario 1 below the user has activated three recovery regions within one subscription region. The total minimum required capacity for three recovery regions is 30 TiB (10 TiB for each recovery region). The total usage across this subscription region is 26 TiB. Since this is number below the 30 TiB total minimum, the user is charged 30 TiB for all three recovery regions.

In scenario 2, the user has activated one recovery region within the NA (North America) subscription region, and another recovery region in the EMEA subscription region. The user has consumed 26 TiB in the NA subscription region, and only 2 TiB in the EMEA subscription region.

Since these two recovery regions are in two separate subscription regions, they must adhere to their own 10 TiB minimum per recovery region. In the NA subscription region, the user has met the 10 TiB minimum and so is charged 26 TiB for usage in the US East Recovery Region. In the EMEA subscription region, the user only has 2 TiB of usage, which is below the 10 TiB minimum. Therefore, the user is charged 10 TiB minimum for Milan Recovery Region in the EMEA subscription region. The user's total charge for both NA and EMEA subscription regions equal 36 TiB.

Scenario	Subscription Region	Recovery Region	Usage per-Recovery Region	Total Usage	Total charged capacity
1	NA	US West	20 TiB	26 TiB	30 TiB
	NA	US East	5 TiB		
	NA	US Central	1 TiB		
2	NA	US East	26 TiB	28 TiB	36 TiB (26 TiB for NA + 10 TiB for EMEA)
	EMEA	Milan	2 TiB		

Consider the following before creating a subscription:

- You cannot cancel, convert, or modify a subscription after you have ordered it.
- You can only change a seller for a subscription if there are no active subscriptions in the region.
- If you require more than 1600 TiB of protected capacity, contact VMware support

#### Procedure

- 1 Log in to VMware Cloud Services using your VMware account at <https://console.cloud.vmware.com>.
- 2 Under **Services > My Services**, click Launch service in the VMware Cloud DR tile.
- 3 From the left navigation, select **Subscriptions**.
- 4 On the Subscriptions page, click the **Create Subscription** button.

- 5 On the Seller and region page, select a subscription region from the drop-down menu.

When you select a subscription region, you see all of the associated AWS regions that your subscription will cover.

- 6 Next, select a seller of record for the subscription.

If the seller of record has already been selected for the initial subscription in this sales region, the seller of record is already selected. Or, if there is already an active term subscription in this region, then seller of record that was chosen for that subscription is already selected and cannot be changed.

---

**Note** You can switch the seller of record if there are no active subscriptions in the subscription region.

---

- 7 On the Quantity and term page, select the quantity and term for the subscription.

You can choose protected capacity and/or any number of protected VMs. If you already have subscriptions, the current terms and VM count is listed.

You can also add VMs for [Chapter 14 Ransomware Recovery](#). For information on how much the ransomware add-on costs, see the VMware Cloud DR [pricing page](#).

- 8 On the Quantity and term page, select a subscription term, either one or three years, and monthly or pay in full and click **Next**.
- 9 On the Summary page, confirm the agreement terms and then click the **Finish**.

## Activate Recovery Region

To start using VMware Cloud DR, activate a recovery region from the Global DR Console.

When you activate a recovery region, you establish the Orchestrator and cloud file system on VMware Cloud on AWS. The recovery region is where you deploy recovery SDDCs, run failover recovery plans, and replicate snapshots to the cloud file system using protection groups.

---

**Note** If a recovery region activation fails, see [Reactivate a Recovery Region](#).

---

### Procedure

- 1 Under My services, click the VMware Cloud DR tile.
- 2 If this is the first time you are activating a region, click the **Pre-Deployment Checklist** button to view important information about setting up VMware Cloud DR. After viewing the checklist, select the agreement check box to acknowledge that you have read and understand the VMware Cloud DR pre-deployment checklist.
- 3 Click the **Activate Recovery Region** button.

- 4 On the Activation page, select the AWS region you want to use for VMware Cloud DR and click **Next**.

The region you select is where the Orchestrator is deployed. When you deploy a cloud file system later, it will deploy to this selected AWS region. When you choose an AWS region, it is associated with a VMware Cloud subscription region. A subscription region covers one or more AWS regions and can cover consumption in any of the AWS regions within the subscription region.

- 5 On the Activation page, under Purchase through select a seller.

Only sellers configured with your organization appear here. If you have active subscriptions in the VMware Cloud subscription region associated with this AWS region, then only that seller is available and is automatically selected. You can change seller for future region deployments only if there are no active subscriptions in the region.

- 6 Review the payment details, which shows the hourly and monthly rate for the service and then click **Next**.

A summary of those subscriptions is displayed here, if you currently have active subscriptions associated with this region.

- 7 On the Summary page, confirm the agreement terms at the bottom of the page, and then click **Finish**. (You can only activate one recovery region at a time.)

#### What to do next

On the Recovery regions page, after the recovery region finishes deploying you now can:

- Activate another recovery region.
- [Create a Subscription](#).
- [Deactivate Recovery Region](#).

## Deactivate Recovery Region

When you no longer need a region for recovery, you can deactivate it.

#### Prerequisites

When you deactivate a recovery region:

- All replicated data and service configuration settings in the selected recovery region is deleted. Complete deletion of the data can take up to ten (10) days.
- All recovered VMs running off any cloud file systems in the region are terminated.
- All recovered VMs in the region are deleted.
- All access to VMware Cloud DR in this recovery region is lost.
- Any active term subscriptions applicable to the region are canceled.
- Billing for the service stops within one (1) hour.

- You cannot re-activate the recovery region until all of your data has been deleted. Full deactivation can take up to 10 days, during which the deactivated region tile remains until the region is completely removed.

#### Procedure

- 1 Log in to VMware Cloud Services using your VMware account at <https://console.cloud.vmware.com>.
- 2 Under **Services**, click Launch service in the VMware Cloud DR tile.
- 3 From the left navigation, select **Region**.
- 4 On the tile select the Actions drop-down menu on the region tile and then select **Deactivate Region**.
- 5 In the Deactivate region confirmation page select all of the check boxes to confirm the region deactivation.
- 6 To deactivate the region permanently, click the **Deactivate** button. It can take up to 10 days before the deactivated region tile disappears from the recovery regions page. If you click a deactivated service tile, a "page not found" error is displayed.

## Reactivate a Recovery Region

If a recovery region fails activation, you can reactivate it.

If a region activation fails:

- Deactivate the recovery region.
- After the region is deactivated, activate the recovery region again.

#### Procedure

- 1 If a recovery region activation fails for any reason, from your organization find the region that failed to activate and click the **Deactivate** button.
- 2 After the region has fully deactivated, click the **Activate** button for the region.

## Invite Users to Your Organization

After you receive access to VMware Cloud DR, invite users to your organization.

As an organization owner, when you [invite users](#) to your organization, you assign them roles which specify privileges that an organization member has over organization assets, and service roles. Service roles give users the permission to use the VMware Cloud DR service.

For more information about VMware Cloud service roles, and how to add them to your users, see [Identity and Access Management](#) and [Edit User Roles](#).

**Note** When you modify a VMware Cloud DR user role in the VMware Cloud console, the changes take approximately 15 minutes to be applied. To apply the changes faster, the user can log out and then log back in to the VMware Cloud console, and then access the VMware Cloud DR service.

## User Roles

When setting up user access, you assign your users roles so that they can perform specific tasks with VMware Cloud DR.

Each role has a specific set of operations associated with it, so when a user is assigned a role, the user can perform all operations associated with that role.

VMware Cloud has two general categories of roles: organization and service. Organization roles provide capabilities for working with the VMware Cloud Services platform, such as adding users and creating API tokens. Service roles provide capabilities related directly to a VMware Cloud service.

The ability to create an API Token requires the following organization role:

- Organization Owner

VMware Cloud DR also requires the following two VMware Cloud service roles:

- VMware Cloud on AWS Administrator
- VMware Cloud on AWS NSX Cloud Admin

## VMware Cloud DR Service Roles

VMware Cloud DR service roles provide users access to specific features.

The following table provides an overview of VMware Cloud DR roles and the features each role permits. Match the user role in each column with the capabilities in each row.

**Note** VMware Cloud DR roles are additive. For example, if you want a user to create snapshots for backup and configure and run recovery plans, you must assign the Recovery Admin and Protection Admin roles to the user account.

The following table lists all roles and operations related to using the VMware Cloud DR UI.

Capability	Recovery Admin	Recovery Tester	Protection Admin	Recovery SDDC Admin	Data Protection Auditor	Orchestrator Admin
Configure API token			✓	✓		✓
Edit Access lists						✓

Capability	Recovery Admin	Recovery Tester	Protection Admin	Recovery SDDC Admin	Data Protection Auditor	Orchestrator Admin
Edit Plans	✓	✓				✓
Plan test	✓	✓				✓
Plan recovery	✓					✓
Activating ransomware recovery services requires: Organization owner, Global Console Admin, and Orchestrator Admin roles						✓
Run plan for ransomware recovery	✓	✓				
Full recovery of VMs during ransomware recovery	✓					
Open ransomware recovery security console	✓ In addition to Recovery Admin, viewing the Carbon Black Cloud console requires at least one Carbon Black Cloud role. For more information about Carbon Black Cloud user roles, see <a href="#">User Roles</a> .					
Deploy a cloud file system			✓			
Create and edit protection groups, set snapshot schedules			✓			✓
VM restore, guest file recovery			✓			✓

Capability	Recovery Admin	Recovery Tester	Protection Admin	Recovery SDDC Admin	Data Protection Auditor	Orchestrator Admin
Deploy the DRaaS Connector, create and edit protected sites, run connectivity check			✓			✓
Deploy/add, edit, delete SDDC				✓		✓
View compliance checks	✓	✓	✓	✓	✓	✓
Reports	✓	✓	✓	✓	✓	✓
View data	✓	✓	✓	✓	✓	✓

The following table describes all roles and operations related to recovery regions and subscriptions within the VMware Cloud DR Global Console.

Capability	Global Console Admin	Deployment Admin (activation)	Deployment Admin (deactivation)	Subscription Admin
View existing deployments and their metadata.	✓	✓	✓	
Create a new VMware Cloud DR deployment in any of the supported regions.	✓	✓		
Delete a partial/failed deployment.	✓		✓	
Delete any of the existing deployments in any of the regions.			✓	
View term subscriptions and pricing	✓			✓
Create new commit subscriptions	✓			✓

## Service Roles and Permitted Operations

VMware Cloud DR service roles provide access to specific features and functionality.

Role	Permitted Operations
Orchestrator Admin	This user role can perform all operations listed in this table, except for creating a subscription.
Data Protection Auditor	<ul style="list-style-type: none"> <li>■ View the UI read-only: Lists, tasks, reports, dialogs (except user management)</li> <li>■ Create PDF of a compliance report and download it</li> </ul> <p><b>Note</b> All other roles include this level of access.</p>
Recovery Admin	<p><b>Recovery plans</b></p> <ul style="list-style-type: none"> <li>■ Create, edit, delete, duplicate recovery plans</li> </ul> <p><b>Test disaster recovery</b></p> <ul style="list-style-type: none"> <li>■ Run a recovery plan to test disaster recovery</li> <li>■ Stop a test disaster recovery (no cleanup)</li> <li>■ Cancel test disaster recovery</li> <li>■ Roll back test disaster recovery</li> <li>■ Retry failed tasks for a completed test disaster recovery</li> <li>■ Retry failed tasks in a step and continue</li> <li>■ Ignore failure and continue without retry</li> <li>■ Continue tasks after user confirmation</li> <li>■ View recovery results</li> </ul> <p><b>Disaster recovery</b></p> <ul style="list-style-type: none"> <li>■ Run recovery plan for disaster recovery</li> <li>■ Preview disaster recovery</li> <li>■ Stop a running task</li> <li>■ Cancel disaster recovery</li> <li>■ Retry failed tasks for a completed disaster recovery operation</li> <li>■ Retry failed tasks in a step and continue</li> <li>■ Ignore failures, continue without retry</li> <li>■ Continue task post user confirmation</li> <li>■ Commit plan after recovery</li> </ul> <p><b>Ransomware recovery</b></p> <ul style="list-style-type: none"> <li>■ Run plan for ransomware recovery</li> <li>■ Full recovery of VMs during ransomware recovery</li> <li>■ Open ransomware recovery security console (also requires at least one Carbon Black Cloud role)</li> </ul>



Role	Permitted Operations
Protection Admin	<p><b>Cloud file system</b></p> <ul style="list-style-type: none"> <li>■ Deploy a cloud file system</li> </ul> <hr/> <p><b>Note</b></p> <p>After you deply the first cloud file system, contact VMware support for assistance deploying more.</p> <p><b>API token</b></p> <ul style="list-style-type: none"> <li>■ Configure API token</li> </ul> <p><b>Protected sites</b></p> <ul style="list-style-type: none"> <li>■ Create, update, delete a protected site</li> <li>■ Add or remove a DRaaS Connector to/from a protected site</li> <li>■ Add or remove a vCenter Server to/from a protected site</li> <li>■ Run a connectivity check for the DRaaS Connector</li> </ul> <p><b>Protection groups</b></p> <ul style="list-style-type: none"> <li>■ Create, edit, delete a protection group</li> <li>■ Activate/deactivate protection group</li> <li>■ Schedule, take, delete snapshots</li> <li>■ Restore, edit, delete a snapshot</li> </ul> <p><b>VMs</b></p> <ul style="list-style-type: none"> <li>■ Restore VM</li> <li>■ Guest file recovery</li> </ul> <p><b>Ransomware recovery</b></p> <ul style="list-style-type: none"> <li>■ Run a plan for ransomware recovery</li> </ul>
Recovery Tester	<p><b>Recovery plan</b></p> <ul style="list-style-type: none"> <li>■ Create, edit, delete, duplicate a recovery plan</li> </ul> <p><b>Test recovery</b></p> <ul style="list-style-type: none"> <li>■ Run a test recovery</li> <li>■ Stop a test recovery</li> <li>■ Cancel a test recovery</li> <li>■ Rollback a test recovery</li> <li>■ Retry failed tasks</li> <li>■ Ignore failed tasks</li> </ul> <p>View and clear alarms</p>

Role	Permitted Operations
Recovery SDDC admin	<b>Recovery SDDCs</b> <ul style="list-style-type: none"> <li>■ Configure an API token</li> <li>■ Deploy or add, edit, delete a recovery SDDC</li> <li>■ Add, rename, or delete a network on an SDDC</li> <li>■ Request a new public IP address</li> <li>■ Rename or delete a public IP address</li> <li>■ Add, remove hosts</li> <li>■ Add, edit, delete NAT rules</li> <li>■ Add, edit, delete new firewall rules</li> </ul> <b>API token</b> <ul style="list-style-type: none"> <li>■ Configure API token</li> </ul>
Global Console Admin	<ul style="list-style-type: none"> <li>■ View existing deployments and their metadata</li> <li>■ Create a new VMware Cloud DR deployment in any of the supported regions</li> <li>■ Delete a partial/failed deployment</li> <li>■ Query offers with all pricing information</li> <li>■ View existing subscriptions and pricing</li> <li>■ Create new commit subscriptions</li> </ul>
Deployment Admin (activation)	<ul style="list-style-type: none"> <li>■ Create a new VMware Cloud DR deployment in any of the supported regions</li> <li>■ View existing deployments and their metadata.</li> </ul>
Deployment Admin (deactivation)	<ul style="list-style-type: none"> <li>■ Delete any existing deployment in any of the regions</li> <li>■ Delete a partial/failed deployment</li> <li>■ View existing deployments and their metadata</li> </ul>
Subscription Admin	<ul style="list-style-type: none"> <li>■ View term subscriptions and pricing</li> <li>■ Create new commit subscriptions</li> </ul>

## Manage Recovery Region

You manage a recovery region to use the VMware Cloud DR UI.

### Procedure

- 1 Log in to VMware Cloud Services using your VMware account at <https://console.cloud.vmware.com>.
- 2 Under **Services > My Services**, click the 'Launch service' link on the VMware Cloud DR tile.
- 3 Click the tile.
- 4 From the left navigation, select **Regions**.
- 5 From a region tile, click the **Manage Region** button.

## Create an API Token

To use VMware Cloud DR, you first must create an API token in the VMware Cloud console.

Before your users access the VMware Cloud DR UI, create an API token to authorize service access for your organization. For instructions on how to create an API token from the VMware Cloud Services console, see [generate an API token](#).

---

**Note** Using Multi-Factor Authentication (MFA) with API tokens is currently not supported with VMware Cloud DR. This limitation applies only to MFA for API tokens (**My account > API Tokens**), and does not apply to your organization authentication policy (**Organization > Authentication Policy > Multi-Factor Authentication**) or your VMware Cloud user account (**My account > Security**).

---

When you create an API token, you define its scope of permissions by assigning specific organization roles and service roles. For VMware Cloud DR, scope the following roles to the API token.

- **Organization Role:** Organization Owner
- **Service Roles:**
  - VMware Cloud on AWS Administrator
  - VMware Cloud on AWS NSX Cloud Admin

The maximum lifespan of a VMware Cloud Services API token is 60 months, after which you must regenerate a new token and configure it inside of VMware Cloud DR. If you do not regenerate a new token when the old one expires, the product features cannot function. The best practice in this case is to create an API token with the longest Time To Live (TTL) possible, to avoid service interruption.

After you create the API token, you can [Add the API Token](#) to the VMware Cloud DR UI.

---

**Important** Your user account must have the Organization Owner role and VMware Cloud Services service roles (Administrator and NSX Cloud Admin) associated with it to create an API token to use with VMware Cloud DR.

---

## Add the API Token

Once you have created an API token, you must add that token to VMware Cloud DR.

---

**Note** The maximum lifespan of a VMware Cloud Services API token is 60 months, after which you must regenerate the API token in the VMware Cloud Console. Then, you can add the new API token in the VMware Cloud DR UI.

---

### Procedure

- 1 From the left navigation, select **Settings**.
- 2 Under API token, click the **API token** button.

- 3 In the **Configure API token** dialog box, enter the API token.
- 4 Click the **Validate** button.
- 5 Click **OK**.

## Change the API Token

Before the current API token expires, you need to regenerate a new token from the VMware Cloud on AWS console, and then replace the existing token in VMware Cloud DR.

The maximum lifespan of a VMware Cloud Services API token is 60 months, after which you must regenerate a new token and configure it inside of VMware Cloud DR.

You can [generate a new API token](#) in the VMware Cloud console. After you generate a new API token, you then add it to VMware Cloud DR.

### Procedure

- 1 From the left navigation, select **Settings**.
- 2 Under API token, click the **API token** button.
- 3 In the **Configure API token** dialog box, click the **Change token** button.
- 4 Enter the new API token in the dialog box.
- 5 Click the **Validate** button.
- 6 Click **OK**.

# Configuring Access to VMware Cloud DR

## 3

To control access to VMware Cloud DR components, you can configure access lists.

Configuring access lists allows you to:

- Control access from the DRaaS Connector to the cloud file system and the Orchestrator.
- Control which users can access the VMware Cloud DR UI, including users who want to [Recover VM Guest Files](#) and download them.
- Harden your VMware Cloud DR environment for [PCI DSS Compliance](#).

Once you enable this setting, only IP addresses or IP address ranges added here can access the service.

## Access Lists

There are two access lists you can configure:

Access List	Description
Connector access list	<p>Specify the public IP addresses and/or IP address ranges for all DRaaS Connectors that can access the Orchestrator and a cloud file system.</p> <p><b>Note</b> Do not enter private IP addresses that are behind a NAT gateway.</p>
Management access list	<p>Specify the public IP addresses and/or IP address ranges for all users you want to allow access to the VMware Cloud DR UI.</p> <p><b>Note</b> Do not enter private IP addresses that are behind a NAT gateway.</p>

## Configure Access to VMware Cloud DR

You can configure access lists to only allow specific IP addresses to access VMware Cloud DR components and UI.

Before you enable this setting, make sure that you compile a list of all allowed IP addresses or IP address ranges of all deployed DRaaS Connectors and all IP addresses to add to the lists. Once you enable this setting, only IP addresses or IP address ranges added here can access the service.

---

**Note** Do not enter private IP addresses that are behind a NAT gateway.

---

### Configure Access.

- 1 From the left navigation, select **Settings**.
- 2 Click the **Security and compliance** button.
- 3 In the **Security and compliance** dialog box, select the **Use access list** option.
- 4 Under Connector access list, enter the public IP addresses and/or IP address ranges for all DRaaS ConnectorDRaaS Connectors.

When you deploy a new DRaaS Connector, or if you already have DRaaS Connector deployed, add the IP addresses here. If you do not know the IP addresses of existing DRaaS Connectors, enter one IP address in the list and the dialog box displays all deployed connectors and their IP addresses at the bottom. IP addresses in the connector access list can also access the VMware Cloud DR UI.

- 5 Next, specify the public IP addresses / IP address ranges of all computers that you want to access to the VMware Cloud DR (sometimes called VCDR) UI.

For example, to allow a specific user's computer to download a VM guest file, enter the user's computer IP address here.

## Security and compliance



## IP address access lists

For improved security, restrict access to VMware Cloud DR in US East (N. Virginia) to specific IP addresses.

☒ Use access lists

## Connector access list

IP addresses in this list can access the cloud file systems, the orchestrator, UI and guest file downloads. All connectors must be in this list.

104.175.5.212

## Management access list

IP addresses in this list can access the VMware Cloud DR UI and guest file downloads.

192.2.22.1

Enter IP addresses or subnets separated by spaces or new lines. Example: 123.145.167.89 123.145.167.0/24. VMware Cloud DR services always have access to the orchestrator.

**IMPORTANT!** You may permanently lose access if you are not in an access list.

## PCI DSS

PCI DSS is an information security standard administered by the Payment Card Industry Security Standards Council applicable to organizations that handle branded credit cards. Please enable the option below if you are protecting workloads that persist or process cardholder data and are subject to PCI DSS. Selecting this option is necessary but not sufficient for PCI DSS compliance. Refer to the [Shared Responsibility Model](#) for supplemental guidance relevant to PCI DSS compliance hardening.

☒ I consent to periodic PCI DSS related scans of VMware Cloud DR in recovery region US East (N. Virginia)

This is relevant to the following controls in PCI DSS version 3.2.1 and 4.0:

- Security vulnerabilities are identified and addressed
- External and internal vulnerabilities scans
- Penetration testing methodology is defined, documented, and implemented

CANCEL

OK

6 Click **OK**.

# VMware Cloud DR Dashboard

# 4

The VMware Cloud DR dashboard gives a high level view of a recovery region, providing protection and recoverability health status and information about storage capacity, deployments, sites, and more.

Dashboard for US East (N. Virginia)

## Welcome to VMware Cloud Disaster Recovery

VMware Cloud DR is VMware's easy-to-use, on-demand disaster recovery service, delivered as SaaS, with cloud economics.



**Quick setup**

- 1 Configure the API token
- 2 Deploy the cloud file system
- 3 Set up a protected site
- 4 Create a protection group
- 5 Add the recovery SDDC
- 6 Create the Recovery plan

### Recovery region summary

<b>Status</b> <ul style="list-style-type: none"><li>Protection</li><li>Recoverability</li></ul>	<b>Cloud file systems</b> <div>1</div> <div>22.7 GiB</div> <div>Calculated every 12h</div>	<b>Protected sites</b> <div>0</div> VMware Cloud <div>2</div> On-premises
<b>Recovery SDDCs</b> <div>0</div>	<b>Protection</b> <div>4</div> Protection groups <div>2 VMs</div> VMs in groups	<b>Recovery</b> <div>6</div> Recovery plans <div>4 VMs - disaster</div> <div>2 VMs - ransomware</div> VMs in plans

### Sites

cloudBackup	22.7 GiB backup size
onPrem	1 vCenter 1 connector
onPrem2	1 vCenter 1 connector

### Topology



Read the following topics next:

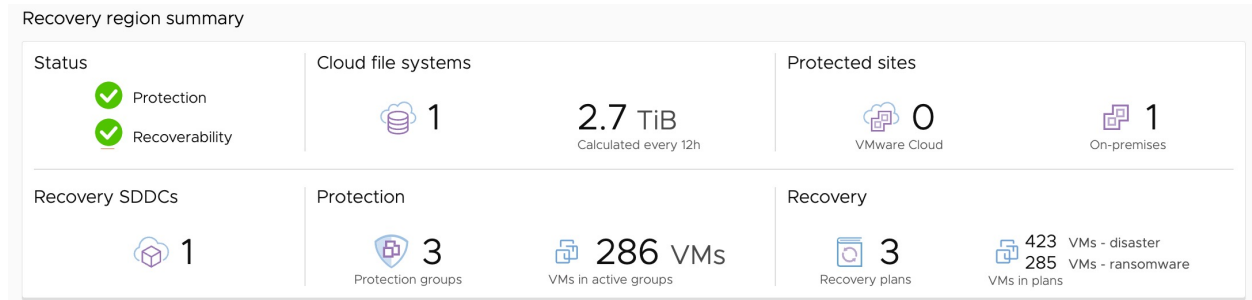
- [Recovery Region Summary](#)
- [Quick Setup Guide](#)



- [The Sites List](#)
- [The Topology Map](#)
- [Running Tasks and Recent Alarms](#)

## Recovery Region Summary

The Recovery region summary panel displays the overall system health, deployed cloud file systems, protected sites, protection groups, recovery plans, and recovery SDDCs.



The summary panel shows the following information:

Section	Description
<b>Status</b>	Overall status for: <ul style="list-style-type: none"> <li>■ <b>Protection.</b> Protected sites, protection groups, snapshots, and cloud file systems.</li> <li>■ <b>Recoverability.</b> Recovery plans, recovery SDDCs, and cloud file systems.</li> </ul> Status can be OK (green), warning (yellow), or red (critical). For more information, see <a href="#">SLA Status</a> .
<b>Cloud file systems</b>	Shows the total number of deployed cloud file systems, and the amount of protected storage capacity across all cloud file systems.
<b>Protected sites</b>	Shows the number of protected sites, both on-premises vSphere sites and protected VMware Cloud on AWS SDDCs.
<b>Recovery SDDCs</b>	Shows the currently deployed recovery SDDCs.
<b>Protection</b>	Shows the total number of protection groups, and the total number of VMs in those groups that are being replicated to a cloud file system.
<b>Recovery</b>	Shows the number of configured recovery plans and the total number of VMs in all plans, also showing how many VMs are in disaster recovery plans and ransomware recovery plans.

## Quick Setup Guide

The VMware Cloud DR dashboard quick setup guide shows you the main tasks you need to perform for disaster and ransomware recovery.

For each main workflow, the quick setup window provides a list of clickable links that lead directly to the feature.



## The Sites List

The Sites list on the dashboard shows all configured sites in your organization: cloud file systems, recovery SDDCs, and protected sites (configured with the DRaaS Connector).

From the dashboard, you see a list of all configured sites in your organization. Click a site name to visit the site configuration.

## Sites



Cloud file system



2.7 TiB backup size



Protected site



1 vCenter  
1 connector



Recovery SDDC



1 host  
10.4 TiB storage

## The Topology Map

The topology map shows a dynamic representation of your VMware Cloud DR deployment.

The topology map shows protected on-premises vSphere sites, recovery SDDCs that are used for recovery to VMware Cloud on AWS, and cloud file systems. The topology map nodes also reflect the health of the given entity as monitored by VMware Cloud DR.








The color of an icon indicates the node's health:

- Green indicates if it is healthy.
- Blue can change to yellow or red, depending on the node's health.
- Gray nodes are currently not reporting health and are likely not participating in any recovery plans.

## Arrows and Icons

The topology map uses different icons and arrows to show connections between the recovery SDDC, on-premises protected sites, VMware Cloud on AWS protected SDDCs, and cloud file systems.

You can highlight a node in the topology map, or highlight an item in a site list, by moving your mouse pointer over them, which highlights the protected site in the topology map and its corresponding system list item. In addition, highlighting directly connected nodes and related edges in the topology (and all others de-emphasized) gives you a better view of systems that are connected to the highlighted system.





Icon	Name
	recovery SDDC
	A VMware Cloud on AWS protected SDDC
	On-premises protected vSphere site
	Cloud file system.
	Dashed lines with an arrowhead in the topology represent single direction replication between the connected entities.
	The bolder, dashed, non-arrowhead edge connecting a cloud file system to an SDDC node represents the live mount connection from the cloud file system to the recovery SDDC. When replication occurs between nodes, the arrowhead edges turn from dashed to solid, and when highlighted, display the current data throughput rate for the data replication.
	A line with slightly longer blue dashes ending in an arrowhead represent a configured recovery from one protected site to a recovery site (such as an SDDC):

## Running Tasks and Recent Alarms



The right side of VMware Cloud DR UI displays status information about running tasks, finished tasks, and recent alarms.

Here you can view currently running tasks, cancel some tasks, view completed tasks, and any alarms that might be important to you.





## Recent alarms

-  SDDC VCHA SDDC will expire in 1 days  
17h ago [x](#)
-  SDDC VCHA SDDC is provisioned and running for 59 days  
17h ago [x](#)
-  SDDC VCHA SDDC is provisioned and running for 58 days  
2d ago [x](#)
-  SDDC VCHA SDDC will expire in 2 days  
2d ago [x](#)

## Running tasks

-  Recovering from ransomware... [R WR windows](#)  
4d ago
-  Recovering from ransomware... [R WR1](#)  
4d ago

## Recently finished tasks

-  Snapshot PG w2-hs3-q0607\_cj\_100 - Half-hourly - 2022-11-14T17:00 UTC  
13m ago
-  Snapshot PG w2-hs3-q0607\_cj\_100 - Half-hourly - 2022-11-14T16:30 UTC  
43m ago
-  Snapshot PG w2-hs3-q0607\_cj\_100 - Half-hourly - 2022-11-14T16:00 UTC  
1h ago
-  Snapshot PG live - dvx38 - Every 8 hours - 2022-11-14T16:00 UTC  
1h ago

For a more detailed list of all tasks and events in the VMware Cloud DRUI, see [Chapter 16 Monitor Events, Tasks, and Alarms](#).

# Deploy a Cloud File System

# 5

To enable storage capacity for your protected sites, deploy a cloud file system.

VMware Cloud DR (sometimes called VCDR) replicates protection group snapshots to the cloud file system for backup. These snapshots are later used for failover operations to a recovery SDDC.

---

**Note** The cloud file system is also sometimes referred to as the Scale-Out Cloud File System, or SCFS.

---

All cloud file systems and all recovery SDDCs must be the same AZ inside one AWS region. This specific AZ is referred to as the recovery AZ, where you deploy recovery SDDCs and add existing SDDCs for recovery operations. For more background information, see [Choose an Availability Zone for Recovery](#).

When you deploy a cloud file system, you make three selections:

- Select a single recovery AZ to use exclusively for recovery operations and snapshot replication. (If a cloud file system has already been deployed, then the recovery AZ is already selected for you.)
- Select an existing SDDC to use as recovery SDDCs in failover operations. A recovery SDDC can only exist in the same AZ that you choose as the recovery AZ.
- Select an existing SDDC that you want to protect with the DRaaS Connector. You can only choose an SDDC that does not exist in the recovery AZ. (On-premises SDDCs are not bound by this constraint.)

## Procedure

- 1 If this is the first time deploying a cloud file system, then click the Deploy cloud file system link in the quick setup panel. (If you want to deploy more than one cloud file system, contact VMware support for assistance.)
- 2 Under Protected sites select any VMware sites you want to protect. Any SDDC you select here must be in a different AZ than the AZ you plan to use for recovery operations.
- 3 Next, under recovery SDDC, you can either deploy a new SDDC or choose existing SDDCs in VMware Cloud on AWS. If you select an existing SDDC, it cannot be in the same AZ as the SDDCs you chose for protection.

- 4 Under Recovery AZ, choose an AZ that is different than the AZ where your protected SDDCs are deployed. If a cloud file system has already been deployed, then the recovery AZ is already selected.

---

**Note** Once you select the availability zone for a cloud file system and recovery SDDCs, you cannot change it.

---

- 5 Enter a name for the new cloud file system.
- 6 Click **Deploy** to deploy the cloud file system.

## Video: What is the Scale-Out Cloud File System?

The Scale-Out Cloud File System (SCFS) enables the efficient storage of backups of protected VMs in cloud storage and allows VMs to be recovered quickly, without requiring data rehydration.

The SCFS, or cloud file system for short, provides cloud storage for snapshot replications so you can quickly and easily fail over recovery plans to a recovery SDDC.

How exactly does the cloud file system work? Watch this video for more details:



([What is the Scale Out Cloud File System \(SCFS\)?](#))

## Choose an Availability Zone for Recovery

The first time you deploy a cloud file system, you must select one AWS availability zone (AZ) for all backup and recovery.

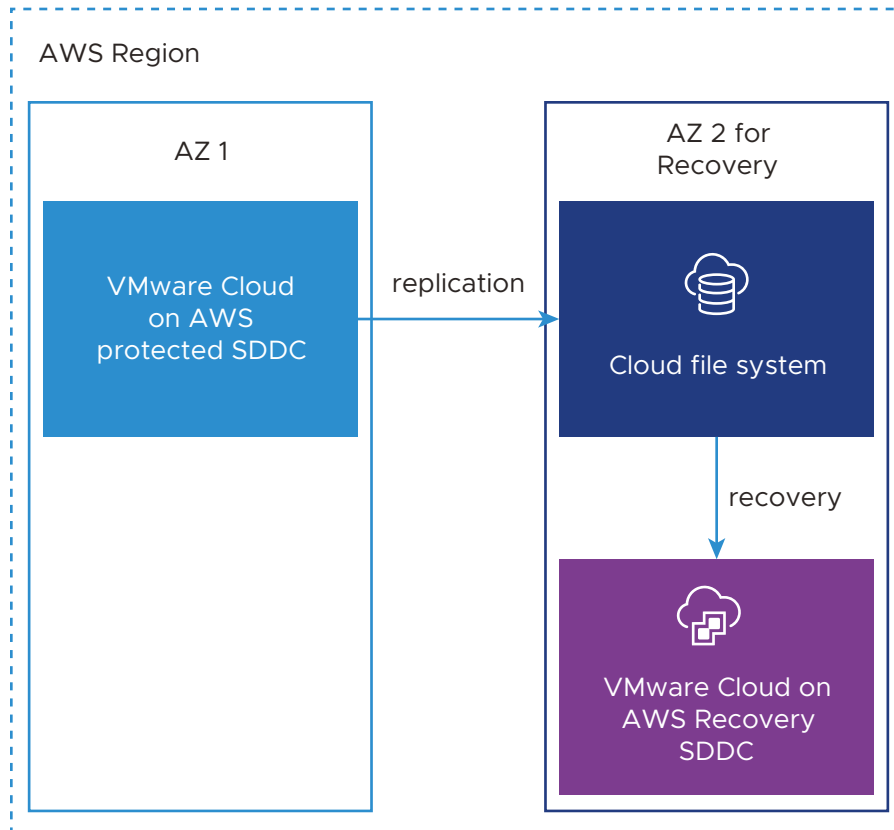
All cloud file systems and all recovery SDDCs must reside the same AZ inside one AWS region. This applies to newly deployed recovery SDDCs and existing SDDCs added for recovery.

Once you select the recovery region and choose an AZ you will not be able to protect SDDCs in that AZ. Your protected SDDCs can be in the same region as your recovery SDDC, but not in the same availability zone.

## Recovery in the Same AWS Region

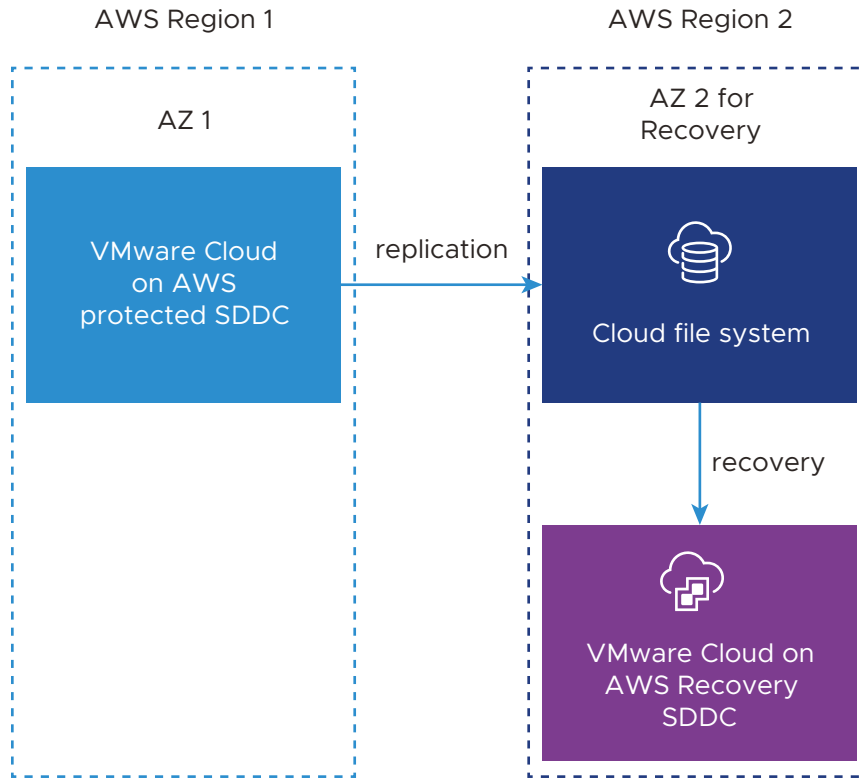
If you want to protect your VMware Cloud on AWS SDDCs using a single AWS region, when you deploy the cloud file system you must select an AZ that is different than the availability zone used for protected SDDCs. This choice is for users who are restricted to one AWS region due to region unavailability or from data residency laws.





## Recovery in a Different AWS Region

If you want to protect your SDDC in one AWS region, then perform backup and recovery to another AWS region, when you deploy the cloud file system, select a different AWS region than the one used for protected SDDCs. This choice is for users who plan to keep their protected data in one AWS region and then backup and recovery to another AWS region.




---

**Note** After first deployment, to increase the number of cloud file systems deployed in your organization, contact VMware Support.

---

## Availability Zone Failure Handling

VMware Cloud DR is designed to handle AWS Availability Zone (AZ) failures.

As part of the service architecture, VMware Cloud DR designates one availability zone (AZ) for recovery. You choose this AZ when you deploy a cloud file system for the first time. This AZ is known as the recovery AZ in that AWS region, and it cannot contain any protected SDDCs. For more information, see [Choose an Availability Zone for Recovery](#).

After the recovery AZ is designated, all other AZs in the region can be used for protected SDDCs for intra-region recovery operations.

In the unlikely event of an AZ failure, the following conditions apply to VMware Cloud DR service components:

- If an AZ where your protected SDDCs are deployed fails, all protected SDDCs in that AZ can be failed over to the recovery SDDC.
- If the recovery AZ where VMware Cloud DR is deployed fails, the Orchestrator and all cloud file systems restart automatically. If a cloud file system has a recovery SDDC attached, then you must redeploy a new or add an existing recovery SDDC after the cloud file system has been restarted.

- If the recovery AZ is down for an extended period of time, contact VMware support for assistance migrating your cloud file systems to a new recovery AZ. During this process, production workloads continue to operate, although they are not protected until all cloud file systems are recovered. All recovery SDDCs must be redeployed (or added) after cloud file system migrations.
- Contact VMware support for assistance with recovering the VMware Cloud DR database in AWS.

## Cloud File System Information

The cloud file system provides cloud storage for snapshot replication that enables failover operations and ransomware recovery on a recovery SDDC.

### Summary

When you select a cloud file system, you can see the following information from the Summary tab:

Cloud File System Information	Description
<b>Details</b>	Shows the name of the cloud file system and the availability zone where you deployed it.
<b>Capacity</b>	<b>Protected capacity.</b> The billable amount of storage capacity used by protected VMs and their snapshots.
<b>Protected sites</b>	Lists the number of protected sites (on-premises or VMware Cloud on AWS SDDC) replicating snapshots to this cloud file system.
<b>Protection</b>	<p>The Protection panel shows the following information:</p> <ul style="list-style-type: none"> <li>■ <b>Protected VMs.</b> The total number of VMs being protected by snapshots. Also displays the percentage of protected VMs out of the total number of VMs in vCenter Server.</li> <li>■ <b>Snapshots.</b> The total number of protection group snapshots and VM snapshots taken.</li> </ul>
<b>Recovery</b>	<p>The Recovery panel shows the following information:</p> <ul style="list-style-type: none"> <li>■ <b>Paired Recovery SDDC.</b> Displays the recovery SDDC that is paired with this cloud file system.</li> <li>■ <b>Live VMs.</b> Current number of VMs running live on the cloud file system.</li> </ul>

## Capacity Utilization and Storage Space Reclamation

On the cloud file system, capacity usage is computed as part of a background storage space reclamation job that runs every 13 hours. When large amount of data gets added to or removed from the cloud file system, storage space reclamation can take a long time, so storage capacity might not get updated in 13 hours. You would see old space usage in the UI, and VMware Cloud DR charges based on space usage information.

For example, if large amount of data gets written to the cloud file system, space reclamation might run longer and you won't be charged for this new data until the space reclamation job is completed.

## Protection Groups

The Protection groups tab for a cloud file system shows all protection groups that are replicating VM snapshots to the selected cloud file system. The list displays protection groups by name and shows their status, schedule, frequency, and other useful information.

cloudEastBackup ☰

Summary Protection groups Performance

Protection groups CREATE PROTECTION GROUP

Protection group	Site	Status	Schedul	Frequency	Quiesc	Size ⓘ	Last snapshot	VMs
tinyPG	onPrem2	✓ Good	Active	Standard	No	377.2 MiB	Jan-06 12:00 am (12...	2 VMs <span>☰</span>
winPG	onPrem2	✓ Good	Active	Standard	No	28.4 GiB	Jan-06 12:10 am (12...	1 VM <span>☰</span>

## Performance

You can view the following performance information about a cloud file system and any VMs running live on it:

Live VM performance	Description
<b>IOPS</b>	Input/Output Operations Per Second (IOPS) measures the number of read and write commands a storage device or medium completes every second for all workloads on the cloud file system datastore.  IOPS also indicates the average KiB size for each read or write operation, measured in Kibibytes (KiB).
<b>Throughput</b>	The data transfer rate to and from the cloud file system datastore, measured in kibibytes (KiB) or Mebibytes (Mibs). Measures the total amount of data that moves along the network path.
<b>Latency</b>	The time it takes for an IOPS request to complete, measured in milliseconds (ms).

Live VM performance	Description
<b>Cache hit rate</b>	<p>A measurement of how many content requests the cloud file system NVMe cache can fulfill successfully, compared to how many requests it receives.</p> <p>A cache hit ratio of 90% and higher means that most of the requests are satisfied by the cache. A value under 80% indicates a working set that is substantially larger than supported by the cache.</p>
<b>Replication throughput</b>	<p>The current data transfer rate of all currently running snapshot tasks replicating to the cloud file system datastore, measured in kibibytes (KiB). Measures the total amount of data that moves along the network path.</p>

# Set Up Protected Sites

## 6

A VMware Cloud DR protected site encompasses vCenter Servers, protection groups, and recovery plans.

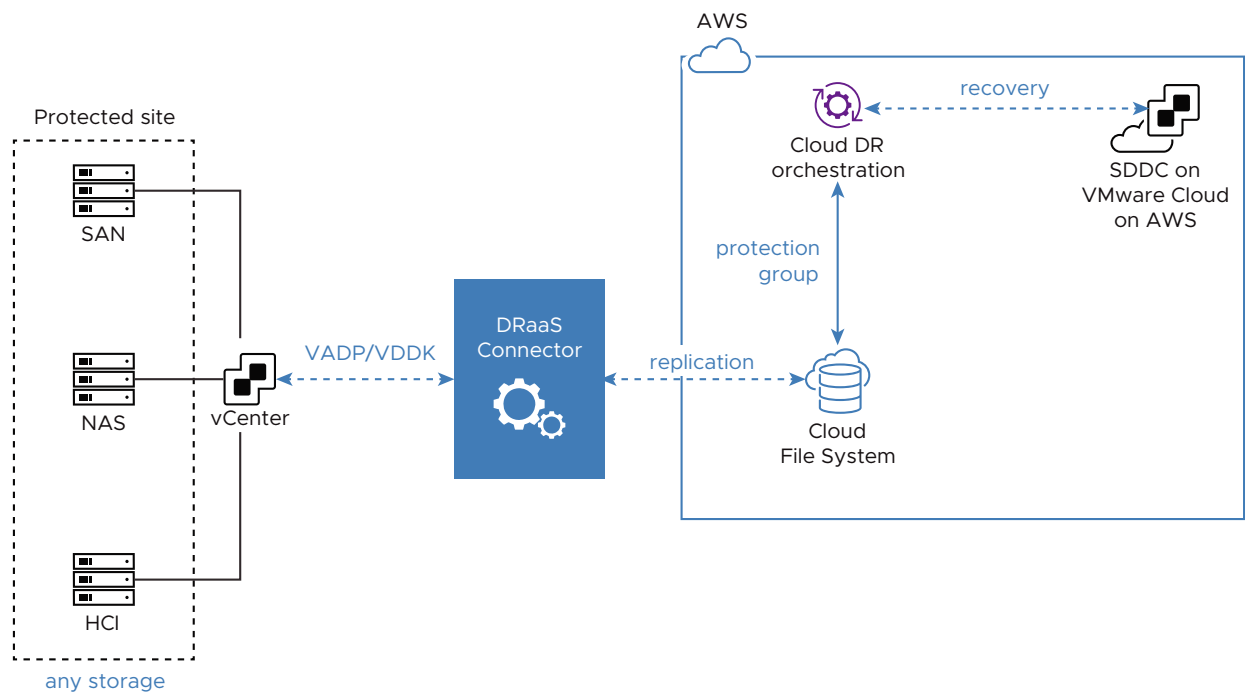
To set up a VMware Cloud DR (sometimes called VCDR) protected site, you create the site and then [Deploy the DRaaS Connector Using vCenter Server UI](#) as a virtual machine into your vSphere environment, either on-premises or a VMware Cloud on AWS SDDC.

After you set up a protected site, you create [Create a Protection Group](#) to replicate snapshots to a cloud file system. You can then use available snapshots from the cloud file system to recover protected VMs into your recovery SDDC using [Chapter 11 Set Up Recovery Plans in VMware Cloud DR](#). Once the protected site is available, you can initiate failback.

---

**Note** When protecting an SDDC using VMware Cloud DR, the recovery SDDC and VMware Cloud DR deployment must be in the same CSP organization as the protected SDDC.

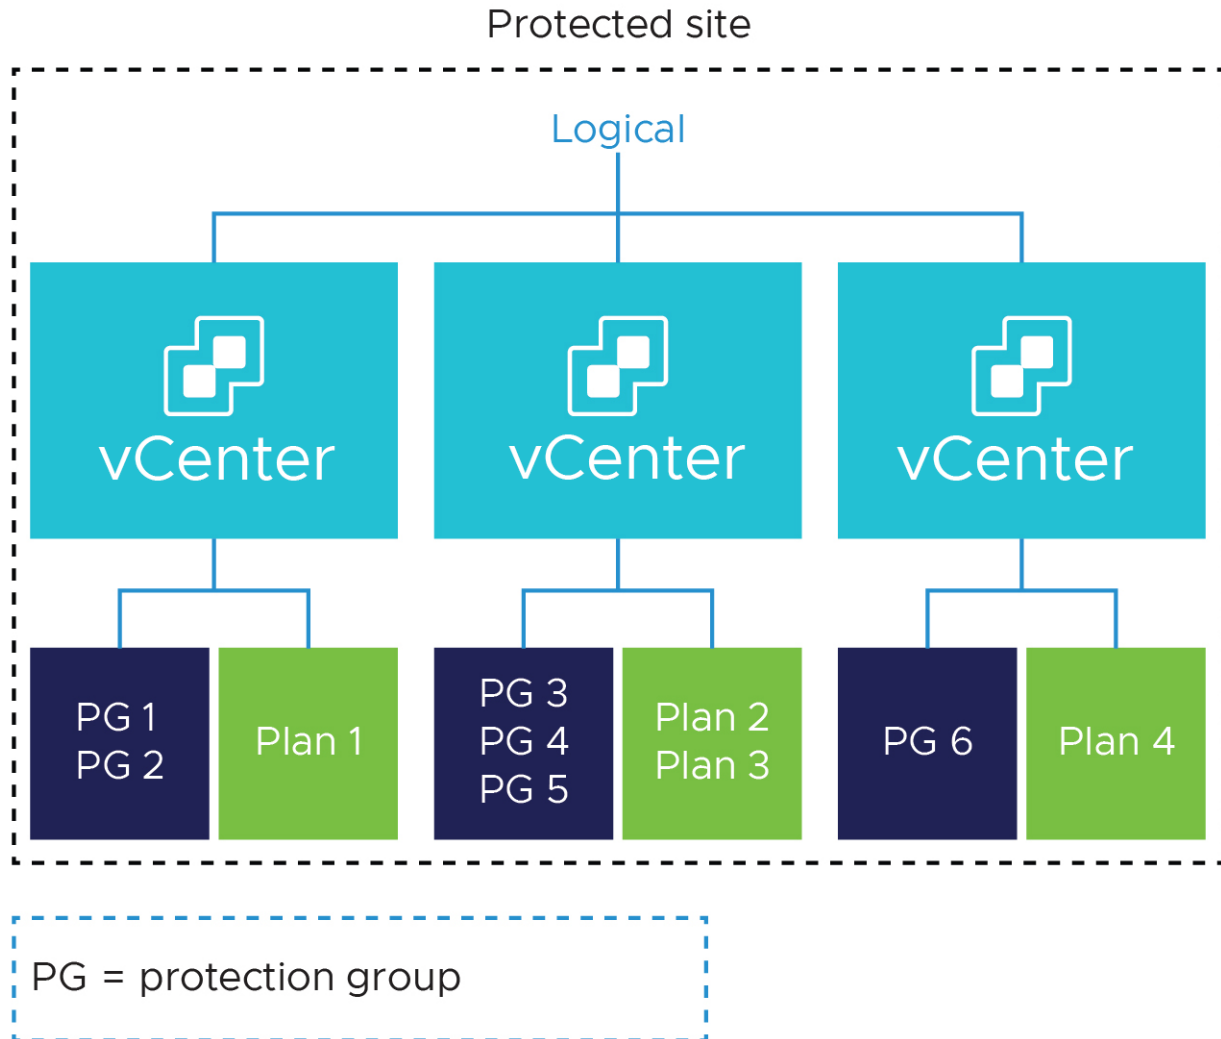
---



## Dimensions of a Protected Site

A VMware Cloud DR protected site encompasses vCenter Servers, protection groups, and recovery plans.

A protected site (on-premises vSphere or an SDDC) includes vCenter Servers which contain the VMs you want to protect. A vCenter Server can only be registered to one protected site, but one protected site can protect multiple vCenter Server. Each vCenter Server can have multiple protection groups and recovery plans.



**Note** For more information about VMware Cloud DR configuration limits, visit the [VMware Configuration Maximums](#) tool.



## Protected Site Setup Considerations

Consider the following suggestions when deploying the DRaaS Connector on your protected site.

---

**Note** These suggestions are not operational scale limits.

---

- Deploy one DRaaS Connector for every 250 VMs total in the protected site's vCenter Server inventory, counting all VMs in vCenter Server, protected or not. If you have 1000 VMs, you do not have to deploy more than four DRaaS Connectors (although there is no harm in deploying additional DRaaS Connectors). You can add connectors as needed. You need not commit to a particular number of connectors up front.
- Deploy only one DRaaS Connector on a single host.
- Deploy at least two connectors per-protected site, for redundancy.
- Sites with more 10,000 VMs might exhibit some responsiveness issues with the VMware Cloud DR UI, such as slow loading of pages or windows when previewing protection group VM membership, creating and editing recovery plans, and during plan compliance checking.
- VMware Cloud DR supports protecting up to 6000 VMs on a site with a single vCenter Server. To protect up to 6000 VMs in a single vCenter Server, you need four separate protected sites, each with its own cloud file system (four cloud file systems).

## Protected SDDC Network Considerations

Before you set up a protected site for an SDDC, you must create the SDDC and have a network segment already configured for it.

Follow these guidelines when configuring a network segment for the DRaaS Connector on the protected site:

- If you are using DHCP for the DRaaS Connector VM, when configuring DHCP from the VMC Console, leave the DNS value empty. Leaving this value empty allows the network to use the default DNS server for the SDDC.
- If you are using a static IP address for the DRaaS Connector VM, log in to the VMC Console and on the **Networking & Security** tab for your SDDC, you can use the DNS service Compute Gateway IP address for the connector VM.
- When setting up the protected site, decide whether you want VMware Cloud DR to create the DRaaS Connector firewall rules, or if you want to create the firewall rules yourself (manually). For more information, see [DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC](#).

## AWS Direct Connect

You have the option of using AWS Direct Connect for connecting your protected site to VMware Cloud DR.

AWS Direct Connect provides a dedicated network connection between your on premises data center and AWS services. With this connection, you can create public virtual interfaces (VIFs) that give you direct access to all public AWS IP addresses, including VMware Cloud DR components.

For more information, see [AWS Direct Connect](#).

Read the following topics next:

- [AWS Direct Connect](#)
- [Set Up Protected SDDC with Recovery in Different Region \(Inter-Region DR\)](#)
- [Set Up Protected SDDC with Recovery in Same Region \(Intra-Region DR\)](#)
- [Create Protected Site for On-premises vSphere](#)
- [Edit or Delete a Protected Site](#)
- [Remove a vCenter Server from a Protected Site](#)
- [Remove a DRaaS Connector from a Protected Site](#)
- [Send Support Bundle](#)

## AWS Direct Connect

VMware Cloud DR supports using Amazon Web Services (AWS) Direct Connect(DX) public or private virtual interfaces (VIFs) for on-premises protected site networks.

AWS Direct Connect provides a dedicated network connection between your on-premises data center and AWS services underlying VMware Cloud DR (also sometimes called VCDR). You can order this connection using your customer-managed AWS account.

AWS Direct Connect allows you to connect protected sites to VMware Cloud DR over the internet. You can target recovery plan failovers to any AWS regions that support Direct Connect.

Direct Connect offers faster speeds and lower latency than with a connection over the public internet, which can increase replication speed to the cloud file system, management traffic, failbacks, and any VMware Cloud DR operation that requires internet connectivity.

Whether you choose to connect through Direct Connect or over the internet, you must configure on-premises firewall rules to open the required ports required to access VMware Cloud DR public IP addresses. For more information, see [System and Network Requirements for the DRaaS Connector](#) and [Service Public IP Addresses](#).

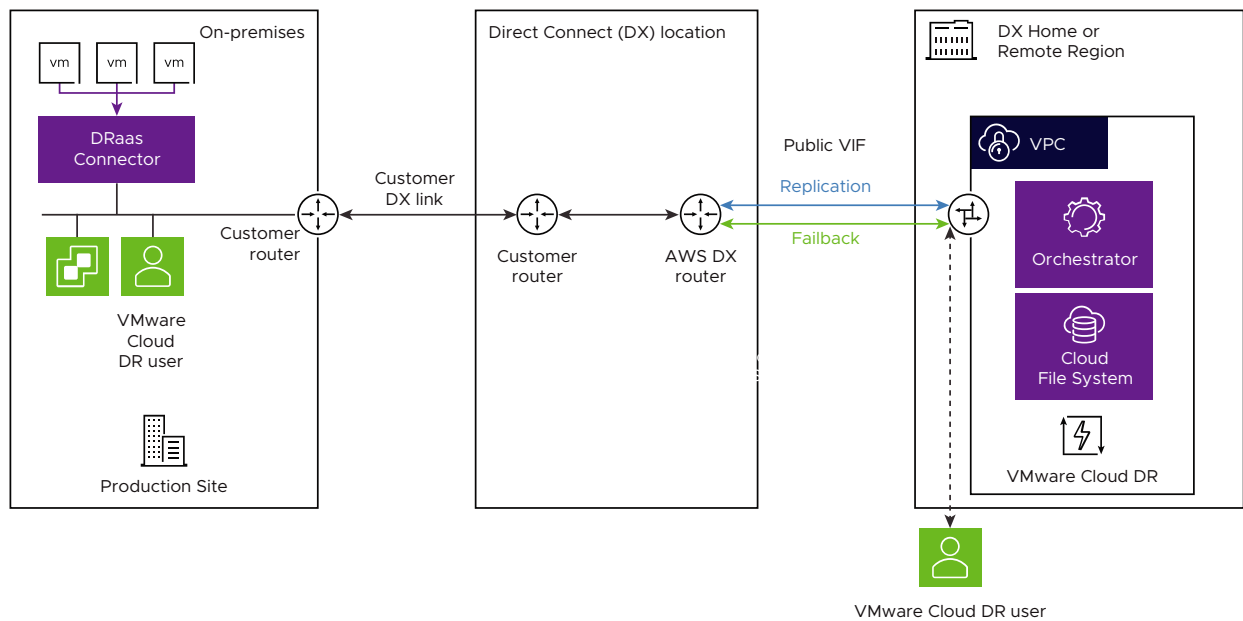
VMware Cloud DR allows multiple VIFs that you can configure to use the same or different Direct Connect connections for redundancy. You can also configure Direct Connect for VIFs that terminate in the same or different AWS region where VMware Cloud DR is deployed.

For more information, see [AWS Direct Connect Resiliency Recommendations](#) and [AWS Direct Connect Quotas](#).

## AWS Direct Connect Public VIF

With Direct Connect, you can create a public VIF that gives you direct access to all public AWS IP addresses, including VMware Cloud DR components.

For information about Direct Connect and setting up a public VIF, see [AWS Direct Connect](#) documentation.

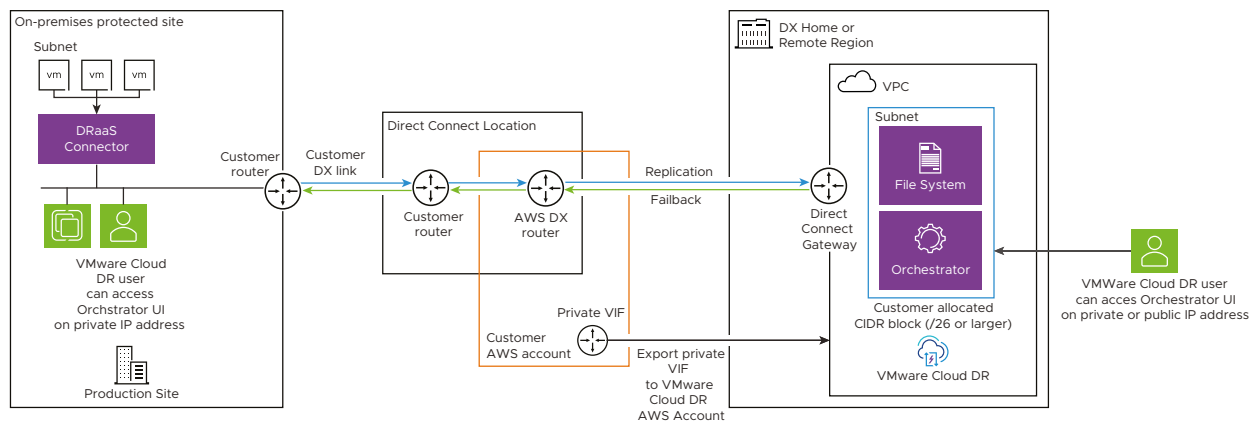


## AWS Direct Connect Private VIF

When your Direct Connect connection has been provisioned, create private virtual interfaces to connect private IP addresses to the VMware Cloud DR Virtual Private Cloud (VPC). From your AWS account, you can create a hosted private VIF using the account shown in the AWS Account ID text box of the Direct Connect page of the **Networking & Security** tab.

For information about Direct Connect and setting up a private VIF, see [AWS Direct Connect](#) documentation.

See [Configure Direct Connect \(Private VIF\)](#) for instructions on configuring Direct Connect private VIFs with VMware Cloud DR.



## Configure Direct Connect (Private VIF)

You can configure AWS Direct Connect for private connections between VMware Cloud DR and your protected sites.

### Prerequisites

Before configuring Direct Connect with VMware Cloud DR, do the following:

- Select a /26 CIDR block within your company's private IP network scheme. This CIDR block must not overlap with other allocated CIDR blocks within your routed on-premises and cloud networking sites. The VMware Cloud DR internal networking uses 172.30.0.0/26, which cannot be used. The allocated CIDR block is part of VMware Cloud DR's Transit VPCs, which host the xENIs for the Orchestrator and cloud file system that are exported over Direct Connect.
- After the VIF is attached, the original /26 CIDR block is split into two equal /27 CIDR blocks that are advertised by the interface. In some cases, you might need to make to your protected site's networking configuration to match advertised prefixes.
- Select an appropriate and valid autonomous system number (ASN) number. VMware Cloud DR uses ASN number 64512, and cannot be used for your side of the Border Gateway Protocol (BGP) connection.
- Obtain your VMware Cloud DR AWS shadow account ID. You can find this account ID and number by navigating to **Settings > Direct Connect**. Your network administrator needs this account number to export your private VIFs to VMware Cloud DR.
- Have your network administrator create a private VIF that uses the allocated CIDR block, VMware Cloud DR shadow account ID, and the allocated BGP ASN number.
- Export your private VIFs to the VMware Cloud DR shadow account ID. Your network administrator performs this task from your AWS account.

Using a private VIF with a VMware Cloud DR protected site is restricted by the following caveats:

- Only a single CIDR block is supported and is shared among all VIFs.
- Do not use the 172.30.0.0/26 CIDR block because it overlaps with CIDR blocks 172.30.16.0/24 and 172.16.0.0/16, which are reserved for use by VMware Cloud DR.
- Multiple protected sites are supported for use with private VIFs, if all protected sites share the same CIDR block and are connecting to the end point of their Private VIF.
- Post-deployment CIDR block changes are not supported.
- You cannot convert a non-private VIF protected site to use a private VIF by yourself. If you are interested in converting a VMware Cloud DR protected site from using native internet or public VIF to private VIF, contact VMware support for assistance.
- VMware Cloud DR cannot determine if a private VIF is being used for a specific protected site.

## Procedure

### 1 Navigate to **Settings > Direct Connect**.

The Direct Connect dialog box shows the VMware Cloud DR AWS shadow account ID. Your IT administrator needs this information to create and export private VIFs to VMware Cloud DR. You cannot configure Direct Connect if you have not exported private VIFs to VMware Cloud DR.

### 2 In the Direct Connect dialog box, click the **Set CIDR block** button.

Snapshot replication routes through a private IP network using IP addresses in the AWS transit VPC CIDR. Do not use the 172.30.0.0/26 CIDR block because it overlaps with CIDR blocks 172.30.16.0/24 and 172.16.0.0/16, which are reserved for use by VMware Cloud DR.

---

**Note** Once you set the CIDR, you cannot change it.

---

### 3 In the Set transit VPC CIDR block dialog box, enter the CIDR block to use with Direct Connect. Select an IP address range that does not conflict with any on-premises network on the protected site that uses Direct Connect.

### 4 Select the check box to confirm that once you set the transit VPC CIDR, it cannot be changed.

### 5 Click **OK**.

### 6 If the connection is successful, the Direct Connect dialog box shows all private VIFs exported to your account.

For each VIF, the dialog box shows the interface name and ID, Direct Connect ID, state (available, unavailable, attaching, or attached), and BGP status (up, down, or unknown).

### 7 To enable a VIF, select the small menu to the right of the VIF row and select **Attach**.

### 8 In the Attach virtual interface confirmation dialog box, select the check box to confirm, and then click **OK**.

## What to do next

After you have established a Direct Connect connection, you can select this connection type when you [Chapter 6 Set Up Protected Sites](#).

---

**Note** After the VIF is attached, the original /26 CIDR block is split into two equal /27 CIDR blocks that are advertised by the interface. In some cases, you might need to make to your protected site's networking configuration to match advertised prefixes.

---

## Unset Direct Connect

If you want to change or edit the CIDR block being used for Direct Connect, you can unset the connection.

After you unset Direct Connect, you must [Configure Direct Connect \(Private VIF\)](#) again to start using it.



Unsetting Direct Connect does not delete a private VIFs that were using it, but they become unusable. To delete the VIFs, see [Delete a Private VIF](#).

### Prerequisites

Unsetting Direct Connect makes the existing connection unusable, which you must reconfigure to start using it again for a protected site.

Before you unset an active Direct Connect connection, you must configure the DRaaS Connectors on the protected site to use the public internet. For more information, see [Service Public IP Addresses](#).

### Procedure

- 1 Navigate to **Settings > Direct Connect**.
- 2 In the Direct Connect dialog box, click **Unset Direct Connect**.

**Direct Connect**

AWS Direct Connect links your internal network directly to an AWS Direct Connect router over a standard network cable. With this connection, you can create virtual interfaces to VMware Cloud DR, bypassing the public Internet, with higher speeds. [Learn more about Direct Connect](#).

AWS shadow account ID: 067055051921

AWS BGP ASN: 64512

On-premises CIDR block for Direct Connect: 10.98.86.128/26

**UNSET DIRECT CONNECT**

**Direct Connect virtual interfaces**

Interface name	Interface ID	Direct Connect ID	State	BGP status
SEA03-DX-N9K02-C...	dxvif-ffgbvum7	dxcon-fgr6o9rt	Available	Down (1)
SEA03-DX-N9K01-C...	dxvif-fg8n5fbz	dxcon-fh5fv2rw	Attached	Up (1)

**CLOSE**

- 3 In the Unset Direct Connect dialog box, enter UNSET DIRECT CONNECT in the field, and then click **Unset Direct Connect**.

## Delete a Private VIF

When you no longer need a private VIF, you can delete it.

**Note** Once you delete a VIF, it is no longer usable and is removed from your AWS account.

### Procedure

- 1 Navigate to **Settings > Direct Connect**.

- 2 In the Direct Connect dialog box, select the small menu to the right of the VIF entry and then click **Delete**.
- 3 In the Delete virtual interface confirmation dialog box, select the check box to confirm, and then click **OK**.
- 4 Click **OK**.

## Switch Between Public Internet and Direct Connect for Protected Sites

If you have a protected site, you might want to switch connectivity from using public Internet to private VIF after configuring Direct Connect

You can switch from public internet to Direct Connect (if configured), and from Direct Connect to public internet. After you switch your protected site connection, check the DRaaS Connector configuration using the DRaaS Connector CLI.

### Procedure

- 1 From the left navigation, select Protected sites, and then select a protected site.
- 2 From the upper-right, click the small menu widget and select **Edit site**.
- 3 In the Edit site dialog box, under Connection to cloud select Use public internet.
- 4 Click **OK**.
- 5 After you switch your protected site connectivity, you need to re-[Chapter 7 Deploy the DRaaS Connector](#) on the protected site.
- 6 After you deploy the new connector, delete the old connector. Then, run the DRaaS Connector CLI `drc show` command.

The output shows a configuration property named `useDirectConnect`.

- When Direct Connect is being used, the `useDirectConnect` property equals `true`.
- When the public internet is being used, the `useDirectConnect` property equals `false`.

For more information about the DRaaS Connector CLI, see [DRaaS Connector Connectivity Check](#).

## Set Up Protected SDDC with Recovery in Different Region (Inter-Region DR)

You can set up a protected SDDC to use a different AWS region for recovery than the region where you replicate snapshots to a cloud file system and perform recovery operations.

### Prerequisites

You can set up a protected SDDC in one AWS region, and perform backup and recovery operations in a different AWS region.

**Procedure**

- 1 From the left navigation, select **Protected sites**.
- 2 Click the **Set up protected site** button.
- 3 In the **Setup protected site** dialog box, under Site types select VMware Cloud on AWS.
- 4 Under Cloud backup, select a cloud file system to use for replicating snapshots from the protected SDDC. If there is already one cloud file system deployed, then it is selected.
- 5 On the next page under Protected SDDC, select an existing SDDC to protect.  
  
The UI indicates the AWS region and availability zone for each SDDC. To perform backup and recovery operations across AWS regions, select an SDDC that is in a different AWS region than the cloud file system.
- 6 On the next page, you have a choice to either allow the system to create firewall rules for the DRaaS Connector (recommended). Or, you can manually create those firewall rules from the VMware Cloud DR UI.  
  
If you are not sure which to select, see [DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC](#) for more information.
- 7 Click **Setup**. When the site is created, it displays under Sites as a protected site.

**What to do next**

Now that you have set up the protected site for your SDDC, you need to [Chapter 7 Deploy the DRaaS Connector](#) on the SDDC.

## Set Up Protected SDDC with Recovery in Same Region (Intra-Region DR)

You can set up a protected VMware Cloud on AWS SDDC to use the same AWS region where you replicate snapshots to a cloud file system and perform recovery operations.

**Prerequisites**

You can set up a protected SDDC in the same AWS region where you plan to deploy a recovery SDDC. Select this configuration If you are constrained by data residency laws, or have access to only one AWS region, and must use the same region for both protection (protected SDDC) and backup and recovery.

If you are using only one AWS region for VMware Cloud DR, when you create a protected SDDC you can select an SDDC that resides in the same region, but different availability zone, than the cloud file system.

**Procedure**

- 1 From the left navigation, select **Protected sites**.
- 2 Click the **Set up protected site** button.

- 3 In the **Setup protected site** dialog box, under Site types select VMware Cloud on AWS.
- 4 Under Cloud backup, select a cloud file system to use with the new protected site. If there is only one cloud file system deployed, then it is already be selected.
- 5 On the next page under Protected SDDC, select an existing SDDC to protect. The UI indicates both the AWS region and availability zone used for each SDDC.

If you want to perform backup and recovery operations in the same AWS region, you can select an SDDC that is in the same AWS region but in a different availability zone as the cloud file system.

- 6 On the next page, you have a choice to either allow the system to create firewall rules for the DRaaS Connector (recommended).

Or, you can manually create those firewall rules from the VMware Cloud DR UI. If you are not sure which to select, see [DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC](#) for more information.

- 7 Click **Setup**. When the site is created, it displays as a protected site.

#### What to do next

Now that you have set up the protected site for your SDDC, you need to [Chapter 7 Deploy the DRaaS Connector](#) on the SDDC.

## Create Protected Site for On-premises vSphere

You can create a protected site for your on-premises vSphere environment.

#### Prerequisites

#### Procedure

- 1 From the left navigation, select **Protected sites**.
- 2 Click the **Set up protected site** button.

- 3 In the **Set up protected site** dialog box, under Site type select **On-premises site**.

Set up protected site

Site type

☒ On-premises site

Protect VMs in customer-managed vSphere environments.

☐ VMware Cloud on AWS

Protect VMs in a VMware Cloud SDDC.

Cloud backup

RWR\_CDVX

Connection to cloud

Select how the connectors will reach out to the cloud file system to replicate and restore snapshots.

☒ Use public internet

Including Direct Connect with public VIF.

☐ Use Direct Connect with private VIF [not configured yet]

Time zone

Select the time zone for protection group schedules.

Filter time zones by country

United States of America ▼

Time zone

Los Angeles

PDT GMT-07 ▼

On-premises site name

Oregon DC

CANCEL

SET UP

- 4 Under Cloud backup, select a cloud file system to use for backups from the protected SDDC. If a cloud file system is already deployed, then it is selected.

- 5 Under Connection to cloud, select either **Use public internet** or **Use Direct Connect with private VIF**.

To use Direct Connect, you must already have configured VMware Cloud DR to use [Configure Direct Connect \(Private VIF\)](#).

- 6 Under Time zone, select a time zone from the drop-down menu, and then click the button to the right to set the time zone for the protected site.
- 7 Under On-premises site name, enter a name for the on-premises site.
- 8 Click **Set up**.

#### What to do next

Next, you can [Chapter 7 Deploy the DRaaS Connector](#) on the protected site.

## Edit or Delete a Protected Site

You can edit a protected site to change the site time zone or name, or you can delete a protected site.

---

**Note** If you delete a protected site, then the VMs on the site are no longer be protected by the DRaaS Connector.

---



---

**Note** Before you delete a protected site, you must first remove its vCenter and all deployed DRaaS Connectors from the site.

---

#### Procedure

- 1 From the left navigation, click **Protected sites** and select a site.
- 2 From the upper-right of the protected sites page, click the **Edit site** button.
- 3 From the drop-down menu, click **Edit site** (or **Delete site**, if you are deleting the site).
- 4 If you are editing the site, you can change the protected site name label, or change the time zone for the site.
- 5 If you are deleting the site, click **Delete site** from the menu.
- 6 Enter the phrase **DELETE SITE** in all upper case letters then click **OK**.

## Remove a vCenter Server from a Protected Site

If you no longer need a vCenter Server on your protected site, you can remove the vCenter Server.

When you remove a vCenter Server from a protected site, the vCenter Server still functions normally inside of vSphere. After you remove the vCenter Server, the VMs on that vCenter are no longer protected by the DRaaS Connector.

#### Procedure

- 1 From the left navigation, select **Protected sites** and then select a site.
- 2 On the protected site page, from the vCenters section click the drop-down menu icon next to the vCenter you want to remove, and then click **Remove vCenter**.
- 3 In the **Remove vCenter confirmation** dialog box, enter the phrase **REMOVE VCENTER** in all upper case letters, and then click **OK**.

## Remove a DRaaS Connector from a Protected Site

If you no longer need a DRaaS Connector on your protected site, you can remove it.

#### Procedure

- 1 From the left navigation, select **Protected sites** and then select a protected site.
- 2 Under the Connectors section and next to the connector you want to remove, select **Remove connector** from the drop-down menu.
- 3 In the **Remove connector confirmation** dialog box, enter the phrase **REMOVE CONNECTOR** in all upper case letters, and then click **OK**.

## Send Support Bundle

To help troubleshoot and diagnose problems, you can send VMware support information about a protected site and the DRaaS Connector, called a support bundle.

When you submit a support bundle, VMware Cloud DR collects support data, creates a support bundle, and sends it to the VMware Support team. Typically, you submit a support bundle only after consulting with VMware support.

#### Procedure

- 1 From the left navigation, select **Protected sites** and then select a protected site.
- 2 Select a protected site.
- 3 From the upper-right of the window, click the drop-down menu and select **Send support bundle**.
- 4 In the dialog box, click **Submit support bundle**.

# Deploy the DRaaS Connector

# 7

After you set up a protected site, you can deploy the DRaaS Connector, a stateless virtual appliance that enables replicating VM snapshot deltas from a protected site to the cloud file system, which are used for recovery operations.

The DRaaS Connector is distributed as virtual appliance in the Open Virtualization Format (OVF) standard. You can deploy the DRaaS Connector in vCenter Server to enable data protection of VMs within on-premises vSphere or VMware Cloud on AWS SDDCs.

The DRaaS Connector can be redeployed if needed at any time without losing backup data. Software upgrades for it are over-the-air and automatic across time. Each connector provides additional replication bandwidth for the site.

Before you deploy the connector, see [System and Network Requirements for the DRaaS Connector](#).

For VMware Cloud on AWS SDDC protected sites, you must deploy the DRaaS Connector on each cluster in the SDDC in order to protect the VMs running on those clusters. If you do not deploy the DRaaS Connector on an SDDC cluster, then you cannot add the VMs from those clusters to protection groups.

Deploying the DRaaS Connector consists of the following tasks:

- [Download the DRaaS Connector OVA from VMware Cloud DR DR UI](#).
- [Deploy the DRaaS Connector Using vCenter Server UI](#).
- [Configure the DRaaS Connector VM from VM Console](#).
- [Register vCenter Server](#).

Read the following topics next:

- [System and Network Requirements for the DRaaS Connector](#)
- [DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC](#)
- [Service Public IP Addresses](#)
- [\(Optional\) Install SSL Certificate Before DRaaS Connector Deployment](#)
- [Download the DRaaS Connector OVA from VMware Cloud DR DR UI](#)
- [Deploy the DRaaS Connector Using vCenter Server UI](#)



- [Configure the DRaaS Connector VM from VM Console](#)
- [Register vCenter Server](#)
- [Refresh vCenter Server User Credentials](#)
- [Power the DRaaS Connector On and Off](#)
- [DRaaS Connector Connectivity Check](#)
- [DRaaS Connector Performance Check](#)

## System and Network Requirements for the DRaaS Connector

The DRaaS Connector has several system and network requirements you need to be aware of before you deploy one.

### vSphere and VMware Cloud on AWS Compatibility

For the most current information on vSphere and VMware Cloud on AWS compatibility with the service, see the [VMware Product Interoperability Matrix](#).

For configuration limits of the service, see [VMware Cloud DR Configuration Maximums](#).

### DRaaS Connector System Requirements

To deploy the DRaaS Connector virtual appliance, make sure that the vSphere host where you deploy it has the following available resources.

Site Resources	Value
VMware vCenter Server	See the <a href="#">VMware Product Interoperability Matrix</a> for the latest supported versions.
CPU	8 GHz (reserved)
RAM	12 GiB (reserved)
Disk	100 GiB virtual disk
Network connectivity	Required

---

**Note** VMware Cloud DR (sometimes called VCDR) does not support an internet proxy server between the DRaaS Connector and the cloud.

---

---

**Note** If the DRaaS Connector becomes unregistered, delete the virtual appliance and re-deploy it using the OVF file. Do not use the vCenter Server datastore browser to register the connector.

---

## DRaaS Connector Deployment Considerations

Consider the following suggestions when deploying the DRaaS Connector on your protected site.

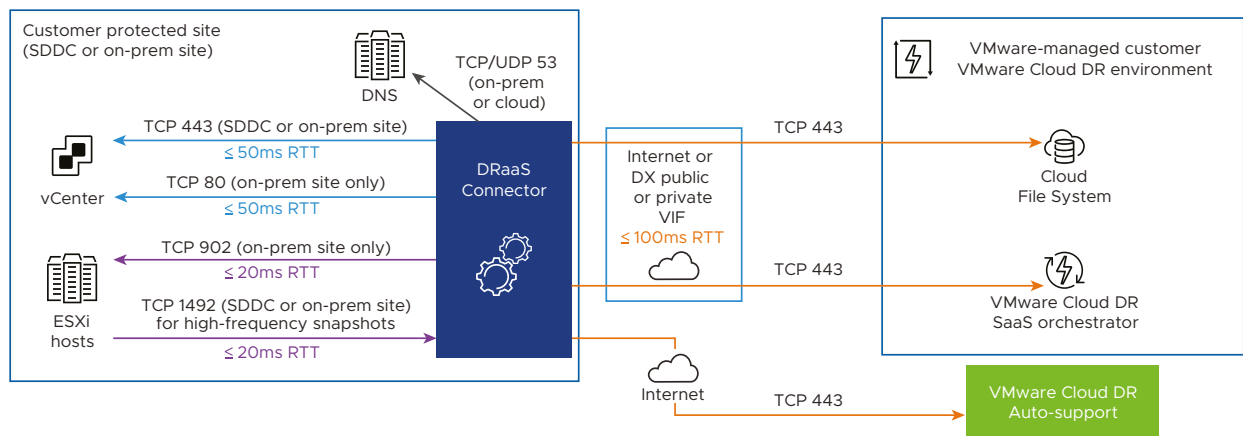
**Note** These suggestions are not operational scale limits.

- Deploy one DRaaS Connector for every 250 VMs total in the protected site's vCenter Server inventory, counting all VMs in vCenter Server, protected or not. If you have 1000 VMs, you do not have to deploy more than four DRaaS Connectors (although there is no harm in deploying additional DRaaS Connectors). You can add connectors as needed. You need not commit to a particular number of connectors up front.
- Deploy only one DRaaS Connector on a single host.
- Deploy at least two connectors per-protected site, for redundancy.
- Sites with more 10,000 VMs might exhibit some responsiveness issues with the VMware Cloud DR UI, such as slow loading of pages or windows when previewing protection group VM membership, creating and editing recovery plans, and during plan compliance checking.
- VMware Cloud DR supports protecting up to 6000 VMs on a site with a single vCenter Server. To protect up to 6000 VMs in a single vCenter Server, you need four separate protected sites, each with its own cloud file system (four cloud file systems).

## DRaaS Connector Networking Requirements

The DRaaS Connector requires the following open outbound ports listed in the diagram in your network to allow connector traffic. See [Service Public IP Addresses](#) for how to find the public IP addresses of the Orchestrator and the cloud file system.

**Note** The DRaaS Connector uses port 443 to connect to the Orchestrator. Previously, if the DRaaS connector was unable to connect to the Orchestrator using port 443, it reverted to port 1759 for the connection. Now, port 1759 can no longer be used for the DRaaS Connector. Ensure that your protected site firewall allows traffic on port 443 for the DRaaS Connector.



**Table 7-1. Open Ports Required for the DRaaS Connector**

Protocol	Port	Source	Destination	Service Description	Classification
<b>Protected site</b>					
TCP	443	DRaaS Connector	vCenter Server (on-premises site or SDDC)	vCenter Server web service	Outbound
TCP	80	DRaaS Connector	vCenter Server (on-premises site only)	vCenter Server web service	Outbound
TCP	902	DRaaS Connector	ESXi Management IP address (on-premises site only)	Reading/writing vdisks	Outbound
TCP	1492	ESXi hosts	DRaaS Connector (For on-premises sites. For VMware Cloud on AWS SDDCs, this outbound rule for ESXi hosts is already configured.)	For high-frequency snapshots, reading/writing vdisks.	Inbound
<b>VMware Cloud DR</b>					
TCP	443	DRaaS Connector	Cloud file system	Encrypted tunnel for data transfers and metadata operations	Outbound
TCP	443	DRaaS Connector	Orchestrator	Management service	Outbound
TCP	443	DRaaS Connector	VMware auto-support server	Support service	Outbound

## DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC

Setting up a protected site for your VMware Cloud on AWS SDDC requires creating firewall rules for the DRaaS Connector.

As you set up your protected sites, you must decide if you want VMware Cloud DR to create the firewall rules needed for the DRaaS Connector (recommended). Or, if you want to create the firewall rules manually.

If you allow VMware Cloud DR to automatically create firewall rules for your protected site, you must create a dedicated network segment to use for the DRaaS Connector on the SDDC. This is recommended as a best practice.

If you wish to [Set Up Protected SDDC with Recovery in Different Region \(Inter-Region DR\)](#) to allow the DRaaS Connector to communicate with your protected SDDC, follow these guidelines:

- *Protected SDDC vCenter Server* outbound on TCP port 443
- *Cloud file system* outbound on TCP port 443
- *Orchestrator* outbound on TCP ports 443
- *VMware Cloud DR auto-support server* outbound on TCP port 443

---

**Note** See [Service Public IP Addresses](#) for how to find VMware Cloud DR public IP addresses.

---

**Note** VMware Cloud DR does not support an internet proxy server between the DRaaS Connector and the cloud.

---

You can open these ports by configuring firewall rules for the SDDC's Compute Gateway as described here: [Add or Modify Compute Gateway Firewall Rules](#).

## Service Public IP Addresses

You can find the public IP addresses of the Orchestrator and the cloud file system in the Deploy connector appliance dialog box.

For customer-managed protected sites, you create firewall rules in your network to allow the DRaaS Connector to communicate with VMware Cloud DR service in AWS (over the internet or through AWS Direct Connect). The Orchestrator and the cloud file system public IP addresses are located in the Deploy connector appliance dialog box for a protected VMware Cloud on AWS SDDC.

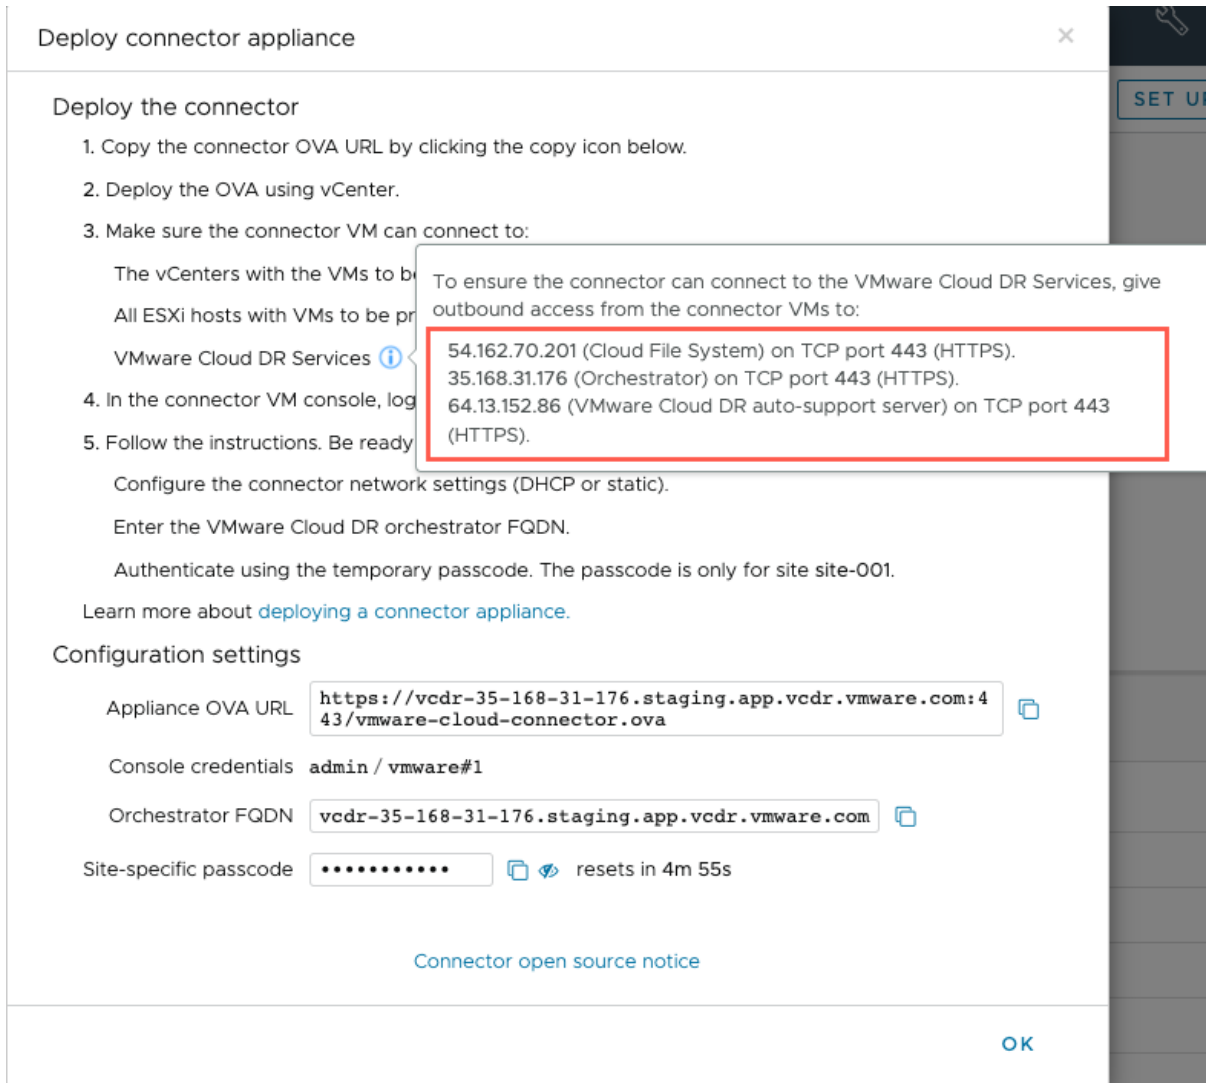
### Procedure

- 1 [Set Up Protected SDDC with Recovery in Different Region \(Inter-Region DR\)](#).
- 2 On the new protected VMware Cloud on AWS protected site, look under Connectors and click the **Deploy** button.
- 3 In the Deploy connector appliance dialog box, on step number 3 click the small information icon and view the tooltip information, which provides the public IP address for the cloud file system and the Orchestrator.

---

**Note** You can also find the Orchestrator IP address (or FQDN) and DRaaS Connector OVA URL by navigating to **Settings > About VMware Cloud DR > .**

---



## (Optional) Install SSL Certificate Before DRaaS Connector Deployment

If you see a warning message about SSL certificates while deploying the DRaaS Connector, you might have to install a new certificate in vCenter Server before proceeding.

When you deploy the DRaaS Connector OVA in vCenter Server 6.7 or newer, if you receive a message stating 'SSL certificate cannot be trusted', you have two options:

- Proceed and click **Yes** to accept the certificate. Or,
- Install in the necessary Certificate Authority root certificate in vCenter Server to enable verification (see instructions below).

**Note** If you do not see this message, you do not have to perform this task.

**Procedure**

- 1 In a browser, go to <https://www.digicert.com/kb/digicert-root-certificates.htm>
- 2 Download the DigiCert TLS RSA SHA256 2020 CA1 global root certificate.  
You can download the certificate directly by going to <https://cacerts.digicert.com/DigiCertTLSRSASHA2562020CA1-1.crt.pem>.
- 3 In your vCenter Server, click **Menu > Administration > Certificate Management**
- 4 Click to add a new Trusted Root Certificate, and use the downloaded certificate.

## Download the DRaaS Connector OVA from VMware Cloud DR DR UI

Setting up a protected site requires that you download the DRaaS Connector connector OVA from the VMware Cloud DR UI.

Using the VMware Cloud DR UI, you can copy the URL to download the OVA into your environment in one of two ways:

- Navigate to **Settings > About VMware Cloud DR**. You can copy the VMware Cloud DR OVA URL here.
- You can get the same connector OVA URL from the **Download connector** dialog box, which is described below.

**Procedure**

- 1 In the VMware Cloud DR UI, click **Sites > Protected sites** and then click the protected site on the left side of the application.
- 2 Under Connectors, click **Deploy**. If this protected site is an SDDC, the **Deploy** button is under Clusters.
- 3 In the **Download connector** dialog box, there is a list of steps that guide you in deploying the connector, as well as the URL to the connector OVA, with an option to download it locally to your system. If the site you are protecting is a VMware Cloud on AWS SDDC, for more information about networking considerations see [DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC](#).
- 4 Click the **Copy** button to copy the download URL. You need this URL when you deploy the OVA in vSphere, after you download the connector.
- 5 Make a note of the Console credentials, which you need to log in to the VM console: `admin/vmware#1`.
- 6 Also copy (or write down) the Orchestrator Fully Qualified Domain Name (FQDN), which you need when you configure the connector in the VM console.
- 7 Click **OK**.

## What to do next

Your next task in setting up a protected site is to [Deploy the DRaaS Connector Using vCenter Server UI](#).

# Deploy the DRaaS Connector Using vCenter Server UI

When you deploy the DRaaS Connector from the vCenter Server UI, you select the host, cluster, and resource pool for the connector.

Before you begin:

- Do not name the DRaaS Connector VM using the same naming conventions you use to name VMs in your vSphere environment. Avoid giving the connector VM a name that might match the VM name pattern you use when you define protection groups.
- If you are deploying the DRaaS Connector to a VMware Cloud on AWS SDDC with more than one cluster, you must choose a cluster to deploy the connector VM on. Each cluster in your SDDC needs the connector VM deployed on it in order for the VMs running there to be added to protection groups and replicated to a cloud backup site.

---

**Note** If you are deploying the DRaaS Connector for a Google Cloud VMware Engine protected site, see [Google Cloud documentation](#) for information about logging in to vSphere from Google Cloud.

---

## Procedure

- 1 In vSphere, select any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host.
- 2 Right-click the object and select **Actions** → **Deploy OVF Template**.
- 3 Click **Next**.
- 4 In the **Deploy OVF Template** dialog box, Step 1, Select an OVF template, paste the connector OVA URL into the URL field. The exact URL to download the connector OVA displays in the **Download Connector** dialog box. For example: `https://<vmware-cloud-dr-ip-address>/cloud-connector.ova`.
- 5 Click **Next**.
- 6 Next, specify a name for the connector. Do not use non-ASCII characters for the connector name. Use a name that is different than the naming conventions you use to name VMs in your vSphere environment, to avoid this VM being included in a snapshot.
- 7 Below the name, choose a location for the connector and then click **Next**.
- 8 Select a compute resource for the connector. If this vSphere is a VMware Cloud on AWS SDDC with more than one cluster, choose a cluster to deploy the connector VM on. For an SDDC, each cluster you want to protect must have the DRaaS Connector VM deployed on it.
- 9 Click **Next**.

- 10 Review the details for your connector deployment, then click **Next** to select storage for the connector VM.
- 11 Select a datastore for the connector and then click **Next**.
- 12 Select the network to use for the connector, and then click **Next** to review the deployment details.
- 13 Click **Finish**. You can now find the connector VM in your vSphere client.
- 14 Click on **Edit settings** and select Virtual Hardware. Enter the CPU and memory reservation as described in [System and Network Requirements for the DRaaS Connector](#).

### What to do next

Now you are ready to [Configure the DRaaS Connector VM from VM Console](#).

## Configure the DRaaS Connector VM from VM Console

Before you configure the DRaaS Connector VM using the VM console, make sure you have all required information by filling out the following worksheet (optional):

Parameters	Value
Console credentials	admin/vmware#1
	<p><b>Note</b> This password will change at the end of the configuration, and you can obtain the new password in the VMware Cloud DR UI.</p>
IP address allocation: Static or DHCP	<p>If you deployed the DRaaS Connector VM on an SDDC, we recommend that you use the default DNS server settings that is associated with your SDDC.</p> <p>If you choose (a) Static, then enter the following</p> <ul style="list-style-type: none"> <li>■ IP address: If you are configuring a static IP address for the connector VM, when asked for the DNS server use the default DNS server created when the SDDC was first deployed. You can find this setting under the SDDC <b>Network and Security &gt; DNS Service IP</b></li> <li>■ Subnet mask</li> <li>■ Gateway</li> <li>■ DNS servers</li> </ul> <p>If you are using a DHCP service, make sure you provide DNS servers that can resolve host names on the protected site internal network.</p> <p><b>Note</b> Public DNS servers are not supported, such as Google DNS servers (8.8.8.8 and 8.8.4.4), which do not work reliably with VMware Cloud DR.</p>
Orchestrator FQDN or IP address	<p>You can obtain the Orchestrator FQDN (or IP address) by navigating to <b>Settings &gt; About VMware Cloud DR</b>. Or, you can also obtain this information from the Deploy Connector appliance dialog box, as described <a href="#">Service Public IP Addresses</a>.</p> <p>If you have <a href="#">AWS Direct Connect</a> configured with a private VIF, then use the Orchestrator IP address</p> <p>.</p>



Parameters	Value
Temporary, site specific passcode	<p>This temporary passcode is used to configure the DRaaS Connector in the VM console.</p> <hr/> <p><b>Important</b> You can obtain a passcode in the Download connector VM dialog box from inside the VMware Cloud DR UI. Make sure that you open the dialog box from the specific site you want the DRaaS Connector to able to connect to.</p> <hr/>
Name to give the connector, as it will appear in the VMware Cloud DR UI.	

**Note** Use the vSphere web console for changing the connector IP address. Do not use an SSH session to configure the connector VM. You must use the vSphere web console if the network does not have DHCP enabled.

#### Procedure

- 1 In the vCenter UI select VMs, and then select the DRaaS Connector VM. Right-click and choose **Power** → **Power on**.

**Note** Always power on and power off the DRaaS Connector from the vCenter on the protected site.

- 2 Under the VM, click **Launch web console**.
- 3 When the console session is open, log in to the connector VM console using the following credentials: `admin/vmware#1`.
- 4 Next, in the Select the network address IP address allocation, either (a) Static or (b) DHCP.

If a DRaaS Connector has already been configured on this site, the current network configuration displays in the console and the current values are set as defaults. To use the existing configuration in all of these steps, press the Enter key.

- 5 If you chose (a) for static IP address allocation: Enter an IP address, Subnet mask, and Gateway. Enter the IP address of one or more DNS servers (up to three supported) that are able to resolve host names on the internal network. To use the existing configuration in all of these steps, press the Enter key to accept the suggested default values.

**Note** Google DNS servers (8.8.8.8 and 8.8.4.4) do not work reliably with VMware Cloud DR. We advise you to use private, non-Google and non-public DNS servers when configuring the connector VM.

- 6 Next, enter the Orchestrator FQDN or IP address. If you have [AWS Direct Connect](#) configured with a private VIF, then you use the Orchestrator IP address.
- 7 Next, enter the DRaaS Connector temporary passcode.
- 8 Enter a name to identify the connector.

## What to do next

After you configure the DRaaS Connector, the console window will return you to the command prompt. If your protected site is an on-premises vSphere, you can now return to the VMware Cloud DR UI to [Register vCenter Server](#) for the protected site. (If your protected site is an SDDC, you do not need to register vCenter.) If your protected site is a Google Cloud VMware Engine private cloud, then see [Register vCenter for Google Cloud VMware Engine Protected Site](#).

## Register vCenter Server

After you deploy the DRaaS Connector to an on-premises site and configure it, you must register it with the vCenter Server.

---

**Note** This task is only required if the protected site is an on-premises vSphere. If the protected site is a VMware Cloud on AWS SDDC, you do not need to register the vCenter Server.

---

Registering vCenter Server with the DRaaS Connector requires vCenter Server credentials (user name and password). You can use two types of users to register the DRaaS Connector with a vCenter Server on a protected site:

- **vCenter Server Administrator user.** The vCenter Server Administrator user role provides sufficient privileges for VMware Cloud DR data protection and DR operations. Using a vCenter Server Administrator user registers the as DRaaS Connector an extension, without storing its credentials.
- **Restricted user.** If you prefer to register the DRaaS Connector with a user that has limited permissions, you can use a provided Python script that creates a restricted user with only the [Privileges for a Restricted vCenter Server User](#) needed for data protection and DR operations. VMware Cloud DR stores credentials for this user. For more information on how to create this user, see [Custom Script for Creating a Restricted vCenter Server User](#). You can also use the DRaaS Connector CLI to create this user manually. For more information, see [Create a Restricted vCenter Server User with the DRaaS Connector CLI](#).

### Procedure

- 1 From the left navigation, select Protected sites and then choose a specific protected site.
- 2 On the protected site page, under vCenters click the **Register vCenter** button.
- 3 In the **Register vCenter** dialog box, enter the vCenter server's IP address.
- 4 Under Authentication, choose one of the two options:
  - **Authenticate with vCenter administrator.** Use this option if you use an existing vCenter Administrator user account with full administrator access.
  - **Authenticate with restricted vCenter user.** Use this option if you want to supply a custom user with only the minimum set of privileges to perform snapshots and failover and/or fallback operations. You can use the [Custom Script for Creating a Restricted vCenter Server User](#) to create this user.

- 5 After selecting an option, enter the vCenter user name and password.
- 6 Click **Register**.

#### What to do next

If you are using a restricted user for vCenter registration, and you change the user's password, you must [Refresh vCenter Server User Credentials](#).

## Custom Script for Creating a Restricted vCenter Server User

You can use a custom Python script to create a restricted vCenter Server user, with optional flags to create a vCenter Server role and apply snapshot, failover, and failback privileges.

This script allows you to create a vCenter Server user with restricted roles that contain the minimum set of [Privileges for a Restricted vCenter Server User](#) required by the DRaaS Connector to perform snapshot replication, failover, and failback operations. Run this script from the DRaaS Connector VM command line.

### Prerequisites

- GOVC installed on the system where you run this script. You can download GOVC here: <https://github.com/vmware/govmomi/tree/master/govc#readme>.
- Set the GOVC\_PATH variable in the script to the path to the 'govc' executable on the system where you run this script. The default is /usr/local/bin/govc.
- Administrator user account and password from the vCenter Server where you are creating this user. You must pass these credentials when you run the script.

### Usage

Create a vCenter Server user and role with a minimum set of privileges needed by the DRaaS Connector. Optional commands appear inside square brackets [ ].

```
vcdr-create-vcenter-user.py [-h]
```

Options:

Option	Description
-h, --help	Displays a list of commands and options.
--vcenter VCENTER	IP address of the vCenter Server where you create this user.
--admin-username admin-username	Username of a user with vCenter Server Administrator privileges. User must have vCenter Server Administrator privileges before you run the CLI commands. If not supplied, the script prompts you for this username.
- [-admin-password admin-password]	Password of a user with vCenter Server Administrator privileges. If not supplied, the script prompts you for this password.

Option	Description
<code>--new-username new-username</code>	Use name for the new user.
<code>[--new-password new-password]</code>	Password for the new user.
<code>--vcenter-role vcenter-role</code>	Name of new or existing vCenter Server role to assign to the new user. You can apply <code>--snapshot-privs</code> and/or <code>--failback-privs</code> to this role.
<code>[--keep-existing-privs]</code>	Keep existing privileges on the provided role. Using this option ensures that you do not unnecessarily remove privileges from the provided role.
<code>[--create-role-only]</code>	Create role with necessary privileges, but do not create a user account or assign permissions.
<code>[--snapshot-privs]</code>	This option adds snapshot and failover privileges to the user role.
<code>[--failback-privs]</code>	This option adds snapshot, failover, and failback privileges to the user role.

## Example

The following example shows how to use the CLI to create a vCenter Server role called `vcdr-bkup-role` and assign the role snapshot and failover privileges (but no failback). This command then creates a vCenter Server user named `vcdr-bkup-user@vsphere.local`, and globally assigns this new user the new role.

```
connector-name>> vcdr-create-vcenter-user.py --vcenter 192.0.170.23 --admin-username
administrator@vsphere.local --new-username vcdr-bkup-user@vsphere.local --vcenter-role vcdr-
bkup-role --snapshot-privs
```

## Script

```
#!/usr/bin/env python

"""
Copyright 2021 VMware, Inc. All rights reserved.

=====
IMPORT! READ THE FOLLOWING
=====

This script can be downloaded and used by customers to create a vCenter user,
role, and permission for use with vCenter registration with the VMware Cloud
Disaster Recover (VCDR) product.

REQUIREMENTS

- Set the GOVC_PATH variable, below, to the path to the 'govc' executable on
  the system where this script will run. The default is '/usr/local/bin/govc'.

- GOVC
```

GOVC must be installed in order for this script to work.  
<https://github.com/vmware/govmomi/tree/master/govc#readme>

#### EXAMPLE

```
$ vcdr-create-vcenter-user.py --vcenter 10.80.15.23 --admin-username
administrator@vsphere.local \
  --new-username vcdr-bkup-user@vsphere.local --vcenter-role vcdr-bkup-role --snapshot-privs
```

The above will create a vCenter role called vcdr-bkup-role with the privileges necessary to perform VM backup (i.e., it will not work for failback). It will also create a vCenter user called vcdr-bkup-user@vsphere.local. Finally, it will globally assign this new user the above role.

#### USAGE

```
$ vcdr-create-vcenter-user.py --help
```

```
usage: vcdr-create-vcenter-user.py [-h] --vcenter VCENTER --admin-username
admin-username --new-username new-username
[--admin-password admin-password]
[--new-password new-password] --vcenter-role
vcenter-role [--keep-existing-privs]
[--create-role-only]
[--snapshot-privs | --failback-privs]
```

Create vCenter user/role/permission with minimal privileges for use with VCDR.

#### optional arguments:

```
-h, --help          show this help message and exit
--vcenter VCENTER   vCenter IP on which to create/update user
--admin-username admin-username
                    Admin username at desired vCenter
--new-username new-username
                    New username to be created in the desired vCenter
--admin-password admin-password
                    Admin password at desired vCenter (will be prompted if
                    not provided)
--new-password new-password
                    New password for the new user in the desired vCenter
                    (will be prompted if not provided)
--vcenter-role vcenter-role
                    Name of new or existing vCenter role to associate to
                    the new user
--keep-existing-privs
                    Keep existing privileges on provided role (i.e., do
                    not prune unnecessary privileges from provided role)
--create-role-only   Create role with necessary privileges but do not
                    create a user account or assign permissions
--snapshot-privs     Create a user with privileges necessary to snapshot
                    VMs and failover (the default)
--failback-privs     Create a user with privileges necessary to snapshot
                    VMs, failover, and failback
```

#### MISC

This script cannot be used to create/update localos users (i.e., the created VCDR user account must be of the form 'name@domain', e.g., 'Administrator@vmware.local').

This script has been tested with python3.

```

"""
__author__ = "VMware, Inc"

from getpass import getpass
import json
import logging
import os
import subprocess
import time

logger = logging.getLogger(__name__)
logger.setLevel(logging.DEBUG)

GOVC_PATH = '/usr/local/bin/govc'
LOG_FILE = '/var/tmp/vcdr-create-vcenter-user'

VCDR_SNAPSHOT_PRIVS = [
    'Datastore.Browse',
    'Datastore.FileManagement',
    'Global.DisableMethods',
    'Global.EnableMethods',
    'System.Anonymous',
    'System.Read',
    'System.View',
    'VirtualMachine.Config.ChangeTracking',
    'VirtualMachine.Config.DiskLease',
    'VirtualMachine.Config.QueryUnownedFiles',
    'VirtualMachine.Config.ReloadFromPath',
    'VirtualMachine.Provisioning.DiskRandomRead',
    'VirtualMachine.Provisioning.GetVmFiles',
    'VirtualMachine.Provisioning.ReadCustSpecs',
    'VirtualMachine.State.CreateSnapshot',
    'VirtualMachine.State.RemoveSnapshot'
]

VCDR_FAILBACK_PRIVS = VCDR_SNAPSHOT_PRIVS + [
    'Datastore.AllocateSpace',
    'Datastore.Config',
    'Datastore.DeleteFile',
    'Datastore.UpdateVirtualMachineFiles',
    'Datastore.UpdateVirtualMachineMetadata',
    'Global.CancelTask',
    'Global.GlobalTag',
    'Global.Licenses',
    'Global.ManageCustomFields',
    'Global.SetCustomField',
    'InventoryService.Tagging.AttachTag',
    'Network.Assign',
    'Resource.AssignVAppToPool',
    'Resource.AssignVMToPool',

```

```

'Sessions.GlobalMessage',
'Sessions.ValidateSession',
'VirtualMachine.Config.AddExistingDisk',
'VirtualMachine.Config.AddNewDisk',
'VirtualMachine.Config.AddRemoveDevice',
'VirtualMachine.Config.AdvancedConfig',
'VirtualMachine.Config.CPUCount',
'VirtualMachine.Config.DiskExtend',
'VirtualMachine.Config.EditDevice',
'VirtualMachine.Config.HostUSBDevice',
'VirtualMachine.Config.ManagedBy',
'VirtualMachine.Config.Memory',
'VirtualMachine.Config.MksControl',
'VirtualMachine.Config.QueryFTCompatibility',
'VirtualMachine.Config.RawDevice',
'VirtualMachine.Config.RemoveDisk',
'VirtualMachine.Config.Rename',
'VirtualMachine.Config.ResetGuestInfo',
'VirtualMachine.Config.Resource',
'VirtualMachine.Config.Settings',
'VirtualMachine.Config.SwapPlacement',
'VirtualMachine.GuestOperations.Execute',
'VirtualMachine.GuestOperations.Modify',
'VirtualMachine.GuestOperations.Query',
'VirtualMachine.Interact.Backup',
'VirtualMachine.Interact.DeviceConnection',
'VirtualMachine.Interact.GuestControl',
'VirtualMachine.Interact.PowerOff',
'VirtualMachine.Interact.PowerOn',
'VirtualMachine.Interact.SetCDMedia',
'VirtualMachine.Interact.SetFloppyMedia',
'VirtualMachine.Interact.ToolsInstall',
'VirtualMachine.Inventory.Create',
'VirtualMachine.Inventory.CreateFromExisting',
'VirtualMachine.Inventory.Move',
'VirtualMachine.Inventory.Register',
'VirtualMachine.Inventory.Unregister',
'VirtualMachine.Provisioning.Customize',
'VirtualMachine.Provisioning.DiskRandomAccess',
'VirtualMachine.Provisioning.FileRandomAccess',
'VirtualMachine.Provisioning.MarkAsVM',
'VirtualMachine.Provisioning.ModifyCustSpecs',
'VirtualMachine.Provisioning.PromoteDisks',
'VirtualMachine.State.RevertToSnapshot'
]

VCDR_LWD_PRIVS = [
    'vSphereDataProtection.Protection',
    'Host.Config.NetService',
    'vSphereDataProtection.Recovery',
    'System.Read',
    'System.Anonymous',
    'Host.Config.Storage',
]

```

```

class GOVC(object):
    def __init__(self, vc_url, vc_username, vc_password, verify_server_cert=False,
exit_on_cmd_failure=True):
        """
        The govc URL scheme defaults to https and the URL path defaults to /sdk.
        This means that:
        1. 'host' is equivalent to 'https://<host>/sdk'.
        2. 'username:password@host' is equivalent to 'https://<username:password@host>/
sdk'. # pragma: allowlist secret

        :param vc_url: URL of ESXi or vCenter instance to connect to
        :param vc_username: vCenter or ESXi admin user's username
        :param vc_password: vCenter or ESXi admin user's password
        :param verify_server_cert: Verify vCenter or ESXi server certificate
        :return: None
        """
        self.vc_url = vc_url
        self.vc_username = vc_username
        self.vc_password = vc_password
        self.verify_server_cert = verify_server_cert
        self.exit_on_cmd_failure = exit_on_cmd_failure
        self.set_envs()

    def _run_govc_cli(self, cmd_args, timeout, json=False):
        """
        Run command.

        :param cmd_args: List of command line args
        :param json: Enable JSON output
        :param timeout: command timeout in secs
        :return: command output
        """
        cmd = [GOVC_PATH]
        cmd.extend(cmd_args)
        if json:
            cmd.insert(2, '-json')

        wait = 0
        logger.debug('Run command %s.', cmd)
        p = subprocess.Popen(cmd, stdout=subprocess.PIPE, stderr=subprocess.PIPE)
        while p.poll() is None and wait < timeout:
            wait += 1
            time.sleep(1)

        if p.poll() is None:
            p.kill()
            rc = 124
        else:
            rc = p.returncode
        stdout, stderr = p.communicate()

        if rc != 0:
            if rc == 124:
                stderr = 'timed out after {} secs, {}'.format(timeout, stderr)
            logger.error('%s failed, rc: %d, stdout: %s, stderr: %s', cmd, rc, stdout, stderr)

```



```

        if self.exit_on_cmd_failure:
            exit(rc)
        else:
            raise subprocess.CalledProcessError(rc, cmd, output=stdout+stderr)

    logger.info(stdout)
    return stdout.strip()

def set_envs(self):
    """
    Set govc login environment variables.

    While using '-u https://user:pass@host/sdk' from the govc command line,      # pragma:
allowlist secret
    there is a bug if the username or the password includes special characters.
    Using environment variables instead confirmed works.
    """
    os.environ['GOVC_URL'] = self.vc_url
    os.environ['GOVC_USERNAME'] = self.vc_username
    os.environ['GOVC_PASSWORD'] = self.vc_password
    os.environ['GOVC_INSECURE'] = 'false' if self.verify_server_cert else 'true'
    logger.debug('govc envs: %s', {k: v for k, v in os.environ.items() if
k.startswith('GOVC') and k != 'GOVC_PASSWORD'})

def ssogroup_create(self, name, description=None, timeout=15, json=False, **kwargs):
    """
    Create SSO group.

    :param name: SSO group name
    :param description: SSO group description
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: Command output
    """
    kwargs.clear()

    cmd_args = ['sso.group.create']
    if description:
        cmd_args.extend(['-d', description])
    cmd_args.append(name)

    return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssogroup_ls(self, timeout=15, json=False, **kwargs):
    """
    List SSO groups.

    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: SSO groups
    """
    kwargs.clear()

    cmd_args = ['sso.group.ls']

```

```

        return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssogroup_rm(self, name, timeout=15, json=False, **kwargs):
    """
    Remove SSO group.

    :param name: SSO group name
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: Command output
    """
    kwargs.clear()

    cmd_args = ['sso.group.rm', name]

    return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssogroup_update(self, name, description='', add_user=None, remove_user=None,
                    timeout=15, json=False, **kwargs):
    """
    Update SSO group.

    :param name: SSO group name
    :param description: SSO group description
    :param add_user: Add user to group
    :param remove_user: Remove user from group
    :param description: SSO group description
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: Command output
    """
    kwargs.clear()

    cmd_args = ['sso.group.update']
    if description:
        cmd_args.extend(['-d', description])
    if add_user is not None:
        cmd_args.extend(['-a', add_user])
    if remove_user is not None:
        cmd_args.extend(['-r', remove_user])
    cmd_args.append(name)

    return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssouser_create(self, name, password=None, description=None,
                  solution=False, act_as_user=True, role='Administrator',
certificate=None,
                  timeout=15, json=False, **kwargs):
    """
    Create SSO user.

    :param name: SSO user name
    :param password: SSO person user password
    :param description: SSO user description
    :param solution: Whether is solution user

```

```

:param act_as_user: ActAsUser role for solution user WSTrust
:param role: Role for solution user (RegularUser|Administrator)
:param certificate: Certificate for solution user
:param timeout: command timeout in secs
:param json: Enable JSON output
:return: Command output
"""
kwargs.clear()

cmd_args = ['sso.user.create']
if description:
    cmd_args.extend(['-d', description])
if solution:
    assert certificate, 'Solution user certificate is required!'
    assert role in ['RegularUser', 'Administrator'], (
        'Invalid solution user role: {}'.format(role))
    if act_as_user:
        cmd_args.append('-A')
    cmd_args.extend(['-R', role, '-C', certificate])
else:
    # Person user.
    assert password, 'SSO user password is required!'
    cmd_args.extend(['-p', password])
cmd_args.append(name)

return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssouser_id(self, name=None, timeout=15, json=False, **kwargs):
    """
    Get SSO user and group IDs.

    :param name: SSO user name
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: SSO user and group IDs
    """
    kwargs.clear()

    cmd_args = ['sso.user.id']
    if name is not None:
        cmd_args.append(name)

    return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssouser_ls(self, solution=False, timeout=15, json=False, **kwargs):
    """
    List SSO users.

    :param solution: List solution users
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: SSO users
    """
    kwargs.clear()

```

```

        cmd_args = ['sso.user.ls']
        if solution:
            cmd_args.append('-s')

        return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssouser_rm(self, name, timeout=15, json=False, **kwargs):
    """
    Remove SSO user.

    :param name: SSO user name
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: SSO user and group IDs
    :return: Command output
    """
    kwargs.clear()

    cmd_args = ['sso.user.rm', name]

    return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def ssouser_update(self, name, password=None, description=None,
                    solution=False, act_as_user=True, role=None, certificate=None,
                    timeout=15, json=False, **kwargs):
    """
    Update SSO user.

    :param name: SSO user name
    :param password: SSO person user password
    :param description: SSO user description
    :param solution: Whether is solution user
    :param act_as_user: ActAsUser role for solution user WSTrust
    :param role: Role for solution user (RegularUser|Administrator)
    :param certificate: Certificate for solution user
    :param timeout: command timeout in secs
    :param json: Enable JSON output
    :return: Command output
    """
    kwargs.clear()

    cmd_args = ['sso.user.update']
    if description is not None:
        cmd_args.extend(['-d', description])
    if solution:
        if act_as_user:
            cmd_args.append('-A')
        if role is not None:
            assert role in ['RegularUser', 'Administrator'], (
                'Invalid solution user role: {}'.format(role))
            cmd_args.extend(['-R', role])
        if certificate is not None:
            cmd_args.extend(['-C', '{}'.format(certificate)])
    else:
        if password is not None:

```

```

        cmd_args.extend(['-p', password])
    cmd_args.append(name)

    return self._run_govc_cli(cmd_args, timeout=timeout, json=json)

def get_sso_user_privileges(self, username, timeout=15, **kwargs):
    """
    Get privileges associated with an SSO user
    :param name: SSO user name
    :param timeout: command timeout in secs
    :return: List of privileges
    """
    kwargs.clear()

    permission_cmd_args = ['permissions.ls']
    permission_output = self._run_govc_cli(permission_cmd_args, timeout=timeout)
    privileges_list = []
    normalized_user = username.lower()
    for line in permission_output.splitlines()[1:]:
        line_split = line.split('/') #formatted [role, entity, username, propagate]
        user = line_split[1].split()[0].lower() # user formatted either <user> or
<domain>\\<user>
        user_split = user.split("\\")
        role = line_split[0].strip()
        constructed_user = user_split[1] + "@" + user_split[0] if len(user_split) == 2
    else user
        if normalized_user == constructed_user:
            role_cmd_args = ['role.ls', role]
            privileges_list.extend(self._run_govc_cli(role_cmd_args,
timeout=timeout).split())
            break

    return privileges_list

def set_sso_user_privileges(self, username, role, entity='/', timeout=15, **kwargs):
    """
    Attach a role to a user
    :param username: SSO user name
    :param role: Role to attach to SSO user
    :param entity: Entity to attach privileges on, default root
    :param timeout: command timeout in secs
    """
    kwargs.clear()

    permission_cmd_args = ["permissions.set"]
    permission_cmd_args.extend(['-principal', username])
    permission_cmd_args.extend(['-role', role])
    permission_cmd_args.extend(['-propagate=true'])
    permission_cmd_args.extend([entity])
    self._run_govc_cli(permission_cmd_args, timeout=timeout)

def role_ls(self, role=None, json=False, timeout=15, **kwargs):
    """
    Get information about roles. Lists names of all roles if a role is
    not provided.

```

```

:param role: name of role to get information about
:param timeout: command timeout in secs
"""
kwargs.clear()

role_cmd_args = ["role.ls"]
if role:
    role_cmd_args.append(role)
return self._run_govc_cli(role_cmd_args, timeout=timeout, json=json)

def role_create(self, role, privileges, timeout=15, **kwargs):
    """
    Get information about roles. Lists names of all roles if a role is
    not provided.
    :param role: name of role to get information about
    :param privileges: list of privileges to attach to role
    :param timeout: command timeout in secs
    """
    kwargs.clear()

    role_cmd_args = ["role.create"]
    role_cmd_args.append(role)
    role_cmd_args.extend(privileges)
    return self._run_govc_cli(role_cmd_args, timeout=timeout)

def role_update(self, role, privileges_to_add=None, privileges_to_remove=None,
timeout=15, **kwargs):
    """
    Get information about roles. Lists names of all roles if a role is
    not provided.
    :param role: name of role to get information about
    :param privileges_to_add: list of privileges to attach to role
    :param privileges_to_remove: list of privileges to remove
    :param timeout: command timeout in secs
    """
    kwargs.clear()

    role_cmd_args = ["role.update"]
    if privileges_to_add:
        role_cmd_args.append("-a")
        role_cmd_args.append(role)
        role_cmd_args.extend(privileges_to_add)
    if privileges_to_remove:
        role_cmd_args.append("-r")
        role_cmd_args.append(role)
        role_cmd_args.extend(privileges_to_remove)
    return self._run_govc_cli(role_cmd_args, timeout=timeout)

class vcdr_create_vcenter_user(object):
    def _get_password(self, account=None, retype=False):
        """
        Prompt for a password.
        """
        prompt = 'Retype password' if retype else 'Password'
        if account is not None:

```

```

        prompt += ' for account %s' % account
    prompt += ': '
    while True:
        password = getpass(prompt=prompt)
        if password:
            return password

def _get_username_and_domain(self, user):
    username = user
    domain = None
    if '@' in user:
        username, domain = user.split('@', 1)
    return username, domain

def _get_user(self, govc_channel, user):
    username, domain = self._get_username_and_domain(user)
    users = govc_channel.ssouser_ls(json=True)
    if users:
        users = json.loads(users)
        for user in users:
            if user['Id']['Name'] == username and user['Id']['Domain'] == domain:
                return user
    return None

def _get_role(self, govc_channel, role_name):
    try:
        role = govc_channel.role_ls(role_name, json=True)
        return json.loads(role) if role else None
    except:
        # Unable to find role.
        pass
    return None

def handle_vcdr_create_vcenter_user(self, vcenter, adminUsername, newUsername,
adminPassword=None, newPassword=None, vcenterRole=None,
                                vcenterEntity=None, snapshotPrivs=False,
failbackPrivs=True, keepExistingPrivs=False, createRoleOnly=False):
    """
    @param vcenter: vCenter IP to create new user on
    @param adminUsername: Admin username at desired vCenter
    @param newUsername: New username to be created in the desired vCenter
    @param adminPassword: Admin password at desired vCenter
    @param newPassword: New password to be attached to the new user in the desired vCenter
    @param vcenterRole: Name of the vCenter role to attach to the new user
    @param vcenterEntity: VCenter entity path to give the new user privileges on, by
default root
    @param snapshotPrivs: Associate privileges necessary for snapshotting and failover
with the role
    @param failbackPrivs: Associate privileges necessary for snapshotting, failover, and
failback with the role
    @param keepExistingPrivs: Keep existing privileges on provided role (i.e., do not
prune unnecessary privileges from provided role)
    @param createRoleOnly: Create role with necessary privileges but do not create a user
account or assign permissions
    """

```

```

try:
    if adminPassword is None:
        adminPassword = self._get_password(account=adminUsername)
    if newPassword is None:
        while True:
            newPassword = self._get_password(account=newUsername)
            newPassword2 = self._get_password(account=newUsername, retype=True)
            if newPassword == newPassword2:
                break
            print('Passwords do not match.')
        govc_channel = GOVC(vcenter, adminUsername, adminPassword,
exit_on_cmd_failure=False)
    except:
        msg = u'Unable to establish connection with vCenter {} with username
{}'.format(vcenter, adminUsername)
        logging.exception(msg)
        print(msg)
        return

# Decide which privileges we care about. TODO: select based on CLI.
if fallbackPrivs:
    target_privs = VCDR_FAILBACK_PRIVS
else:
    target_privs = VCDR_SNAPSHOT_PRIVS

# Check the privileges of the Admin, and assign only the permissible privileges.
admin_role = self._get_role(govc_channel, 'Admin')
target_privs = list(set(admin_role['Privilege']) & set(target_privs + VCDR_LWD_PRIVS))

# Create user if it doesn't already exist.
if not createRoleOnly:
    existing_user = False
    try:
        existing_user = self._get_user(govc_channel, newUsername)
        newUsernameWithoutDomain, _ = self._get_username_and_domain(newUsername)
        if not existing_user:
            govc_channel.ssouser_create(newUsernameWithoutDomain, newPassword)
        else:
            govc_channel.ssouser_update(newUsernameWithoutDomain, newPassword)
    except:
        msg = u'Unable to {} user {} on vCenter {}'.format("update" if existing_user
else "create", newUsername, vcenter)
        logging.exception(msg)
        print(msg)
        return

# Create role if it doesn't already exist and add appropriate privs.
try:
    existing_role = self._get_role(govc_channel, vcenterRole)
    if existing_role:
        govc_channel.role_update(vcenterRole, privileges_to_add=target_privs)
    else:
        govc_channel.role_create(vcenterRole, target_privs)
except:
    msg = u'Unable to create/update role {} on vCenter {}'.format(vcenterRole,

```



```

vcenter)
    logging.exception(msg)
    print(msg)
    return

    # Remove extraneous privs from role.
    if not keepExistingPrivs:
        try:
            role = self._get_role(govc_channel, vcenterRole)
            privs = role['Privilege']
            extra_privs = list(set(privs) - set(target_privs))
            if extra_privs:
                govc_channel.role_update(vcenterRole, privileges_to_remove=extra_privs)
        except:
            msg = u'Unable to remove extraneous privs from role {} on vCenter
{}'.format(vcenterRole, vcenter)
            logging.exception(msg)
            print(msg)
            return

    # Update the permissions.
    if not createRoleOnly:
        try:
            vcenterEntity = vcenterEntity or '/'
            govc_channel.set_sso_user_privileges(newUsername, vcenterRole, vcenterEntity)
        except:
            msg = u'Unable to create permission on {} for {}/{} on vCenter
{}'.format(vcenterEntity, newUsername, vcenterRole, vcenter)
            logging.exception(msg)
            print(msg)
            return

    return

if __name__ == '__main__':
    import argparse

    logging.basicConfig(level=logging.DEBUG, filename=LOG_FILE, filemode='a')
    logging.info('Starting script')

    parser = argparse.ArgumentParser(description='Create vCenter user/role/permission with
minimal privileges for use with VCDR.')
    parser.add_argument('--vcenter',
                        required=True,
                        help='vCenter IP on which to create/update user')
    parser.add_argument('--admin-username',
                        required=True,
                        metavar='admin-username',
                        help='Admin username at desired vCenter')
    parser.add_argument('--new-username',
                        required=True,
                        metavar='new-username',
                        help='New username to be created in the desired vCenter')
    parser.add_argument('--admin-password',

```

```

        metavar='admin-password',
        help='Admin password at desired vCenter (will be prompted if not
provided)')
    parser.add_argument('--new-password',
        metavar='new-password',
        help='New password for the new user in the desired vCenter (will be
prompted if not provided)')
    parser.add_argument('--vcenter-role',
        required=True,
        metavar='vcenter-role',
        help='Name of new or existing vCenter role to associate to the new
user')
    parser.add_argument('--keep-existing-privs',
        action='store_true',
        help='Keep existing privileges on provided role (i.e., do not prune
unnecessary privileges from provided role)')
    parser.add_argument('--create-role-only',
        action='store_true',
        help='Create role with necessary privileges but do not create a user
account or assign permissions')

    group = parser.add_mutually_exclusive_group()
    group.add_argument('--snapshot-privs',
        action='store_true',
        help='Create a user with privileges necessary to snapshot VMs and
failover (the default)')
    group.add_argument('--failback-privs',
        action='store_true',
        help='Create a user with privileges necessary to snapshot VMs,
failover, and failback')

    args = vars(parser.parse_args())

    vcenter = args.get('vcenter')
    admin_username = args.get('admin_username')
    new_username = args.get('new_username')
    admin_password = args.get('admin_password')
    new_password = args.get('new_password')
    vcenter_role = args.get('vcenter_role')
    vcenter_entity = '/'
    snapshot_privs = args.get('snapshot_privs')
    failback_privs = args.get('failback_privs')
    keep_existing_privs = args.get('keep_existing_privs', False)
    create_role_only = args.get('create_role_only', False)

    worker = vcdr_create_vcenter_user()
    worker.handle_vcdr_create_vcenter_user(vcenter, admin_username, new_username,
adminPassword=admin_password, newPassword=new_password,
        vcenterRole=vcenter_role,
vcenterEntity=vcenter_entity, snapshotPrivs=snapshot_privs,
        failbackPrivs=failback_privs,
keepExistingPrivs=keep_existing_privs, createRoleOnly=create_role_only)

```

```
logging.info('Ending script')  
  
exit(0)
```

## Privileges for a Restricted vCenter Server User

When you create a restricted vCenter Server user for the DRaaS Connector, apply the following privileges to the vCenter Server role to allow snapshot replication, failback, and failover.

CLI Option	Privileges
<code>--snapshot-privs</code>	<b>Datastore.Browse</b> <b>Datastore.FileManagement</b> <b>Global.DisableMethods</b> <b>Global.EnableMethods</b> <b>System.Anonymous</b> <b>System.Read</b> <b>System.View</b> <b>VirtualMachine.Config.ChangeTracking</b> <b>VirtualMachine.Config.DiskLease</b> <b>VirtualMachine.Config.QueryUnownedFiles</b> <b>VirtualMachine.Config.ReloadFromPath</b> <b>VirtualMachine.Provisioning.DiskRandomRead</b> <b>VirtualMachine.Provisioning.GetVmFiles</b> <b>VirtualMachine.Provisioning.ReadCustSpecs</b> <b>VirtualMachine.State.CreateSnapshot</b> <b>VirtualMachine.State.RemoveSnapshot</b> For high-frequency snapshots: <b>vSphereDataProtection.Protection</b> <b>Host.Config.NetService</b> <b>vSphereDataProtection.Recovery</b> <b>System.Read</b> <b>System.Anonymous</b> <b>Host.Config.Storage</b>
<code>--failback-privs</code>	This option applies all privileges listed above for <code>--snapshot-privs</code> , plus all of the following privileges: <b>Datastore.AllocateSpace</b> <b>Datastore.Config</b> <b>Datastore.DeleteFile</b> <b>Datastore.UpdateVirtualMachineFiles</b> <b>Datastore.UpdateVirtualMachineMetadata</b> <b>Global.CancelTask</b> <b>Global.GlobalTag</b> <b>Global.Licenses</b> <b>Global.ManageCustomFields</b> <b>Global.SetCustomField</b> <b>InventoryService.Tagging.AttachTag</b> <b>Network.Assign</b> <b>Resource.AssignVAppToPool</b> <b>Resource.AssignVMToPool</b> <b>Sessions.GlobalMessage</b> <b>Sessions.ValidateSession</b> <b>VirtualMachine.Config.AddExistingDisk</b> <b>VirtualMachine.Config.AddNewDisk</b> <b>VirtualMachine.Config.AddRemoveDevice</b> <b>VirtualMachine.Config.AdvancedConfig</b>

CLI Option	Privileges
	VirtualMachine.Config.CPUCount
	VirtualMachine.Config.DiskExtend
	VirtualMachine.Config.EditDevice
	VirtualMachine.Config.HostUSBDevice
	VirtualMachine.Config.ManagedBy
	VirtualMachine.Config.Memory
	VirtualMachine.Config.MksControl
	VirtualMachine.Config.QueryFTCompatibility
	VirtualMachine.Config.RawDevice
	VirtualMachine.Config.RemoveDisk
	VirtualMachine.Config.Rename
	VirtualMachine.Config.ResetGuestInfo
	VirtualMachine.Config.Resource
	VirtualMachine.Config.Settings
	VirtualMachine.Config.SwapPlacement
	VirtualMachine.GuestOperations.Execute
	VirtualMachine.GuestOperations.Modify
	VirtualMachine.GuestOperations.Query
	VirtualMachine.Interact.Backup
	VirtualMachine.Interact.DeviceConnection
	VirtualMachine.Interact.GuestControl
	VirtualMachine.Interact.PowerOff
	VirtualMachine.Interact.PowerOn
	VirtualMachine.Interact.SetCDMedia
	VirtualMachine.Interact.SetFloppyMedia
	VirtualMachine.Interact.ToolsInstall
	VirtualMachine.Inventory.Create
	VirtualMachine.Inventory.CreateFromExisting
	VirtualMachine.Inventory.Move
	VirtualMachine.Inventory.Register
	VirtualMachine.Inventory.Unregister
	VirtualMachine.Provisioning.Customize
	VirtualMachine.Provisioning.DiskRandomAccess
	VirtualMachine.Provisioning.FileRandomAccess
	VirtualMachine.Provisioning.MarkAsVM
	VirtualMachine.Provisioning.ModifyCustSpecs
	VirtualMachine.Provisioning.PromoteDisks
	VirtualMachine.State.RevertToSnapshot

## Create a Restricted vCenter Server User with the DRaaS Connector CLI

You can create a restricted vCenter Server user with the DRaaS Connector CLI, to ensure that the connector can only access vCenter Server with a minimal set of privileges.

You can [Create a Restricted vCenter Server User](#) using the DRaaS Connector CLI, or you can [Custom Script for Creating a Restricted vCenter Server User](#) provided by VMware Cloud DR that creates a user, and then applies specific roles with only the privileges required for DR operations.

Once you create the user, you can then [Register vCenter Server](#) with the DRaaS Connector, using this new user to authenticate with vCenter Server.

Before you use this feature, consider the following:

- Manually creating a user can be error prone.
- Determine when the password for the provided account expires and manually change the password before its expiration.
- Manually update the user account to provide privileges to perform failback, so that failback privileges are only available when you run a failback recovery plan.

## DRaaS Connector CLI Commands for Creating a Restricted vCenter Server User

You can use the DRaaS Connector CLI to create a vCenter Server user, with optional flags to create a vCenter Server role and apply snapshot, failover, and failback privileges.

### Usage

```
drc create-vcenter-user --option
```

Optional commands are listed inside square brackets [ ].

Option	Description
-h, --help	Displays a list of commands and options.
--vcenter VCENTER	IP address of the vCenter Server where you want to create this user.
--admin-username admin-username	Username of a user with vCenter Server Administrator privileges. User must vCenter Server Administrator privileges before you run the CLI commands. If not supplied, the script prompts you for this username.
- [-admin-password admin-password]	Password of a user with vCenter Server Administrator privileges. If not supplied, the script prompts you for this password.
--new-username new-username	Username for the new user.
[--new-password new-password]	Password for the new user.
--vcenter-role vcenter-role	Name of new or existing vCenter Server role to associate to the new user. You can apply --snapshot-privs and/or --failback-privs to this role.
[--snapshot-privs]	Adds snapshot and failover privileges to the user role.
[--failback-privs]	Adds snapshot, failover, and failback privileges to the user role.

## Create a Restricted vCenter Server User

Use the DRaaS Connector CLI to create a custom user that can be used to register vCenter Server with the DRaaS Connector.

### Prerequisites

When you manually create a vCenter Server user, make sure that the "propagate to children" permission is applied to the inventory root with a role providing the necessary privileges.

When you run this command, it creates a role with the minimal privileges (based on `--snapshot-privs` or `--failback-privs`), and it creates a "propagate to children" permission on the inventory root, associating this user with this role.

### Procedure

- 1 Open an SSH connection to the DRaaS Connector. Or, you can use the VM console to connect with the DRaaS Connector VM.

You can obtain the DRaaS Connector IP address for an on-premises protected site in the VMware Cloud DR UI under **Sites > Protected Sites**.

To obtain the connector admin password, go to **Sites > Protected Sites > Show password**.

- 2 Run the following command:

```
connector-name>> drc create-vcenter-user --vcenter 192.0.2.0
--admin-username administrator@vsphere.local --admin-password drcSmeck$2!
--new-username drc_user@vmware.com --new-password drcShlep$2!
--vcenter-role draasrole --snapshot-privs --failback-privs
```

### What to do next

---

**Note** If you omit any of the password flags when running this command, you are interactively prompted to supply the password after you run the command.

---

Now you can register vCenter Server with the DRaaS Connector and authenticate with this new user.

## Register vCenter Server Using Restricted vCenter Server User

Once you have created a restricted vCenter Server user, you can now register vCenter Server with the DRaaS Connector using the restricted user account.

### Prerequisites

You only have to register one DRaaS Connector with vCenter Server. Other DRaaS Connectors at the site automatically update to be aware of the newly registered vCenter Server.

**Procedure**

- 1 SSH to the DRaaS Connector virtual appliance.

To SSH, you can obtain the DRaaS Connector IP address for an on-premises protected site in the VMware Cloud DR UI under **Sites > Protected Sites**.

To obtain the connector admin password, go to Protected Sites, select the site and under the Connectors tile select the more options menu and click **Show password**.

- 2 Run the following command:

```
connector-name>> drc register-vcenter --vcenter 10.80.11.235 --username
drc_user@vsphere.local --password drcShlep$2!
```

**Example**

```
connector-name>> drc register-vcenter --vcenter 10.80.11.235 --username
drc_user@vsphere.local --password drcShlep$2!
```

**What to do next**

You can now start creating protection groups on your protected on-premises vSphere site.

**Reregister vCenter Server**

If you update the custom user password, or create a different restricted user, you must reregister vCenter Server using the DRaaS Connector CLI.

**Prerequisites**

You only have to register one DRaaS Connector with vCenter Server. Other DRaaS Connector at the site are automatically updated to be aware of the newly registered vCenter Server.

**Procedure**

- 1 SSH to the DRaaS Connector virtual appliance.

You can obtain the DRaaS Connector IP address for an on-premises protected site in the VMware Cloud DR UI under **Sites > Protected Sites**.

To obtain the connector admin password, go to Protected Sites, select the site and under the Connectors tile select the more options menu and click **Show password**.

- 2 Run the following command, adding a `--reregister` flag to the `register-vcenter` command:

```
connector-name>> drc register-vcenter --reregister --vcenter 192.10.11.2 --username
drc_user@vsphere.local --password drcShlep$2!
```

**DRaaS Connector CLI Commands for Registering/Re-registering vCenter Server**

The DRaaS Connector CLI enables you to register or reregister vCenter Server with a restricted user.



## Register vCenter Command

This command registers vCenter Server with the DRaaS Connector.

```
drc register-vcenter
```

Options	Description
--vcenter	IP address of the vCenter Server to register with the DRaaS Connector.
--username	The user name of the restricted user you created for registering vCenter Server with the DRaaS Connector.
--password	The password for the user.

## Reregister vCenter Command

This command reregisters vCenter Server with the DRaaS Connector and associates it with the given user. This command is for users who previously registered vCenter Server with the DRaaS Connector, but now want to use a different user to reregister the DRaaS Connector with vCenter Server. Or, if you change a restricted user's password, you must reregister it.

```
drc reregister-vcenter
```

Options	Description
--vcenter	IP address of the vCenter Server to reregister with the DRaaS Connector.
--username	The user name of the new user you created to reregister the DRaaS Connector with vCenter Server.
--password	The password for the user.

## Refresh vCenter Server User Credentials

If you registered the DRaaS Connector with vCenter Server using a restricted user and then update the user's password, refresh the user credentials after making the update.

If you used a vCenter Server Administrator user account to register the DRaaS Connector with vCenter Server, you do not have to refresh the Administrator user's credentials.

### Procedure

- 1 From the left navigation, select Protected sites, and then choose a protected site.
- 2 On the protected site page, under vCenter click the small menu and click **Refresh Credentials**.
- 3 In the **Refresh credentials** dialog box, enter the username and password. If the user does not have sufficient [Privileges for a Restricted vCenter Server User](#), an error displays.
- 4 Click the **Refresh** button.

## Power the DRaaS Connector On and Off

You can power the DRaaS Connector on and off from the vCenter Server UI.

You can power the DRaaS Connector virtual appliance on and off from the protected site vCenter Server.

### Procedure

- 1 On the protected site vCenter Server, select the DRaaS Connector virtual appliance, right-click and select **Power** → **Power on**.

---

**Note** Always power on and power off the DRaaS Connector from the vCenter Server that the connector is protecting.

---

- 2 To power it off, select the DRaaS Connector virtual appliance, right-click and select **Power** → **Power off**.

## DRaaS Connector Connectivity Check

The DRaaS Connector CLI allows you to check network connectivity between the DRaaS Connector and the Orchestrator, cloud file system, and auto-support server, and the protected site vCenter Server and ESXi hosts.

### DRaaS Connector Connectivity Check CLI

The connectivity checker can:

- Determine if the connector can communicate successfully with the Orchestrator, the cloud file system, and autosupport server (cloud scope).
- Determine if the connector can communicate successfully with registered vCenter Servers and its ESXi hosts on the protected site (local scope).
- Obtain the IP address or FQDN, and port number of unreachable entities, and error information if an entity is unreachable.

To check connectivity from the DRaaS Connector to other entities in your environment, run the following DRaaS Connector CLI command on the virtual appliance:

```
drc network test
```

You can use the following options with the command:

Option	Description
<code>--scope [all   cloud   local]</code>	<p>Sets the scope of the check, which can be one of three options:</p> <ul style="list-style-type: none"> <li>■ <code>all</code>: Checks connectivity to both local and cloud entities.</li> <li>■ <code>cloud</code>: Checks connectivity to the Orchestrator, the cloud file system, and the autosupport server.</li> <li>■ <code>local</code>: Checks connectivity to vCenter Server or ESXi hosts on the protected site where you deployed the DRaaS Connector.</li> </ul>
<code>--orchestrator</code>	<p>Fully Qualified Domain Name (FQDN) of the Orchestrator. You must supply a password for authentication.</p> <p>This option is required only if the DRaaS Connector is not configured.</p>
<code>--vcenter [vcenter server IP address/FQDN]</code>	<p>For a local scope check, use this option to check connectivity to a registered vCenter Server using its IP address or FQDN.</p> <p>You must add the <code>--vcenter-user</code> option, which requires password authentication after you run the command.</p>
<code>--vcenter-username</code>	<p>When performing a local connectivity to a registered vCenter Server, you must supply the vCenter Server username. The command prompts you for the password of this user.</p>
<code>--output-format [JSON   default]</code>	<p>Choose the output format (optional):</p> <ul style="list-style-type: none"> <li>■ <code>json</code>: Format output in the JSON format.</li> <li>■ <code>default</code>: Print output using the default formatter.</li> </ul>

## Check Connectivity to All Entities

SSH to the DRaaS Connector virtual appliance and run this command to check connectivity from the DRaaS Connector to all entities (cloud and local) in your environment:

```
drc network test --scope all
```

In this example output, all VMware Cloud DR cloud entities are reachable by the DRaaS Connector, while the protected site vCenter Servers are unreachable. A comment section for the unreachable vCenter Servers indicates the cause of failure.

```
connector-name>># drc network test --scope all
FQDN/IP                               Port  Cloud  EntityType
Status      Comment
192.168.139.207      80    False  VCENTER
UNREACHABLE Cannot complete login due to an incorrect user name or password.
192.168.139.207      443   False  VCENTER
UNREACHABLE Cannot complete login due to an incorrect user name or password.
35.84.127.181        443   True   SCFS
REACHABLE    --
autosupport.datrium.com 443   True   SUPPORT
```

```

REACHABLE    --
vcdr-56-188-244-16.staging.app.vcdr.vmware.com      443  True    ORCHESTRATOR
REACHABLE    --
vcdr-56-188-244-16.staging.app.vcdr.vmware.com      443  True    ORCHESTRATOR
REACHABLE    --

```

## Check Connectivity to Cloud Entities

Run this command to check connectivity from the DRaaS Connector to the Orchestrator, the cloud file system, and the Auto-support server:

```
connector-name>># drc network test --scope cloud
```

This example output shows that all VMware Cloud DR entities are reachable by the DRaaS Connector:

```

drc network test --scope cloud
FQDN/IP                               Port  Cloud  EntityType
Status      Comment
35.84.127.181                               443  True   SCFS
REACHABLE    --
autosupport.datrrium.com                   443  True   SUPPORT
REACHABLE    --
vcdr-52-184-288-53.staging.app.vcdr.vmware.com 443  True   ORCHESTRATOR
REACHABLE    --
vcdr-52-184-288-53.staging.app.vcdr.vmware.com 443  True   ORCHESTRATOR
REACHABLE    --

```

## Check Connectivity to Local Entities

Run this command to check connectivity from the DRaaS Connector to a protected site vCenter Server and its ESXi hosts:

```
connector-name>># drc network test --scope local
```

This example output shows that all local entities are reachable by the DRaaS Connector:

```

drc network test --scope local
FQDN/IP                               Port  Cloud  EntityType
Status      Comment
192.168.139.207                               80  False  VCENTER
REACHABLE    --
192.168.139.207                               443  False  VCENTER
REACHABLE    --
192.168.139.211                               902  False  ESXI
REACHABLE    --

```

## Check Connectivity to the Orchestrator and Cloud Entities

Run the following command if the DRaaS Connector is not configured but you want to check connectivity from the DRaaS Connector to the Orchestrator, and other cloud entities. You must provide the Orchestrator IP address or FQDN.

The CLI prompts you for the Orchestrator password.

```
drc network test --scope cloud --orchestrator vcdr-192-168-21-0.staging.app.vcdr.vmware.com
Enter the site-specific Orchestrator password:
```

**Note** For more information on getting the Orchestrator IP address or FQDN and passcode, see [Service Public IP Addresses](#).

This example output shows that all VMware Cloud DR entities are reachable:

FQDN/IP	Port	Cloud	EntityType
Status	Comment		
35.84.127.181	443	True	SCFS
REACHABLE	--		
autosupport.datrium.com	443	True	SUPPORT
REACHABLE	--		
vcdr-54-184-246-36.staging.app.vcdr.vmware.com	443	True	ORCHESTRATOR
REACHABLE	--		
vcdr-54-184-246-36.staging.app.vcdr.vmware.com	443	True	ORCHESTRATOR
REACHABLE	--		

## Check Connectivity to vCenter Server

Run the following command to check connectivity from the DRaaS Connector to a registered vCenter Server and its ESXi hosts. You must provide the registered vCenter Server IP address or FQDN and a vCenter Server username. The CLI prompts you for the vCenter Server user password.

```
drc network test --vcenter 192.168.139.207 --vcenter-username Administrator@datrium.local --
scope local
Enter password for vCenter login:
```

This example output shows that all local vCenter Server entities are reachable:

FQDN/IP	Port	Cloud	EntityType	Status	Comment
192.168.139.207	80	False	VCENTER	REACHABLE	--
192.168.139.207	443	False	VCENTER	REACHABLE	--
192.168.139.211	902	False	ESXI	REACHABLE	--

## DRaaS Connector Performance Check

You can use the DRaaS Connector CLI to view the performance of snapshot replication and restore operations.

You can run the `drc perf` command during snapshot replication or failback operations to display logical and physical throughput between the DRaaS Connector and the cloud file system during these two operations. The CLI also displays how much time it takes to perform snapshot replication or failback operation.

To run a performance check for the DRaaS Connector, run the following command:

```
drc
perf

        usage:
        -h, --

help

        --output-format
OUTPUT_FORMAT
```

You can use the following options with this command:

Option	Description
-h, --help	Displays help information about the performance check CLI.
--output-format OUTPUT_FORMAT	Choose the output format (optional): <ul style="list-style-type: none"> <li>■ json: Format output in the JSON format.</li> <li>■ default: Print output using the default formatter.</li> </ul>

When you run the `drc perf -h` command with the help option, the output displays the following information to describe the results.

**Table 7-2. Backup (Snapshot Replication)**

Category	Description
<b>Throughput</b>	
<b>Physical throughput</b>	The rate at which compressed data is successfully transferred between the DRaaS Connector and a cloud file system during snapshot replication, measured in megabits per second (Mbps).
<b>Logical throughput</b>	The rate at which raw data is successfully transferred between a DRaaS Connector and a cloud file system during snapshot replication, measured in megabits per second (Mbps).
<b>Connector Time Taken (Percentage)</b>	
<b>Read</b>	Percentage of time spent by copy sectors in Read operations.
<b>Compress</b>	Percentage of time spent by copy sectors in compress operations.
<b>Write</b>	Percentage of time spent by copy sectors in write-wait operations.

Table 7-2. Backup (Snapshot Replication) (continued)

Category	Description
<b>SCFS Time Taken (Percentage)</b>	
<b>Write</b>	Percentage of time spent by the cloud file system in write operations.
<b>Uncompress</b>	Percentage of time spent by the cloud file system in uncompress operations.

Table 7-3. Restore (Failback Operation)

Category	Description
<b>Throughput</b>	
<b>Physical</b>	The rate at which compressed data is successfully transferred between the cloud file system and the protected site during a restore operation, measured in megabits per second (Mbps).
<b>Logical</b>	The rate at which raw data is successfully transferred between the cloud file system and the protected site during a restore operation, measured megabits per second (Mbps).
<b>Connector Time Taken (Percentage)</b>	
<b>Read</b>	Percentage of time spent by copy sectors in Read operations.
<b>Uncompress</b>	Percentage of time spent by copy sectors in uncompress operations
<b>Write</b>	Percentage of time spent by copy sectors in write-wait operations.
<b>SCFS Time Taken</b>	
<b>Read</b>	Percentage of time spent by the cloud file system in read operations.
<b>Compress</b>	Percentage of time spent by the cloud file system in compress operations.

For example, if you run the `drc perf` command during snapshot replication, you see the following output:

```
# drc perf
2022-09-21T20:09:35.054934 UTC
----- Backup -----
----- Restore -----
-----

Throughput (Mbps)   Connector Time Taken (%)   SCFS Time Taken (%) |Throughput (Mbps)
Connector Time Taken (%)   SCFS Time Taken
(%)
```

Physical Read Compress	Logical Uncompress Write	Read	Compress Read	Write	Write	Uncompress	Physical	Logical
-----	-----	-----	-----	-----	-----	-----		-----
-----	-----							
-----								
--	--	--	--	--	--	--		--
--	--	--	--	--	--	--		--
--								
--	--	--	--	--	--	--		--
--	--	--	--	--	--	--		--
--								
--	--	--	--	--	--	--		--
--	--	--	--	--	--	--		--
--								
--	19.1	26.8	42.9	11.7	0.0	0.9	0.1	
--	--	--	--	--	--	--	--	
--								
--	--	--	--	--	--	--	--	
--	--	--	--	--	--	--	--	
--								
--	24.8	26.8	1.6	0.1	0.0	0.3	0.1	
--	--	--	--	--	--	--	--	
--								
--	--	--	--	--	--	--	--	
--	--	--	--	--	--	--	--	
--								
--	53.5	78.9	4.4	0.2	0.0	0.1	0.2	
--	--	--	--	--	--	--	--	
--								
--	64.2	115.8	16.6	0.6	0.0	0.1	0.3	
--	--	--	--	--	--	--	--	
--								
--	1.4	154.4	23.7	0.5	0.0	0.2	0.4	
--	--	--	--	--	--	--	--	
--								
--	0.1	206.4	28.4	0.1	0.0	0.2	0.5	
--	--	--	--	--	--	--	--	

For example, if you run the `drc perf` command during a failback operation, you see something like this:

```
# drc perf
2022-08-25T07:55:45.050222 UTC
----- Backup -----
----- Restore -----
-----
```



Throughput (Mbps)		Connector Time Taken (%)			SCFS Time Taken (%)			Throughput (Mbps)	
		Connector Time Taken (%)			SCFS Time Taken				

# Using Google Cloud VMware Engine as a Protected Site



VMware Cloud DR supports protecting a Google Cloud VMware Engine private cloud as a protected site.

Before you begin setting up the Google Cloud VMware Engine private cloud as a protected site, you must do three things:

- Elevate the credentials of the Google Cloud VMware Engine user required to register vCenter with the DRaaS Connector.
- Obtain login credentials for that user.
- Obtain the IP address of the vCenter Server on Google Cloud VMware Engine private cloud.

For information about getting this information, see [Before You Add Google Cloud VMware Engine as a Protected Site](#).

Once you obtain this information, you can then perform these tasks to set up the protected site:

- [Set Up Protected Site for Google Cloud VMware Engine](#).
- [Download the DRaaS Connector OVA from VMware Cloud DR DR UI](#).
- [Deploy the DRaaS Connector Using vCenter Server UI](#).
- [Configure the DRaaS Connector VM from VM Console](#).
- [Register vCenter for Google Cloud VMware Engine Protected Site](#).

## Known Issue

The following issue applies when using Google Cloud VMware Engine private cloud as a protected site:

- **vCenter registration failure.** vCenter registration failed when using the vCenter Administrator user on Google Cloud VMware Engine private cloud protected site because the Google Cloud VMware Engine `CloudOwner@gve.local` user did not have its privileges elevated.

---

**Workaround** To ensure that vCenter registration is successful, you must elevate the privileges of the Google Cloud VMware Engine private cloud `CloudOwner@gve.local` user, and then [Register vCenter for Google Cloud VMware Engine Protected Site](#) selecting the vCenter Administrator user option. To elevate the Google Cloud user privileges, see [Before You Add Google Cloud VMware Engine as a Protected Site](#).

---

Read the following topics next:

- [Before You Add Google Cloud VMware Engine as a Protected Site](#)
- [Set Up Protected Site for Google Cloud VMware Engine](#)
- [Download the DRaaS Connector OVA from VMware Cloud DR DR UI](#)
- [Deploy the DRaaS Connector Using vCenter Server UI](#)
- [Configure the DRaaS Connector VM from VM Console](#)
- [Register vCenter for Google Cloud VMware Engine Protected Site](#)

## Before You Add Google Cloud VMware Engine as a Protected Site

Before you can set up a Google Cloud VMware Engine private cloud as a protected site, you must perform a minor task and obtain key information from Google Cloud.

Before setting up a Google Cloud VMware Engine private cloud as a protected site, do the following:

- Elevate the privileges of the `CloudOwner@gve.local` user in Google Cloud. (For more information, see [Elevating VMware Engine privileges](#).)
- Obtain the login credentials for the `CloudOwner@gve.local` user in Google Cloud.
- Obtain the IP address of the vCenter Server on Google Cloud VMware Engine private cloud.

As a prerequisite, if you are using either static IP assignment or DHCP on your Google Cloud VMware Engine, you must include a Google Cloud VMware Engine internal DNS server when you [Configure the DRaaS Connector VM from VM Console](#).

### Procedure

- 1 Log in to your account on Google Cloud.

- 2 In the Google Cloud VMware Engine site, click **Resources** on the left navigation.
- 3 Under Private Clouds, click the private cloud you want to protect with VMware Cloud DR.
- 4 From the **Summary** tab, in the Basic Info section under Status, click **View** under vCenter login info. Copy the credentials of the `CloudOwner@gve.local` user, which you will use to register vCenter with the DRaaS Connector.
- 5 Next, in the Danger zone section, click **Elevate** to elevate the privileges of the `CloudOwner@gve.local` user. If you do not elevate the privileges for this user, then you cannot register vCenter with the DRaaS Connector.
- 6 Last, select the vSphere Management Network tab to obtain the IP address of vCenter server on Google Cloud. Copy the IP address of the vCenter Server Appliance.
- 7 Exit Google Cloud.

#### What to do next

With this information, you can now [Set Up Protected Site for Google Cloud VMware Engine](#).

## Set Up Protected Site for Google Cloud VMware Engine

When you set up Google Cloud VMware Engine private cloud vCenter as a protected site, you choose to set it up as an on-premises protected site.

In this task, you will set up your Google Cloud VMware Engine private cloud vCenter as an on-premises protected site.

#### Procedure

- 1 From the left navigation, select **Protected sites**.
- 2 Click the **Set up protected site** button.
- 3 In the **Setup protected site** dialog box, under Site types select On-premises site.
- 4 Under Connection to cloud, select public internet.
- 5 Under Cloud backup, select a cloud file system to use for backups from the protected SDDC. If a cloud file system is already deployed, then it is selected.
- 6 Select a time zone from the drop-down menu, and then click the button to the right to set the time zone for the protected site.
- 7 Enter a name for the on-premises site.
- 8 Click **OK**.

#### What to do next

Next, you are ready to [Download the DRaaS Connector OVA from VMware Cloud DR DR UI](#)

## Download the DRaaS Connector OVA from VMware Cloud DR DR UI

Setting up a protected site requires that you download the DRaaS Connector connector OVA from the VMware Cloud DR UI.

Using the VMware Cloud DR UI, you can copy the URL to download the OVA into your environment in one of two ways:

- Navigate to **Settings > About VMware Cloud DR**. You can copy the VMware Cloud DR OVA URL here.
- You can get the same connector OVA URL from the **Download connector** dialog box, which is described below.

### Procedure

- 1 In the VMware Cloud DR UI, click **Sites > Protected sites** and then click the protected site on the left side of the application.
- 2 Under Connectors, click **Deploy**. If this protected site is an SDDC, the **Deploy** button is under Clusters.
- 3 In the **Download connector** dialog box, there is a list of steps that guide you in deploying the connector, as well as the URL to the connector OVA, with an option to download it locally to your system. If the site you are protecting is a VMware Cloud on AWS SDDC, for more information about networking considerations see [DRaaS Connector Firewall Rules for a VMware Cloud on AWS Protected SDDC](#).
- 4 Click the **Copy** button to copy the download URL. You need this URL when you deploy the OVA in vSphere, after you download the connector.
- 5 Make a note of the Console credentials, which you need to log in to the VM console: `admin/vmware#1`.
- 6 Also copy (or write down) the Orchestrator Fully Qualified Domain Name (FQDN), which you need when you configure the connector in the VM console.
- 7 Click **OK**.

### What to do next

Your next task in setting up a protected site is to [Deploy the DRaaS Connector Using vCenter Server UI](#).

## Deploy the DRaaS Connector Using vCenter Server UI

When you deploy the DRaaS Connector from the vCenter Server UI, you select the host, cluster, and resource pool for the connector.

Before you begin:

- Do not name the DRaaS Connector VM using the same naming conventions you use to name VMs in your vSphere environment. Avoid giving the connector VM a name that might match the VM name pattern you use when you define protection groups.
- If you are deploying the DRaaS Connector to a VMware Cloud on AWS SDDC with more than one cluster, you must choose a cluster to deploy the connector VM on. Each cluster in your SDDC needs the connector VM deployed on it in order for the VMs running there to be added to protection groups and replicated to a cloud backup site.

---

**Note** If you are deploying the DRaaS Connector for a Google Cloud VMware Engine protected site, see [Google Cloud documentation](#) for information about logging in to vSphere from Google Cloud.

---

#### Procedure

- 1 In vSphere, select any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host.
- 2 Right-click the object and select **Actions** → **Deploy OVF Template**.
- 3 Click **Next**.
- 4 In the **Deploy OVF Template** dialog box, Step 1, Select an OVF template, paste the connector OVA URL into the URL field. The exact URL to download the connector OVA displays in the **Download Connector** dialog box. For example: `https://<vmware-cloud-dr-ip-address>/cloud-connector.ova`.
- 5 Click **Next**.
- 6 Next, specify a name for the connector. Do not use non-ASCII characters for the connector name. Use a name that is different than the naming conventions you use to name VMs in your vSphere environment, to avoid this VM being included in a snapshot.
- 7 Below the name, choose a location for the connector and then click **Next**.
- 8 Select a compute resource for the connector. If this vSphere is a VMware Cloud on AWS SDDC with more than one cluster, choose a cluster to deploy the connector VM on. For an SDDC, each cluster you want to protect must have the DRaaS Connector VM deployed on it.
- 9 Click **Next**.
- 10 Review the details for your connector deployment, then click **Next** to select storage for the connector VM.
- 11 Select a datastore for the connector and then click **Next**.
- 12 Select the network to use for the connector, and then click **Next** to review the deployment details.
- 13 Click **Finish**. You can now find the connector VM in your vSphere client.

- 14 Click on **Edit settings** and select Virtual Hardware. Enter the CPU and memory reservation as described in [System and Network Requirements for the DRaaS Connector](#).

### What to do next

Now you are ready to [Configure the DRaaS Connector VM from VM Console](#).

## Configure the DRaaS Connector VM from VM Console

Before you configure the DRaaS Connector VM using the VM console, make sure you have all required information by filling out the following worksheet (optional):

Parameters	Value
Console credentials	admin/vmware#1
<p><b>Note</b> This password will change at the end of the configuration, and you can obtain the new password in the VMware Cloud DR UI.</p>	
IP address allocation: Static or DHCP	<p>If you deployed the DRaaS Connector VM on an SDDC, we recommend that you use the default DNS server settings that is associated with your SDDC.</p> <p>If you choose (a) Static, then enter the following</p> <ul style="list-style-type: none"> <li>■ IP address: If you are configuring a static IP address for the connector VM, when asked for the DNS server use the default DNS server created when the SDDC was first deployed. You can find this setting under the SDDC <b>Network and Security &gt; DNS Service IP</b></li> <li>■ Subnet mask</li> <li>■ Gateway</li> <li>■ DNS servers</li> </ul> <p>If you are using a DHCP service, make sure you provide DNS servers that can resolve host names on the protected site internal network.</p> <p><b>Note</b> Public DNS servers are not supported, such as Google DNS servers (8.8.8.8 and 8.8.4.4), which do not work reliably with VMware Cloud DR.</p>
Orchestrator FQDN or IP address	<p>You can obtain the Orchestrator FQDN (or IP address) by navigating to <b>Settings &gt; About VMware Cloud DR</b>. Or, you can also obtain this information from the Deploy Connector appliance dialog box, as described <a href="#">Service Public IP Addresses</a>.</p> <p>If you have <a href="#">AWS Direct Connect</a> configured with a private VIF, then use the Orchestrator IP address</p> <p>.</p>

Parameters	Value
Temporary, site specific passcode	<p>This temporary passcode is used to configure the DRaaS Connector in the VM console.</p> <hr/> <p><b>Important</b> You can obtain a passcode in the Download connector VM dialog box from inside the VMware Cloud DR UI. Make sure that you open the dialog box from the specific site you want the DRaaS Connector to able to connect to.</p> <hr/>
Name to give the connector, as it will appear in the VMware Cloud DR UI.	

**Note** Use the vSphere web console for changing the connector IP address. Do not use an SSH session to configure the connector VM. You must use the vSphere web console if the network does not have DHCP enabled.

#### Procedure

- 1 In the vCenter UI select VMs, and then select the DRaaS Connector VM. Right-click and choose **Power** → **Power on**.

**Note** Always power on and power off the DRaaS Connector from the vCenter on the protected site.

- 2 Under the VM, click **Launch web console**.
- 3 When the console session is open, log in to the connector VM console using the following credentials: admin/vmware#1.
- 4 Next, in the Select the network address IP address allocation, either (a) Static or (b) DHCP.

If a DRaaS Connector has already been configured on this site, the current network configuration displays in the console and the current values are set as defaults. To use the existing configuration in all of these steps, press the Enter key.

- 5 If you chose (a) for static IP address allocation: Enter an IP address, Subnet mask, and Gateway. Enter the IP address of one or more DNS servers (up to three supported) that are able to resolve host names on the internal network. To use the existing configuration in all of these steps, press the Enter key to accept the suggested default values.

**Note** Google DNS servers (8.8.8.8 and 8.8.4.4) do not work reliably with VMware Cloud DR. We advise you to use private, non-Google and non-public DNS servers when configuring the connector VM.

- 6 Next, enter the Orchestrator FQDN or IP address. If you have [AWS Direct Connect](#) configured with a private VIF, then you use the Orchestrator IP address.
- 7 Next, enter the DRaaS Connector temporary passcode.
- 8 Enter a name to identify the connector.



### What to do next

After you configure the DRaaS Connector, the console window will return you to the command prompt. If your protected site is an on-premises vSphere, you can now return to the VMware Cloud DR UI to [Register vCenter Server](#) for the protected site. (If your protected site is an SDDC, you do not need to register vCenter.) If your protected site is a Google Cloud VMware Engine private cloud, then see [Register vCenter for Google Cloud VMware Engine Protected Site](#).

## Register vCenter for Google Cloud VMware Engine Protected Site

After you deploy the DRaaS Connector to your Google Cloud VMware Engine private cloud vSphere and configure it, you must register it with vCenter.

### Prerequisites

For this task, you need the the login credentials for the `CloudOwner@gve.local` user in Google Cloud and the IP address of the vCenter Server on Google Cloud VMware Engine private cloud. You must have also elevated the privileges of the `CloudOwner@gve.local` user in Google Cloud. For more information, see [Before You Add Google Cloud VMware Engine as a Protected Site](#).

### Procedure

- 1 In the VMware Cloud DR UI, from the left navigation select Protected sites, and then choose the Google Cloud VMware Engine protected site.
- 2 On the protected site page, under vCenters click the **Register vCenter** button.
- 3 In the **Register vCenter** dialog box, enter the vCenter server's IP address.
- 4 Under Authentication, select Authenticate with vCenter administrator.
- 5 Enter the vCenter user name and password for the Google Cloud `CloudOwner@gve.local` user.
- 6 Click **Register**.

# Using Protection Groups

# 9

Use protection groups to create recurring VM snapshots and replicate them to a cloud file system, so they can later be used for disaster and ransomware recovery.

After snapshots replicate to a cloud file system, you can use those snapshots in recovery plans for disaster and ransomware recovery. When you configure a recovery plan, you can select protection groups that have scheduled replication to a cloud file system, and when you start a plan you can select snapshots for recovery.

## Snapshot Frequency

VMware Cloud DR provides two types of snapshots for protection groups, based on snapshot frequency:

Snapshot Frequency	Description
Standard-frequency snapshots	<p>Schedule recurring snapshots as frequent as every 4 hours.</p> <p>Standard-frequency snapshots also allow you to select the quiesce option. This option quiesces the guest operating of the system of a VM before taking a snapshot. Quiescing pauses or alters the state of running processes on the VM to guarantee a consistent state of any applications running at the time a snapshot is taken. The quiesce option is only available for standard-frequency snapshots, and only for powered on VMs with VMware Tools installed and running.</p> <p>Linux VMs also need pre-thaw and post-freeze scripts installed. For an example script, see <a href="#">Enabling Quiescing for Linux VMs</a>.</p> <p>To know if your VMs are compatible with snapshot quiescing, <a href="#">Run a Quiesce Compatibility Check</a>.</p>
High-frequency snapshots	<p>Schedule recurring snapshots as frequent as every 30 minutes, which requires that the on-premises protected site is running vSphere 7.0 Update 3 or the protected VMware Cloud on AWS SDDC is running version 1.16.</p> <p>If you are not sure if the hosts on your protected site are compatible with high-frequency snapshots see, <a href="#">Run a Host Compatibility Check for High-Frequency Snapshots</a>.</p> <p>To convert standard-frequency snapshots to high-frequency snapshots, open a standard-frequency snapshot and select the high-frequency snapshot option. Once you switch a protection group from standard-frequency snapshots to high-frequency snapshots, you cannot revert back.</p>

## App-consistent Snapshots with Quiescing

For VMs with VMware Tools installed, you can create protection groups that take quiesced snapshots. Quiescing pauses or alters the state of running processes on the VM to guarantee a consistent state of any applications running at the time a snapshot is taken. So when you restore the VM, you recover applications to the state they were in at the time the snapshot was taken.

---

**Note** High-frequency snapshots do not support quiescing.

---

Requirements for quiescing:

- VM is powered on.
- VMware Tools installed and running. VMware Tools requires Windows Volume Shadow Copy Service (VSS) or protection groups cannot take quiesced snapshots. Windows VMs require VMware Tools version 10.x and above.
- Linux VMs only: Pre-freeze and post-thaw scripts installed on the VM. VMware Tools must be version 10.2 or above.

## Dynamic Group Membership

A protection group query dynamically defines the protection group membership at the time a snapshot is taken. Protection groups provide three types of queries:

- **VM name pattern.** A VM name pattern is a string of characters that matches the names of VMs in your vSphere inventory, either for inclusion or exclusion in the protection group snapshot. Any VMs that match the pattern specified become included (or explicitly excluded from) the protection group for snapshots.
- **Folders.** You can add VM folders that are present in your vSphere inventory to a protection group, so that all VMs in those folders are included in snapshots. Folder selection does not include sub folders. To include sub folders, select them manually.

---

**Note** Protection groups do not support folder-based snapshots for VMs that are a part of vApp.

---

- **vSphere tags.** Use tags to define protection group membership. Any VMs that match the tags you specify are included in the protection group snapshot. You can select any tags defined in vSphere on the protected site. For successful failover operations, ensure that the selected tags also exist on the target Recovery SDDC vCenter, or the compliance check display warnings, and the failback operation fails.

---

**Note** Creating, deleting, and assigning vSphere tags on VMs are not immediately visible to protection groups. For example, if you create a tag and associate it with 10 VMs, a protection group might not immediately show the VMs associated with this tag. Typically, it can take up to 15 minutes for vSphere tags to appear in protection groups, but usually is much faster.

---



---

**Note** Protection groups do not support queries based on vSphere tags on protected sites running vSphere 6.0. If your protected site is running vSphere 6.0, you cannot create a tag query.

---

Before a protection group takes a snapshot, VMs that match any name patterns evaluate first, and then combine with any defined folder or tag queries.

If you use an exclusion name pattern in your query, it is possible that other Folder or tag queries defined in protection group might override the previously excluded name pattern. For example, if one of your queries excludes a VM by name, if that same VM lives inside a folder you have selected in a Folders query, that VM is included in the snapshot.

To verify the snapshot before a scheduled job, [Take a Manual Snapshot](#).

---

**Note** To create a VM name pattern to match VMs that have name using the following special characters, they must be escaped (prefixed with \ ) in the specified name pattern: ? , \* .

---

## Changing VM Protection Group Association

You might want to change a VM's protection group membership so it belongs to a second, new protection group. You can add a VM to a new protection group by defining the new protection group's queries (name, tag, folder) to include the VM.

Once a VM has been added to a new protection group and the first snapshot of that VM is taken in the new group, if the VM is currently a member of **both** protections groups and at least one snapshot of the VM exists in the original group, you can fail over and restore from any snapshot in the original or new protection group.

VMware Cloud DR tracks VM snapshots at the VM level, creating incremental backups of a VM even if it is added to a new protection group, for as long as there are snapshots of the VM available in the first or second protection group.

## Snapshot Retention

A protection group snapshot retention policy defines how long snapshots remain on the cloud file system. You can set retention for any duration of hours, days, weeks, months, or years.

---

**Note** A 90 day retention schedule might result in higher storage capacity consumption, which is charged as described on the [VMware Cloud DR pricing](#) page.

---

Read the following topics next:

- [Protection Group Caveats](#)
- [Video: Configuring Protection Group Policies](#)
- [Create a Protection Group](#)
- [Snapshots](#)
- [Edit a Protection Group](#)
- [Delete a Protection Group](#)
- [Throttle Replication](#)
- [Export VM-to-Protection Group Mappings](#)
- [Unprotect a VM](#)
- [Replication Progress Statistics](#)

## Protection Group Caveats

VMware Cloud DR protection groups are bound by specific caveats and restrictions.

---

**Note** These considerations apply to both standard and high-frequency snapshots. For caveats specific to high-frequency snapshots, see [Caveats for High-Frequency Snapshots](#).

---

- VMware Cloud DR does not support using protection groups for VMs with:
  - Shared disk clusters. Do not configure a protection group for any VMs that reside on a shared disk in a shared disk cluster.
  - VMs that use Integrated Drive Electronics (IDE) disks for storage
  - Independent disks and Raw Device Mappings (RDM)
  - Encrypted VMs
- **The members of a single protection group must share the same vCenter.** In other words, you cannot create a protection group that contains VMs from two different vCenters.
- **You cannot protect VMs that belong to two different vCenters but share the same VM instanceUUID.** An instanceUUID is a unique identifier that is used by vCenter Server to uniquely identify a Virtual Machine within a vCenter Server. If two VMs have the same instanceUUID, you cannot protect those VMs with VMware Cloud DR.
- **VMware Transit Connect not supported.** You cannot use VMware Transit Connect to send replication data traffic from your protected site to VMware Cloud DR.
- **Cannot clone DRaaS Connector VMs.** VMware Cloud DR does not support cloning DRaaS Connector VMs.
- **Datasets not supported for snapshot replication, single VM recovery, and failover/failback for vSphere 8 VMs.** VMware Cloud DR does not support capturing datasets in snapshot replication or VM recovery operations. The VMware vSphere 8.0 release provides a new datasets feature that only works on VMs with virtual hardware version 20 (hardware version 20 is new to vSphere 8.0).
- **Event logged when replicating a VM with datasets (not supported).** During snapshot replication, any dataset information in VMs is deleted, and an event is logged in VMware Cloud DR. If you attempt to recover, fail over, or fail back a VM with datasets, the missing dataset information is ignored. For example, if you attempt to fail over a VM with datasets, you see an event like this: "Datasets file of VM 'VM-DB' is ignored in snapshot 'PG-snapshot Daily Backups'..
- **VMware Cloud DR does not support protecting the same VM with two cloud file systems.** If a VM is being protected with high-frequency snapshots replicating to two cloud file systems, the replication task will fail with a 'PeAlreadyClaimed' error. If a VM is being protected with

standard-frequency snapshots replicating to two cloud file systems, the replication task will fail with no error given. This scenario can also occur if a VM belongs to one vCenter, and the vCenter is registered to two separate protected sites, and each protected site is replicating snapshots of that VM to different cloud file systems.

- **Only five snapshot replication tasks can be run concurrently across all protection groups (on all protected sites).** If more than five snapshot replication tasks are started at the same time, the first five tasks are started, and the remaining snapshot tasks are queued until the first five complete, and continues until all other tasks are completed.
- **VMware Cloud DR does not snapshot VM templates.** If you include a VM template in a protection group, the snapshot does not capture the template.
- **No changing vSphere configuration during snapshot replication.** Do not make any vSphere environmental or configuration changes when a snapshot is in progress, or you can experience accuracy issues with the snapshot.

## Video: Configuring Protection Group Policies

When you configure a protection group, you set the policies that define what VMs you want to snapshot and replicate to the cloud file system, the frequency at which you replicate the snapshots, and how long you want to retain them.

This short video illustrates the three main configurings you define to set your protection group policies:



(Setting protection group policies)

## Create a Protection Group

When you create a protection group, you configure the settings needed to enable successful disaster and ransomware recovery.

When you create a protection group, you select a snapshot frequency, choose whether to quiesce VMs before taking snapshots, define group membership (VMs) using dynamic queries, set the snapshot schedule, and configure the snapshot retention policy.

When you save the protection group, it begins taking snapshots of group members according to its snapshot schedule, with each snapshot replicating to a cloud file system. You can then use snapshots in recovery plans for disaster recovery and ransomware recovery operations.

A protection group only supports one snapshot type: standard-frequency OR high-frequency. You cannot mix snapshot frequency types in a protection group. Once a VM has a high-frequency snapshot taken of it, all subsequent snapshots taken of that VM are high-frequency, even if the protection group is configured to standard-frequency.

---

**Note** Do not configure a protection group for any VMs that reside on a shared disk in a shared disk cluster.

---

For more information, see [Chapter 9 Using Protection Groups](#).

Before you can create a protection group, you must first:

- [Chapter 5 Deploy a Cloud File System](#).
- [Chapter 6 Set Up Protected Sites](#).
- [Chapter 7 Deploy the DRaaS Connector](#).

#### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 From the upper-right, click the **Create protection group** button.
- 3 In the **Create protection group for site** dialog box, under Protection group name enter a name for the group.
- 4 Next, select a protected site and a vCenter instance you want to take snapshots from.
- 5 After you select a protected site and vCenter, the protection group automatically checks the protected site for compatibility with the high-frequency snapshot feature. The compatibility check has three possible outcomes:
  - If the protected site vCenter and hosts are compatible, then the Use high-frequency snapshots option is selected and you can start defining protection group membership (step 6).
  - If the protected site vCenter and hosts are not compatible, then this protection group cannot take high-frequency snapshots. In this case, you can leave the Use high-frequency snapshot option deselected. Or, you can upgrade the vCenter Server and ESXi hosts on the protected site.
  - If the protected site you select is "mixed," it means some hosts are compatible with high-frequency snapshots and some are not. In this case, you can [Run a Host Compatibility Check for High-Frequency Snapshots](#) to find hosts and VMs that are not compatible.
- 6 To use quiesced snapshots, deselect the High-frequency snapshots option and then select the Quiesce option. If you are not sure if your protected site VMs are compatible, you can [Run a Quiesce Compatibility Check](#).
- 7 Next, define the dynamic membership of the protection group using VM name patterns, tags and/or VM folders in a vCenter query. The protection groups will contain all VMs that match the queries.



- 8 Under Group membership, click the **VM name pattern**, **Tags**, or **Folder** buttons.
- 9 For a **VM Name Pattern** query, enter a VM name pattern that is evaluated before a snapshot is taken. You can also enter a name pattern in the query in the Excluding text box for exclusion. (If there is already one vCenter query, then from the Add vCenter query drop-down menu select **VM name pattern**.)
- 10 For a folder query, click the **Select folders** button. In the **vCenter folders** dialog box, search the list of folders in your vCenter and click one to add it. (If there is already one vCenter query, then use the Add vCenter query drop-down menu and select Folders.) Folder selection does not include sub folders, so you must select subfolders specifically.

---

**Note** Protection groups do not support folder-based snapshots for VMs that are a part of vApp.

---

- 11 For a tag query, click the **Tags** button. In the vSphere tags dialog box, select tags to define protection group membership. Any VMs with selected tags are included in the protection group snapshots. Click **OK** when you finish.
- 12 Click **Preview VMs** to see VMs that match the queries.
- 13 Click **Next**.
- 14 To set the snapshot schedule, click **New Schedule**.
- 15 On the Schedule page, edit the snapshot schedule. A protection group can have a maximum of 10 schedules associated with it. Select snapshot frequency based on the following intervals:
  - Standard-frequency snapshots allow you to set snapshots schedules as frequently as every 4, 6, 8, or 12 hours, daily, weekly, or monthly.
  - With high-frequency snapshots, you can set snapshots schedules every 30 minutes, every hour, 2, 4, 8, and 12 hours, and daily, weekly, or monthly.

---

**Note** Protection group snapshots can exhibit errors if your environment is running other VM backup software that runs at the same time as the VMware Cloud DR snapshot schedule. If a VM is currently being backed up by VMware Cloud DR by a protection group snapshot, and the same VM is also being backed up by other software that uses vSphere Storage APIs for Data Protection (VADP), the snapshot is still taken, but that VM is not included in the snapshot.

---

- 16 Next, set the snapshot retention for this protection group. You can select any number and then select the duration (hours, days, weeks, months, years).

For effective ransomware recovery, create a snapshot retention schedule that is at least 90 days. A 60 day retention schedule might result in higher storage capacity consumption, which is charged as described on the [VMware Cloud DR pricing](#) page.

- 17 Optionally, you can add a custom name for the snapshot schedule. Click once in the default schedule name field above the schedule, type a custom a name, and then press Enter on your keyboard.
- 18 When you have finished adding snapshot scheduled and retention, click **Finish**. When the protection group appears in the Protection groups list, you can add it to a recovery plan for testing and disaster and ransomware operations.

#### What to do next

After you create a protection group, you can now check the protection group's [Protection Group Health Status](#) and make sure the replication schedule is active.

## VM Name Pattern

To define dynamic group membership when configuring a protection group, you can use VM name patterns (for inclusion and exclusion).

These VM name patterns are used at the time of taking snapshots to determine protection group VM membership.

A VM name pattern is a string of characters. VMware Cloud DR uses glob syntax matching for VM name patterns. For example, the pattern `windows[1-5]` matches the set of virtual machines with names in the sequence of `windows1..windows5`. The search is case-insensitive.

---

**Important** A protection group evaluates VM name patterns first, and then combines those results with folders queries before the protection group takes a snapshot. Keep in mind that if you use an exclusion name pattern in your vCenter query for a protection group, it is possible that other folder queries defined in protection group might override the previously excluded name pattern.

---



---

**Note** If you want to create a VM name pattern to match VMs that are named using the following special characters, they must be escaped (prefixed with `\`) in the specified name pattern: `?` , `*` . .

---

The following table provides a brief description of the special characters that you can use in a virtual machine name pattern.

Virtual Machine Name Patterns – Special Characters	Description
<p>*</p>	<p>A single asterisk matches any string of zero or more characters, excluding the slash character ( / ).</p> <p>For example:</p> <p><code>windows*</code></p> <p>matches all virtual machines with names that start with the string <code>windows</code>.</p>
<p>?</p>	<p>A question mark matches any single character, excluding the slash character / .</p> <p>For example, the pattern:</p> <p><code>linux?</code></p> <p>matches any virtual machine with a name that starts with the string <code>linux</code> followed by any single character.</p>
<p>[ ]</p>	<p>Square brackets enclose a set of characters.</p> <p>The brackets specify a match of exactly one character out of the set. The character set can include one or more ranges of characters.</p> <p>For example:</p> <p><code>linux[0-9]</code></p> <p>matches anything in the datastore root that begins with the string <code>linux</code> and ends with a single digit.</p>

There are two types of virtual machine name patterns you can set in the protection group configuration:

- Inclusion patterns. VMs to be selected. You must specify at least one inclusion pattern.
- Exclusion patterns. VMs to be excluded from the results of inclusion pattern selection. An exclusion pattern begins with the exclamation character !.

The order of patterns in a pattern list does not matter. A pattern list uses a comma , as a delimiter.

## Protection Group Health Status

Protection group health status indicates if a replication task was successful or encountered failures.

You can view the status of a protection group from the Protection groups list. A protection group can have one of the following health statuses:

Status	Meaning
OK	Protection group snapshot (and all VMs) successfully replicated to the target cloud file system.
Warning	<p>Snapshot replication task encountered some failures. A warning can mean that some or all of the VMs in the snapshot did not successfully replicate to the target cloud file system.</p> <p>If you see a warning status for a protection group, select the <b>Events</b> tab and select the Protection filter to look for the event.</p>
Critical	<p>Snapshot replication task encountered failures and no VMs were snapshotted.</p> <p>If you see a critical status for a protection group, select the <b>Events</b> tab and select the Protection filter to look for the event.</p> <p>For those snapshot jobs that experienced errors, you can view <a href="#">Snapshot Events and Logs</a> to help troubleshoot any issues.</p>

You can view protection group health and snapshot schedule (active or inactive) from the protection group list:

Protection groups <span>CREATE PROTECTION GROUP</span>									
Protection group	Site	Cloud file system	Status	Schedule	Frequency	Quiesce	Size	Last snapshot	VMs
1 VM	tfw-vcsa1	RWR_CDVX	Good	Active	Standard	No	613.6 MiB	Oct-03 12:38 am (12h a...	1 VM
100 linux vms	tfw-vcsa1a	RWR_CDVX	Good	Active	Standard	No	182.6 MiB	Oct-03 12:23 am (12h a...	101 VMs
100 linux VMs	tfw-vcsa1a	RWR_CDVX	Good	Active	Standard	No	327.2 GiB	Oct-03 12:44 am (12h a...	101 VMs
100 windows	tfw-vcsa1a	RWR_CDVX	Critical	Active	High	No	7.0 TiB	Oct-03 12:46 am (12h a...	102 VMs
c4-w2-hs3-q0812...	tfw-vcsa1a	RWR_CDVX	Good	Active	Standard	No	4.8 GiB	Oct-03 12:01 am (13h a...	1 VM

## Snapshots

You can schedule recurring VM snapshots to the cloud file system, which are later used to failover VMs to a recovery SDDC for disaster and ransomware recovery.

You can view snapshots for a [Create a Protection Group](#) by navigating to **Protection Groups** view and then selecting one of the snapshot in the group. A protection group can take two types of snapshots, standard-frequency or high-frequency.

**Note** A protection group can have only one snapshot type.

- With [Standard-frequency Snapshots](#), you can set snapshot schedules as frequently as every 4, 6, 8, or 12 hours, or Daily, Weekly, or Monthly. You can also choose to quiesce standard-frequency snapshots. To find out if you can quiesce standard-frequency snapshots, see [Run a Quiesce Compatibility Check](#).

- **High-frequency Snapshots** provide the ability to set snapshot schedules as frequently as 30 minutes, hourly, every 2, 4, 6, 8, or 12 hours, or Daily, Weekly, or Monthly. To determine if you can take high-frequency snapshots, see [Run a Host Compatibility Check for High-Frequency Snapshots](#).

## Snapshots Details

For every snapshot, you can view the following details:

Detail	Description
Name	Name of the VM included in the snapshot.
Origin vCenter	Protected site vCenter where the VM snapshot resides.
Snapshot type	Either standard-frequency or high-frequency.
Transferred bytes	The logical size of transferred data from the protected site to the cloud file system in a snapshot. The actual physical size of transferred data might be smaller due to compression.
Change rate	<p>The logical VM data change rate since the previous snapshot.</p> <hr/> <p><b>Note</b> For the first snapshot of a VM after a product upgrade, 'Change Rate' is not reported.</p> <hr/> <p>Change rate can be useful for detecting anomalous snapshot behaviors. For example, a very high change rate outside of normal expectations might indicate an unplanned encryption event related to a ransomware attack.</p>

For Transferred bytes, these two caveats apply:

- **Transferred Bytes and VM changes during snapshot initial seeding.** When using standard -frequency snapshots, any failures or interruptions at the time of initial seeding can impact 'Transferred bytes' as reported in the UI, so the value might be lower or larger than the logical size of the VM. When using high-frequency snapshots, if there are changes on the source VM while initial seeding is in progress, then the reported Transferred bytes might be larger than the logical size of the source VM.
- **Transferred Bytes for snapshot replication different depending on VM disk provisioning.** When a new disk is added to a VM that is being protected by high-frequency snapshots, the Transferred Bytes value for the snapshot task depends upon the type of datastore and provisioning of the disk. For example, if the disk is on vSphere Virtual Machine File System (VMFS) datastore and is provisioned “lazy zeroed,” then Transferred Bytes counts only the bytes that are allocated on the disk. If the disk is provisioned “eager zeroed” then all the bytes are accurately reports in Transferred Bytes. (For more information, see [About Virtual Disk Provisioning Policies](#).) When increasing the size of disk, Transferred Bytes might not show the extended size as Transferred Bytes, depending on the method used to increase the disk capacity.

## Standard-frequency Snapshots


With standard-frequency snapshots, you can schedule recurring snapshots as frequent as every four hours.

### Standard-frequency snapshots and VADP

If your environment is running backup software that uses VMware APIs for Data Protection (VADP), then that software's processes might conflict with replication as follows:

- If a VM is being backed up another backup solution, and VMware Cloud DR starts a snapshot of the same VM after the other backup software started the backup, then the VMware Cloud DR snapshot might fail if the backup software finishes the backup before finishing replication of the VMware Cloud DR snapshot. In this case, even though the snapshot failed, the next scheduled snapshot will resume from where it failed.
- If a VM snapshot is being replicated by VMware Cloud DR, and another backup solution (non-VMware Cloud DR) starts a backup of the same VM after VMware Cloud DR started replication, then the backup might fail if VMware Cloud DR finishes the replication of snapshot before the backup software finishes the backup.

For example:



### App-consistent Snapshots with Quiescing

Standard-frequency snapshots also allow you to quiesce the guest operating of the system of a VM before taking a snapshot. Quiescing pauses or alters the state of running processes on the VM to guarantee a consistent state of any applications running at the time a snapshot is taken.

To convert standard-frequency snapshots to [High-frequency Snapshots](#), open a standard frequency snapshot and select the high-frequency snapshot option. Once you switch a protection group from standard frequency snapshots to high-frequency snapshots, you cannot revert back. If you are not sure if the hosts on your protected site are compatible with high-frequency snapshots see, [Run a Host Compatibility Check for High-Frequency Snapshots](#).

---

**Note** High-frequency snapshots do not support quiescing.

---

Requirements for quiescing:

- VM is powered on.
- VMware Tools installed and running. VMware Tools requires Windows Volume Shadow Copy Service (VSS) or protection groups cannot take quiesced snapshots.
- Linux VMs only: Pre-freeze and post-thaw scripts installed on the VM. VMware Tools must be version 10.2 or later.

The quiesce option is only available for standard-frequency snapshots, and only for powered on VMs with VMware Tools installed and running. To know if your VMs are compatible with snapshot quiescing, [Run a Quiesce Compatibility Check](#). Linux VMs also need pre-thaw and post-freeze scripts installed to enable quiescing. For an example script, see [Enabling Quiescing for Linux VMs](#).

## Limitation for Standard-frequency Snapshots

Protecting a VM with VMware Tools ApplInfo plug-in enabled is not supported. If you are protecting a VM with VMware Tools installed with ApplInfo enabled, collected application data is stored in the VMX file of the VM, causing it to expand to >55 KB. Because VMware Cloud DR does not allow backup of VMX files that are >55 KB, snapshot replication fails. To avoid this situation, deactivate the VMware Tools ApplInfo plug-in on the VM.

## Run a Quiesce Compatibility Check

You can run a quiesce compatibility check to determine if VMs are compatible with quiesced snapshots.

Requirements for quiescing:

- VM is powered on.
- VMware Tools installed and running. VMware Tools requires Windows Volume Shadow Copy Service (VSS) or protection groups cannot take quiesced snapshots. Windows VMs require VMware Tools version 10.x and above.
- Linux VMs only: Pre-freeze and post-thaw scripts installed on the VM. VMware Tools must be version 10.2 or above.

---

**Note** VMs with high-frequency snapshots cannot be quiesced.

---

### Procedure

- 1 From the left navigation, select **Protection groups**.

- 2 From the Protection groups page, click **Create protection group**. Or, select an existing protection group and click the **Edit group** button.
- 3 In the dialog box, enter a name for the protection group.
- 4 Next, select a protected site and a vCenter. The dialog box also displays the cloud file system associated with the protected site, which is where snapshots replicate to.
- 5 After you select a protected site and vCenter, the system checks the vCenter for software compatibility with high-frequency snapshots. If the protected site is compatible, the High-frequency option is selected. Deselect the High-frequency option.
- 6 Next, select the Quiesce option.
- 7 To check if the VMs on the protected site have VMware Tools installed to enabled quiesced snapshots, create any type of query (name pattern, tags, folders).
- 8 Click **Quiesce compatibility check**.

When the check finishes, there are three possible states for each VM:

- Question mark icon. VMware Tools is not running. Quiescing not supported.
- Warning icon. VMware Tools is out of date. Quiescing might fail.
- Green light icon. VMware Tools is installed and running. Quiescing supported.

---

**Note** The quiesce compatibility check does not check if a VM has a high-frequency snapshot associated with it. If a VM already has a high-frequency snapshot taken, then that VM cannot be quiesced, even if it has VMware Tools installed. The quiesce compatibility check also does not check if Linux VMs have pre-freeze and post-thaw scripts installed, which is required for quiescing Linux VMs.

---

- 9 If the check passes, you can proceed in creating the protection group.

## Enabling Quiescing for Linux VMs

To enable quiescing for Linux VM snapshots, you must create custom quiescing scripts to run pre-freeze and post-thaw commands.

### Prerequisites

Prerequisites for custom quiescing scripts on Linux VMs are as follows:

- The scripts have to be created in the `/etc/vmware-tools/backupScripts.d` directory on Linux VMs. (This directory does not exist by default, so you must create it.)
- The directory can contain one script or multiple scripts that are executed in sequence. The filenames of the scripts affect the execution order (for example, `10-webapp.sh`, then `20-database.sh`).
- Each script must be able to handle `freeze`, `freezeFail` and `thaw` arguments passed by VMware Tools during the different phases.
- VMware Tools version 10.2 or higher must be installed on the VMs.



## Example Script

This is an example script written to support quiescing for a Linux VM running a PostgreSQL database.

```
#!/bin/sh
if [[ $1 == "freeze" ]]
then
    # set log directory
    log="/var/log/vpostgres_backup.log"
    # set and log start date
    today=`date +%Y/%m/%d\ %H:%M:%S`
    echo "${today}: Start of creation consistent state" >> ${log}
    # execute freeze command.
    # This command can be modified as per the database command
    cmd="echo \"SELECT pg_start_backup('${today}', true);\n\" | sudo -i -u postgres psql >>
${log} 2>&1"
    eval ${cmd}
    # set and log end date
    today=`date +%Y/%m/%d\ %H:%M:%S`
    echo "${today}: Finished freeze script" >> ${log}
elif [[ $1 == "thaw" ]]
then
    echo "This section is executed when the Snapshot is removed"
    log="/var/log/vpostgres_backup.log"
    # set and log start date
    today=`date +%Y/%m/%d\ %H:%M:%S`
    echo "${today}: Release of backup" >> ${log}
    # execute release command
    cmd="echo \"SELECT pg_stop_backup();\n\" | sudo -i -u postgres psql >> ${log} 2>&1"
    eval ${cmd}
    # set and log end date
    today=`date +%Y/%m/%d\ %H:%M:%S`
    echo "${today}: Finished thaw script" >> ${log}
elif [[ $1 == "freezeFail" ]]
then
    echo "This section is executed when the Quiescing Fails."
else
    echo "No argument was provided"
fi
```

## High-frequency Snapshots

High-frequency snapshots allow you to schedule recurring snapshots as frequently as every 30 minutes, and as a preview feature, you can schedule snapshots every 15 minutes.

To enable high-frequency snapshots, your protected site must be running vSphere 7.0 Update 3 or higher, and your protected SDDCs must be running version 1.16 or higher. For the latest release information on vSphere 7.0 Update 3, see these product [release notes](#).

Leveraging a unique data capture technology that is different than VMware's vStorage APIs for Data Protection (VADP) and VMware vSphere replication, high-frequency snapshots create a consistent snapshot of a protected workload by extracting only the delta between two consecutive snapshots and only applies the changes to new snapshots.

High-frequency snapshots are also not dependent on VM-level snapshots to insure consistency, are not bound by snapshot hierarchy on disks, and are not affected by the impact of consolidating snapshots at the end of the sync process. These capabilities enable high-frequency snapshots to take as many as 48 snapshots per-day (every 30 minutes).

## Caveats for High-Frequency Snapshots

High-frequency snapshots have several caveats and limitations.

- **Interoperability between high-frequency snapshots and VADP on same VM.** If a VM is protected using VMware Cloud DR and enabled for high-frequency snapshots, there might be potential interruptions in snapshot replication if the same VM is also being backed up by a third party backup solution that uses VMware APIs for Data Protection (VADP). When the third party backup solution creates or deletes a VADP backup at the same time VMware Cloud DR is replicating a high-frequency snapshot, this snapshot task pauses and retries after a few seconds. VMware Cloud DR will continue the snapshot replication from the point of interruption.

---

**Best Practice** If you are using both a third party backup solution that uses VADP and VMware Cloud DR snapshots to protect the same VM, first create a standard-frequency snapshot of the VM for initial seeding. Then, convert the protection group to high-frequency snapshots and resume the snapshot schedule. Interruptions from third party backup solution that uses VADP are less likely to require a full ingest by VMware Cloud DR during steady state replications, as compared to during initial seeding.

---

- **VMware Tools ApplInfo plug-in is not supported.** If you are protecting a VM with VMware Tools installed with ApplInfo enabled, collected application data is stored in the VMX file of the VM, causing it to expand to >55 KB. Because VMware Cloud DR does not allow backup of VMX files that are >55 KB, snapshot replication fails. To avoid this situation, deactivate the ApplInfo plug-in VMware Tools for the VM.
- **Interoperability limitation between VMware Cloud DR high-frequency snapshots and VMware HCX.** You cannot use VMware HCX to perform a bulk migration or a replication assisted vMotion (RAV) for set of VMs on a VMware Cloud DR protected SDDC, if those VMs are also being replicated using high-frequency snapshots. However, HCX bulk migrations and RAV are supported if the VMs being migrated to (or from) a VMware Cloud DR protected site are not being replicated with high-frequency snapshots.
- **Software Requirement.** High-frequency snapshots are only supported on protected sites running vSphere 7.0 Update 3 or higher and protected SDDCs running version 1.16 or higher. For the latest release information on vSphere 7.0 Update 3, see these product [Release Notes](#).

- **vSphere Virtual Machine Encryption not supported.** High-frequency snapshots use a temporary external snapshot file to capture information that gets overwritten after a snapshot is taken. This external snapshot file is not protected by vSphere Virtual Machine Encryption and thus impacts data confidentiality of captured disk blocks. For these reasons, VMs with vSphere encryption enabled are not supported for high-frequency snapshots.
- **Special characters not supported.** High-frequency snapshots are not supported in vSphere environments that use non-ASCII or special characters in inventory object names (such as VMs, VM templates, hosts, clusters, networks, datastores).
- **High-frequency snapshots cannot be enabled on protection groups if there are vSphere snapshot types present.** Protection Groups with VMs where existing vSphere (non-VMware Cloud DR) snapshots are present cannot be enabled for high-frequency snapshots. This limitation is because Change Block Tracking (CBT) cannot be enabled when vSphere snapshots exist at the time high-frequency snapshots are enabled. However, vSphere snapshots can be taken after the first high-frequency snapshot completes. If you have VMs with existing snapshots not taken by VMware Cloud DR, you can delete the existing snapshots on disks attached to the VM, or consolidate them so none are present at the time a snapshot is taken. Once you have deleted or consolidated the existing vSphere snapshots, high-frequency snapshots can be enabled and new vSphere Snapshots can be created for a VM.
- **Restoring a deleted VM fails when restore destination site is non-compatible with high-frequency snapshots.** If you attempt to restore or fail back a deleted VM from high-frequency snapshot, and none of the hosts on the cluster are compatible with high-frequency snapshots (restore destination must be on vSphere 7.0 Update 3), the VM is created but is not fully restored. In this situation, the VM is unusable. If this situation occurs, make sure that at least one host in the cluster's resource pool is compatible with high-frequency snapshots, and then retry the VM restore operation. In this case, you can only perform a single VM restore. You cannot retry the failback.
- **Protection group can only have one snapshot type.** A protection group can only be configured for one snapshot type: standard-frequency, high-frequency, or quiesced. You cannot mix snapshot types in a protection group.
- **VMs must be part of a cluster.** The enablement of high-frequency snapshots of VMs and their disks is only supported for VMs that are part of a cluster, not on a standalone host. If you have standalone hosts, place them in a cluster. You can then add VMs on the host to protection groups and schedule high-frequency snapshots.
- **Converting protection groups from high-frequency snapshots to standard-frequency snapshots not supported.** Once a protection group is enabled for high-frequency snapshots, you cannot revert back the protection group to take standard-frequency snapshots.
- **VMs with mixed snapshot types.** If a VM belongs to a protection group configured for either standard-frequency snapshots or quiesced snapshots, and the same VM is captured in another snapshot from a different protection group configured for high-frequency snapshots, all subsequent snapshots for that VM are captured as high-frequency snapshots.

- **'Unprotect' a VM from a high-frequency snapshot protection group.** If you have a VM that previously belonged to a protection group with high-frequency snapshots enabled and you want to 'unprotect' the VM, see [Deactivate High-frequency Snapshots from VMs](#).
- **VMs using VMware vSphere Virtual Volumes (vVols) not supported.** VMs that use vVols are not supported for high-frequency snapshots.
- **Not supported with vSphere Replication and array-based replication using VMware Site Recovery Manager (SRM).** High-frequency snapshots do not support replicating VMs using vSphere Replication or array-based replication with VMware Site Recovery Manager (SRM), if those VMs are also being protected by VMware Cloud DR high-frequency snapshots.
- **Increased VM memory overhead for vSphere Distributed Resource Scheduler (DRS)-enabled clusters (on-premises protected sites only).** Using high frequency snapshots increases the VM memory overhead for protected VMs for on-premises protected vSphere sites. If you have manually adjusted the VM memory overhead setting, then using high-frequency snapshots might require a change to the manual setting. For example, for a VM with large number of disks, the memory overhead specified might be insufficient. If needed, contact VMware support to adjust this setting.
- **If vSphereBackupNFC is set manually, it must be set to use a routable network interface.** If you change the vSphere backup network setting vSphereBackupNFC, ensure that it is set to use a routable network interface, or high-frequency snapshots won't work.

## Run a Host Compatibility Check for High-Frequency Snapshots

If you are not sure if all of the hosts and VMs on a protected site are compatible with high-frequency snapshots, you can run a host compatibility check.


When you create a protection group and select a protected site and vSphere instance, the system runs a compatibility check to determine if the hosts on the site are compatible with high-frequency snapshots.

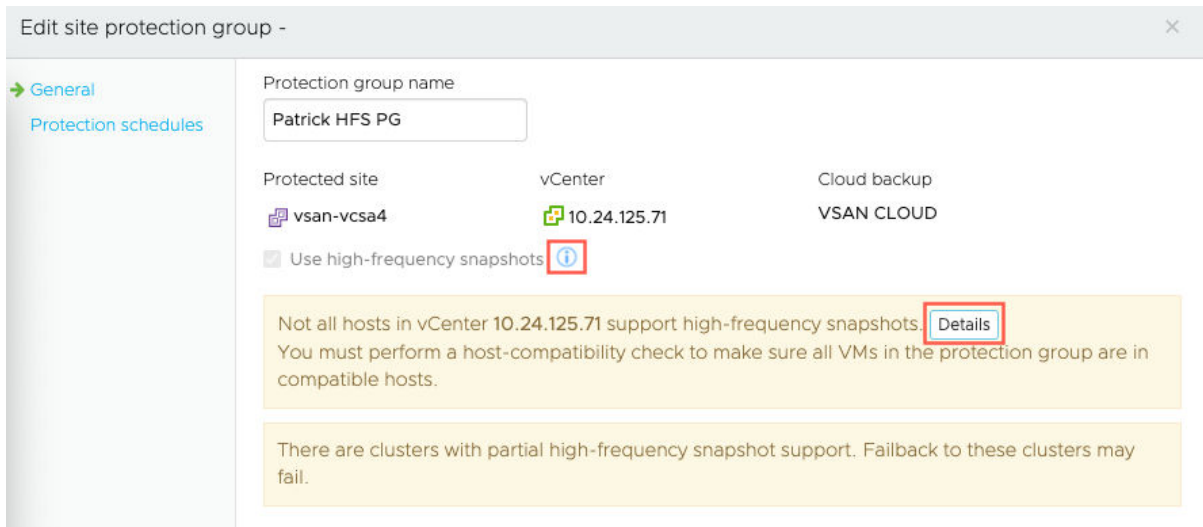
If the compatibility check fails, it means either that the cluster does not support high-frequency snapshots, or the cluster is "mixed." A mixed cluster has some hosts running vCenter Server 7.0 Update 3 and ESXi 7.0 Update 3, and some that are not.

In this situation, you can run a host compatibility check to identify any non-compatible VMs that you might want to migrate to a compatible host.

### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 From the upper-right of the Protection groups page, click **Create protection group**.
- 3 In the **Create protection group for site** dialog box, enter a name for the protection group.
- 4 Next, select a protected site and a vCenter. The dialog box also displays the cloud file system associated with the protected site, which is where snapshots replicate to.

- 5 After you select a protected site and vCenter, the system checks the vCenter for software compatibility with high-frequency snapshots. The software compatibility check has three possible outcomes:
  - If the protected site vCenter and hosts are compatible, then the Use high-frequency snapshots option is selected, and you can start defining protection group membership.
  - If the protected site vCenter and hosts are not compatible, then high-frequency snapshots cannot be enabled for the selected vCenter.
  - If the protected site is "mixed," then some hosts are compatible with high-frequency snapshots and some are not. In this case, continue the steps in this task to run a host compatibility check and preview VMs on the hosts.
- 6 If the selected vCenter's cluster mixed, the Use high-frequency snapshots option is deselected. Select the option, and then to see which hosts are compatible click **Details**. (The Details button only appears if the cluster is mixed and the option is selected.) You can also click the Information icon (  ) next to the Use high-frequency snapshots option.




- 7 In the High-frequency snapshots dialog box, hover over the **Some hosts** link to see which hosts on the protected site are compatible or not. When you are finished, click **OK**.

High-frequency snapshots

High-frequency snapshots allow you to schedule snapshot as often as every 30 minutes.

**Partial support in vCenter 10.24.125.71**

Cluster ▲	Supported in
Mix Cluster	 <b>Some hosts</b>

Incompatible hosts  
n1394.datrium.com

Compatible hosts  
n1522.datrium.com

**Requirements and downloads**

High-frequency snapshots have strict requirements at the vCenter, host, and VM levels. Enabling high-frequency snapshots requires using specific versions of ESXi and vCenter Server, which can only be downloaded here:

vCenter Server 7.0.2 hfs	ESXi 7.02 hfs
<a href="#">ISO image</a>	<a href="#">ISO image</a>
<a href="#">Appliance OVA</a>	<a href="#">Offline bundle</a>
<a href="#">Appliance update bundle</a>	

**IMPORTANT:** VMs with existing vSphere snapshots or VMware Cloud DR standard-frequency snapshots will **not** be included when snapshotting a high-frequency protection group

OK

- 8 To run the host compatibility check, create an all-encompassing VM name pattern query using \*, and then click **Host compatibility check**.

Edit site protection group -

General

Protection schedules

Protection group name

Patrick HFS PG

Protected site

vsan-vcsa4

vCenter

10.24.125.71

Cloud backup

VSAN CLOUD

☒ Use high-frequency snapshots

Not all hosts in vCenter 10.24.125.71 support high-frequency snapshots. You must perform a host-compatibility check to make sure all VMs in the protection group are in compatible hosts.

There are clusters with partial high-frequency snapshot support. Failback to these clusters may fail.

VM name pattern

VM name pattern

\*

Excluding

Use \* as wildcard, comma to separate patterns. VM pattern tips

Add vCenter query ▾

Preview VMs

Host compatibility check

The queries are evaluated with vCenter whenever a new snapshot is taken.

Cancel

Back


Next

Finish

- 9 In a mixed cluster, the compatibility check fails and displays a list all VMs on hosts that are not compatible with high-frequency snapshots.






















High-frequency host compatibility check

Compatibility check failed!



There are VMs in hosts that do not support high-frequency snapshots. They will **not** be included when snapshotting this protection group.

VMs in incompatible hosts

VM name	Cluster	Datastore	State
 da-drc-stg-dvx23-00	Mix Cluster	dvx79-Datastore1	On
 da-drc-stg-dvx23-01	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_01	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_02	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_03	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_04	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_05	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_06	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_07	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_08	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_09	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_10	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_11	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_12	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_13	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_14	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_15	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_16	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_17	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_18	Mix Cluster	dvx79-Tags-ds	On
 n1394_tag_19	Mix Cluster	dvx79-Tags-ds	On

22 virtual machines

**IMPORTANT:** VMs with existing vSphere snapshots or VMware Cloud DR standard-frequency snapshots will not be included when snapshotting a high-frequency protection group.

Cancel



**Note** Any VMs that have existing vSphere snapshots or VMware Cloud DR standard-frequency snapshots also cannot be captured in high-frequency snapshots.

10 Click **OK**.

## Preview: 15 Minute RPO with High-frequency Snapshots

As a preview feature, you can now configure protection groups with high-frequency snapshots with up to 200 VMs to achieve 15 minute RPO.

With this preview feature, you can configure two half-hourly snapshot schedules for the release, with each schedule staggered every 15 minutes.

For 15 minute RPO, configure your protection group snapshot schedule as follows:

- One half-hourly snapshot schedule set to start each hour at the start of the hour 00 and one to start at 30 minutes past the hour.
- A second half-hourly snapshot schedule to start at 15 and 45 minutes past the hour.
- For snapshot retention, under 'Keep snapshots for' select 3 days.

**Note** You can ignore the warning about keeping snapshots for 90 days.

Configuring these two schedules results with snapshots being taken every 15 minutes.

Edit site protection group -

General  
Protection schedules

### Protection schedules

*Half-hourly*

Take snapshots: Half-hourly (dropdown)  
On: :00 and :30 (dropdown)  
Keep snapshots for: 3 (input) days (dropdown) [X]

*Half-hourly-2*

Take snapshots: Half-hourly (dropdown)  
On: :15 and :45 (dropdown)  
Keep snapshots for: 3 (input) days (dropdown) [X]

To improve the chance of recovering your data in case of a ransomware attack, a schedule with a retention of at least 90 days (or 12 weeks or 3 months) is recommended.

**NEW SCHEDULE**

These schedules will result in times between snapshots under 30 minutes. This is not supported and may result in skipped snapshots and a breached SLA.

## Caveats and Restrictons

- 15 minute RPO is only available for high-frequency snapshots.
- Up to 200 VMs per cloud file system are supported.
- 15 minute RPO does not work on a first time snapshot of a VM. The VMs in the protection group must already have snapshots taken before attempting 15 minute RPO.

- VMs with a high change rate and large snapshot deltas might not achieve a 15 minute RPO. In such cases, snapshots could take longer than 15 minutes to complete.
- 15 minute RPO can be affected if other snapshot jobs are running at the same time.

#### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 From the upper-right, click the **Create protection group** button.
- 3 In the **Create protection group for site** dialog box, enter a name for the group.
- 4 Next, select a protected site and a vCenter instance you want to take snapshots from.
- 5 After you select a protected site and vCenter, the protection group automatically checks the protected site for compatibility with the high-frequency snapshot feature. Select high-frequency snapshots.
- 6 Next, define the dynamic membership of the protection group using VM name patterns, tags and/or VM folders in a vCenter query. The protection groups will contain all VMs that match the queries. Under Group membership, click the **VM name pattern**, **Tags**, or **Folder** buttons.
  - For a **VM Name Pattern** query, enter a VM name pattern that is evaluated before a snapshot is taken. You can also enter a name pattern in the query in the Excluding text box for exclusion. (If there is already one vCenter query, then from the Add vCenter query drop-down menu select **VM name pattern**.)
  - For a folder query, click the **Select folders** button. In the **vCenter folders** dialog box, search the list of folders in your vCenter and click one to add it. (If there is already one vCenter query, then use the Add vCenter query drop-down menu and select Folders.) Folder selection does not include sub folders, so you must select subfolders specifically.
  - For a tag query, click the **Tags** button. In the vSphere tags dialog box, select tags to define protection group membership. Any VMs with selected tags are included in the protection group snapshots. Click **OK** when you finish.
- 7 Click **Preview VMs** to see VMs that match the queries.
- 8 Click **Next**.
- 9 To set the snapshot schedule, click **New Schedule**.
- 10 On the Schedule page, edit the snapshot schedule. Select snapshot frequency based on the following intervals:
 

Configure one half-hourly snapshot schedule set to start each hour at the start of the hour 00 and one to start at 30 minutes past the hour.

Configure a second half-hourly snapshot schedule to start at 15 and 45 minutes past the hour.
- 11 Next, set the snapshot retention for this protection group. You can select any number and then select the duration (hours, days, weeks, months, years).
 

For 15 minute RPO, a three day retention schedule is recommended.

- 12 Optionally, you can add a custom name for the snapshot schedule. Click once in the default schedule name field above the schedule, enter a custom name, and then press Enter on your keyboard.
- 13 When you have finished adding snapshot scheduled and retention, click **Finish**. When the protection group appears in the Protection groups list, you can add it to a recovery plan for testing and disaster and ransomware operations.

## Deactivate High-frequency Snapshots from VMs

You can use the DRaaS Connector CLI to deactivate high-frequency snapshots from VMs.

You might have a VM that belongs to a protection group that replicates high-frequency snapshots, and then at a later date, you want to deactivate high-frequency snapshots from the VM.

---

**Note** If you want to convert high-frequency snapshots to standard-frequency snapshots, contact VMware support. The task described in this topic does not complete the complete conversion from high-frequency snapshots to standard-frequency snapshots.

---

You can use the `drc deactivate-hfs` CLI commands to detach high-frequency snapshots from a single VM or multiple VMs on a protected site.

---

**Note** Once you deactivate a high-frequency snapshot from a VM, the next high-frequency snapshot taken of that VM requires a full seeding process.

---

### Usage

To deactivate high-frequency snapshots from a VM If you know the VM instanceID, run this command:

```
drc deactivate-hfs -h
usage: drc deactivate-hfs
        [-h] --vcenter-ip VCENTERIP
        [--vm-instance-uuid VMINSTANCEUUID]
        [--vendor-id VENDORID]
        [--drc-agent-svc-id DRCAGENTSVCID]
        [--cloud-file-system-id CLOUDFILESYSTEMID]
        [--name-pattern NAMEPATTERN]
        [--output-format OUTPUT_FORMAT]
```

---

**Note** Options inside of square brackets [ ] are optional, unless otherwise noted.

---

You can use the following options with this command:

Option	Description
<code>--vcenter-ip VCENTERIP</code>	IP address of the protected site vCenter where you want to deactivate high-frequency snapshots from a VM.
<code>--vm-instance-uuid VMINSTANCEUUID</code>	<p>A VM <code>instanceUUID</code> is a unique identifier that associated a VM with a specific vCenter.</p> <p>For information about finding a VM instance UUID using the VMware Managed Object Browser or Powershell, view the information <a href="#">here</a>.</p>
<code>--vendor-id VENDORID</code>	<p>You can obtain the vendor ID for a VM in the URL of VMware Cloud DR combined with the cloud file system ID. In VMware Cloud DR, select the cloud file system where the VM snapshots replicate to. For example:</p> <pre>vcd.r.vmware.com/#/backup-sites/summary/ r43zf5c1-5172-4af9-8925-10b3959df2k1</pre> <p>The last string on the URL is the cloud file system ID. To get the VM vendor ID, do these two things:</p> <ul style="list-style-type: none"> <li>■ Add <code>.dvx.</code> to the <code>com</code> of the URL.</li> <li>■ Add the cloud file system ID after <code>.dvx.</code></li> </ul> <p>For example:</p> <pre>com.vmware.vcd.r.dvx.r43zf5c1-5172-4af9-8925-10 b3959df2k1</pre>
<code>--drc-agent-svc-id DRCAGENTSVCID</code>	This command is reserved for support usage. Do not use this option unless requested to do so by VMware Support.
<code>--cloud-file-system-id CLOUDFILESYSTEMID</code>	<p>The ID of the VMware Cloud DR cloud file system. This option is required if the protected site is a VMware Cloud on AWS SDDC. This option is not required for on-premises protected sites.</p> <p>You can find the cloud file system ID by selecting it in VMware Cloud DR and looking at the URL:</p> <pre>vcd.r.vmware.com/#/backup-sites/summary/ d43zf5c1-5172-4af9-8925-10b3959df2k1</pre> <p>In the example, the cloud file system ID is:</p> <pre>d43zf5c1-5172-4af9-8925-10b3959df2k1</pre>

Option	Description
<code>--name-pattern NAMEPATTERN</code>	<p>Name patterns dynamically define protection group membership. You can also use name patterns with the CLI to determine which VMs you want to deactivate high-frequency snapshots from.</p> <p>If no name pattern is given in to deactivate multiple VMs, you are asked to supply one.</p> <p>For more information, see <a href="#">VM Name Pattern</a>.</p> <p>For example, if you want to deactivate all VMs from high-frequency snapshots, the name pattern is:</p> <pre>--name-pattern '.*'</pre> <p>Other examples:</p> <p>Deactivate high-frequency snapshots from VMs that have a name starting with 'h':</p> <pre>'h.*'</pre> <p>Deactivate high-frequency snapshots from VMs that have a name ending with 'st':</p> <pre>--name-pattern '.*st'</pre> <p>Deactivate high-frequency snapshots from VMs that have a name containing 'oo':</p> <pre>--name-pattern '.*oo.*'</pre>
<code>--output-format OUTPUT_FORMAT</code>	<p>Choose the output format (optional):</p> <ul style="list-style-type: none"> <li>■ <code>json</code>: Format output in the JSON format.</li> <li>■ <code>default</code>: Print output using the default formatter.</li> </ul>

For example, to deactivate high-frequency snapshots from a single VM using the VM instance UUID:

```
connector-name>># drc deactivate-hfs --vcenter-ip 192.168.226.228 --vm-instance-uuid
5017f5c8-5588-c638-e252-3734d5e00b6c
```

The results:

```
result: Deactivate VM vm-21 succeeded with vendorId: com.vmware.vcdr.dvx.31336361-2301-1018-
abd2-667ca9d91a2e
```

This example command deactivates high-frequency snapshots from multiple VMs based on the name pattern `'.*tiny-core.*'`:

```
drc deactivate-hfs --vcenter-ip 192.168.226.228 --name-pattern '.*tiny-core.*'
```

The results:

```
result: All VMs deactivated successfully
total VMs found: 9, deactivated: 8, deactivate failed: 0, unmatched: 1, excluded: 1
```

## Error Handling

The detach high-frequency snapshots CLI provides verbose error messaging to help you understand when some VMs fail to deactivate their high-frequency snapshots. In some cases, a VM fails to have its high-frequency snapshots deactivated because the VM only has standard frequency snapshots associated with it.

This example indicates that several of the VMs that matched the name pattern likely do not have high-frequency snapshots associated with them:

```
[vcd_r_00:50:56:8b:4b:2a:~]# drc deactivate-hfs --vcenter-ip 192.168.226.228 --name-pattern
'.*'
result: Deactivate-HFS failed for vm: tiny-core-scsi_3 result: Deactivate VM vm-19 failed
with vendorId: com.vmware.vcdr.dvx.deactivateHFS. VM was not HFS activated.
Deactivate-HFS failed for vm: tiny-core-scsi_1 result: Deactivate VM vm-21 failed with
vendorId: com.vmware.vcdr.dvx.deactivateHFS. VM was not HFS activated.
There are VM(s) that failed deactivate-hfs, please verify those VMs aren't unprotected
total VMs found: 9, deactivated: 6, deactivate failed: 2, unmatched: 0, excluded: 1
```

This example shows you the results if you enter a name pattern that does not match any VMs:

```
[vcd_r_00:50:56:8b:4b:2a:~]# drc deactivate-hfs --vcenter-ip 192.168.226.228 --name-pattern
'test'
result: No matching VMs. Please try again with a valid name-pattern. No VMs were deactivated
total VMs found: 9, deactivated: 0, deactivate failed: 0, unmatched: 9, excluded: 1
```

This example shows the results of an incomplete command, where you only entered the vCenter IP address and nothing else:

```
[vcd_r_00:50:56:8b:4b:2a:~]# drc deactivate-hfs --vcenter-ip 192.168.226.228
result: Please pass in the name pattern of the VMs that are targeted to get hfs deactivated
```

## Take a Manual Snapshot

In addition to taking snapshots on a recurring schedule, you can also take manual snapshots of group member VMs.

The manual snapshot includes the current members of the protection group based on the vSphere queries defined for the group (VM name pattern and/or folders). The snapshot includes all VMs that match any of the query criteria at the time when you take the snapshot.

---

**Note** If another scheduled or manual snapshot is in progress, you cannot take a manual snapshot of that protection group until the current snapshot finishes.

---

**Important** To view a completed snapshot, reload the snapshots list on the protection group page. To verify that a snapshot job has started, you can view the **Monitor > Events > Protection** filter to view the snapshot job start event.

---

#### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 From the list, select a protection group to take a snapshot of.
- 3 In the protection group page, from the drop-down menu on the far right, click **Take snapshot**.
- 4 In the **Take manual snapshot** dialog box, select a retention time the snapshot, which determines how long the snapshot is retained before it is automatically deleted. For example, you can choose Forever (is never deleted), for a specific period of time (for example, 12 days), or a specific date on which the snapshot is deleted.
- 5 Click **Take snapshot**.

---

**Note** If a snapshot job is already in progress, then you see an error stating that the system failed to take the snapshot. In this case, you must wait until the current snapshot job is finished to take the manual snapshot.

---

- 6 Refresh the list to see the new snapshot. Click the snapshot to see its contents.

## Restore an Individual VM

Once a protection group has taken snapshots of the VMs on your protected site, you can restore individual VMs from a snapshot.

VMware Cloud DR restores the VM to the same state it was in when the snapshot was taken, including its vCenter location, configuration, and data.

An example of needing to restore a VM is during a failed software upgrade attempt or when something was accidentally deleted or uninstalled from a virtual machine.

---

**Note** This section shows you how to restore one VM at a time. For multi-VM restore automation, see [Chapter 11 Set Up Recovery Plans in VMware Cloud DR](#).

---

A few prerequisites before you can restore a VM:

- There must be at least one snapshot taken of a protection group to restore a VM from the group.
- Make sure that the VM you are restoring is powered OFF in vSphere before you restore it.

## Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 Select a protection group from the list.
- 3 Click one of the snapshots associated with the group.
- 4 In the list of VMs in the snapshot, click the **Restore** button next to the VM you want to restore.
- 5 In the **Restore VM** dialog box, you see information that describes the VM and the location to which it will be restored on the protected site.
- 6 Click **OK** to restore the VM.

## Guest File Recovery

VMware Cloud DR snapshots allow you to download guest files of individual VMs to recover those files to a safe site.

Guest files downloaded as ZIP packages, which you can unzip and manually restore them to the destination of your choice.

Guest files are downloaded to the browser host where guest file recovery is being performed. For example, you can use guest file recovery to find a clean guest file from an older but known good snapshot of a VM.

You can run a recovery plan for ransomware and when a VM is in validation, you can open a browser on the VM running in the recovery SDDC and download imported guest files from a more recent snapshot directly into the running VM. After clean guest files have been downloaded to the VM, you can complete security analysis and recover the VM back to production.

Every guest file download is also available as a link from the **Monitor > Tasks** list that you can send to other users. The download link expires after six hours. The user on the local system must have file-level permissions to unzip the package.

---

**Note** If you are using [access lists](#) for VMware Cloud DR, only IP addresses listed in the Management access list can download a guest file for recovery.

---

**Best Practice** For security reasons, guest file download links are valid for six hours after the task completes. Any attempt to access an expired link results in an HTTP 403 forbidden error. Download links for this feature use enterprise grade encryption at the source and only allow SSL-based connections for download. Each download link contains both proof of identity and means of authentication, so anyone with the link can download the file. Share these links only at the discretion of the backup administrator. As a security best practice, use great caution while sharing these links.

---

Guest file recovery supports the following file systems:

- **Windows:** NTFS and FAT32.



- **Linux:** Ext3 and Ext4.

Guest file recovery does not support the following technologies:

- Windows dynamic volumes.
- Linux VMs that use Logical Volume Manager (LVM)
- Linux VMs formatting with the XFS file system.
- Microsoft Storage Spaces.

Current caveats for guest file recovery:

- You can run one guest file download at a time.
- Maximum path length of the download file directory = 255 characters.
- Maximum number of files or folder paths per ZIP package = 25. A folder path that contains multiple files is only counted as one item in the ZIP package out of a maximum of 25.
- Maximum individual file size allowable for download = 40 GB. This means that any given file in a download package cannot be larger than 40 GB.
- Maximum ZIP package export size = 100 GB.
- Windows OS unzip utility. Currently, restoring guest files does not support using the Windows OS default unzip utility in the File Explorer. Use 7ZIP or WinRAR utility on Windows systems for guest file restore operations.

## Recovering Guest Files on a Recovery SDDC

If you are restoring VM guest files or folders directly on a VM in a recovery SDDC on VMware Cloud on AWS, you must first configure access to the cloud file system S3 bucket in AWS.

The cloud file system, where protection group snapshots are stored, uses S3 as a repository of recovery points.

You have two options for configuring access for guest file restore from the cloud file system into a VM on a recovery SDDC;

- Use an S3 endpoint in a linked customer account. Create the endpoint in a linked account, and then add VMC firewall rules, described here: [Access an S3 Bucket Using an S3 Endpoint](#). Or
- Use an internet gateway to access your S3 bucket. This method also requires deactivating the S3 option on the connected Amazon VPC for your SDDC. For information, see [Access an S3 Bucket Using the Internet Gateway](#).

## Recover VM Guest Files

You can recover virtual machine guest files from VMware Cloud DR snapshots, which you can then manually restore to any destination of your choice.

You can recover guest files from three different locations:

- The virtual machines list.

- A snapshot inside a protection group.
- The **Other** menu during [Chapter 14 Ransomware Recovery](#), when VMs are in the validation state.

---

**Note** For supported file systems on Windows and Linux, and for a full list of caveats and limitations with guest file recovery, see [Guest File Recovery](#).

---

#### Procedure

- 1 From the left navigation, select **Virtual Machines**. Or, when a VM is in validation during ransomware recovery, from the **Other** menu.

If you are accessing guest file recovery during ransomware recovery, select **Recover guest files** from the **Other** menu after you start a VM in validation.

---

**Note** If you have more than one cloud file system, select one from the upper left of the list. If you have deployed only one cloud file system, it is already selected.

---

- 2 In the search box at the top you can search for specific VMs using [VM Name Pattern](#) (except for exclusion patterns). By default, the search box uses the \* wildcard to search for VMs across all snapshots.
- 3 Decide which VM to recover files from, and then click the **Recover guest files** button next to the VM.
- 4 In the **Recover guest files** dialog box, the most recent snapshot of the VM is selected. (If you select a VM snapshot from a protection group list, then that snapshot is the one selected.) To select a different snapshot, click the left or right arrow, or click the **Use different snapshot button**.

- 5 In the **Recover guest files** dialog box, you can select files from the list of Available files and folders. Click the down arrow to download individual files or folders. When you click the down arrow after selecting a file or folder, the ZIP package downloads immediately.

Recover guest files for win2012-vm3

Snapshot



win-server - Daily - 2021-09-24T19:30 UTC

Sep-24 12:30 pm (2h ago)

Guest file system path

/ SCSI 0:0-Partition-1 (C:) / System Volume Information /

Available files and folders

File name	Type	Size	Modified	
MountPointManagerRemoteDatabase	File	--	Sep-24 12:26 pm	
tracking.log	Log file	20.0 KiB	Jun-29 04:47 am	

Selected files and folders

Click the arrow on any file or folder to add it to the selected file and folder list for ZIP archive creation.


2 folders and files loaded

0 files and folders selected

Create ZIP file

Close

- 6 To download multiple files, click the right arrow to select multiple files. After you have selected files for download, click the **Create ZIP file** button to start the download.



- 7 After the ZIP file downloads, click the **Close** button. Your user must have file-level permissions to unzip the package.
- 8 To access a URL for the guest file package, select **Monitor > Tasks**.
- 9 Find the guest file download task in the list. To filter the list, select the Protection filter.
- 10 From the menu at the far right of the task entry, right-click the download icon and select **Copy Link Address**.

## Snapshot Retention

When you create a protection group, you configure how long to retain snapshots on the cloud file system.

With protection groups, you can retain snapshots for a short time (one hour) or a long time (several years), depending upon your data retention policy.

As a best practice for ransomware recovery, configure snapshot schedules with a retention of at least 90 days. A 90 day retention schedule might result in higher storage capacity consumption, which is charged as described on the [VMware Cloud DR pricing](#) page.

You can set snapshot retention when you [Create a Protection Group](#).

---

**Note** VMware Cloud DR deletes all snapshots that are past the time of expiration, except the last snapshot. This final snapshot is retained in order to make sure recovery is still possible, and so VMware Cloud DR can perform an incremental sync of the VMs if the schedule is activated again.

---

## Snapshot Retention for Running Ransomware Recovery Plan

When you start a ransomware recovery plan, VMware Cloud DR pauses all snapshot expiration for snapshots taken prior to starting the plan. Existing snapshots are deleted upon expiration, regardless of protection group retention policy, until the plan is ended. Any subsequent snapshots taken since the starting of the plan will expire and be deleted according to the configured retention policy.

When the plan is ended, snapshots expiration will resume according to the defined protection group retention policy.

## Delete Snapshots

For a VMware Cloud DR protection group, you can delete one or more snapshots at a time.

### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 From the list, select a protection group.
- 3 Select one or more of the snapshots by clicking the check box to the left of its name.
- 4 Click the **Delete** button.
- 5 In the Delete snapshot dialog box, click **Delete** to remove the snapshots.

## Edit Snapshot Name

After a snapshot is taken, you can edit its name.

### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 From the list, select a protection group.
- 3 Select a snapshot by clicking the check box to the left of its name.
- 4 Click the **Edit** button.
- 5 In the **Edit snapshot** dialog box, under Snapshot name, enter a new name.
- 6 Click **OK**.

## Edit Snapshot Schedules

You can edit one or more snapshots at a time to change their snapshot retention schedule.

**Procedure**

- 1 From the left navigation, select **Protection Groups**.
- 2 From the list, select a protection group.
- 3 Select one or more snapshots by clicking the checkbox next to its name.
- 4 Click the **Edit** button.
- 5 In the **Edit snapshot** dialog box, if you have selected multiple snapshots, the title indicates the number of snapshots whose schedules you are changing. You have three options for snapshot retention: Forever (indefinite retention), For a specific amount of time, or Until a specific date.

---

**Note** For effective ransomware recovery, create a snapshot retention schedule that is at least 90 days. A 90 day retention schedule might result in higher storage capacity consumption, which is charged as described on the VMware Cloud DR pricing page.

---

- 6 Click **OK**.

## Deactivate Snapshot Schedule

To stop a protection group from replicating its snapshots to a backup site, you can delete its snapshot schedules.

When you deactivate a snapshot schedule, the already replicated snapshots are retained on the cloud file system for as long as the configured retention time.

VMware Cloud DR deletes all snapshots that are past the time of expiration, except the last snapshot. This final snapshot is retained in order to make sure recovery is still possible, and so VMware Cloud DR can perform an incremental sync of the VMs if the schedule is activated again.

If you delete a protection group, then all snapshots are deleted from the cloud file system.

**Procedure**

- 1 From the left navigation, select **Protection groups**.
- 2 Select a protection group.
- 3 In the **Edit site protection group** dialog box, select Protection group schedules from the left side.
- 4 Click the small X to the right of each schedule to delete the snapshot schedule.

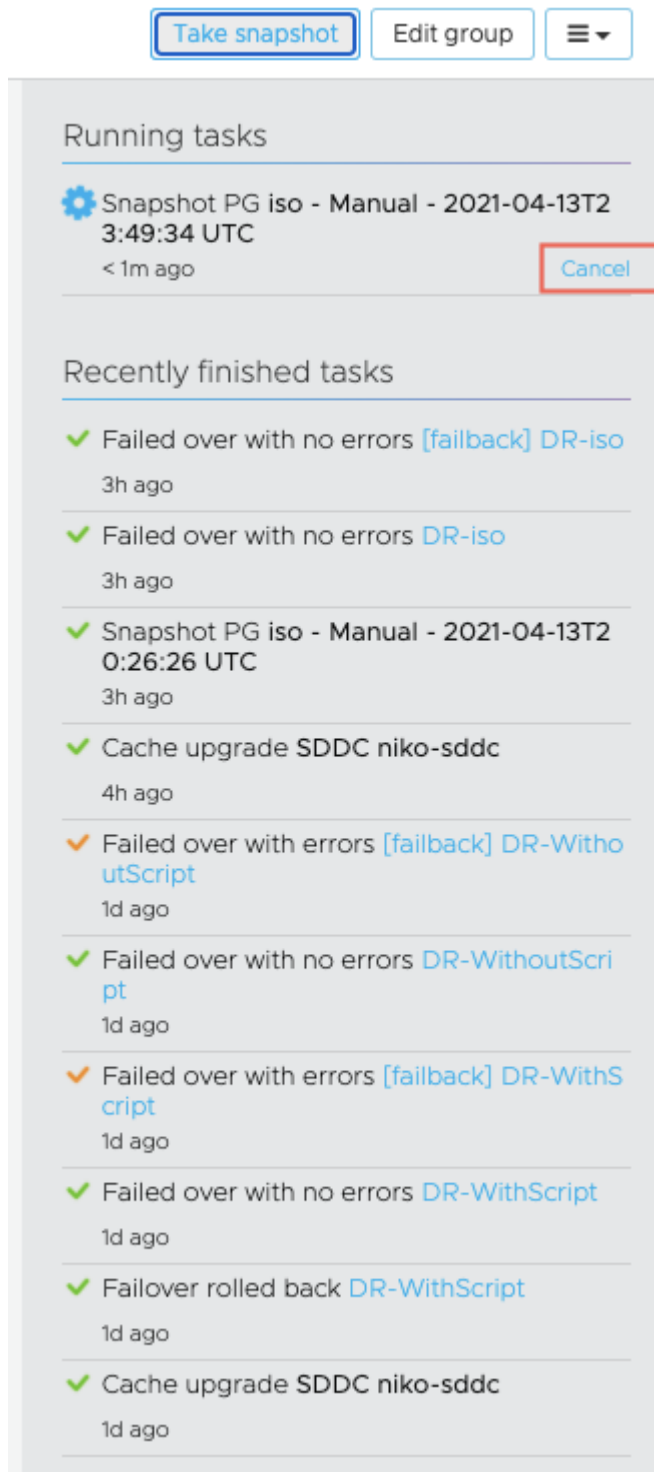
## Cancel a Snapshot Task

You can cancel a snapshot task from the running Tasks list on the right side of the VMware Cloud DR UI.

Canceling a snapshot is only possible while the snapshot task is running. Once a snapshot task completes, it cannot be canceled.

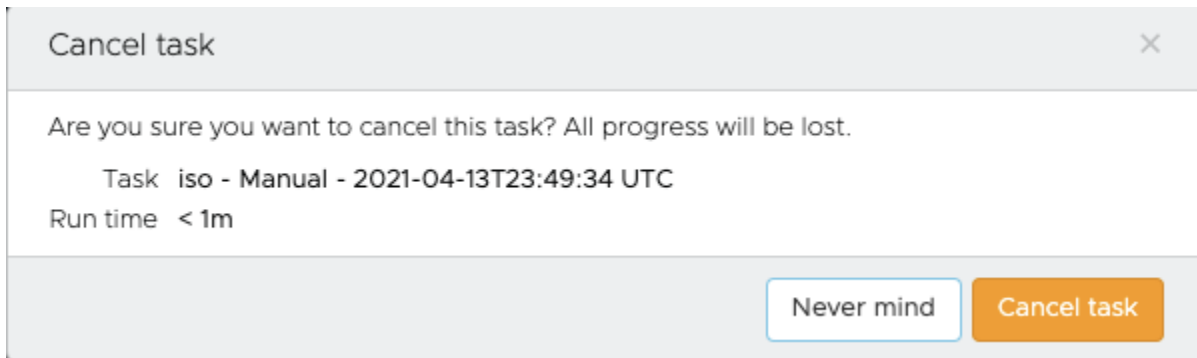
## Procedure

- 1 On the right side of the VMware Cloud DR UI, locate the running task for the snapshot task and click the Cancel link below it.



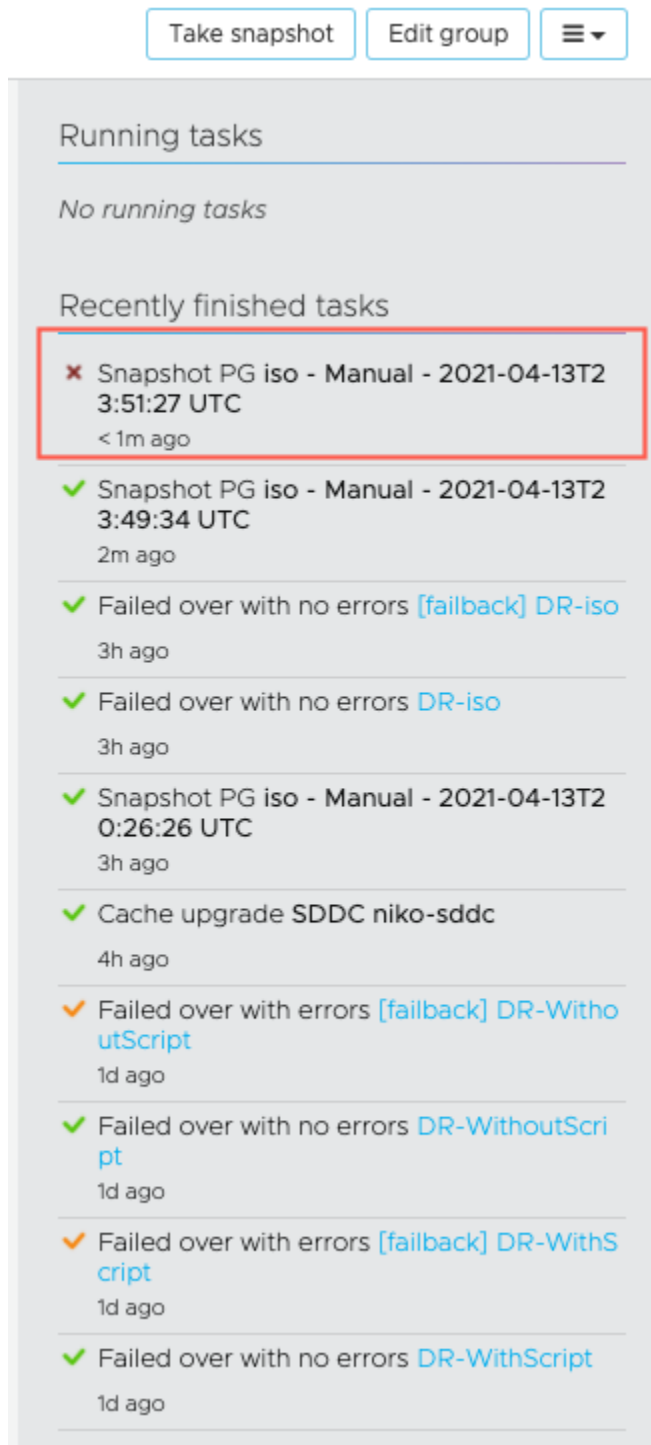
- 2 Click the Cancel link under the snapshot task.

- 3 In the **Cancel task** dialog box, click **Cancel task**.





- 4 The canceled snapshot task now appears in the Recently finished tasks list.



## Snapshot Events and Logs

Snapshot events and logs include information about snapshot jobs.

## Snapshot Events

Snapshot event information includes:

- Severity of the snapshot event (Info, Warning, or Error).
- Time the snapshot was taken.
- Start and finish times of the snapshot.
- System on which the snapshot was taken.
- If a snapshot has expired.
- If a snapshot is the last snapshot in the protection group schedule.

To view snapshot events for a protection group, from the left navigation select **Protection groups**. Then, select a protection group and then click the **Events** tab.

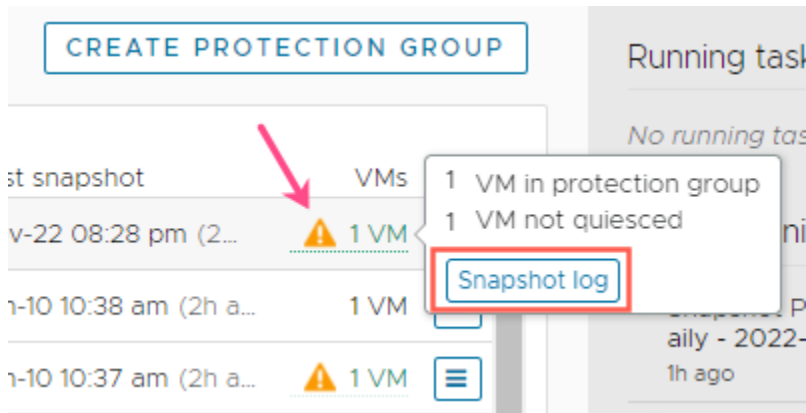
The screenshot shows the VMware Cloud Disaster Recovery web interface. The left navigation pane is expanded, showing 'Protection groups' selected. The main content area displays the 'Events' tab for the protection group 'centos-2TB-vm1'. A table lists various events with columns for Severity, Description, Target, Timestamp, and User name. The events include warnings about app-consistent snapshots and information about task completions and new snapshots.

Severity	Description	Target	Timestamp	User name
Warning	Failing to use app-consistent snapshot 1 out of 1 VMs in group 'centos-2TB-vm1'	centos-2TB-vm1	Jan-09 11:08 pm	
Info	Task drc-71d91c32-71ac-11ec-a529-0a2330b5d125 for PG centos-2TB-vm1 completed s...	centos-2TB-vm1	Jan-09 04:30 pm	
Info	New snapshot 'centos-2TB-vm1 - Daily - 2022-01-10T00:30 UTC'. Includes 1 VMs and 0...	centos-2TB-vm1	Jan-09 04:30 pm	
Warning	Failed to use app-consistent snapshot 1 out of 1 VMs in snapshot 'centos-2TB-vm1 - Da...	centos-2TB-vm1	Jan-09 04:30 pm	
Warning	Quiesced Snapshot was not attempted for VM centos-2TB-vm1 - Daily - 2022-01-10T0...	centos-2TB-vm1	Jan-09 04:30 pm	
Info	DRC Task drc-71d91c32-71ac-11ec-a529-0a2330b5d125 started for PG centos-2TB-vm1.	centos-2TB-vm1	Jan-09 04:30 pm	
Info	Task drc-476f92ce-70e3-11ec-b3da-0a2330b5d125 for PG centos-2TB-vm1 completed...	centos-2TB-vm1	Jan-08 04:30 pm	
Info	New snapshot 'centos-2TB-vm1 - Daily - 2022-01-09T00:30 UTC'. Includes 1 VMs and ...	centos-2TB-vm1	Jan-08 04:30 pm	
Warning	Failing to use app-consistent snapshot 1 out of 1 VMs in group 'centos-2TB-vm1'	centos-2TB-vm1	Jan-08 04:30 pm	

## Snapshot Logs

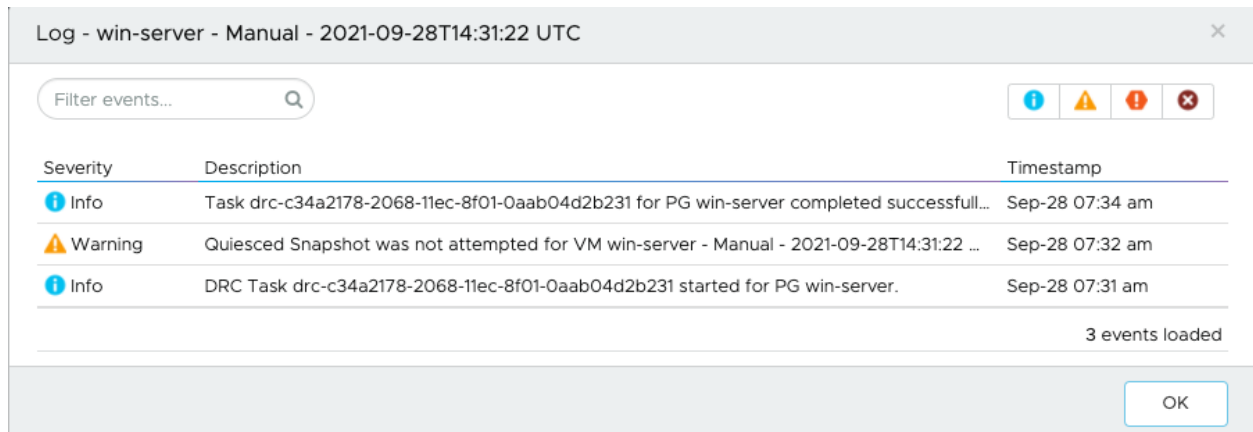
You can view a snapshot log in the protection group page by selecting a protection group then selecting the **Events** tab. A snapshot log shows information related to the most recent protection group snapshot. Snapshot logs also show any related events that occurred at the time of the task.

To see the snapshot log, move your pointer over the status icon at the right side of the event, and then click the **Snapshot log** button.



You can also view snapshot logs from **Monitor > Events**, if you select the Protection filter, then select a snapshot event. The small panel opens at the bottom of the page and shows the **Snapshot log** button for all completed snapshot tasks.

A snapshot log also indicates if your protection group has either high-frequency or quiesced snapshots configured, but only some of the VMs were captured with those features. For example, you have the quiesce option selected in a protection group with 10 VMs, and one of the VMs does not have VMware Tools installed (required for quiesced snapshots). The snapshot task generates a log indicating that one VM was not quiesced.



## Edit a Protection Group

Edit protection groups to change group membership parameters, change the snapshot schedule, or provide a custom name for the schedule.

### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 To the right of the protection group you want to edit, from the drop-down menu on the far right, click **Edit**.
- 3 You can now edit the protection group's VM name pattern, tag, or folder queries.

- 4 When you click **Next**, you can also edit the snapshot schedule and retention policy for the protection group.
- 5 Optionally, you can add a custom name for the snapshot schedule. Click once in the default schedule name field above the schedule, type a custom a name, and then press Enter on your keyboard.
- 6 Click **Save**.

## Delete a Protection Group

When you no longer need a protection group, you can delete it.

---

**Caution** When you delete a protection group, you also delete ALL of the snapshots associated with it, forever.

---

### Procedure

- 1 From the left navigation, select **Protection groups**.
- 2 Identify a protection group to delete, from the drop-down menu on the far right, click **Delete**.
- 3 In the **Delete protection group** dialog box, verify that you want to delete this group and all the associated snapshots in it and enter the phrase **DELETE GROUP AND SNAPSHOTS** in all upper case letters.

---

**Note** When you click **OK**, all snapshots in the group are deleted permanently and are not retrievable.

---

- 4 Click **OK** to delete the selected protection group and associated snapshots.

## Throttle Replication

You can control the speed at which a protection group replicates snapshots to a cloud file system by changing the replication bandwidth throttle settings.

For optimum performance, set throttling to at least 100 Mbps. (The lowest setting is 20 Mbps.) You can also turn off the throttle setting.

### Procedure

- 1 From the left navigation select **Protected sites**, and then select the site you want to throttle replication for.
- 2 From the upper-right of the page, select the drop-down menu and click **Throttle replication**.
- 3 In the **Throttle replication** dialog box, select Throttle to and then enter a numerical value in the Mbps text box to set the desired replication speed.
- 4 Click **OK**.

## Export VM-to-Protection Group Mappings

You can export a list of all VM to protection group mappings to the CSV file format.

For each VM you are protecting, you can export the following information:

- VM name.
- Protected site name where the VM is running.
- Protected site vCenter IP address.
- Number of protection groups the VM belongs to.
- Names of all protection groups that the VM belongs to.
- The pathname of the VM virtual machine configuration file (.vmx) on the protected site vCenter.

---

**Note** Exporting protection group membership to CSV can take up to one minute per-2500 VMs. Do not close the browser window during export, or the export will be canceled.

---

### Procedure

- 1 From the left navigation, select Virtual Machines.
- 2 Click the **Export CSV** button.
- 3 When the file is finished downloading, you can open the file to view all VM to protection group relationships.

## Unprotect a VM

If for any reason you do not want a VM to be included any longer in a protection group, you can 'unprotect' a VM.

When you include a VM in a protection group and begin snapshot replication to the cloud file system, the VM is considered 'protected' by VMware Cloud DR.

Having a protected VM ensures that in the event of a disaster, or a ransomware recovery attack, you can recover the VM from those snapshots.

To unprotect a VM, you can remove it from any protection groups it is a member of by doing the following:

- If the protection group is using a VM name pattern, change the name pattern so it does not include the VM name.
- If the protection group is using tag queries, remove those tags from VM you want to unprotect.
- If the protection group is using folder queries, manually move the VM to a folder not configured in the protection group.
- Delete all snapshots of the VM from all protection groups it was previously included in.

- If the VM was being protected by high-frequency snapshots, you must [Deactivate High-frequency Snapshots from VMs](#).

After these tasks are completed, allow roughly two to three days for VMware Cloud DR space reclamation processes to completely remove the remaining VM snapshots. This process can take longer if there are a large number of snapshots to delete.

## Replication Progress Statistics

You can monitor the real time progress of snapshot replication tasks from the dashboard.

Replication progress shows you the percentage of data being transferred and the rate of data transferred ("logical throughput") during and after the task.

## Replication Progress in the Task List

You can monitor the progress of snapshot replication tasks from the Tasks list and view the following information:

---

**Note** Replication statistics only measure replication from a protected site to a cloud file system. These statistics do not represent failover or failback progress.

---

Metric	Description
Start time	Time when the task started.
End time/Progress	During the task, you can see the percentage of the total task completed. When the task is finished, it shows the time of completion.
Logical throughput	The progress during snapshot replication from a protected site to a cloud file system, since the start of the task. Measured in Megabits (Mbps) and Gigabits per second (Gbps). This measurement only displays while the task is in progress.

---

**Note** This view does not show progress for single VM restore tasks

---

## Replication Statistics on Protected Site Page

The protected sites page shows replication statistics for both replication throughput and throttle (if configured), and also shows throughput for restore operations (during the operation).

### Cloud backup target



prasanna-scfs-1

Backup 778 Mbps

Throttle None

The protected site page also shows the current [Throttle Replication](#) maximum, if configured. You can also view replication transfer rate in the Topology pane.

If no snapshots are replicating to a cloud file system, then the throughput value is empty.

## Replication Throughput in the Dashboard Topology Pane

The Topology map shows The rate of data being transferred to and from the protected site to a cloud file system, measured in Mbps and Gbps.

### Topology



# Deploy a Recovery SDDC

# 10

You can deploy a VMware Cloud on AWS recovery SDDC and configure it as a site for both disaster and ransomware recovery.

In addition to deploying a new recovery SDDC in VMware Cloud DR (sometimes called VCDR), you can also add an existing SDDC for recovery. For more information, see [Add an Existing SDDC for Recovery](#). Once you deploy a VMware Cloud on AWS in a recovery region, it cannot be used in another recovery region.

You can deploy a recovery SDDCs using i3en, i3n, and i4i hosts.

Deploying a recovery SDDC is restricted by the following caveats:

- VMware Cloud on AWS recovery SDDCs are restricted by VMware Cloud on AWS configuration limits. For example, if you want to know the currently supported maximum number of hosts per cluster, the maximum number of SDDCs per region, or the maximum number of SDDCs per Organization, and more, see [VMware Configuration Maximums](#).
- Stretched clusters are not supported for recovery SDDCs.
- If you need PCI hardening for your recovery SDDC, contact VMware Support for assistance.

---

**Important** For more information about VMware Cloud DR configuration limits, visit the [VMware Configuration Maximums](#) tool.

---

Read the following topics next:

- [Before You Start Using an SDDC for Recovery](#)
- [Deploy a New recovery SDDC](#)
- [Add an Existing SDDC for Recovery](#)
- [Delete a recovery SDDC](#)
- [Daily Usage Email Reminder](#)
- [Adding and Removing Hosts](#)
- [Adding, Attaching, Deleting Clusters](#)
- [Add a Network to a recovery SDDC](#)
- [Request Public IP Addresses](#)



- [Add NAT Rules](#)
- [Add a Firewall Rule to the recovery SDDC Network](#)
- [Create a Firewall Rule for Public IP Addresses Accessing vCenter](#)
- [Access Recovery SDDC on VMware Cloud on AWS](#)

## Before You Start Using an SDDC for Recovery

Before you deploy a new or add an existing recovery SDDC, it is important to understand how it functions and its limitations.

VMware Cloud DR leverages the VMware Cloud on AWS Software Defined Data Center ("SDDC") as a disaster recovery site, which you can use if disaster strikes (or for testing) and you have to fail over your protected vCenter to the cloud.

### Create Firewall Rule for Public IP Addresses Accessing vCenter

As a general best practice, create a firewall rule for all Public IP addresses (or IP address ranges) that require access to the vCenter on your recovery SDDC. For more information, see [Create a Firewall Rule for Public IP Addresses Accessing vCenter](#).

### Recovery SDDC Default Firewall Rules

Do not change the recovery SDDC default firewall rules. Changing the default firewall rules could interrupt access from the SDDC to the cloud file system or Orchestrator components. By default, when your recovery SDDC is deployed its network will contain a set of pre-configured firewall rules which begin with the "CloudDR-SystemRule-" prefix.

Do not delete these firewall rules. recovery SDDC firewall rules with this prefix cannot be edited or deleted in the VMware Cloud DR UI, but these rules can be edited and deleted in the VMware Cloud on AWS console. So, do not change or delete any of these SDDC firewall rules.

### Recovery SDDC Datastore

When you deploy a recovery SDDC, the system automatically creates an NFS datastore named 'ds01'.

This NFS datastore is created exclusively to expose VM backups to the recovery SDDC to facilitate disaster recovery, so never use it as general-purpose storage. Do not use the vSphere Client, vSphere APIs, or any other means to create and power on VMs directly on this NFS datastore, except through VMware Cloud DR.

### CloudDR-Proxy VM

During deployment of a recovery SDDC, or while attaching an existing SDDC to VMware Cloud DR, a virtual machine named "CloudDR-Proxy" is deployed into the Compute-ResourcePool in your SDDC.

The CloudDR-Proxy VM is critical to the operations of VMware Cloud DR and facilitates the launching and orchestration of failover and failback operations when a recovery plan is run. The CloudDR-Proxy VM manages data movement during failover and failback. During failback, the CloudDR-Proxy VM copies back the changes that occurred to the virtual machines while they were running in the recovery SDDC after a failover.

Do not modify this VM, power it down, or apply any network configurations, or failover and failback operations may be adversely affected.

**Note** If you want to use route filtering with your recovery SDDC, you must configure route aggregation to advertise routes for the /26 CIDR of “sddc-cloud-dr-proxy-network” subnet to the cloud file system, as described in [Aggregate and Filter Routes to Uplinks](#).

For the CloudDR-Proxy VM, VMware Cloud DR automatically creates the following infrastructure-level Allow firewall rules in the distributed firewall in vCenter on the recovery SDDC:

**Table 10-1. Allow firewall rules created for the CloudDR-Proxy VM**

Source	Destination	Firewall Rule
Cloud file system	CloudDR-Proxy VM	Allow
CloudDR-Proxy VM	Cloud file system	Allow
CloudDR-Proxy VM	vCenter Server	Allow

The CloudDR-Proxy VM is managed automatically by VMware Cloud DR and requires no setup or management by the customer.

**Note** The subnet you choose for the CloudDR-Proxy VM when deploying a recovery SDDC must not overlap with the SDDC management subnet, the subnet used for a policy based VPN, or the linked AWS account VPC subnet.

## Recovery SDDC and Cluster Names

Cluster names are given automatically when you create a cluster, you cannot provide a custom name for clusters you add to your recovery SDDC.

Clusters added to a recovery SDDC follow this naming pattern: 'Cluster-<x>-<y>'. For example, the first cluster for your first SDDC is named: Cluster-1-1. If you add another cluster to the same SDDC, the new cluster name is 'Cluster-1-2', and so on.

If you tear down the first SDDC, the cluster names in any recovery plan change and appear with an asterisk (\*) for the SDDC name, such as 'Cluster-\*-1' and Cluster-\*-2. When the second SDDC is deployed, the UI displays the proper cluster names in the recovery plans with the incremented SDDC name.

This same behavior applies to cluster names in plan compliance reports. If the SDDC is currently deployed, then the cluster names appear with the correct name in the report, such as 'Cluster-1-2'. If the SDDC is deleted at the time the report runs, then the cluster name uses the asterisk, such as 'Cluster-\*-2'.

## Your AWS Account and the Recovery SDDC

When you deploy a recovery SDDC, you connect it to an AWS account belonging to you (also called the 'customer AWS account').

Before you can deploy a new recovery SDDC, your AWS account must be linked to your VMware Cloud organization. The purpose of this account is to provide a network connection from customer AWS services to your VMware Cloud on AWS SDDCs. For more information, see [Deploy an SDDC from the VMC Console](#) and [AWS Account Linking](#).

Your AWS account must also have a subnet created in your AWS Virtual Private Cloud (VPC) in the same AWS region and availability zone where you deployed a [Chapter 5 Deploy a Cloud File System](#). As a best practice, create a subnet in every AWS Availability Zone (AZ) you want to use before you deploy a cloud file system and recovery SDDC. The size of the subnet must be /26. For more information, see [Work with VPCs and subnets](#).

For the full list of VMware Cloud on AWS networking requirements for the customer AWS account, see [Deploying and Managing a Software Defined Data Center](#).

## Recovery SDDC Deployment Restrictions

Before you deploy a recovery SDDC, there are several restrictions you need to know.

- When you create or add a recovery SDDC, a new /26 subnet is created for the CloudDR-Proxy VMs (one proxy VM per-cluster). The CloudDR-Proxy VM subnet is connected to the Compute Gateway and is separate from the management subnet used by the recovery SDDC. You can either connect to the /26 range or you can enter a new subnet.
- The subnet you choose for the CloudDR-Proxy VMs cannot overlap with the SDDC management subnet, or if you have a policy-based VPN configured on an existing SDDC, the remote network cannot overlap with the /26 subnet. The CloudDR-Proxy VM subnet also must not overlap with the AWS VPC CIDR for the AWS account that is linked to your VMware Cloud organization.
- Each recovery SDDC you deploy (after the first deployment) must have a cloud file system site associated with it. For more information, see [Chapter 5 Deploy a Cloud File System](#).
- If you deploy a single-host recovery SDDC (also known as a 'starter' Recovery SDDC), it is deleted after 30 days and all data on the recovery SDDC is lost. For this reason a single host Recovery SDDC is only intended for testing purposes, not for use in production.
- You can add a recovery SDDC attached to an SDDC Group. However, you might fail to attach an existing SDDC that is already a member of SDDC Group.

- You can scale-up a single host recovery SDDC into a three or more 'multi-host' recovery SDDCs and retain all your data. A multi-host recovery SDDC with three or more hosts is not time bound but is subject to recurring costs. A multi-host Recovery SDDC also provides data protection and production-level SLAs.
- You can remove hosts from the recovery SDDC if the number of hosts in your SDDC cluster remains above the 2-host minimum. You cannot scale down a 2-host Recovery SDDC. Ensure that you have sufficient capacity in your cluster to hold the workload VMs that will be evacuated from the hosts that you remove.
- Two host (I3 type) recovery SDDC deployments are not supported with VMware Cloud SDDC version 1.15.
- Do not change the user credentials for the NSX Cloud Admin account.
- If you have enabled an [authentication policy](#) for your VMware Cloud organization, to either block or allow specific IP addresses, make sure that you do not accidentally block or disallow the Orchestrator and cloud file system IP addresses. To find these IP addresses, see [Service Public IP Addresses](#).

## Maintaining recovery SDDC Settings

Once you have deployed your recovery SDDC, do not change any of the following settings or configurations in the VMware Cloud DR UI:

- **Recovery SDDC default firewall rules.** Changing the default recovery SDDC firewall can interrupt access to the cloud file system or Orchestrator components. By default, when you deploy your SDDC, its network contains a set of pre-configured firewall rules which begin with the "CloudDR-SystemRule-" prefix. Do not delete these firewall rules. Keep in mind that you cannot edit or delete SDDC Firewall rules using this prefix in the VMware Cloud DR UI, but these rules can be edited and deleted in the VMC Console. Make sure you do not delete any of these SDDC firewall rules.
- **Do not rename your recovery SDDC once you have deployed it.** Once you name and deploy your Recovery SDDC, do not change its name.
- **Network configuration on the CloudDR-Proxy VM.** Do not make any changes to the network configuration on the CloudDR-Proxy VM.

## Deploy a New recovery SDDC

You can deploy a new recovery SDDC to use for disaster recovery and ransomware recovery.

When you deploy a new recovery SDDC, it connects to an AWS account belonging to you (also called the 'customer AWS account') that is linked to your VMware Cloud organization. If you have not linked your AWS account with your organization, see [AWS Account Linking](#).

VMware Cloud DR supports recovery SDDCs using I3en, I3n, and I4i hosts. If you want to deploy an I4i recovery SDDC from the VMware Cloud DR UI, contact support to enable this capability.

Each recovery SDDC you deploy must have a cloud file system associated with it. For more information, see [Chapter 5 Deploy a Cloud File System](#).

---

**Note** VMware Cloud DR does not currently support stretched clusters for recovery SDDC.

---

#### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 From the upper-right of the page, click the **Add recovery SDDC** button.
- 3 In the **Add recovery SDDC** dialog box, select **Deploy new SDDC**.

---

**Note** If no cloud file systems are available, then you must [Chapter 5 Deploy a Cloud File System](#) first.

---

- 4 Next, under Cloud file system, select a cloud file system to pair with the recovery SDDC. The recovery SDDC is deployed in the same AWS availability zone as the selected cloud file system.
- 5 Click Next, and in the next page of the dialog box, enter the following settings:

Setting	Description
SDDC Name	Enter a name for the recovery SDDC.
Deployment type: Single Host	<p>A single-host SDDC is designed for testing purposes and can be deployed for the following reasons:</p> <ul style="list-style-type: none"> <li>■ A single-host recovery SDDC can be used for building and testing your recovery plans, after which you can delete the recovery SDDC to reduce costs. A single host SDDC does not provide any data protection, does not offer production level Service Level Agreements (SLAs), and are automatically torn down in 60 days. Use a single-host deployment for testing purposes and not for production use.</li> <li>■ A single host SDDC can also be deployed on-demand, where you deploy the SDDC after a disaster event or for DR or ransomware testing.</li> <li>■ You can scale up a single host SDDC, but you cannot scale down to one host after you have scaled up.</li> <li>■ If you want a permanent recovery SDDC, select the multi-host option and select the numbers of hosts to deploy.</li> <li>■ A single host recovery SDDC can use either I3 or i4i host type. If you want to deploy the SDDC with I3en hosts, you must deploy a multi-host recovery SDDC.</li> </ul> <p>Scroll down in the dialog box to the Host type section, where you can see options for a multi-host deployment:</p>

Setting	Description
Deployment type: Multi-host	<p>If you select Multi-host option (two or more), you can select I3, I3en, or I4i hosts.</p> <ul style="list-style-type: none"> <li>■ The I3 host type is the default host type. I3 hosts have 36 cores, 512 GiB RAM, and 10.37 TiB raw storage capacity per host. You can deploy two, three, or four hosts of the I3 host type.</li> <li>■ The I3en host type is optimized for data-intensive workloads. I3en hosts have 96 logical cores, 768 GiB RAM, and 45.84 TiB raw storage capacity per host.</li> <li>■ The I4i host type provides up to 128 logical cores, 1024 GiB of RAM, and 30 TiB raw storage capacity per host.</li> </ul> <p><b>Note</b> If you want to add a recovery SDDC with I4i hosts, contact VMware support for assistance.</p> <p>You can remove hosts from the recovery SDDC if the number of hosts in your SDDC cluster remains above the 2-host minimum. You cannot scale down a 2-host Recovery SDDC. Ensure that you have sufficient capacity in your cluster to hold the workload VMs that are evacuated from the hosts that you remove.</p> <p>You must use the VMware Cloud on AWS UI to scale down an SDDC.</p>
Number of hosts	<p>For multi-host deployments you can select two, three, or four hosts. Adding a host increases the available storage capacity and costs.</p>
Management subnet	<p>Enter a subnet for the management network of the recovery SDDC. This private subnet range (RFC 1918) is used for vCenter Server, NSX Manager, and ESXi hosts.</p> <ul style="list-style-type: none"> <li>■ Select a range that does not conflict with other networks that you want to connect to this recovery SDDC.</li> <li>■ Minimum CIDR sizes: /23 for up to 27 hosts, /20 for up to 251 hosts, /16 for up to 4091 hosts. Reserved CIDRs: 10.0.0.0/15, 172.30.0.0/16. Enter a CIDR block size of either /16 or /20.</li> </ul>
Compute subnet	<p>Enter a gateway logical network for the recovery SDDC.</p> <p>This private subnet range is for the logical network that the workload VMs use. This network supports a maximum of 1000 MAC addresses, so using a /22 range or smaller is recommended.</p> <p>The SDDC management subnet you use here cannot overlap with the subnet used for the VMware Cloud DR proxy VMs, which is /26 . If you have configured a policy based VPN on the SDDC, that remote network cannot overlap with the subnet.</p> <p>(For more information, see <a href="#">Selecting IP subnets and Connectivity for your SDDC.</a>)</p> <p>Only one logical network is created by default. More networks can be created in the VMware Cloud DR UI after the recovery SDDC is deployed. For more information, see <a href="#">Add a Network to a recovery SDDC</a></p>

Setting	Description
Compute network name	Give the compute subnet a name.
Proxy subnet	The subnet range for the VMware <a href="#">CloudDR-Proxy VM</a> . When you create or add a recovery SDDC, a new /26 subnet is created for the VMware Cloud DR CloudDR-Proxy VM (one proxy VM per-cluster). The CloudDR-Proxy VM subnet is connected to the Compute Gateway and is separate from the management subnet used by the recovery SDDC. You can either connect to the /26 range or you can enter a new subnet.

The subnet for the [Cloud-DR Proxy VM](#) cannot overlap with the SDDC management subnet. If you have a policy based VPN configured on the SDDC, the remote network cannot overlap with the /26 subnet.

- 6 Click **Next**. Under AWS Settings, you see your linked AWS account information. If you have more than one AWS account linked to your organization, you can select the account you want to use for this SDDC.
- 7 Next, select a subnet for the SDDC. If the availability zone where the recovery SDDC is being deployed has more than one subnet configured, you can select the subnet from the drop-down menu, next to the name of the VPC where the SDDC will be deployed.
- 8 Next, under Seller choose either VMware or AWS, if you have bought from more than one seller. If you have used only one seller, it is selected here. For more information, see [Choose a Purchase Option](#).
- 9 When you are ready to deploy the recovery SDDC, enter the phrase **DEPLOY SDDC** in all uppercase letters in the confirmation field and then click **Deploy**.

#### What to do next

Once you deploy a recovery SDDC, follow these guidelines to ensure a consistent SDDC configuration: [Maintaining SDDC Settings](#).

## Add an Existing SDDC for Recovery

You can leverage an existing SDDC to use for recovery operations.

When you add a VMware Cloud on AWS SDDC, a cloud file system is attached to it. Both the SDDC and cloud file system must be in the same AWS availability zone (AZ). If no cloud file systems are available, then you can [Chapter 5 Deploy a Cloud File System](#).

---

**Note** The recovery AZ, where all cloud file systems and recovery SDDCs exist, must support I3, I3en, and I4i host types.

---



---

**Note** Adding an existing VMware Cloud on AWS SDDC to VMware Cloud DR is a permanent operation. When you add an SDDC, it mounts the cloud file system that cannot be unmounted.

---

## Prerequisites

When adding an existing VMware Cloud on AWS SDDC for recovery, the SDDC:

- Must be version 1.12 or higher.
- Cannot have stretched clusters.
- Cannot have a management subnet that is different than the subnet used by recovery SDDCs already deployed in the current recovery region.

To allow access from VMware Cloud DR to an existing VMware Cloud on AWS SDDC vCenter Server, see [Allowing Access to vCenter for an Existing SDDC](#).

## Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 In the upper-right of the page, click the **Add recovery SDDC** button.
- 3 In the **Add recovery SDDC** dialog box, select **Attach existing SDDC**.
- 4 Next, select a cloud file system and then click **Next**.
- 5 In the **Attach existing SDDC** dialog box, select an SDDC. If an SDDC is not selectable, the red text explains the reason. For example, Different region means the SDDC you want to add cannot be deployed because it is not in the same region as the selected cloud file system. Click **Add**.
- 6 In the next page of the dialog box, you can enter a subnet for the VMware Cloud DR proxy VMs, which is used for communication between the proxy VM and the cloud file system.  
  
When you add a recovery SDDC, a new /26 subnet is created for the Cloud-DRProxy VM (one proxy VM per-cluster). The proxy VM subnet is connected to the Compute Gateway and is separate from the management subnet used by the recovery SDDC. You can either connect to the /26 range or you can enter a new subnet.  
  
The subnet for the Cloud-DR Proxy VM cannot overlap with the SDDC management subnet. If you have configured a policy based VPN on the SDDC, that remote network cannot overlap with the /26 subnet.
- 7 Finally, enter the phrase **ATTACH SDDC** in all uppercase letters in the confirmation text box, and then click **Attach**. This process can take up to 30 minutes.

## What to do next

When you add an SDDC with multiple clusters, the default cluster is attached first. All additional clusters are attached next, and their progress can be monitored at **Sites > recovery SDDC**. If any clusters fail to attach, you can attach them manually. For more information, see [Attach a Cluster](#).

## Create Firewall Rule for PCI-Hardened Recovery SDDC

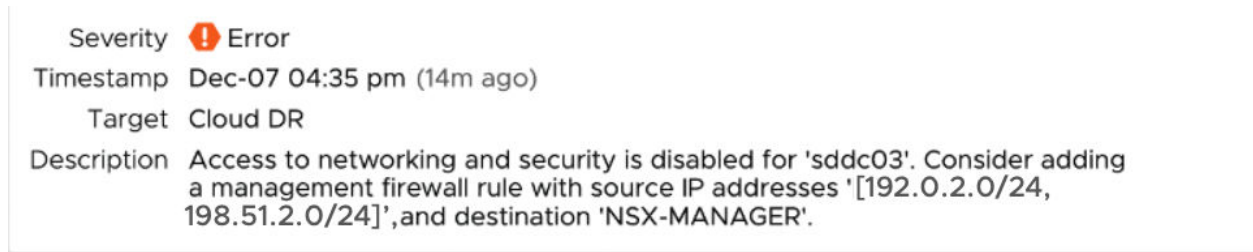
If the SDDC you want to add for recovery is PCI-hardened, you must create a firewall rule to allow the cloud file system to connect with the SDDC.



A PCI-hardened SDDC is configured with limited network access. To allow the VMware Cloud DR cloud file system to connect to a PCI-hardened SDDC for recovery, you must create a management gateway firewall rule that allows the cloud file system access the NSX-MANAGER on the SDDC.

Before you can create this firewall rule, you need the two IP addresses of the cloud file system. You can obtain the cloud file system IP addresses by first attempting to add the PCI-hardened SDDC to VMware Cloud DR, which will fail. When this operation fails, an event is generated that lists the IP addresses of the cloud file system.

For example, you can see the two cloud file system IP addresses in this event:



You will add these two IP addresses to the firewall rule on the SDDC. Once the firewall rule is created to allow the cloud file system to access the PCI-hardened SDDC, you can [Add an Existing SDDC for Recovery](#).

## Prerequisites

### Procedure

- 1 Log in to VMware Cloud Services at <https://console.cloud.vmware.com>.
- 2 Click **Inventory > SDDCs**, then pick the recovery SDDC card and click **View Details**.
- 3 Click the **Networking & Security** tab.
- 4 On the **SDDC Networking at Security** tab, select Gateway Firewall, and then click the **Management Gateway** tab.
- 5 Click the **Add Rule** button.
- 6 Enter a name for the rule, such as **CloudDR-Access-NSX**.
- 7 Click the pencil icon in the Source field in the firewall rule.
- 8 In the **Set Source** dialog box, select the User Defined Groups option. You need to create a group because the cloud file system has two IP addresses.
- 9 Click the **Add Group** button.
- 10 Enter a group name, such as **CloudDR-Access**.
- 11 Under Compute Members, click the **Set** button.
- 12 Enter both cloud file system IP addresses separated by a space.
- 13 Click **Save**, and then click **Apply**.

- 14 When you return to the firewall rule table, click the small pencil icon in the Destination field.
- 15 In the **Set Destination** dialog box, select the System Defined Groups option.
- 16 Next, select the NSX Manager destination, and then click **Apply**.
- 17 In the firewall rule table, click the small pencil icon in the Services field.
- 18 Select both services listed in the field (HTTPS and ICMP).
- 19 When the rule is configured, click the **Publish** button.

#### What to do next

You can now [Add an Existing SDDC for Recovery](#).

## Delete a recovery SDDC

When you want to stop using a recovery SDDC (and stop accruing costs), you can delete it.

Deleting a recovery SDDC terminates all running workloads, deletes all the VMs, and destroys all SDDC data and configuration settings, including public IP addresses.

If the SDDC is included in a recovery plan, all plan compliance checks related to the deleted recovery SDDC are skipped. The plan compliance check will be healthy, but the plan cannot be run until the new recovery SDDC is created or added.

When you delete a recovery SDDC, the following event is logged:

`SDDCDeleteFromVMCConsoleSucceededEvent.`

#### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 From the small menu next to the **Open vCenter** button, click **Delete recovery SDDC**.
- 3 In the dialog box, enter the phrase **DELETE SDDC INCLUDING VMS AND SETTINGS** in all upper case letters to confirm and then click **OK**.

## Daily Usage Email Reminder

Because deploying a recovery SDDC accrues costs, VMware Cloud DR sends a daily email reminder informing you that you have a recovery SDDC deployed, how many hours it has been running, how many hosts are associated with it, and when it expires (if applicable).

The email reminder is automatically sent to all email addresses you have configured for alerts in **Settings > Email alerts**.

When this email reminder is sent, an event is triggered that sends a daily email to specify the name of the SDDC and the number of days remaining before it expires.

## Adding and Removing Hosts

You can add hosts to clusters in your recovery SDDC to expand your recovery capacity.

Host types include:

- **I3.** I3 hosts have 36 cores, 512 GiB RAM, and 10.37 TiB raw storage capacity per host. Supported for single to four host clusters
- **I3en.** The I3en host type is optimized for data-intensive workloads. I3en hosts have 96 logical cores, 768 GiB RAM, and 45.84 TiB raw storage capacity per host. Single-host Recovery SDDCs cannot contain the I3en host type.
- **I4i.** The I4i host type provides up to 128 logical cores, 1024 GiB of RAM, and 30 TiB raw storage capacity per host.

The number of hosts on a recovery SDDC cluster has the following restrictions:

- If there is one existing host in the cluster, you can add exactly two hosts (no other option available).
- You cannot revert back from a two or three host cluster down to a one host cluster.
- If you want to scale up a two host deployment, you can expand the recovery SDDC by adding a cluster to it, rather than adding more hosts to the two host primary cluster.
- A single host cluster supports I3 or I3en hosts only.

For more information on VMware Cloud DR limits, see <https://configmax.vmware.com/home>.

---

**Note** Adding a host increases both the available storage capacity and costs.

---

**Note** If your cluster has 3 or more hosts, and during operation your data expands beyond the capacity of the configured vSAN, VMware's Elastic DRS automatically adds a host if the recovery SDDC free space drops below 25%. Hosts are never removed automatically, but can be removed in the VMware Cloud DR UI all the way down to three hosts. Keep in mind also that you cannot remove hosts from 2 host recovery SDDC.

---

## Add a Host

Add hosts to clusters in your recovery SDDC to expand your recovery capacity.

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.

---

**Note** If a cluster only has one host, then the **Delete** button is deactivated. If there is one existing host, you can add exactly 2 hosts (no other option available). If there are 3 or more hosts, you can add up multiple hosts up to the supported limits.

---

- 3 The dialog box shows the current number of hosts in the cluster and the estimated capacity. Enter the number of hosts you want to add. The dialog box then shows the projected storage capacity for the number of hosts to be added.
- 4 To add the hosts, enter the phrase **ADD HOSTS** in the text field in all upper case letters and then click **OK**.

## Remove a Host

When you no longer need a host in your cluster, you can remove the host.

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Under Clusters, next to the cluster you want to remove hosts from, select **Remove hosts**. If a cluster has two hosts, then this menu item is dimmed. You cannot remove hosts from a two host cluster.
- 4 In the **Remove hosts** dialog box, enter the number of hosts to remove. When you enter a number, the system estimates the amount of removed capacity that is available when you remove the hosts.
- 5 Enter the phrase **REMOVE HOSTS** in all upper case letters in the text field to confirm.
- 6 Click **OK**.

## Adding, Attaching, Deleting Clusters

You can add or attach clusters to your recovery SDDC, with each additional cluster created in the same availability zone as the recovery SDDC.

## Caveats for Clusters on recovery SDDCs

Consider these caveats when adding a cluster to your recovery SDDC:

- Each additional cluster you add increases the costs of your recovery SDDC.
- You cannot delete the default cluster that was created when your recovery SDDC was first created (displays as 'cluster1' in VMware Cloud DR).
- Before you delete a cluster, make sure you check for any potential breakage to mappings in your recovery plans.

## Add a Cluster

You can add a cluster to your recovery SDDC to expand its failover capacity.

### Prerequisites

When you first add a cluster to a recovery SDDC, you can only add a two, three, or four host (i3 and i3en) cluster. Deploying a four host cluster takes roughly one hour.

---

**Note** Adding a cluster to a recovery SDDC increases costs.

---

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Expand the Clusters section, and then click **Add cluster**.
- 4 In the **Add cluster** dialog box, select the number of hosts to add to this cluster (minimum of two hosts).
- 5 Next, select the host type you want to add to this cluster, either I3 or I3en for clusters with three or more hosts. If you selected a two host cluster, you can only deploy it using I3 hosts.
- 6 In the Confirm section of the dialog box, enter the phrase **ADD CLUSTER** in all upper case letters.
- 7 Click **OK**.

## Attach a Cluster

Clusters added to a recovery SDDC from the VMware Cloud on AWS must be attached before they can be used for recovery operations.

### Prerequisites

After you deploy a recovery SDDC (or add an existing one), if you add a cluster to the SDDC from VMware Cloud on AWS, the new cluster displays in VMware Cloud DR. However, before you can use this cluster for recovery operations, you must attach it.

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Expand the Clusters section, and then look for the new cluster.
- 4 To the right of the cluster, click **Attach**.

## Delete a Cluster

When you no longer need a cluster in a recovery SDDC, you can delete it.

## Prerequisites

You cannot delete the default cluster that gets created when you create your recovery SDDC (displays as 'cluster1' in the VMware Cloud DR UI).

---

**Note** Only perform this operation from the VMware Cloud DR UI (not from the VMware Cloud Services console or vSphere Client).

---

## Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Expand the Clusters section and then from the small drop-down menu on the right select **Delete cluster**.
- 4 In the **Delete cluster** dialog box, under the Confirm section, enter the phrase **DELETE CLUSTER** in all upper case letters.
- 5 When you are ready to delete the cluster, click **Delete cluster**.

## Add a Network to a recovery SDDC

Use VMware Cloud DR to create a routed network segment for your recovery SDDC.

If some of your VMs have to connect over logical networks beyond the default network that was created when you first deployed the recovery SDDC, you can use VMware Cloud DR to change the default network.

## Procedure

- 1 From the left navigation, select recovery SDDCs.
- 2 Select a recovery SDDC.
- 3 From the drop-down menu next to the **Open vCenter** button, click the icon and from the menu select **Add network**.
- 4 In the **Add network** dialog box, enter the following information:
  - **Network name.** Any name you want to give the network.
  - **Gateway:** The IP address of the gateway. For example: 192.0.2.1
  - **Bits:** Add an IP prefix for the gateway. For example: 27.
  - **Optional:** DHCP. You can activate DHCP for the new network. In these text boxes, add the gateway IP address range.
- 5 Click **OK**.
- 6 Once the network is created, you can rename it if you wish by clicking **Rename**.

## Request Public IP Addresses

You can request public IP addresses for your recovery SDDC network for the VMware Cloud DR UI.

Once you have public IP addresses, you can create NAT rules to map the public IP addresses to private IP addresses assigned for your VMs.

---

**Important** Using public IP addresses increases costs. You can release public IP addresses that are not in use.

---

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Under Public IPs, click **Request new IP**.
- 4 In the **Request public IP** dialog box, enter an optional label for the public IP address. Labeling the IP addresses is not required but can be useful for identifying the application or VM that the public IP address is being used for.
- 5 Click **OK**.
- 6 The new public IP address appears under Public IP addresses on the Failover SDDC page. You can rename the IP address label at any time by clicking **Rename**.
- 7 To delete the public IP address, click **Delete**.

## Add NAT Rules

You can configure inbound Network Address Translation (NAT) rules for mapping internet traffic from a public-facing IP address to a private IP address in the recovery SDDC's compute network.

---

**Note** With this feature you can map public IP addresses to private IP addresses and all services that communicate over whichever port is assigned to the private IP address. If you want to map specific ports or services, contact VMware Support.

---

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Under NAT rules click **Add new rule**.
- 4 In the **Add new rule** dialog box, enter the following information:
  - **Rule name.** Enter a name for the NAT rule.
  - **Public IP address.** Enter a public IP address.
  - **Private IP address.** Enter the internal IP address to map the public IP address to.

- 5 Click **OK**.
- 6 After the NAT rule is created, click **Edit** if you want to rename the rule, or to change the public or internal IP addresses.

## Add a Firewall Rule to the recovery SDDC Network

You can add firewall rules that allow or block specific traffic to and from your recovery SDDC network.

You can add firewall rules to allow or deny all connections or specific IP addresses or a range of IP addresses. You can also configure specific services and ports to be allowed or blocked in the rules.

---

**Important** If there is a firewall rule that you are unable to create using the VMware Cloud DR UI, contact VMware support for assistance.

---

**Note** For information about finding VMware Cloud DR public IP addresses, see [Service Public IP Addresses](#).

---

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 Under Firewall rules, click **Add new firewall rule**.
- 4 In the **Add firewall rule** dialog box, enter a name for the rule.
- 5 Under Source, select Any to specify traffic from any source, or select IP address to enter individual IP addresses or an IP address range.
- 6 Under Destination, select Any to specify all traffic to any destination, or select an IP address to enter individual IP addresses or an IP address range.
- 7 Under Services, select Any if you want to create a rule for any service on this connection. Or, select a specific service from the drop-down menu and enter the ports on which those services run.
- 8 Under Action, select either Allow or Drop to permit or deny the specified connection.
- 9 Click **OK**.

## Recovery SDDC Firewall Rules Naming Convention

Firewall rules of a recovery SDDC use a specific naming convention for both system default firewall rules and firewall rules created in VMware Cloud DR.



For example, when you deploy your recovery SDDC, its network contains a set of pre-configured firewall rules which begin with the "CloudDR-SystemRule-" prefix. Do not delete these firewall rules.

---

**Note** recovery SDDC Firewall rules with this prefix cannot be edited or deleted in VMware Cloud DR, but these rules can be edited and deleted in the VMware Cloud on AWS console.

---

When you create a firewall rule in VMware Cloud DR, those rules begin with the "CloudDR-UserRule-" prefix. Firewall rules with this prefix can be edited and deleted by VMware Cloud DR users that have the SDDC admin service role granted to them.

Finally, it is possible to create some recovery SDDC firewall rules in the VMware Cloud on AWS console that do not appear in the VMware Cloud DR UI. So when looking at firewall rules, be aware that the VMware Cloud DR UI does not display all possible types of recovery SDDC firewall rules that you can create in the VMware Cloud on AWS console.

## Create a Firewall Rule for Public IP Addresses Accessing vCenter


After you deploy a recovery SDDC, you can create a firewall rule to allow an SDDC vCenter access from public IP addresses in your network.

To allow access to your recovery SDDC by public IP addresses (or IP address ranges) in your network, create a firewall rule in the VMware Cloud on AWS console.

### Prerequisites

For this task, make sure you have compiled a list of the public IP addresses that you want to allow access to your recovery SDDC.

### Procedure

- 1 Log in to VMware Cloud console at <https://console.cloud.vmware.com>.
- 2 Select the VMware Cloud on AWS on AWS tile, and then select a recovery SDDC.
- 3 On the **Networking & Security** tab, select **Gateway Firewall**.
- 4 On the **GATEWAY FIREWALL** page, select **Management Gateway**.
- 5 To add a rule, click **ADD RULE** and give the new rule a name.
- 6 Click the pencil icon (  ) in the Source field of the rule and select **User defined groups**.
- 7 Click the **Add group** button, and enter a name for the group.
- 8 To add your public IP addresses (or ranges) to this group, click **Set members** and then enter your public IP addresses in the middle field.
- 9 Click **Apply**.

- 10 Click the pencil icon in the **Destination** field, and then set the rule Destination to your recovery SDDC vCenter IP address.
- 11 Next, click the pencil icon in the **Services** field, and add the following services to the firewall rule:
  - HTTPS
  - ICMP ALL
  - SSO
- 12 Ensure that the firewall rule is enabled, and then click **Publish** to save and activate this rule.

## Access Recovery SDDC on VMware Cloud on AWS

You can access a recovery SDDC from the VMware Cloud DR UI.

### Prerequisites

When you access a recovery SDDC on VMware Cloud on AWS from the VMware Cloud DR, the UI provides for convenience the default user account credentials associated with a new SDDC in VMware Cloud on AWS, named `cloudadmin@vmc.local`.

For security reasons, if you change the `cloudadmin@vmc.local` user's password in the VMware Cloud console, those changes are not updated in VMware Cloud DR.

### Procedure

- 1 From the left navigation, select **Recovery SDDCs**.
- 2 Select a recovery SDDC.
- 3 In the upper-right of the SDDC information, click the **Open vCenter** button.
- 4 In the **Open vCenter** dialog box, click **Open vCenter**. You can now view your recovery SDDC.

# Set Up Recovery Plans in VMware Cloud DR

# 11

Set up a recovery plan to define the configuration and the orchestration steps required for successful disaster or ransomware recovery.

After you [Create a Recovery Plan](#) and [Configure Recovery Plans](#) a recovery plan, you can run it as a [Chapter 12 Running a Recovery Plan for Failover](#) or test failover, or you can run the plan for ransomware recovery.

---

**Note** Topics in this section discuss running a recovery plan for disaster recovery. For information about running a recovery plan for ransomware recovery, see [Chapter 14 Ransomware Recovery](#).

---

When a recovery plan for disaster recovery starts, it performs all of its recovery steps until completion. A recovery plan configured for disaster recovery can also continue to run to a specific point in the process and wait for user input. Or you can configure a plan to stop and wait for a specified time limit, and then continue until the next stop, or to completion. You can also launch a [Configure Script VM](#) as a recovery step.

For each active recovery plan, VMware Cloud DR performs a compliance check and reports on environmental changes, such as vSphere misconfigurations or network outages. Compliance checks also give you an opportunity to fix issues and restore plan integrity before disaster occurs.

VMware Cloud DR can maintain multiple plans of different types, and multiple plans can be in various stages of execution at any given time, even concurrently.

## vSphere Tags and Recovery Plans

Make sure that all vSphere tags that you are using for VMs on a protected site also exists on the recovery SDDC before you initiate a failover. The failover process associates vSphere tags with recovered VMs that were associated with the VM on the original protected site. For this reason, the tags and their associated categories must be pre-configured on the recovery SDDC for successful failover and failback. For more information about tags and failover operations, see [VM Tags and Tag Categories](#).

Recovery plan compliance checks indicate any missing tags on the recovery SDDC, so before you run a recovery plan, view the plan compliance report. Also, recovery plan [Alerts](#) can be configured to send an email if a recovery plan has missing tags on the selected recovery SDDC.

Read the following topics next:

- ['Just-in-Time' DR](#)
- [VM Customization During Failover/Failback](#)
- [Create a Recovery Plan](#)
- [Configure Recovery Plans](#)
- [Recovery Plan Compliance Checks](#)
- [Viewing Recovery Plans](#)
- [Activating/Deactivating Recovery Plans](#)

## 'Just-in-Time' DR

If your organization has a limited budget, it is possible to use VMware Cloud DR recovery plans in a 'Just-in-time' mode.

'Just-in-time' is a type of configuration where you only deploy a recovery SDDC just at the time of disaster. Because deploying an SDDC costs money, you might only want to deploy (and start paying for) the SDDC at the time of a disaster strikes. A just-in-time deployment is also useful for occasional DR testing activities.

Using recovery plan in VMware Cloud DR (sometimes called VCDR) for a 'Just-in-time' DR has some restrictions:

- A recovery SDDC must be created before you can fully configure and run the plan.
- Once the required recovery plans are created, you can delete the recovery SDDC to save on costs
- Creating the recovery SDDC is a manual task and can take up to a couple of hours to complete, so factor this into your overall RTO planning for 'Just-in-time' disaster scenarios.
- After the SDDC is deleted, you do not have to deactivate the plan. You have the option to deactivate the plan, but once deactivated no more compliance checks are run on the plan.
- When a new SDDC is created to use for the original recovery plan, then you must manually reconfigure all of its resources, networks, folders, vSphere tags.
- For compliance checks and the failover or test to work properly, the new SDDC must have an identical cluster geometry, network layout, folder structure, and tags as the previous SDDC used in the plan, or the failover will not work.

## VM Customization During Failover/Failback

When running a recovery plan for failover or failback, VMware Cloud DR might customize a Windows VM's guest OS based up the plan configuration.

Because VM snapshots are not suitable to run immediately after a failover or failback operation, VMware Cloud DR might need to take an extra snapshot of a Windows VM, customize the guest OS, and restart the VM.

Customizations that might occur to a Windows VM during failover or failback include:

- [IP Address Mapping](#) to set IP addresses to the values defined in the plan mappings.
- [Virtual Networks Mapping](#) changes to accommodate network mapping rules.
- [Compute Resources Mapping](#) changes to ensure sufficient system resources to run the VM.

## Create a Recovery Plan

The first step in orchestrating disaster and ransomware recovery is to create a recovery plan.

### Prerequisites

Before you create a recovery plan, do the following:

- [Chapter 5 Deploy a Cloud File System](#).
- [Chapter 10 Deploy a Recovery SDDC](#).
- [Chapter 6 Set Up Protected Sites](#).
- [Chapter 7 Deploy the DRaaS Connector](#) on the protected site.
- [Create a Protection Group](#) for snapshot replication.

---

**Note** You might have created a recovery plan with the option named 'recovery SDDC deployed in case of disaster'. If you have a plan with this option enabled, the plan will not run successfully, as this feature was removed. You must edit and reconfigure any plans that were created with that option selected. As a workaround, ensure that both the source protected site and target recovery SDDC referenced in the plan are running. Then, edit the plan and on the first page select the 'Existing recovery SDDC' option. Next, reconfigure all of the plan mappings and then save the plan. Then, we recommend that you test the plan for disaster preparedness.

---

### Procedure

- 1 From the left navigation, select **Recovery plans**.
- 2 Click the **Create Recovery Plan** button in the upper-right.
- 3 In the **Recovery Plan** wizard, enter a name for the plan and a description.  
 Write a good description, since this information appears when you generate a recovery plan report. A good description can help others on your team understand what VMs and applications this plan protects.
- 4 In the Recovery site section, the **Existing Recovery SDDC** is selected.
- 5 When you click **Next**, you can start mapping all components of a protected vCenter site.

What to do next

[Configure Recovery Plans.](#)

## Configure Recovery Plans

Configuring recovery plans require defining where to move your protected data when you run the plan for disaster or ransomware recovery.

Specifically, the recovery plan defines where protected resources move to on the recovery site, such as protection groups, VMs, files, vCenters, all vCenter folders, compute resources, virtual networks, and IP addresses (individual or ranges). In addition, you can optionally configure the Test Site operating environment differently for failover exercises.

## Selecting Sites

Configuring a recovery plan requires selecting a protected site and a recovery SDDC.

### Protected Site Selection

On the Site page you define the protected site you want to fail over to a recovery SDDC using this recovery plan. In some cases, you might have multiple protected sites. The Protected Site drop-down menu displays a set of sites known to VMware Cloud DR that you can use for this recovery plan.

### Recovery Site Selection

One of the most important configurations of a recovery plan is the recovery site defined in the plan, which is the recovery SDDC, where you recover the protected site to recovery from a disaster or a ransomware attack.

Protection groups replicate snapshots to a cloud file system in the same AWS region as the recovery SDDC. These backups are instantly available to the recovery SDDC, resulting in instant VM power-on, once the SDDC is available and configured.

You have the option to failover to a recovery SDDC in a different Availability Zone from the Availability Zone of a protected SDDC. For example, you want to guarantee that backups never leave the origin region, which might result in lower costs and better replication bandwidth and RPO.

## Select Protection Groups

Select protection groups for your recovery plans that have regularly scheduled snapshots that you can use for disaster recovery or ransomware recovery operations.

Before you configure a recovery plan and run it, first [Create a Protection Group](#) and schedule snapshot replication to a cloud file system. When a failover occurs, you can use VM snapshots and restore them on the recovery SDDC. Snapshots can also be used when starting VMs in ransomware recovery. When you run the plan, you choose which snapshot to use for failover or ransomware recovery.

In the Groups page of the recovery plan wizard, you select the protection groups you want to include for failover. These selections ensure that the plan has sufficient information to recover all VMs and files from the selected protection groups. The plan uses snapshots of these protection groups for failover when you run the plan. Protection group selection affects a set of automatic compliance checks run for this plan.

A warning displays If the selected protection groups do not have scheduled replication configured for the backup site.

---

**Note** Protection groups that you add to a plan can only originate from the same protected site.

---

## Configure Plan Mappings

Recovery plan mappings are an important part of disaster recovery, as they define where the VMs you want to recover are located (a specific vCenter on a protected site), and the recovery SDDC where you want the VMs to recover.

You must instruct VMware Cloud DR how and when to restart recovered VMs and where to place them from a compute resource pool and network perspective. Specifically, you configure mappings for the following items in a recovery plan for disaster recovery:

- vCenter and vCenter folders
- Compute resources
- Virtual networks
- IP addresses (optional)

Recovery plan mappings have a one to one relationship, which means if an item exists on the protected site, it also must exist on the failover site (the recovery SDDC). Fix any missing items on the failover before you run the recovery plan.

### vCenter Mapping

Mapping vCenters in a recovery plan consists of selecting vCenters registered to the protected site.

You must map all source vCenters that contain VMs protected by protection groups to a target vCenter on a recovery SDDC.

Every recovery plan has vCenter mappings for the following three objects:

- vCenter folders
- Compute resources
- Virtual networks

### vCenter Folders Object Inventory Mapping

The recovery plan displays a subset of the vCenter object inventory item for both the protected and recovery vCenters. You must map source vCenter object nodes that contain protected VMs, which are displayed in the UI with blue text. All other mappings are optional.

To add a mapping, select the source vCenter node and the corresponding target vCenter node indicating where the source VMs are recovered. Then, click **Add**. Complete this step for each mapping, and then click **OK** when finished.

---

**Note** If your VMs on the protected vSphere site have tags associated with them, make sure that the same sets of tags and tag categories also exist on the target site of the plan (the recovery SDDC).

---

**Important** Avoid having other VMs in target folders on the recovery SDDC, because name conflicts can arise when registering VMs with vCenter.

---

## Compute Resources Mapping

Another important mapping for a recovery plan is defining which vCenter computing resources are used for failover.

Mappable vCenter compute objects:

- **Clusters.** If the cluster contains VMs, the cluster icon is highlighted in blue text to indicate a required plan mapping. (This color scheme applies to all required mappings.)
- **Resource pools**
- **Standalone hosts** (not in a cluster). Note that a standalone host can only be mapped to another standalone host.

---

**Note** Regarding vCenter cluster names, "Cluster-1-<clusterIndex>" represents the name of the initial cluster when the recovery SDDC was first created.

---

If the SDDC that your clusters belong to is deleted, then any plans with mappings to clusters on that SDDC displays the target cluster names with an asterisk. For example, "Cluster-\*-<clusterIndex>".

Plan compliance reports indicate an error when clusters are mapped to a deleted SDDC, or if there is a mapping to a deleted cluster.

## Virtual Networks Mapping

With virtual network mappings, you map protected site networks to networks on the recovery SDDC.

Before you map virtual networks in a recovery plan, make sure that you create a network segment for each network on the protected site that your VMs on the site are connected to. For more information, see [Create or Modify a Network Segment](#) and [VMware Cloud DR Networking Best Practices](#).

When you configure network mappings in the plan, you can either use the same mappings for failover and for testing, or you can map to a different network for test failovers. For example, if you are running a pilot light deployment using production workloads, you can create separate network for testing, so you can run a test fail over to a sandbox environment.



To use a different network for testing, deselect the 'Same for test and failover' check box. Then, you can set different network mappings for failover and test failover operations.

## IP Address Mapping

VMware Cloud DR IP address mappings determine how a VM IP address is assigned when you perform a failover to a recovery SDDC.

When you recover a VM from one site to another, you must instruct VMware Cloud DR which IP addresses to use for the recovered VMs.

You configure IP address mappings for VMs installed with Linux or Windows guest OS. VMs configured for IP address mapping displays with a target IP, target subnet mask, target gateways, and target DNS servers.

Before you map IP addresses in a recovery plan, read the following information:

- To map IP addresses for Windows VMs, the system drive of the VMs must be mapped to `c:\`. The mapped `c:\` drive cannot be dynamic volume. The drive must be a basic disk.
- VMware Tools must be installed on the guest OS to ensure successful IP address mapping. For Linux OS, you can also use Open VM Tools (open-vm-tools).
- Only IPv4 is supported for protection plan IP address mapping. Windows VMs can have an IPv6 address configured on the VM, but ONLY IPv4 addresses can be mapped. Linux VMs cannot have an IPv6 address configured on any network interface or IP mapping will not be performed.
- IP address mapping does not alter the type of address assignment. During the recovery process, if a VM interface has been assigned an IP address statically, it cannot be switched to DHCP. Conversely, an interface initially configured with DHCP cannot be converted to static. This feature is designed to maintain the consistency and reliability of the recovery process.

### Individual IP Address Mapping

The following figure illustrates the IP address mapping page, which has the following fields:

- Optional rule description
- Source and target IP addresses
- Source and target subnet masks
- Source and target gateways
- Source and target DNS servers

Edit IP address mapping rule

Rule description

windows\_static\_vlan3

(optional)

Source

IP addresses

10.3.0.160 10.3.0.161 10.3.0.162 10.3.0.163 10.3.0.164  
10.3.0.165 10.3.0.166 10.3.0.167 10.3.0.168 10.3.0.169  
10.3.0.170 10.3.0.171 10.3.0.172 10.3.0.173 10.3.0.174  
10.3.0.175 10.3.0.176 10.3.0.177 10.3.0.178 10.3.0.179  
10.3.0.180 10.3.0.181

Source subnet mask

255.255.0.0

Example: 255.255.255.0

Source gateways

10.3.0.1

Up to two IP addresses, separated by spaces.

Source DNS servers

10.126.0.18

Up to two IP addresses, separated by spaces.

Target

IP addresses

10.3.0.182 10.3.0.183 10.3.0.184 10.3.0.185 10.3.0.186  
10.3.0.187 10.3.0.188 10.3.0.189 10.3.0.190 10.3.0.191  
10.3.0.192 10.3.0.193 10.3.0.194 10.3.0.195 10.3.0.196  
10.3.0.197 10.3.0.198 10.3.0.199 10.3.0.200 10.3.0.201  
10.3.0.202 10.3.0.203

Target subnet mask

255.255.0.0

Example: 255.255.255.0

Target gateways

10.3.0.1

Up to two IP addresses, separated by spaces.

Target DNS servers

10.126.0.19

Up to two IP addresses, separated by spaces.

Cancel

OK

Entries for individual IP addresses must be separated either by white spaces or new lines.

Entries for gateways and DNS servers must be separated by white spaces. If multiple IP addresses are specified, they will be matched in the specified order from source to target.

To configure IP address mapping, you must enter:

- The Rule description field text (optional)
- Source and target IP addresses
- Source and target gateways
- Source and target DNS servers

### IP Address Range Mapping

You can configure IP address ranges for your recovery plan rather than individual IP addresses.

You can use IP ranges in a recovery plan by selecting Range from Range/IP addresses in configuration wizard.

**Note** Only IPv4 is supported for protection plan IP address mapping. Windows VMs can have an IPv6 address configured on the VM, but only IPv4 addresses can be mapped. Linux VMs cannot have an IPv6 address configured on any network interface or IP mapping is not performed.

The IP address mapping page has the following text boxes: optional rule description, source and target IP range prefixes/bits, source and target subnet masks, source and target gateways, and source and target DNS servers. Entries for gateways and DNS servers must be separated by white spaces.

The Range prefix text box provides an IP address within a range of IP addresses for mapping for both source and target. The Bits text box defines the available range of IP addresses to be mapped.

The following table describes the available CIDR Prefix values for Bits text box:

CIDR Prefix	Dotted Decimal Notation	# Node addresses	# of Traditional Class Networks
/13	255.248.0.0	512 K	8 B or 2048 C class
/14	255.252.0.0	256 K	4 B or 1024 C class

CIDR Prefix	Dotted Decimal Notation	# Node addresses	# of Traditional Class Networks
/15	255.254.0.0	128 K	2 B or 512 class
/16	255.255.0.0	64 K	1 B or 256 class
/17	255.255.128.0	32 K	128 C class
/18	255.255.192.0	16 K	64 C class
/19	255.255.224.0	8 K	32 C class
/20	255.255.240.0	4 K	16 C class
/21	255.255.248.0	2 K	8 C class
/22	255.255.252.0	1 K	4 C class
/23	255.255.254.0	512	2 C class
/24	255.255.255.0	256	1 C class
/25	255.255.255.128	128	½ C class
/26	255.255.255.192	64	¼ C class
/27	255.255.255.224	32	⅛ C class

For example:

- Range prefix = 10.116.1.50
- Bits = /24

This means the available range of IP addresses to be mapped is 10.116.1.0 through 10.116.1.255.

The Bits (CIDR Prefix) specified can be a smaller range within the defined subnet in your environment. For example, you can define the subnet as follows:

- [network: 10.116.0.0/20,
- netmask: 255.255.240.0,
- gateway: 10.116.0.1,
- range: 10.116.0.0 - 10.116.15.255]

In this example, providing a bits value of /24 with Range prefix of 10.116.1.0 allows you to provide a smaller range of IP addresses to be mapped within that subnet. The subnet mask value provided is used when the IP addresses are mapped.

Limitations when mapping IP address ranges:

- You can provide a bits value that is smaller than the subnet mask size (CIDR prefix). For instance, if the subnet is a /20 you can define a CIDR prefix (bits) that provides a smaller IP range (/21, /22) for the range mapping.

- You cannot do the reverse. If the subnet is a /20, you cannot enter a CIDR prefix (bits) that provides a greater IP range (/19, /18) for the range mapping. If attempted, the UI displays an error.

To configure IP address range mapping, enter:

- Text description (optional)
- Source (protected site) and target (recovery site) ranges expressed in CIDR notation
- Source and target subnet masks
- Source and target gateways
- Source and target DNS servers

### Supported IP Address Mapping Combinations

Recovery plans support several IP address mapping combinations.

**Note** VMware Cloud DR does not support IPv6 remapping, and IPv6 addresses cannot be entered in the plan wizard.

Mapping	Configuration Support
DHCP	VMware Cloud DR supports DHCP mappings for Linux VMs.
Static IP addresses - Linux	VMware Cloud DR supports IP address mapping for Linux VMs, with adapters that have exactly one IPv4 and optionally one link local unicast IPv6. Only IPv4 addresses can be mapped. IPv6 configurations are preserved.
Static IP addresses - Windows	VMware Cloud DR supports IP address mapping for Windows VMs. Single IPv4 address per network adapter is supported. If present on the interface, IPv6 configurations are preserved.
Multiple Gateways	VMware Cloud DR supports mapping of multiple gateways per adapter for supported versions of Windows and Linux VMs.
Multiple Adapters	DHCP or static mappings of IPv4 addresses are supported per adapter for Windows and Linux VMs.
Multiple VM recovery	VMware Cloud DR supports IP address mapping of static and DHCP IPv4 addresses for multiple Linux and Windows VMs. IP address mappings can be specified in the plan wizard as individual IP addresses or IP address ranges.
Mapping rules do not apply to recovered VMs	If IP address mapping rules added to the plan wizard do not match any recovered VMs, remapping is skipped.

### OS Support for Static IP Address Mapping

Recovery plans support both Linux and Windows OS IP address mapping.

#### Windows IP address Mapping

VMware Cloud DR supports IP mapping of commonly used Windows versions:

- Windows 10
- Windows 2016

- Windows 2012 R2

Multiple network adapters per VM are supported. Single IPv4 address per network interface is supported. If present on the interface, IPv6 configurations are preserved.

If the network interface has no statically configured IPv4 addresses, no IP mapping is performed, even if matching IP address is found on the source.

### Linux IP Address Mapping

VMware Cloud DR supports IP address mapping for a VM if all NICs have exactly one IPv4, and optionally one link local unicast IPv6.

Also, since IP address mapping for Linux depends on vSphere guest customization, the source VM machine hostname must meet the naming requirements from [vSphere Customization Spec](#). Otherwise, IP address mapping is skipped.

Supported Linux versions:

- CentOS 7.0-1406
- CentOS 7.3-1611
- CentOS 7.5-1804
- RHEL 6 minimal

## Configure Script VM

You can add custom scripts to run on a dedicated VM during plan execution as a recovery step.

In the **Script VM > Recovery Steps** page of the recovery plan wizard, you can configure either Recover steps in the following ways:

- **Recovery Steps > Step Type > Recover Protection Groups | Individual VMs | All remaining VMs, files and groups > Pre/Post-recover actions for each VM > Run script in the Script VM.** These recovery step types allow you to execute a script on the script VM once per-recovered VM. You can configure each recover step to be run on the script VM before and/or after failover or failback operations.
- **Recovery Steps > Step Type > Other Actions > Add Action > Run script in the Script VM.** This recovery step executes a script on the script VM once, and can be placed before, in-between, or after other recovery steps. You can also add multiple other actions to run a script in a single recovery plan. For example, you can add one recover step action to launch a script before recovery is performed, in the middle of the VMs being recovered, and then another script launch after the VMs are recovered.

For more information on script VMs and recovery steps, see [Configure Recovery Steps](#).

Any script must be accessible to the script VM, since script execution is performed on the script VM.

Both Windows (Powershell) and Linux (Python) are supported for the script VM guest OS. You can use a script VM for testing and for real recovery operations.

Script VM restrictions:

- A recovery plan only runs on a single script VM.
- Scripts are run on the VM that you designate as the Script VM. Scripts are not run on other recovered VMs.
- VMware Tools must be installed and running on the script VM.
- The script VM must be available in the environment before the first step requiring a script call. The Script VM can be recovered as the first steps of a plan, or it must already be running on the target recovery SDDC. For example, do one of the following:
  - Deploy the script VM on both the protected site and the recovery SDDC. The script is called from both locations during failover and failback. Or,
  - Deploy the script VM on the protected site and add it to a protection group. Then, the script VM must be failed over as part of the recovery plan, and it needs a separate recovery step where the script VM is recovered before other VMs are recovered.

---

**Note** A script VM configured to run on multiple VMs is run sequentially on each target VM in the batch. Depending on the number of VMs targeted by the script VM, it can take longer by slowing the recovery completion process.

---

#### Procedure

- 1 From the **Recovery Plan > Script VM** page, select the Run script VM option.
- 2 Enter the VM name on which the script is run and the vCenter where the VM is hosted. You can select to run the same script for both real and test failovers, or you can use different scripts for real failovers and test failovers.
- 3 Click **Next** in the recovery plan Wizard.
- 4 On the Recovery steps page, click **Add step**.
- 5 In the Add step dialog box, you can choose either one of the Recover steps, or Other.
  - Select **Other Actions > Add Action > Run script in the Script VM** to run a script on the script VM once, and can be placed before, in-between, or after other recovery steps. Or:
  - Select **Recover Steps > Step Type > Recover Protection Groups | Individual VMs | All remaining VMs, files and groups > Pre/Post-recover actions for each VM > Run script in the Script VM** execute a script on the script VM once per-recovered VM.
- 6 When you are finished, click **OK**.

## Configure Recovery Steps

Recovery steps in a recovery plan dictate what actions the plan takes during a failover and the specific order in which those actions occur when the plan is running.

Recover steps in a plan consist largely of restoring individual VMs or VMs contained in protection group snapshots, and copying and restoring files and vSphere groups to the target recovery site.

With recovery steps, you specify scripts to run before and/or after a VM is powered on during a failover. You can also configure how you want the running plan to handle errors it encounters during failover, and also require user input on some steps before the plan continues running. Recover step types:

Step type	Description
<b>Recover protection groups</b>	Allows you to recover any of the protection groups associated with the plan. This step recovers and promotes each protection group, registers all VMs in a protection group snapshot with vCenter, then customizes and powers on the VMs. You can also choose to run a script VM before or after recovery of all VMs in the protection group.
<b>Recover individual VMs</b>	Allows you to recover individual VMs (you can select more than one) in the order that you specify. This type of step recovers then promotes each VM, registers each with vCenter, then customizes them and powers them on. You can also choose to run a script VM before or after recovery of an individual VM.  <b>Note</b> All VMs in the protection group referenced by the plan are recovered when the plan is run, in addition to any individual VMs you configure to restore here.
<b>Recover all remaining VMs, files, and groups</b>	Allows you to recover everything else referenced in plan mappings, in addition to any selected protection groups or individual VMs. This step recovers and promotes all remaining protection group and VMs and other files, registers the VMs with vCenter, then customizes and powers them on. You can also choose to run a script VM before or after recovery of all VMs.
<b>Other actions</b>	This step type allows you to define other steps while the plan is running: <ul style="list-style-type: none"> <li>■ <b>Wait for a fixed amount of time.</b> Pause the plan execution for a specified duration of time.</li> <li>■ <b>Wait for user input.</b> Requires that a user enter text in the running Recovery Plan to confirm a step before the plan continues executing.</li> <li>■ <b>Run script.</b> Run a script using the script VM configured for the plan.</li> </ul>

## Select Power Actions

For each recovery step in a plan, you must select a power action which determines if you want VMs to be powered on or off after recovery.

A recovery step for protection groups or VMs require one of these three power actions:

- Power on only VMs that were powered-on when the snapshot was taken.
- Power on all recovered VMs.
- Do not power on VMs.

## Protection Groups, Snapshots, and VM Power State

When a protection group takes a snapshot of its member VMs, it captures the power state of VMs in the group, either on or off. If the VM is powered on when a snapshot is taken, then after failover the VM is powered on when it is restored.

Conversely, If the VM is powered off at the time the snapshot was taken, the VM is powered off when it is restored.

VMs that are powered off when the snapshot is taken are not able to be powered on until after the storage vMotion of that VM to the SDDC completes.



To be sure, if your VMs must be powered on and ready for use immediately after recovery, you can override that default behavior when you set the recovery power state in your recovery plan to be on.

## Select Pre-recover and Post-recover Actions

In addition to configuring a power action for recovered VMs, you can also select actions that occur before and after a VM is powered on.

- Pre-recover action:
  - **Run a script in the script VM.** This pre-recover action requires entering the script path and any custom parameters.
- Post-recover actions:
  - **Wait for VM IP address.** Wait for the VM's IP address to be assigned before moving to the next step in the plan.
  - **Wait for VMware Tools.** Wait for VMware Tools to launch before continuing to the next step in the plan.
  - **Run script in the Script VM.** This requires entering the script path and any custom parameters.

## Script Configuration for Pre- or Post-Recover Actions

If you select to run a script as a pre-recover or post-recover action, you must configure script settings. This script VM is independent of the VMs that you recover as part of the plan.

To configure script execution, you must identify both the script and the script VM.

- You can only specify one virtual machine for script execution. The name of this virtual machine must be unique in its vCenter context.
- You must identify the script by its location on the script VM and by execution requirements. See Recovery steps.

There are two types of script execution:

- A pre-recover action script runs before powering on a recovered virtual machine.
- A post-recover action script runs after the recovered VM has been powered on. Post-action scripts can be paused for a certain amount of time to allow IP address configuration on recovered virtual machines, and can be paused to allow VMware Tools installation on recovered virtual machines.

## Script Actions

---

**Note** This action runs a script on a VM within the context of a plan execution recovery step. The script action takes an absolute path to the script on the script VM and a list of parameters that you can specify.

---

For Powershell scripts, only the absolute path to the 'powershell.exe' can be in the script path, and the Powershell script must be set in the parameters.

For example:

Configure

Execute a file **before** each VM is recovered on the Script VM. It can be any executable file, including PowerShell, Python, bash, and batch files. Will fail if the script returns a non-zero value.

Full path to script file

Examples: c:\scripts\post\_boot\_script.ps1 or /usr/bin/script.sh

Custom parameters

Example: --priority 1 --force

Failover and test parameters

Rollback and clean up parameters

The actual VM name and the new path are automatically passed to the script.

Timeout

 seconds

The action will fail if the script times out.

Cancel

OK

The timeout value (measured in seconds) is the amount of time to wait for this action before returning a failure on timeout. If the script takes more than 300 seconds, it will fail.

The script execution action returns an exit code for the script, where a non-zero exit code means failure, and an exit code of zero means success. At the time of recovery execution, you must supply the script VM credentials so that it is possible to run the script in the script VM remotely.

## Failback and Rollback Script Actions

Script actions have both a forward (failback) and backward (rollback) execution:

- If the script was run in the forward direction, a “`--failover`” flag is added to the parameters list so the script can distinguish between directions.
- If the script was run in the reverse direction, a “`--rollback`” flag is added to the parameter list.

## Configure Plan for Ransomware Recovery

You can configure a recovery plan for ransomware recovery with integrated vulnerability and security analysis.

In a recovery plan, select the 'Activate ransomware recovery' option to use the plan for ransomware recovery or ransomware recovery tests. Running a plan for a ransomware recovery test is the same as running a ransomware recovery plan, except that a ransomware recovery test has no option to recover VMs to a protected site.

When you activate ransomware recovery in a plan, VMs in the plan are charged for ransomware recovery for VMware Cloud DR. For more information on costs, see <https://www.vmware.com/products/cloud-disaster-recovery.html> and click the **Pricing** tab.

---

**Note** For more details on creating and setting up a recovery plan, see [Configure Recovery Plans](#).

---

When you click either the **Ransomware Recovery** or **Ransomware Test** buttons in the recovery plans list, you make VMs in the plan available for ransomware recovery.

---

**Note** If you configure test network mappings for a recovery plan, and the plan is activated for ransomware recovery, the plan will use the test network mapping by default when you run the plan.

---

Configuring ransomware for a recovery plan requires choosing from the following options:

**Edit plan - DB-backup-nolan**

**General**  
 Sites  
 Groups  
 ✓ vCenters  
 ✓ vCenter folders  
 ✓ Compute resources  
 ✓ Virtual networks  
 ✓ IP addresses  
 ✓ Script VM  
 ✓ Recovery steps  
 ➔ **Ransomware**  
 ✓ Alerts

**Ransomware**

☒ **Activate ransomware recovery**  
 Allow starting the ransomware test and recovery workflows with this plan.

VMs in this plan will be charged the ransomware add-on.  
[See pricing page.](#)

Confirm you understand the following:

☒ **VMs in this plan will generate the additional ransomware-add-on charge.**

**Security and vulnerability analysis during ransomware recovery**

☒ **Use integrated analysis**  
 Install sensors as VMs are restored in the recovery SDDC to perform security and vulnerability analysis.

☒ **Pause when starting a VM to manually remove production security sensors**  
 If your VMs have sensors from a security solution, such as Carbon Black, you should uninstall them when starting validation. This is to avoid polluting your production security solution with alerts occurring in the isolated recovery SDDC.

☐ **Do not use integrated analysis**  
 Use other tools to test for ransomware. VMs will start in fully isolated mode.

**CANCEL** **< BACK** **NEXT >** **FINISH**

- **Activate ransomware recovery.** Enable ransomware recovery for the plan. When you save the plan, you start being charged for ransomware recovery for all VMs protected by a recovery plan.
- **Use integrated analysis.** Enable integrated security and vulnerability analysis for VMs in the plan. When you run the plan and start VMs for recovery, VMware Cloud DR installs a security sensor on VMs, which enables ransomware analysis. For Linux VMs, you must install the security sensor manually, and uninstall any existing security sensors from the VM. For more information, see [Run Plan and Install Linux Launcher and Sensor](#). Integrated analysis does not trigger any additional charges.
- **Pause when starting a VM to manually remove production security sensors.** Pause when starting VMs during ransomware recovery so you can remove any production sensors or security software, which might interfere with VMware Cloud DR integrated analysis and impact the isolated recovery environment of the recovery SDDC. For more information, see [Uninstalling Sensors](#).

- **Do not use integrated analysis.** If you want to use your own security tools for ransomware recovery on your recovery SDDC, select this option. When selected, no VMware Cloud DR security sensors are installed when you start VMs during ransomware recovery.

## Alerts

As a part of recovery plan configuration, you can send email notifications when specific plan-related events occur.

You can select email addresses to send alerts to specific users when a recovery plan is out of compliance.

For information about configuring alert email addresses, see [Configure Email Alerts](#).

## Continuous Compliance Alerts

You can choose to send alerts when a recovery plan configuration or mapping is out of compliance. You must fix compliance errors before you run a plan as a failover or test failover, or the operation might fail. Compliance checks scan plans as follows:

- Every compliance check.
- Compliance warnings.
- Compliance errors.
- Once per week.
- When compliance check results changed.

## Recovery Plan Runtime Alerts

Recovery plans support email alerts for plan continuous compliance and during failover runtime:

- Failover execution status changes.
- Waiting for user input.
- Failover finished. Waiting for user commit.

## Recovery Plan Compliance Checks

Compliance checks scan recovery plans to verify and validate all configurations and mappings between a protected site and a recovery SDDC.

Each recovery plan provides continuous compliance checks to ensure that all plan configurations and mappings are valid before you run a plan. Compliance checks run regularly every 30 minutes.

Compliance checks also make sure that the specified protection groups are live on the protected site and are replicating successfully to the target recovery SDDC. A plan can become out of compliance if any of its conditions become violated because of environmental or plan configuration changes.

Failing compliance checks transition a recovery plan into a degraded health state of warning or critical. When a plan fails a compliance check, an email is sent to the recipients configured in the VMware Cloud DR settings. Health checks run on a per-plan basis. Some plans can have an OK health status, while others can be in degraded states at the same time.

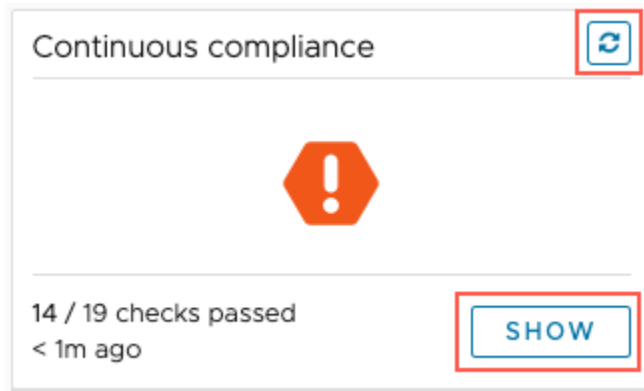
As a recovery plan is created, tested for compliance, and run, VMware Cloud DR maintains detailed logs of all the actions that were performed, the plans that were completed, and the compliance checks on those plans. You can generate and download these reports as a PDF or have them emailed on an automated schedule.

Compliance state for a plan does not restrict failover or test operations. Even if a plan is non-compliant in some areas, you can still run as a failover or test failover plan.

As a best practice, check your recovery plan compliance before you run it. You can also set email [Alerts](#) in the plan to notify you if a recovery plan is out of compliance.

**Note** Only run a recovery plan if the compliance check is green. If a recovery plan compliance check is not green, the plan will likely not run successfully.

To check plan compliance, select a recovery plan, and then check the compliance box:



To view plan compliance, click the **Show** button. To run the compliance report on demand, click the small refresh button at upper-right of the Continuous compliance box.

You can view the Compliance column in the list of recovery plans to quickly identify plans that are out of compliance:

Plans						
Plan	Status	Protected site	Recovery site	Groups	Compliance	
<input type="radio"/> [failback] plan1	Inactive	Prasanna-SDDC	→ site1	1	✓	1d ago
<input type="radio"/> [failback] plan2	✓ Failover committed	Prasanna-SDDC	→ site1	1	✓	1d ago
<input type="radio"/> plan1	✓ Failover committed	site1	→ Prasanna-SDDC	1	✓	2d ago
<input type="radio"/> plan2	✓ Failover committed	site1	→ Prasanna-SDDC	1	✓	2d ago
<input type="radio"/> pra-centos-4-plan	✓ Failover committed	site1	→ Prasanna-SDDC	1	!	5h ago

Recovery plan compliance reports check for the following configurations and mappings:

## Protected Site

Connection to source site	Checks the presence and availability of all source vCenters referenced in the plan and checks connectivity of the vCenter.
Replication health	Checks the relationship between source and destination replication to make sure replication is possible.
Datastores exist on source site	Checks for the existence and health of all source datastores defined in the plan. Datastore capacity is checked as part of snapshot replication health.
Protected groups replication schedule	Checks the snapshot retention for protection groups referenced by the plan. If snapshot retention for protection groups in the plan is less than 90 days, this item is flagged.
Networks exist on source site	Checks the existence of all source DNS servers and gateways listed in plan.
Resource pools exist on source site	Checks the existence of all source resource pools listed in the plan. If a cluster resource pool is mapped for either source or destination, the compliance check ensures that the Distributed Resource Scheduler (DRS) is activated.
Folders exist on source site	Checks for the existence of all source folders listed in the plan.

Protection groups are healthy	<p>Each recovery plan compliance verifies the existence and live status of protection groups referenced by the plan on the source site.</p> <p>Recovery plan compliance checks indicate if any VMs in the snapshot were not snapshotted, according to the protection group configuration. For example, if you configure a protection group to quiesce all VMs before taking a snapshot, and some VMs were not quiesced, then the recovery plan compliance report indicates the mismatch.</p> <p>Protection group health compliance is also evaluated by these factors:</p> <ul style="list-style-type: none"> <li>■ If the most recent snapshot fails, has no VMs, is missing some VMs, or has expired.</li> <li>■ If the most recent snapshot was not quiesced, and was configured to quiesce.</li> <li>■ If the most recent snapshot was taken on schedule or not.</li> <li>■ If a protection group snapshot schedule is expired.</li> <li>■ If a snapshot is empty (has no VMs), which might indicate that the queries defined in the protection group are not capturing any VMs.</li> <li>■ If a snapshot has any warnings during the operation, which indicates not all VMs were captured as configured in the protection group.</li> </ul>
Snapshot parity	Checks to make sure that at least one snapshot exists for every VM referenced in the plan.
Clusters exist in source site and cluster name	<p>Checks if all clusters referenced in the plan exist on the destination site, and checks to verify the name of the cluster as configured in the plan.</p> <p>If the cluster name changes or is deleted, then the plan is out of compliance. To be compliant, the name in the plan must match both the source and destination cluster name.</p>



## Recovery Site

Connection to failover site	<p>Checks to ensure that there is network connectivity to the recovery site (a recovery SDDC), and checks connectivity to all destination vCenters referenced in the plan..</p> <p>If you delete a recovery SDDC, then all compliance checks for that deleted SDDC are skipped. The plan compliance check will be healthy, but the plan cannot be run until the new recovery SDDC is added.</p> <p>Without having a recovery SDDC in place, a positive compliance report does not indicate DR readiness. It only means that the source site and protection group configuration is healthy. In order to ensure full compliance and DR readiness, you must deploy and fully configure all resources in an on-demand recovery SDDC and run a new compliance report to confirm DR readiness.</p>
vCenter server registered in failover site	<p>Checks to ensure that the vCenter server on the recovery SDDC is registered with VMware Cloud DR.</p> <p>The compliance check also verifies all vCenter authentication access and version compatibility, to make sure that source and destination vCenters are both accessible and version-compatible. Additionally, the compliance check ensures that a sufficient number of network ports are available in the vSwitch.</p>
Datastores exist on failover site	<p>Checks for the existence and health of all destination datastores defined in the plan. Datastore capacity is checked as part of snapshot replication health.</p> <p>Also checks that ds01 datastore is in a healthy state by checking for its existence, its maintenance mode status, free space statistics, and mounting and accessibility from hosts.</p>
Protection groups can be recovered in failover site	<p>Checks to ensure that all protection groups listed in the plan can successfully be failed over. Also ensures that at least one snapshot exists for every VM referenced in the plan.</p>
Networks exist on failover site	<p>Checks the existence of all destination DNS servers and gateways listed in plan. If you have defined IP address mappings, compliance checks report the presence of VMs without vSphere tools installed as warnings.</p>
Resource pools exist on failover site	<p>Checks the existence of all destination resource pools listed in the plan. If a cluster resource pool is mapped for either source or destination, the compliance check ensures that the Distributed Resource Scheduler (DRS) is activated.</p>
Folders exist on recovery site	<p>Checks for the existence of all destination folders listed in the plan.</p>

Tags in protection group queries exist in recovery vCenter	Checks to ensure that any vSphere tags and tag categories associated with your protected VMs also exist on the destination recovery SDDC.
Clusters exist in recovery site and cluster name	<p>Checks if all clusters referenced in the plan exist on the destination site and checks to verify the name of the cluster as configured in the plan.</p> <p>If the cluster name changes or is deleted, then the plan is out of compliance. To be compliant, the name in the plan must match both the source and destination cluster name.</p>
VMs can be restored in recovery vCenter	<p>Checks to ensure that there are sufficient snapshots that can be used to recover VMs in the plan.</p> <p>For all VMs referenced in the plan, the compliance check also verifies:</p> <ul style="list-style-type: none"> <li>■ The existence of VMs referenced by name in workflow steps on the source - VMs that are recovered outside of their Protection Groups.</li> <li>■ The existence of VMs referenced by name as targets for executing scripts.</li> <li>■ All VMs referenced by name are unique.</li> </ul>

## Orchestration

IP address mapping	If you have defined IP address mappings in a plan, compliance checks report the presence of VMs without vSphere tools installed as warnings.
Recovery steps	This compliance check ensures that all recovery steps can be run successfully without interference with other recovery steps. For example, it checks to determine if a named VM in a recovery step plan is actually in the protection group and can be recovered, and checks if elements and resources and items required to perform the steps can be found.
Script server recovered before script actions	<p>Checks that the script server is configured to be recovered before any scripts are launched.</p> <p>This compliance check also verifies the user-supplied credentials required to run any custom scripts specified in the plan.</p>

## Ransomware Recovery

Item	Compliance Check Description
VMware Tools	Checks to see if the correct version of VMware Tools (11.2) is installed on the VM. If VMware Tools is missing or at a lower version, this item is flagged.
Carbon Black integration status	Checks the connection to Carbon Black Cloud Integrated security and vulnerability analysis. If VMware Cloud DR cannot connect with Carbon Black Cloud, then this item is flagged.
Carbon Black workload appliance health	Checks the health of the Carbon Black Cloud workload appliance.
Protection group snapshot retention	Checks the snapshot retention for protection groups referenced by the plan. If snapshot retention for protection groups in the plan is less than 90 days, this item is flagged.

## Other Checks

VMC site health	Checks that VMware Cloud on AWS is healthy and reachable
VMC refresh token validity	Checks that the refresh token used for CLI operations is valid.
VMC proxy is running and reachable	Checks that the <a href="#">CloudDR-Proxy VM</a> is running and reachable.
VMC folder structure for file recovery is valid	Check that the folder structure on VMware Cloud on AWS is valid.

## View Compliance Checks

You can view the results of a compliance check of a recovery plan before you run it to make sure recovery will be successful.

### Procedure

- 1 From the left navigation, select **Recovery plans**.
- 2 Select a recovery plan.
- 3 In the Compliance pane, click the **Show** button.
- 4 In the Continuous compliance dialog box, click **Create PDF report** to generate a PDF of the compliance report.

## Viewing Recovery Plans

The recovery plans view shows currently defined plans along with plan summary information: current plan status, protected and recovery sites, and the last run compliance check results.

Each recovery plan protects a single vSphere protected site and a single vCenter within that site. A recovery plan maps a protected site to a recovery site (an SDDC on VMware Cloud on AWS) that can take over workload operations following plan completion.

DR plans

Plans							Create plan	Edit	Duplicate	Delete
Plan ^	Status	Protected site		Recovery site	Groups	Compliance				
<input type="radio"/> [failback] centos-vm1-plan	✓ Failover committed	byo-prasanna-...	→	dvx27-site	1	✓ 16d ago				
<input type="radio"/> [failback] win-server-plan	✓ Failover committed	byo-prasanna-...	→	dvx27-site	1	✓ 9d ago				
<input type="radio"/> [failback] win2012-vm2-pla	✓ Failover committed	byo-prasanna-...	→	dvx27-site	1	✓ 16d ago				
<input type="radio"/> [failback] win2012-vm3-pla	✓ Failover committed	byo-prasanna-...	→	dvx27-site	1	✓ 10d ago				
<input type="radio"/> centos-vm1-plan	✓ Failover committed	dvx27-site	→	byo-prasanna-...	1	✓ 16d ago				
<input type="radio"/> centos-vm13-plan	ⓘ Ready	dvx27-site	→	byo-prasanna-...	1	✓ 21m ago				
<input type="radio"/> win-server-plan	✓ Failover committed	dvx27-site	→	byo-prasanna-...	1	✓ 9d ago				
<input type="radio"/> win2012-vm2-plan	✓ Failover committed	dvx27-site	→	byo-prasanna-...	1	✓ 16d ago				
<input type="radio"/> win2012-vm3-plan	✓ Failover committed	dvx27-site	→	byo-prasanna-...	1	✓ 13d ago				

## Activating/Deactivating Recovery Plans

Recovery plans used for disaster recovery can be in either an active or deactivated state.

When a recovery plan is active, you can run the plan for failover, failback, or ransomware. New plans are active by default. A plan is automatically deactivated upon committing a successful failover or failback. You can explicitly re-activate a previously deactivated recovery plan by clicking **Activate plan**.

To run either a recovery or test recovery, the recovery plan must be active. Activating the plan also triggers continuous compliance checks for the plan. When a plan is inactive, then no compliance checks are run against it. The maximum number of active recovery plans that can be checked for compliance is 15.

The Active plan state is indicated in the Status column of the recovery plan list. Depending on whether the Test site is configured, an active plan can have either Ready or Ready (not testable) status. A recovery plan's inactive state is also indicated in the Status column when viewing a list of plans under the recovery plan view.

# Running a Recovery Plan for Failover

# 12

You can run a recovery plan for failover or test failover in disaster recovery situations.

Running a recovery plan for failover creates a runtime representation of the recovery steps defined in the plan, combined with other information available only when the plan starts running, such as the snapshot selection and the plan's failover operations.

You can also run a recovery plan for ransomware recovery. For more information, see [Run a Ransomware Recovery Plan](#).

Plan recovery steps apply to the plan itself and control the failover workflow. For example, a planned failover creates a workflow of operations based on the recovery steps defined in the plan. When a plan's recovery steps are run on the source site (power off VMs, replicate the last snapshot) and destination site (recover VMs in the predefined order).

An unplanned failover creates a different workflow based on the same recovery steps defined in the plan.

When a plan has finished executing and all of the steps in the running plan workflow have completed, you must explicitly [commit a failover](#) or [Rollback and Acknowledge a Failback Plan](#) in order for the plan to return to a ready state.

---

**Note** The most current version of VMware virtual hardware is 20, but VMware Cloud on AWS only supports up to version 19. If you have a VM in your environment running virtual hardware version 20, and you fail over the VM to VMware Cloud on AWS, VMware Cloud DR will automatically downgrade those VMs to virtual hardware version 19.

To ensure that protected and recovered VMs have the same virtual hardware versions, and to avoid the automatic version downgrade, you can set the default hardware version for new VMs by following the instructions [here](#).

---

**Note** VMware Cloud DR does not support recovering VMs to VMware Cloud on AWS SDDC with NFS-mounted external datastores including Amazon FSx for NetApp datastores, Cloud Control Volumes or VMware Cloud Flex Storage.

---

**Before you run a recovery plan for failover**

Before you run a recovery plan for failover, you must define the protected site, create protection groups, and make sure that your Recovery SDDC contains the same vSphere tags that exist on the protected site.

- **A protected site defined.** You must deploy and configure the DRaaS Connector on the vSphere protected site. For more information, see [Download the DRaaS Connector OVA from VMware Cloud DR DR UI](#).
- **Protection groups created with snapshot schedules configured.** To run or test a plan, you must configure protection group snapshots replication to a cloud backup site that is used for failover.
- **Tags present on the target Recovery SDDC.** For successful failover operations, you must ensure that any vSphere tags and tag categories associated with your protected VMs also exist on the target Recovery SDDC vCenter, or the compliance check displays warnings, and failback will not operate successfully

Read the following topics next:

- [How a Failover Recovery Plan Runs](#)
- [Configure VM Storage for Failover](#)
- [Run a Failover Plan](#)
- [Migrate a VM from the Cloud File System to the SDDC Datastore](#)
- [System Behavior During Failover](#)
- [Recovery Plan States](#)
- [User Input During Failover](#)
- [Failover Error Handling](#)
- [Failover Completion](#)
- [Running a Test Failover Recovery Plan](#)
- [Test Failover Example](#)

## How a Failover Recovery Plan Runs

You can run a recovery plan for failover immediately after a real life disaster event, or run it as a test failover before a real disaster occurs.

You can run a recovery plan for failover in the following ways:

- **Failover.** A failover operation is run following a disaster event when the source site is no longer available. The failover operation orchestrates on the destination site based on previously replicated snapshots. When failing over to a recovery SDDC, VMs that belong

to the protection groups defined in your recovery plan are recovered to the vCenter on the recovery SDDC. With a failover, you have the option to commit, cancel and rollback, or terminate. When a plan is committed, you can then create a failback plan to be used later during failback once the source site is recovered.

- **Test failover.** A test failover operation is similar to regular failover operation, but runs in the context of its own test execution environment only if the recovery plan has been configured to have separate configuration mappings for failovers and test failovers. Another difference is that by default, a test failover stops execution with every encountered error. In an actual failover, the default behavior is to continue running the plan even, if errors are found during recovery. You have the option to override the default behavior for both failover types by changing the default runtime settings prior to starting the failover operation. Finally, with a test failover your only option is to clean up the test plan. There is no failback function or capability to fail workloads back to the original protected site. For more information, see [Running a Test Failover Recovery Plan](#).

## Configure VM Storage for Failover

When you run a recovery plan for disaster recovery, you have the option to run failed-over VMs live on the cloud file system, or you can have VMs fully migrated to the vSAN datastore on the recovery SDDC.

VM storage configuration for a plan provides two options for failed-over VMs:

- **Run VMs live on the cloud file system.** After failover, VMs run live directly on the cloud file system, which offers a faster failover time for better RTO. Another benefit of running VMs on the cloud file system is that subsequent failback operations are also faster, resulting in less downtime. Some VMs recovered on the cloud file system might require performance that is better suited to vSAN. After a recovery plan operation completes, you can selectively Storage vMotion workloads to the vSAN datastore to improve performance. However, this will cause a longer failback process for those VMs, so do not Storage vMotion those VMs back to the cloud file system.
- Another benefit of using the cloud file system for disaster recovery operations is that it you will likely require fewer, and potentially less expensive, host types to operating during disaster recovery. You only have to size and scale your SDDC for CPU and memory to avoid adding hosts to meet requirements for vSAN capacity, which is often the constraint for sizing of an SDDC.
- **Full storage migration to recovery SDDC.** With this option enabled, VMware Cloud DR performs a full Storage vMotion migration from the staging datastore to the SDDC vSAN datastore as the final step of running a plan.
- This option increases RTO, as the plan cannot be committed or finished until all Storage vMotion operations are complete. At scale, this can take hours or days. Without committing a successful failover plan, even with all VMs up and running, you cannot then run a failback operation until the initial Storage vMotion is complete. Also during failback operations, there

will be a longer failback outage to recover workloads that have been migrated to vSAN. Fully migrated VMs provide higher IOPS performance, which is suitable for VMs that require higher performance, such as database VMs. This option might require more hosts on the cluster, depending on the size of the VMs.

**Best practice** If you have some workloads that you want to recovered on vSAN for performance reasons, create separate protection groups and recovery plans for those VMs. When you run a recovery plan and select the target datastore for recovery, all VMs in the plan and all associated protection groups recover to the same datastore. You cannot selectively recover VMs to one datastore or the other when running a plan.

**Best practice** Do not delete an SDDC from VMware Cloud if VMs are running live on the cloud file system. When deleting an SDDC from the VMware Cloud UI, ensure that there are no live VMs on the cloud file system. Deleting the SDDC will cause subsequent failovers to fail.

Failover - pno-testers ×

✓ Compliance check

✓ Snapshots

✓ Runtime settings

→ VM Storage

Preview

Confirmation

VM Storage

Select the storage for the failed-over VMs.

VM storage		Failover VM power-on speed	Failover and failback speed	Estimated I/O performance	Resources that might require additional hosts
<input type="radio"/> Run VMs live on cloud file system	<span>i</span>	●●●●●	●●●●●	●●●●●	CPU and memory
<input type="radio"/> Full storage migration to recovery SDDC	<span>i</span>	●●●●●	●●●●●	●●●●● migrating migrated	CPU, memory, and storage

CANCEL

< BACK

NEXT >

START FAILOVER

## Caveats for Running VMs Live on the Cloud File System

The following restrictions apply when running VMs live on the cloud file system:

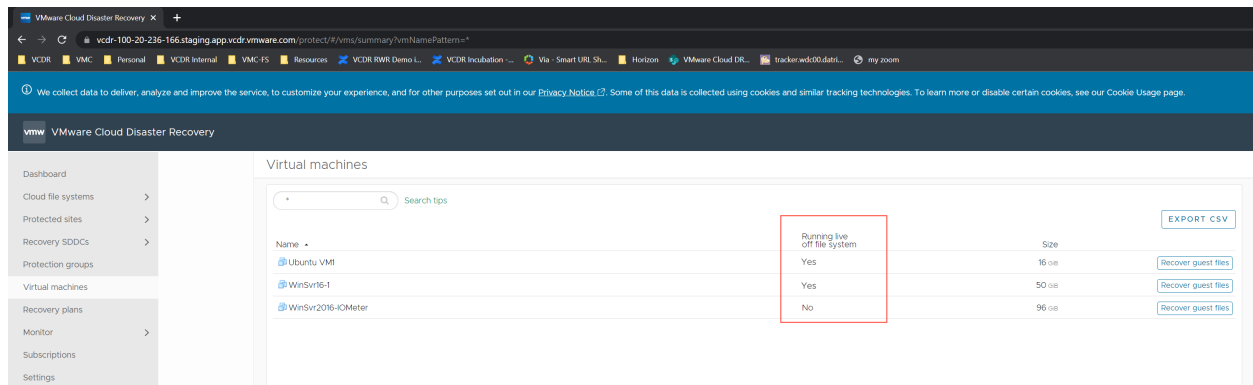
- **Configuration changes made to a VM running live on the cloud file system will not persist if VM is migrated to the vSAN datastore and then failed back.** When a VM is failed over to the cloud file system, any changes made to the VM such as renaming the VM, CPU, memory updates, are retained on the VM after failback. If the VM is migrated to the recovery SDDC vSAN datastore, those configuration changes will not persist if the VM is failed back to the protected site. In this situation, failback reverts the VM to the original configurations at the time of the first failover.



- **Other backup software might increase the VMDK file size of a VM running live on the cloud file system, causing failback to fail.** If a VM running live on the cloud file system is also being backed up by other software, the other software might increase the VM VMDK file size past an allowable size, which could cause failback of that VM to fail.
- **Manually performing a storage vMotion of a VM from a vSAN datastore to the cloud file system is not allowed.** Manual vMotion a VM from a vSAN datastore to the cloud file system is not supported.

## VMs List Indicates VMs Running Live on a Cloud File System


The Virtual Machines list shows all protected VMs, with a column to indicate if the VM is running live on the cloud file system:




Name	Running live off file system	Size
Ubuntu VM	Yes	16 GB
WinSvr16-1	Yes	50 GB
WinSvr2016-iCMeter	No	96 GB

## The Cloud File System Datastore and the SDDC Datastore

When you fail over VMs using the Run VMs live on the cloud file system option, you are running those VMs on the recovery SDDC datastore named 'ds01'.



When you fail over VMs using the 'Full storage migration to SDDC' option, VMs are initially staged on the 'ds01' datastore, and then are migrated to the SDDC vSAN datastore named 'WorkloadDatastore'.



If you start by running VMs live on the cloud file system, but decide later that you want faster performance, you can migrate the VM to the WorkloadDatastore on the recovery SDDC.

Currently, you can only migrate a complete VM from the cloud file system ds01 datastore to the Workload datastore. If you only migrate some disks of the VMs but not all of them, then the VM cannot be failed back to the protected site.

For more information, see [Migrate a VM from the Cloud File System to the SDDC Datastore](#).

## Run a Failover Plan

When you run a recovery plan for failover, a running instance of the plan recovery steps launches and the plan continues until a pause for user input or upon encountering an error (if configured), or if you cancel, roll back, or end the plan.

When you run a recovery plan for failover, a running instance of the plan recovery steps launches and the plan continues to completion, until a pause for user input step takes place, upon encountering an error (if configured), or if the user cancels, rolls back or terminates the plan.

---

**Note** For successful failover operations, ensure that any vSphere tags and tag categories associated with your protected VMs also exist on the target recovery SDDC, or the failback will not succeed.

---

## Procedure

- 1 From the left navigation, select **Recovery plans**.
- 2 In the list of plans, click the plan that you want to run. The plan you select must be in the Ready state.
- 3 Next, click the **DR Failover** button.
- 4 In the Compliance check page of the **Failover** wizard, you can view the Compliance check for the plan. Click **Next**.

---

**Warning** Running a failover when the plan compliance check is not all green will likely result in a failover failure.

---

- 5 In the Snapshots page, you can verify that the failover plan is using the snapshot you want when it fails over. By default, the most recent snapshot is selected, but if you want to select a different snapshot, click **Use different snapshot**.
- 6 In the **Select protection group snapshot** dialog box, select a snapshot you want to use for the failover operation, and then click **OK**.
- 7 Click **Next**. In the Runtime settings page, under Error Handling select one of the following two options:
  - **Ignore all errors.** Select 'Ignore all errors' to run the failover in unattended mode and to allow the failover operation to continue running, even when it encounters errors. The system automatically ignores all errors by default. You can still fix those errors at the end of the failover operation if the failover completes with partial success, by clicking **Retry all errors**.
  - **Stop on every error.** Select this option to run the failover in an attended running mode. This mode instructs the plan to stop on every error and waits for you to click Retry or Ignore and continue. This option is useful if you are running this plan as a test failover.
- 8 Click **Next**, and on the VM Storage page, select how you want the VMs stored once they are failed over:
  - **Run VMs live on the cloud file system.** After failover, VMs run live directly on the cloud file system, which offers a faster failover time for better RTO. Another benefit of running VMs on the cloud file system is that subsequent failback operations are also faster, resulting in less downtime. Some VMs recovered on the cloud file system might require performance that is better suited to vSAN. After a recovery plan operation completes, you can selectively Storage vMotion workloads to the vSAN datastore to improve performance. However, this will cause a longer failback process for those VMs, so do not Storage vMotion those VMs back to the cloud file system.

- Another benefit of using the cloud file system for disaster recovery operations is that it you will likely require fewer, and potentially less expensive, host types to operating during disaster recovery. You only have to size and scale your SDDC for CPU and memory to avoid adding hosts to meet requirements for vSAN capacity, which is often the constraint for sizing of an SDDC.
  - **Full storage migration to recovery SDDC.** With this option enabled, VMware Cloud DR performs a full Storage vMotion migration from the staging datastore to the SDDC vSAN datastore as the final step of running a plan.
  - This option increases RTO, as the plan cannot be committed or finished until all Storage vMotion operations are complete. At scale, this can take hours or days. Without committing a successful failover plan, even with all VMs up and running, you cannot then run a failback operation until the initial Storage vMotion is complete. Also during failback operations, there will be a longer failback outage to recover workloads that have been migrated to vSAN. Fully migrated VMs provide higher IOPS performance, which is suitable for VMs that require higher performance, such as database VMs. This option might require more hosts on the cluster, depending on the size of the VMs.
- 9 Click **Next**, and in the Preview page, you can view the steps that are taken when you finally run the plan. You defined these steps in the Recovery Plan Recovery steps page. To achieve low RTO, VMs are first recovered on the staging datastore. This recovery involves no data copy. VMs are powered on using the stored backups directly.
- If you have selected 'Full storage migration to recovery SDDC', VMware Cloud DR adds automatic migration tasks to the plan, which will run and must complete prior to plan completion and commit.
- 10 Click **Next**, and in the Confirmation page, to run this failover plan, enter **FAILOVER** in all upper case letters in the confirmation text box.
- 11 Click **Finish** to run the failover operation.

## Results

You can monitor the failover process in the VMware Cloud DR UI by clicking the plan to view its details. (You can also monitor the process in the recovery SDDC). After failover, once the VMs have been powered on, they are either migrated by Storage vMotion to the recovery SDDC vSAN datastore or migrated to the cloud file system.

After a failover operation finishes, you must commit the failover to make the effects permanent. When you commit a completed failover plan, the plan transitions to the committed state. You cannot start a failback to the source site until the plan is committed.

Until the completed failover operation is explicitly committed by an administrator, it can be rolled back (even following a successful completion). But after you Commit a plan, there is no rollback.

For more information, see [#unique\\_158](#).

## Migrate a VM from the Cloud File System to the SDDC Datastore

For workloads that require higher random IOPs, such as database VMs, you can use Storage vMotion to migrate those VMs from the cloud file system datastore to the recovery SDDC vSAN datastore.

After running a recovery plan with VMs using the Run VMs live on the cloud file system option, VMs in the plan are running on the datastore named 'ds01' on the SDDC. You can migrate those VMs to the SDDC datastore named 'WorkloadDatastore' using Storage vMotion on the vSphere Client on the SDDC.

---

**Note** You must migrate the entire VM and all its files and disks, or the VM cannot be failed back to the protected site.

---

### Procedure

- 1 Log in to the vSphere Client on the recovery SDDC, and expand the SDDC and cluster where the VM is running.
- 2 Select the VM and make sure that this VM is running on the ds01 datastore (the cloud file system datastore).
- 3 Right-click and select **Migrate**.
- 4 For migration type, select 'Change storage only' and then click **Next**.
- 5 In the table of datastores list, select WorkloadDatastore (the SDDC vSAN datastore) as the destination for the migration. Click **Next**.
- 6 Retain the default settings and click **Finish**.

## System Behavior During Failover

When you run a recovery plan for failover, the Orchestrator accurately moves your VMs to a failover site as defined in your plan and to optimize RTO.

### No Overwriting of Existing VMs

To avoid undesirable side effects of executing potentially erroneous plans, a running plan never overwrites VMs that exist on the destination datastore. If another VM is already present on the destination datastore at the exact datastore path matching the path of the recovering VM, VMware Cloud DR does not attempt to recover that VM.

Such VM recoveries are flagged as 'failed' during failover or test failover and the existing VMs are preserved. To make automatic recovery of these types of VMs possible, the conflicting VMs must be explicitly deleted from the destination datastore before running a failover operation.

## Batching

VMware Cloud DR recovers VMs in fixed size batches, also called substeps. VM batching is done to:

- Recover VMs concurrently to improve RTO
- Fine-grained retry on encountering errors
- Control the load on external components

All VMs in a batch are recovered concurrently, improving overall RTO. Recovering individual VMs can involve many different stages, such as retrieving remote snapshots if selected snapshots are not available on the recovery site, customizing IP addresses, reconfiguring VM to reflect failover mappings, powering-on VMs, and other configurations. Parallelizing stages of a failover plan across a set of all VMs in a batch improves the overall throughput and reduces RTO.

Improving error handling is another reason for failover batching. VMware Cloud DR supports retry of VM recovery on transient errors or following error remediation. When the “stop on all errors” setting is configured in a plan, the running plan will stop following a failed batch running with some VMs encountering errors.

Batch running is atomic in that the execution stops after all VMs in a batch have reached a terminal state, either successful recovery, or an error. Upon addressing the error condition, the failed batch can be retried. Similarly, VMware Cloud DR can automatically retry the operation of the last batch upon observing transient errors (for example, transient network connectivity problems).

As part of retry, VMware Cloud DR rolls back the batch operation and then restarts it. Batching reduces the throwaway work for large plans by limiting the rollback to the failing batch only.

Batching limits concurrency imposed on other external components involved in recovery. For example, batching limits the number of concurrent requests issued to the vCenter Server on the recovery site. When failover involves fetching remote snapshots or performing Storage vMotion, batching will naturally limit concurrency for these operations resulting in better overall system throughput.

## Skip VMs Not Registered with vCenter

If a VM is not registered with vCenter on the protected site, it will not be automatically recovered and registered with vCenter on the recovery SDDC.

## VM Tags and Tag Categories

The failover process associates vSphere tags with recovered VMs that were associated with the VM on the original protected site. However, the tags and their associated categories must be pre-configured on the Recovery SDDC for successful failover and fallback.

When you fail over VMs with tags, be aware of two possible environmental situations:

- A) Tags are present on both the protected site vCenter configuration and on the recovery SDDC.

- B) Tags are present on the protected site vCenter configuration, but the tags do not exist on the Recovery SDDC.

During a recovery plan compliance check, the system scans every VM in the protection group to make sure all tags associated with all VMs in the PG are available on the recovery SDDC.

If the category and the tags present on a VM do not exist on the recovery SDDC, then they will be flagged by the compliance checks as errors.

When you perform a failover in these two scenarios, you might have to perform extra steps before you can commit or failback the recovery plan.

In scenario (A), all categories and tags have been created on the recovery SDDC. After failover, each VM is started on the recovery SDDC and tagged with the same tags it had on the source site. In this situation, no extra action is required before committing or failing back the recovery plan.

In Scenario (B), where some tags are missing (and compliance check was failing, failover will proceed, but it will complete with errors.

Hence, before you try to committing the recovery plan, or failing back, after the failover completes with errors you manually have to create the missing tags on the recovery SDDC from the VMware Cloud on AWS Console. Then, you can proceed to commit the plan or run a failback operation with the plan.

## Migration Limits During Failover/Failback

During failover or failback, VMs in the recovery plan are migrated to the 'WorkloadDatastore' in vSphere. When vSphere migration limits are reached, the failover or failback tasks might report 'Resources currently in use by other operations. Waiting'.

For information, see [Limits on Simultaneous Migrations](#).

You can choose to bypass Storage vMotion migration to the recovery SDDC and run failed over VMs live on the cloud file system. This failover runtime setting uses cloud backup as highly available (HA) storage and runs recovered VMs directly from the cloud file system. With this option, failover is faster and there is no dependency on SDDC hosts for storage capacity. For more information, see [Configure VM Storage for Failover](#).

Keep in mind these restrictions when using this feature:

- When you select this option and perform a failover, the VMware Cloud DR software cannot be upgraded until you fail back the VMs.
- You cannot run two failover recovery plans at the same time that 1) both share some of the same VMs and 2) where both plans have storage migration set to 'Run VMs and files live on the cloud filesystem'. To avoid this, run one recovery plan at a time when you select this option.
- If you choose to 'Run VMs and files live on the cloud filesystem' when you fail over a recovery plan, make sure that if you want to fail back those VMs to a protected site, do not Storage vMotion those VMs. If you Storage vMotion the VMs, then you cannot fail them back.

## Recovery Plan States

Once you define a recovery plan, it can exist in one of several states:

Plan State	Description
Ready	The plan is ready to be run as a failover or as a test failover operation, or the plan can be deactivated.
Ready (not testable)	The plan is ready to be run as a failover operation, but no test site or test mappings are configured. This plan can be run as a failover operation, but it cannot be run as a test failover.
Incomplete (testable)	The plan is not fully configured for failover, but can be run as a test failover.
No recovery site	Plan has no recovery site defined.
Inactive	The plan has been deactivated and cannot be run until it is re-activated. Compliance checks are not run on a plan in this state.
Failover rolled back	A failover operation was successfully rolled back.
Rolled back with errors	A test failover was rolled back, but some errors were encountered during rollback.
Finished with no errors	A failover operation has completed successfully.
Finished with errors	A failover has completed but with some errors. Investigate any errors to fix anything that did not go as expected.
Test finished with no errors	A test failover has completed successfully.
Test cleaned up	A test failover has been cleaned up successfully.
Failed over with no errors	A failover operation completed successfully with no errors.
Test finished with errors	A test failover completed but with errors.
Test cleaned up with errors	A test failover was cleaned up but with errors.
Failover stopped	A failover has been stopped. Once a failover is stopped, it cannot be rolled back or run in forward direction.
Test stopped	A test failover has been stopped. Once a failover is stopped, it cannot be rolled back or run in forward direction.
Failover committed	A failover operation has been committed (cannot be rolled back).

During plan execution, a plan can exist in one of the following states:

Running Plan State	Description
Failing over	A failover operation has been run and is in the process of performing the plan's recovery steps on the target recovery site.
Retrying	Plan operation is in the process of restarting.
Testing	A failover test has been run and is in the process of performing the plan's recovery steps in a test environment.
Canceling failover	A failover operation is in the process of being canceled.



Running Plan State	Description
Rolling back failover	A failover operation is in the process of being rolled back.
Cleaning up test	A test failover is being cleaned up.
Waiting for user input	A failover operation has been started and is running, but is currently stopped as it waits for user input to proceed.
Failing over... Error	A failover operation encountered an error.
Testing ... Error	A test failover encountered an error.
Terminating	A failover operation is in the process of being stopped.
Terminating test	A test failover operation is in the process of being stopped.
Rolling back failover	A failover or failover test is in the process of being rolled back.

## User Input During Failover

Depending on how you configured the recovery plan for failover, when the plan launches and performs all failover steps defined in the plan.

When the failover begins running, the plan moves into the failover state. You can observe the running plan's progress from the plan's detail page.

You can perform the following operations during failover:

### Manual Intervention During Failover

While a plan runs, you can perform the following operations:

- Wait for user input
- Cancel/Cancel and Rollback
- Terminate

### Wait for User Input

A failover operation stops automatically if a plan's recovery steps are configured to stop the running plan and wait for user input before proceeding.

### Cancel and Rollback

A running failover can also be canceled upon completion (but only before you either commit or acknowledge the plan completion). Canceling and rolling back a running failover operation involves reversing the current running direction and rolling back the already completed steps, starting with the last run step until the step sequence stops.

When the failover completes, you can choose to Rollback. During a Rollback, each recovery step defines specific actions for both forward and reverse operations. Actions run in reverse direction cancel out actions that ran in the forward direction (for example, power-off or power-on, VM delete or VM create). A successful rollback implies the elimination of all side effects of a partially or fully completed workflow.

## Terminate

If the running plan is stuck for some reason and cannot make any progress in either forward or reverse direction, the terminate operation forces the failover or test failover to stop. Once a failover is terminated, it cannot be rolled-back or run in forward direction.

Terminate is a powerful, permanent operation you can use only in situations when rollback cannot make any further progress because of errors, or when it is for some reason desirable to retain the side effects of a partially completed operation. The failover runtime environment must be manually cleaned up to avoid conflicts with future failover operations.


## Failover Error Handling

During a failover, you can instruct a recovery plan how to handle errors in the recovery steps of the plan.

### Stop on Every Error - Retry or Ignore and Continue


If you configured a recovery plan to stop when it encounters an error ('Stop on every error'), the supported error handling actions are:

- **Retry.** Re-attempts the step, after some manual intervention to amend the error. The failed substeps are rolled back and re-ran, which can result in success or a repeated error.
- **Ignore and continue.** Continue the failover after a stop. For example, if the error is not critical and the error can be fixed later. This operation skips the failures from the failed substep and continues with the failover operation.



### Retry All Errors

If your plan is configured to 'Ignore all errors' when a failover operation is running, the workflow ends up in a partially completed successful state. To amend those errors, 'Retry all errors' inspects the failover from the start and retries all failed sub-steps. The type of retry is done in the same manner in when it runs the 'Retry' operation.



## Retry Events

When you retry failover steps that initially failed during running plan operation, VMware Cloud DR logs specific events to indicate which recovery steps are being retried.

**Note** For more information on events, see [Chapter 16 Monitor Events, Tasks, and Alarms](#).

Retry event message	Severity Level	Description
"Retry recovery of [count] VMs."	Info	Generated when the recovery of a batch of VMs is retried.
"Retry recovery of VM [VM name]."	Info	Generated when the recovery of an individual VM is retried.

## Ignore Events

When you run a failover or test failover recovery plan, you have the option to 'ignore all errors' when the plan is run. If any errors are encountered when the plan is running, VMware Cloud DR generates the following system events to inform you which recovery steps were ignored.

Ignore event message	Severity Level	Description
"Ignore failed recovery of [count] VMs and continue."	Info	Generated when the recovery of a batch of VMs fails during plan operation and is ignored (but can later be recovered manually).
"Ignore failed recovery of VM [VM name]."	Info	Generated when the recovery of an individual VM fails and is ignored (but can later be recovered manually).

## Failover Completion

When a failover completes, you can commit the plan, which causes the recovery site to become the active site, and for the plan to transition to the 'Failover committed' state, where it is deactivated and stops all compliance checks.

### Commit a Failover

When a recovery plan failover completes with no errors, commit the failover.

When you commit a failover recovery plan, its effects become permanent and the plan is moved to the Failover committed state. When you commit a plan, the plan transitions to the Failover committed state. Use caution when committing a failover.

If you roll back a failover (or clean up a test), or terminate any recovery plan execution, you need to explicitly acknowledge the plan completion to transition the plan into Ready state and make it available for future failover operations.

Until you explicitly commit the failover operation, it can still be rolled back (even following a successful completion). A rollback operation causes the failback recovery plan to run the plan's failover steps in the reverse order.

After you commit a failover plan, you cannot rollback the plan.

---

**Note** VMware Cloud DR service software cannot be upgraded if your environment has any uncommitted plans.

---

- 1 From the left navigation, click **Recovery plans**.
- 2 In the list of plans, select the recovery plan that was run as a failover.
- 3 In the **Commit** dialog box, enter a note for the failover.

Commit - Users

Commit will finish the failover task. Further actions for this task, like rollback or retry, will not be available anymore.

### Failover results

Failover status	Failed over with no errors	Errors	None
Time to recovery	3 minutes, 17 seconds		

### Failover notes

OK

### Failback plan

☐ Create a failback plan
 The failback plan reverses source and destination, and the corresponding mappings.

### Confirmation

Type COMMIT FAILOVER to confirm.

COMMIT FAILOVER

Cancel

Commit

- 4 Under Failback plan, select the Create a failback plan option, which creates a duplicate failback plan that reverses the order of steps in the plan. In the future, you can use the duplicate plan for failback.
- 5 Next, under Default datastore, select a datastore to use for the failback plan.


VMware Cloud DR will attempt to failback VMs to their original datastores. For example, if at the time of failback the original VMs still exist on the failover target, VMware Cloud DR preserves the datastores and folders of all its VMDKs and the VMs.

However, if there is no datastore with the same name on the failback site, or if the VMs no longer exist (or the site is new), VMware Cloud DR recovers the VMs to the datastore you select here.

- 6 In the conformation section, enter `COMMIT FAILOVER` in all capital letters, and then click the **Commit** button.

## Activate the Plan to Run It Again

To run the plan again, click **Activate**.



Once the plan has been activated, you can run the plan again for failover.

## Post-commit Actions

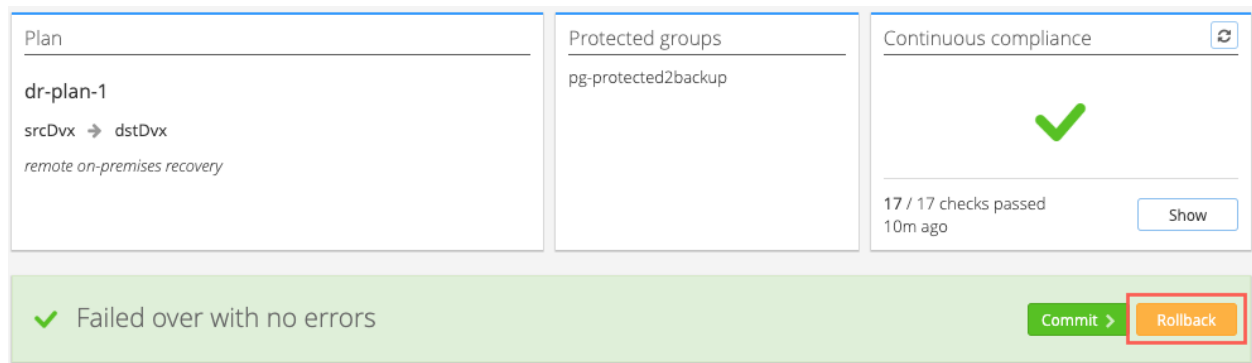
After the recovery plan is committed, you can run the plan again.

If you are preparing to run the recovery plan again for failover, on the protected vSphere site power off all VMs but leave them in place before you run the plan.

## Roll Back and Acknowledge a Failover

If you are not satisfied with a failover, you can perform a rollback operation, which causes the plan to run its failover steps in the reverse order, as specified in the plan.

For example:



After you perform rollback, you must acknowledge the operation.

Acknowledge - dr-plan-1

The test report will be emailed to @datrium.com.

Test results

Test status

Test cleaned up

Test errors

None

Time to recovery

9 seconds

Clean up errors

--

Test notes

Test completed as planned.

Cancel

Acknowledge

After you acknowledge the rollback, the recovery plan is placed back into the Ready state so it can be re-run:

Plan details [Summary](#) [Reports](#)

Plan

dr-plan-1

srcDvx → dstDvx

remote on-premises recovery

Protected groups

pg-protected2backup

Continuous compliance

✓

17 / 17 checks passed

2h ago

Show

✓ Ready

Failover

Test plan

Disable plan

## Running a Test Failover Recovery Plan

To test a recovery plan for disaster recovery, you can run the plan as a test failover.

A test failover runs in the context of its own test failover environment on the recovery SDDC, specified by the recovery plan's test mapping rules. The results of the test failover do not permanently affect a recovery SDDC. Test failovers cannot be failed back. VMs that are recovered for testing will be destroyed at the end of the disaster recovery test during the clean-up process.

The Orchestrator runs continuous compliance checks on a recovery plan to ensure that the plan is compliant with its own steps and mappings. But compliance checks alone do not catch all possible errors that can occur during a failover.

Run a test failover recovery plan on a periodic basis, so that you can determine if your plan works as you intended, to make sure that all IP address mappings work correctly on the VM guest OS, all your vCenter folders and settings can be recovered on the target, all scripts run correctly and in the proper order, and so on. Test failovers can also be used with periodic compliance reports to satisfy your company's disaster recovery preparedness auditing policy.

A test failover stops on the first failure by default. You can override all other default behaviors using custom options. For unattended plan runs, you can configure a test failover to run to completion while ignoring all errors.

Test failover operations give you the option of performing a full Storage vMotion from the staging datastore to the SDDC datastore to simulate a real failover, or to leave VMs on the staging datastore to cut down on the failover time, and to allow you to test and debug your failover faster. With this more cost effective preview feature, the SDDC can be substantially smaller in size because VMs are kept on the cloud file system datastore, eliminating the vSAN storage capacity constraints, which can incur costs.

Once you have fully run the failover test and have checked the VMs in the vCenter, you must clean up the test. Click the **Cleanup** button to roll back all test side effects on test completion.

## Creating an Optional Isolated Test Environment

When you configure a recovery plan, you can specify separate operating environments within the recovery SDDC, depending on whether the plan is run as an actual failover or failover test.

If you are performing a test failover, you can leverage this capability by testing with an isolated set of pre-configured networks, resource pools and/or folders on the recovery SDDC to ensure that there no impact on production. Isolating network segments for testing, for example, may help avoid issues like duplicate IP addresses if production failover networks are routed networks.

## Test Failover Example

An example of a test failover illustrates how a test failover works.

In this example, a failover operation is using a recovery plan that:

- Orchestrates the test failover to a recovery SDDC, for a protection group called 'Users' that regularly replicates snapshots to the cloud file system.



- Returns the VMs to the initial power state after failover. When the VMs recover, they are powered on or off based on their power state when the snapshot was taken. For example, if the VM was powered off when the selected snapshot was taken, then the VM will be powered off after the test failover is complete.

This example illustrates the main steps in running a recovery plan as a test failover:

- Select snapshots
- Define runtime settings
- Preview plan steps and confirm test operation
- Clean up test plan
- Acknowledge test plan cleanup

## Select Snapshots

When running a recovery plan as a test failover, select snapshots to use for the test.

- 1 From the left navigation, select **Recovery plans**.

- 2 In the list of recovery plans, click the recovery plan you want to test.

In this example, we are selecting a recovery plan called User DR, which contains VMs for end user computer systems. The Status indicates Ready, which means the plan is ready to run as a failover or as a test failover (or, the plan can be deactivated).

- 3 In the Plan details page, to test a plan click the **DR Failover Test** button.

- 4 You see the Snapshot page of the **Test plan** wizard.

This page allows you to select a snapshot to use for the test failover. By default, the Test plan wizard selects the most recent snapshot taken of the protection group VMs. If you want to select an older snapshot, click the **Use different snapshot** button.

- 5 In the **Select protection group snapshot** dialog box, you can select older snapshots, depending on the disaster recovery scenario you want to test.

- 6 Click **OK** to select the selected snapshot. Click **Next** to continue.

## Define Runtime Settings

When you run a test failover, you define the recovery plan runtime settings.

- 1 In the Runtime settings page of the wizard, you set two test failover parameters: error handling and test storage migration.

- 2 Under Error handling, select one of the following options:

- **Ignore all errors.** To test a failover plan quickly and see the results, you can choose to ignore all errors and run the plan unattended.

- **Stop on every error.** This option stops the test execution on the first failure, and then requires user intervention to resolve each error before continuing. This option is useful if your plan definition is complex and you want to troubleshoot errors in failover as they occur during plan operation.
- 3 Under the Test storage migration section, choose one of the following options:
- **Run VMs live on the cloud file system.** After failover, VMs run live directly on the cloud file system, which offers a faster failover time for better RTO. Another benefit of running VMs on the cloud file system is that subsequent failback operations are also faster, resulting in less downtime. Some VMs recovered on the cloud file system might require performance that is better suited to vSAN. After a recovery plan operation completes, you can selectively Storage vMotion workloads to the vSAN datastore to improve performance. However, this will cause a longer failback process for those VMs, so do not Storage vMotion those VMs back to the cloud file system.
  - Another benefit of using the cloud file system for disaster recovery operations is that it you will likely require fewer, and potentially less expensive, host types to operating during disaster recovery. You only have to size and scale your SDDC for CPU and memory to avoid adding hosts to meet requirements for vSAN capacity, which is often the constraint for sizing of an SDDC.
  - **Full storage migration to recovery SDDC.** With this option enabled, VMware Cloud DR performs a full Storage vMotion migration from the staging datastore to the SDDC vSAN datastore as the final step of running a plan.
  - This option increases RTO, as the plan cannot be committed or finished until all Storage vMotion operations are complete. At scale, this can take hours or days. Without committing a successful failover plan, even with all VMs up and running, you cannot then run a failback operation until the initial Storage vMotion is complete. Also during failback operations, there will be a longer failback outage to recover workloads that have been migrated to vSAN. Fully migrated VMs provide higher IOPS performance, which is suitable for VMs that require higher performance, such as database VMs. This option might require more hosts on the cluster, depending on the size of the VMs.

---

**Note** Test failover plans cannot be failed back.

---

- 4 Click **Next**.

## Preview Steps and Run the Test Plan

The last tasks in running a test failover recovery plan are to review the steps of the test plan and then run it.

- 1 In the Preview page of the **Test plan** wizard, review the steps that the plan will take when it is run, and then click **Next**.
- 2 In the Confirmation page, type the words `TEST PLAN` and then click the **Run test** button.

When the plan starts, an email alert is sent to users configured for notification in the recovery plan. You can watch the progress of the test failover from the Tasks list on the right side of the VMware Cloud DR UI. You can also observe the progress in vCenter on the recovery SDDC.

## Clean Up and Acknowledge Test Plan

If you are satisfied with the results of the test recovery plan operation, you can clean up the plan and acknowledge the results.

Cleaning up a test plan reverses the plan's instructions and undoes all of the failover operations and results by unregistering and deleting VMs on the test recovery site.

In this example, the test recovery plan finished with no errors. To view more information about what happened during the failover, you can expand each step of the plan operation.

- 1 After you review the test plan failover, click the **Clean up** button.
- 2 In the Clean up dialog box, review the details, and then enter **CLEAN UP TEST**. Click the **Clean up** button to initiate the test plan cleanup.
- 3 After cleanup, you need to explicitly acknowledge that a test failover ran successfully and you want to tear down the test failover. Click the **Acknowledge** button.
- 4 In the **Acknowledge** dialog box, review the clean up operation details, enter an optional note, and then click the **Acknowledge** button.

After you acknowledge a test failover, VMware sends an email to users configured for notification in the recovery plan, along with a PDF report describing the test failover, including a summary of the plan, the plan configuration, and logs for runtime, failover, and any errors.

You now have the option to run the plan again as a test failover, or as a regular failover operation.

# Run a Failback Recovery Plan

# 13

You can run a recovery plan to fail back from a Recovery SDDC to a protected site.

Failback from an SDDC returns only changed data. There is no rehydration, and the data remains in its native compressed and deduplicated form.

---

**Note** VMware Cloud does not support failback of a VM that has its disk geometry changed after failover.

---

You can run the failback recovery plan by clicking the **Failover from VMC** button.

Failback from a recovery SDDC runs several steps, including the following:

- VMs are powered off on the recovery SDDC.
- The last VM snapshot is taken following powering off the VM. The differences between the VM state at the time of recovery and failback are then applied to the snapshot used for recovery to construct a VM backup on the cloud file system for subsequent retrieval.
- These VM backups are then retrieved to a protected site system using a general forever incremental protocol.
- VMs are recovered to a protected site.
- Upon successful recovery, VMs are automatically deleted from the recovery SDDC.

Once a failback recovery plan is created from duplicating the plan and reversing its steps, the new failback plan operates the same way as any other plan. You can edit the plan to change the destination site to point to a new protected site. Or, you can change the vCenter mapping if the failback target site has more than one protected site.

You can also use a new protected site for failback, if the proper mappings are configured, but in this case incremental recovery is not possible. However, if VMware Cloud DR can find a VM with the same instance UUID, then an incremental recovery is performed. If VMware Cloud DR cannot find the same instance UUID for a VM, then a full recovery is initiated.

## General Caveats for Failback

Failback operations have the following restrictions:

- You cannot add an individual VM recovery step in failback plans. However, you can [Restore an Individual VM](#) individually from a protection group snapshot.

- VMware Cloud DR cannot failback any newly-created VMs post-recovery in the SDDC that match the PG name pattern or folder criteria defined in a recovery plan. This means that any new VMs that were created after recovery and that match PG name patterns in the recovery plan are not included when you perform a failback operation to a new or restored protected site. In this situation, an error is generated indicating that the new VMs were not failed back.
- VMware Cloud DR does not support recovering VMs to VMware Cloud on AWS SDDC with NFS-mounted external datastores including Amazon FSx for NetApp datastores, Cloud Control Volumes or VMware Cloud Flex Storage.

Read the following topics next:

- [Failback Process](#)
- [Failback Example](#)
- [Preparing for Failback](#)
- [Set the Plan Datastore](#)
- [Run a Failback Recovery Plan](#)
- [Cleaning up Failback Source](#)
- [Performing Repeated Failbacks with Same Snapshot Issue](#)

## Failback Process

The failback process consists of three general stages: undo, catchup, and completion.

### Power Off and Undo Stage

In the Undo stage, VMs on the failback target are restored to the state that matches snapshots used at recovery time.

The Undo stage of a failback ensures that all VMs protected by a recovery plan power off and their state matches the state of the snapshot selected for failover to a recovery SDDC. The Undo stage is necessary because VMs on the protected site might not be powered off prior to failover (for example, if the site became disconnected). These VMs continue to run and accumulate any arbitrary changes that diverge with the authoritative VM state in the recovery SDDC.

---

**Note** If VMs are missing or have diverged significantly from their failover snapshots, the Undo Stage can take a substantial amount of time. To avoid prolonged interruptions of service, recovered VMs in the recovery SDDC remain in powered-on state during this stage.

---

### Catchup Stage

In the Catchup stage, VM changes that have occurred while running on a recovery SDDC following failover are applied to the VMs on the failback target site

The Catchup stage of a failback operation applies incremental changes of all modified VM state in the recovery SDDC to the VM copies on the failback target. The duration of this stage depends on the change rate and the amount of time VMs are running on the Recovery SDDC.

During this stage, VMs remain in powered off state on both the failback target and on the recovery SDDC. Because of service downtime during this stage, it might be desirable to schedule it during the maintenance window. For example, you can edit the failback plan and inserting a wait-for-user prompt step prior to running the catchup.

## Completion Stage

After the failback recovery plan completes and is committed, it becomes inactive.

---

**Note** You can delete the deactivated failback plan if it is no longer needed, but if you delete the plan then the plan's compliance report will also be deleted. If you want to delete the plan, download its compliance report first.

---

## Failback Example

When you run a failback recovery plan protecting a single Windows VM, the plan moves through several steps.

- **Step 1.** Powers off a VM on the failback target to make sure that VM changes are applied safely.
- **Step 2.** Performs the Undo Stage of failback. Because in this case the recovered VM was still present on the target site, this stage is short. At the end of this step, VM contents match the snapshot used at failover time.
- **Step 3.** Powers off the VM in the recovery SDDC in preparation for Catchup.
- **Step 4.** Synthesizes the last snapshot that contains all changes incurred by the VM while executing in the recovery SDDC following the failover.
- **Step 5.** These changes are captured in the VM snapshot.
- **Step 6.** Customizes the VM for a failback target site. Specifically, VM IP addresses and other network parameters inside the guest are adjusted for the target site in accordance with the mapping rules.
- **Step 7.** Captures customization changes in a snapshot.
- **Step 8.** Runs Catchup. All VM changes are transferred to the target site and applied to the VM. At the end of this step, VM contents on the target site match the snapshot taken in Step 7.
- **Step 9.** Completes VM transformation to match the execution environment of the target site. VM networks are adjusted according to the mapping rules and Ethernet adapters are reconfigured to reflect the virtual networks. VM is powered on and a protection schedule is re-activated.

- **Final step.** On successful completion of the failback, VM is deleted from the recovery SDDC. This step is skipped if failback does not complete successfully preserve VMs on the SDDC.

Following a successful execution, a failback plan must be committed, like any other plan.

A running failback can be terminated at any point during execution. Manual cleanup is necessary following the terminate operation. If failback is terminated, all VMs remain present in the recovery SDDC and can be powered on to resume service.

## Preparing for Failback

To prepare for a failback, create the failback recovery plan, which is based on a committed failover plan.

A failback recovery plan is created in one of two ways:

- When a failover recovery plan is committed, you have the option of creating a duplicate plan with its steps in reverse order. This operation makes a copy of the original plan and adds [FAILOVER] to the name of the duplicated plan. For more information, see [Commit a Failover](#).
- You can duplicate the plan and manually reverse its steps and rename the plan to indicate it is a failover plan. To duplicate a plan, select it and then click the **Duplicate** button.

Once you create a failback plan, you can edit it like any other recovery plan. For example, if the original protected site cannot be recovered following a disaster event, you can adjust a failback plan adjusted to use another on-premises site as its target.

To adjust the plan, add a protected site following a normal process and then select the newly added site as a failover target for the failback plan.

## Set the Plan Datastore

Before you run a failback recovery plan, you must first set the plan's datastore configuration.

A failback recovery plan has a special page of the plan wizard for setting a default failback datastore. In situations where the original datastore from the protected site is not longer available, you need to choose a specific datastore to use for failback.

Typically, VMs fail back to their original datastores. For example, if at the time of failback the original VM still exists on the failover target, VMware Cloud DR preserves the datastores and folders of all its VMDKs, and the VM itself.

If the VM no longer exists (or the site is new), VMware Cloud DR recovers this VM to the datastore you define in the plan.

If the target vCenter Server is no longer available, you can select any other vCenter associated with a failover site while configuring the failback plan. You can edit a newly created failback plan and select a default datastore. You can edit all other failback plan mappings similar to mappings in a regular failover plan.

## Run a Failback Recovery Plan

You can run a failback recovery plan to restore a protected site.

### Procedure

- 1 From the left navigation, select **Recovery plans**.
- 2 Select the recovery plan you want to fail back.
- 3 To run a failback recovery plan, you must activate it. From the list of plans, select the failback plan and click the **Activate** button. If the failback plan has already been activated, then the **Failover from VMC** button is enabled.
- 4 Click the **Failover from VMC** button.
- 5 In the Compliance check page, review the compliance information, and then click **Next**.
- 6 In the Runtime settings page, notice that error handling is deactivated, because during a failback operation, all errors are ignored and cannot be retried. Click **Next**.
- 7 In the Preview page, you can review the steps the plan will complete when the failback is initiated. Click **Next**.
- 8 After the failback finishes, click the **Commit** button.

## Cleaning up Failback Source

After a successful failback, you do not need to perform cleanup on the failback source site.

However, if the failback partially completes, fails, or is ended early, you must delete from disk the stale VMs and files left behind on datastore ds01 in the 'Staging' folder on the failback source site (the recovery SDDC). Any files left behind might also be located in the WorkloadDatastore.

All VMs that remain present in the failback source site (the recovery SDDC) can be powered on to resume service. This applies only to VMs that were part of the failed failback plan that did not get processed. The recovery SDDC might have other VMs running either locally or from other failover plans.

During failback, VMware Cloud DR creates and retains some snapshots to ensure the safety of the VM state when it was running in the cloud, in case we need to recover the data even after failback has been committed. These snapshots are deleted at the time of the next failover. To have these snapshots deleted sooner, contact support.

## Performing Repeated Failbacks with Same Snapshot Issue

If you use the same snapshot for repeated failback or restore operations, it can trigger a longer, non-incremental restore.

This issue occurs with both high-frequency snapshots and standard-frequency snapshots. A full restore transfers the entirety of the VM contents from cloud snapshots, while an incremental restore only transfers the blocks which are different from the target snapshot.



Incremental restore is more rapid, and preferred where possible, but some scenarios make a full restore required.

## Standard-frequency Snapshots

For users whose protected site is a VMware Cloud on AWS SDDC with a protection group using standard-frequency snapshots, this issue can occur due to the loss of change tracking information during a failback or restore operation. This issue does not affect on-premises protected sites.

If your protected site is an VMware Cloud on AWS SDDC and you have performed a prior failover-failback for a protection group leveraging standard frequency snapshots, then any subsequent restores to a snapshot taken before the prior failback requires a full, non-incremental restore.

For example, in the following scenario:

```
snapshot 1  
  
snapshot 2  
  
[...]  
  
snapshot n  
  
failover-failback to snap n
```

Any subsequent restores to snapshots 1-n after this point require a full restore.

### **Workaround:**

Only perform failback with snapshots taken after previous VM restore operations.

## High-frequency Snapshots

This issue can occur when using high-frequency snapshots on any protected site type (on-premises or an SDDC).

If you are using high-frequency snapshots in a protection group, performing multiple consecutive VM restore operations (failback or single VM restore) performed without intervening snapshots results in a full, non-incremental restore.

### **Workaround:**

After performing a failback or VM recovery, take a snapshot before performing failback operations are attempted. This snapshot does not have to be used for failback or restore operations.

For example, the following scenario would result in a full restore:

```
snapshot 1  
  
failover-failback to snap 1  
  
failover-failback to snap 1
```

This procedure would restore the VM incrementally:

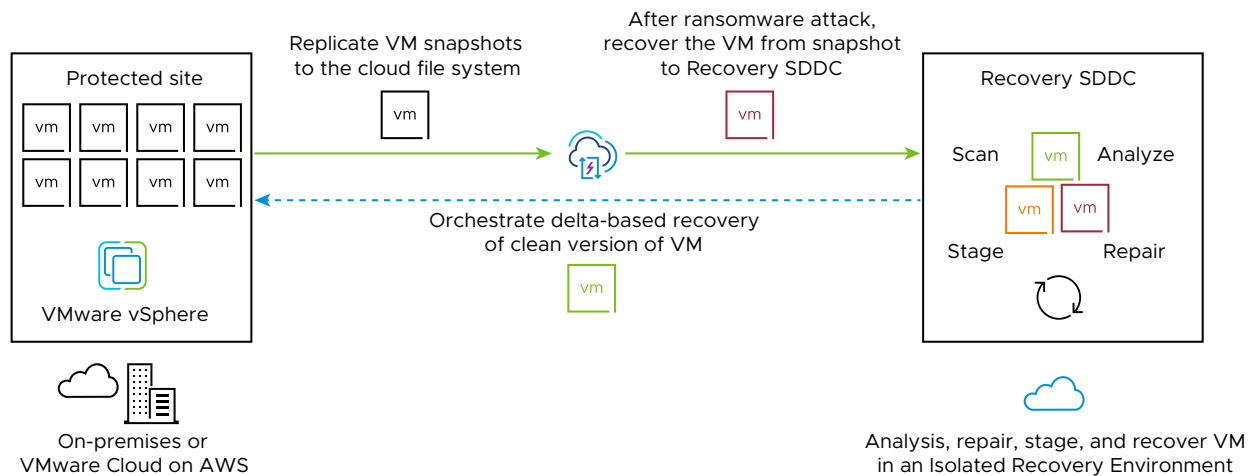
```
snapshot 1  
  
failover-failback to snap 1  
  
snapshot 2  
  
failover-failback to snap 1
```

# Ransomware Recovery


# 14

VMware Ransomware Recovery provides an isolated recovery environment (IRE) on a VMware Cloud recovery SDDC that allows you to inspect, analyze, and recover infected VMs before restoring them to a production environment.

Because VM snapshots are likely to be infected after a ransomware attack, you can use the recovery SDDC as a network-restricted IRE to perform analysis, remediation, and validation of compromised VMs without the risk of infecting other workloads in production. Once you validate clean VMs, you can recover them back to production.



The journey of VMs in a recovery plan for ransomware recovery go through these four states: backup, validation, staged, and recovered.



**Note** Uninstall any existing third party security software on VMs before running them in a recovery plan for ransomware, including any pre-existing Carbon Black security sensors.

Ransomware recovery states are defined as follows:

- In the **In backup** state, VMs have been replicated to a cloud file system and are available in a running recovery plan for validation on the recovery SDDC. You can choose to recover VMs based on the snapshot history of the protection groups to which the VM belongs.
- In the **In validation** state, VMs are moved to a recovery SDDC and powered on. Security sensors are installed on VMs so vulnerability and behavioral analysis and malware signature scanning can begin. You can patch VMs in the validation stage to address discovered vulnerabilities. You can also remove malware using security tools when VMs are in this stage. (You can also start VMs on the recovery SDDC with no sensors installed, if you want to use your own security software for analysis.)
- In the **Staged** state, you can start recovering VMs. VMs in this state have been validated and powered off. Staged VMs can be recovered to the protected site or returned to validation to produce a better recovery candidate.
- In the **Recovered** state, VMs have been recovered to the original protected site, or a different protected site.

Read the following topics next:

- [Ransomware Recovery and Disaster Recovery](#)
- [The Isolated Recovery Environment \(IRE\)](#)
- [Ransomware Recovery Workflow and Tasks](#)
- [Video: Ransomware Recovery Product Demo](#)
- [Activate Ransomware Recovery Services](#)
- [Access to Carbon Black Cloud](#)
- [Change Country for Ransomware Data Analysis](#)
- [Network Isolation Levels](#)
- [Video: Network Isolation for Ransomware Recovery](#)
- [Configure Plan for Ransomware Recovery](#)
- [Run a Ransomware Recovery Plan](#)
- [Start VMs in Recovery SDDC](#)
- [Search for VMs in a Plan](#)
- [Try a Different Snapshot](#)
- [Snapshot Timeline: Change Rate and Entropy Rate](#)
- [Integrated Security and Vulnerability Analysis](#)
- [Discard or Detach VMs](#)
- [Open Security Console](#)

- [Guest File Recovery](#)
- [Copy IP Address for VM Access](#)
- [Badging Snapshots](#)
- [User Notes](#)
- [Restart Validation from the Staging Snapshots on recovery SDDC](#)
- [Restart Validation from Protected Site Snapshots](#)
- [Power Off and Stage VMs](#)
- [Recover VMs in Protected Site](#)
- [Recover VMs in Other Protected Site](#)
- [End Ransomware Recovery](#)
- [Ransomware Events](#)
- [Manual Sensor Installation](#)
- [Uninstalling Sensors](#)
- [Create a Custom Network Isolation Level](#)

## Ransomware Recovery and Disaster Recovery

Ransomware recovery has similarities with disaster recovery, but there are key differences that require unique automation workflows to guarantee the security of production workloads, shorten recovery times, and reduce data loss.

For instance, during ransomware recovery you cannot be certain that snapshots are not infected without validating the snapshots. Unlike disaster recovery, the most recent snapshots are likely compromised in a ransomware attack. Ransomware recovery must be always performed under the assumption that malware is embedded in the snapshot data.

Snapshots must be either validated as free of infection, or malware must be removed during the ransomware recovery process to avoid reintroducing ransomware into a production environment. Because snapshots are potentially infected, they should not be directly recovered to the production environment. Instead, snapshots must be initially restored into the Isolated Recovery Environment (IRE) for security analysis.

The ransomware recovery workflow offers several preconfigured VM [Network Isolation Levels](#) that can be changed with the push of a button. To prevent malware lateral movement in the IRE, each selected VM snapshot is powered on in a quarantine isolation level by default. Ransomware recovery often involves multiple recovery iterations to identify a malware-free backup or to "clean" from a chosen snapshot (malware removal). To avoid reinfection, you restore snapshots to the production environment only following validation in the recovery workflow.

You can use a recovery plan for both disaster and ransomware recovery, but enabling ransomware recovery requires explicit [Configure Plan for Ransomware Recovery](#).

The following table provides a comparison of the main differences between ransomware recovery and disaster recovery:

	Disaster Recovery	Ransomware Recovery
Objective	Recovery of business operations with minimal downtime and data loss.	Recovery of compromised workloads with minimal loss, while providing data integrity and security assurance.
Recovery site	The recovery SDDC in VMware Cloud on AWS. Workloads can be made externally available immediately upon recovery.	VMs are validated and cleaned (malware or other malicious software removed) in the IRE in the recovery SDDC prior to restoring them to the protected site to avoid reinfection of production workloads. Workloads are made externally available only following security validation.
Snapshot retention	Limited snapshot retention time to accommodate the disaster recovery use case, with the latest snapshot being the primary recovery candidate.	Longer snapshot retention. Ransomware 'dwell time' (duration between infection and manifestation) can range from weeks to months. VMware recommends at least 3 months retention for ransomware recovery.
Recovery Point Objective (RPO)	Disaster recovery RPO is configured through snapshot schedules based on business needs.	Recent snapshots are likely encrypted and/or infected and might not be suitable for recover, which can result in higher RPO. Early detection improves RPO.
Recovery Time Objective (RTO)	Instant power-on in the recovery SDDC.	Higher RTO due to security validation in the IRE prior to recovery to the protected site.
Orchestration	Full automation with recovery plans that allow for unattended recovery. A recovery plan recovers all protected VMs to a recovery site.	Iterative recovery requires control over ransomware recovery workflow state transitions. Partial recovery of subsets of protected VMs is possible.
Tools	VMware Cloud Disaster Recovery Orchestrator.	In the addition to using the VMware Cloud Disaster Recovery Orchestrator, requires backup validation and remediation with integrated or third party security tools. Preconfigured network isolation levels available for use during validation.

## The Isolated Recovery Environment (IRE)

With VMware Ransomware Recovery, a recovery SDDC is transformed into an on-demand, cloud-based (IRE) with predefined network isolation levels.

The recovery SDDC provides a network-restricted IRE on VMware Cloud on AWS that does not require building an environment from scratch and patching together different tools and hardware. You can use predefined [Network Isolation Levels](#), or you can [Create a Custom Network Isolation Level](#) to match your security needs.

After a ransomware attack, you can launch a recovery plan and select VMs from a deep snapshot history to be placed into an IRE for forensic analysis and validation. When you start VMs in validation on the IRE, VMware Ransomware Recovery provides [Integrated Security and Vulnerability Analysis](#) that analyzes each VM in recovery for suspicious OS behaviors, malware file signatures, and known vulnerabilities.

When you have succeeded in finding clean VMs, you can orchestrate those VMs back to a protected production site.

## Avoid Private Connections Between Production Environments and the IRE

In order to ensure the security and integrity of the IRE, when validating VMs on the IRE, do not establish any form of private network connectivity between your production environment and the IRE.

Specifically, do not:

- Use Direct Connect to connect on-prem infrastructure to the IRE.
- Use a VPN to connect on-prem or cloud infrastructure to the IRE.
- Stretch L2 networks from on-prem to the IRE.
- Add the IRE SDDC to an SDDC group that enables private connectivity to other SDDCs and native VPCs.

## Ransomware Recovery Workflow and Tasks

The ransomware recovery workflow includes several sets of tasks based on ransomware recovery states:

---

**Note** During ransomware recovery, VMware Cloud DR can process a maximum of 50 VMs at a time. This limit applies to starting VMs, staging VMs, recovering VMs, powering off VMs, changing network isolations, selecting new snapshots, and so on. For example, if 25 VMs are currently being started on the recovery SDDC, you can only operate on 25 more VMs until the other 25 VMs have finished starting.

---

Task	Actions	Ransomware Recovery state
Prepare for ransomware recovery	<p><a href="#">Activate Ransomware Recovery Services</a> for ransomware recovery and optionally enable integrated vulnerability and behavior scanning.</p> <p><a href="#">Create a Protection Group</a>. A protection group replicates snapshots on a regular schedule to a cloud file system. VMs in the group are included in a recovery plan.</p> <p><a href="#">Create a Recovery Plan</a>. When you create a recovery plan for ransomware, you configure many settings, such as selecting protection groups and <a href="#">Configure Plan for Ransomware Recovery</a>.</p>	N/A
Start Recovery Plan for Ransomware Recovery	<a href="#">Run a Ransomware Recovery Plan</a> for ransomware recovery.	In backup
Start VMs on the Recovery SDDC	<p>When you start VMs in validation, it is considered an 'iteration'. Every time you change snapshots of VMs in validation, it is a new iteration. You can iterate VMs in validation as many times as you want.</p> <p>When you start VMs in validation, the following behaviors occur:</p> <ul style="list-style-type: none"> <li>■ VMware Cloud DR uses Live Mount to instantly power-on the selected VM snapshot on the recovery SDDC.</li> <li>■ A security sensor is automatically installed on Windows VMs, if the recovery plan is enabled for integrated security and vulnerability analysis, and if VMware Tools version 11.2 or later is installed on the VM.</li> </ul> <p>For Linux VMs, VMware Tools version 11.2 is also required, but you must manually install the Linux sensor. For more information, see <a href="#">Manual Sensor Installation</a>.</p> <p>Any pre-existing third party or Carbon Black sensors should be uninstalled before proceeding with ransomware validation. For more information, see <a href="#">Uninstalling Sensors</a>.</p> <p>When you start VMs for ransomware recovery, VMware Cloud DR begins <a href="#">Integrated Security and Vulnerability Analysis</a> of the VMs.</p> <hr/> <p><b>Note</b> Currently, malware signature scans for Linux VMs do not report progress in the UI. The scan still occurs, but the progress indicator remains "in progress" even after the scan is finished.</p>	In validation



Task	Actions	Ransomware Recovery state
Iterate security analysis and remediation	<p><a href="#">Snapshot Timeline: Change Rate and Entropy Rate</a> to analyze snapshot change rate and entropy rate.</p> <p>You can <a href="#">Try a Different Snapshot</a> for VMs, for example, so you can find a snapshot which has the least amount of entropy and a higher rate of compression.</p> <p><a href="#">Badging Snapshots</a> for VMs you know are clean or infected.</p> <p>Monitor security events and alerts generated by integrated behavioral analysis and malware scanning.</p> <p><a href="#">Integrated Security and Vulnerability Analysis</a> from an integrated vulnerability scan.</p> <p>Manually patch vulnerabilities and remove malware.</p> <p><a href="#">Guest File Recovery</a>. Recover individual files and directories from a VM snapshot, if you need to replace damaged files with the original one from an earlier date.</p> <p>Change the <a href="#">Network Isolation Levels</a> for VMs running on the recovery SDDC.</p> <p><a href="#">Discard or Detach VMs</a>, if you validate VMs from different snapshots.</p> <p><a href="#">Copy IP Address for VM Access</a> of a VM and open in vCenter.</p> <p><a href="#">Open Security Console</a> for threat hunting and remediation.</p>	In validation / In backup (During analysis and remediation, VMs can be moved between in the backup and in validation states.)
Power off and stage validated VMs	<p>When you <a href="#">Power Off and Stage VMs</a>, VMware Cloud DR takes a snapshot of the VMs, which are used for recovery to a protected site.</p> <p><a href="#">Restart Validation from Protected Site Snapshots</a>. Restart validation iteration for VMs with the same or different snapshots.</p> <p><a href="#">Badging Snapshots</a> that you know are clean or infected.</p>	Staged
Recover VMs	<p>You can <a href="#">Recover VMs in Protected Site</a> where it originated.</p> <p>You can also <a href="#">Recover VMs in Other Protected Site</a>, if you have other protected sites configured.</p>	Recovered
End Recovery	<p>When you are finished validating and recovering VMs, you can <a href="#">End Ransomware Recovery</a> by stopping the recovery plan.</p>	N/A

## Video: Ransomware Recovery Product Demo

Ransomware recovery with VMware Cloud DR gives you the tools to fight back against attacks with an isolated recovery environment (IRE) and a deep history of immutable VM snapshots.

Watch this product demo video to understand how to recover successfully from a ransomware attack using VMware Cloud DR ransomware recovery.



(Ransomware Recovery with VMware Cloud DR)

## Activate Ransomware Recovery Services

To use ransomware recovery with integrated security and vulnerability analysis, you must first activate ransomware recovery services.

Enable integrated security and vulnerability analysis in your recovery plans to recover from a ransomware attack (or to test one). When you run a plan for ransomware recovery, integrated vulnerability and behavioral analysis and malware signature scanning begins.

---

**Note** Performing this task requires that your user has the Organization owner role. If you are a MSP partner using the CPN (Cloud Provider Network) console, you must activate Ransomware Recovery Services in each tenant organization by a user with Provider Admin permissions.

---

### Allowing Activation of NSX-T Advanced Firewall

VMware NSX Advanced Firewall Advanced Firewall is required to enable [Network Isolation Levels](#) levels. NSX Advanced Firewall Advanced Firewall is an on-demand, chargeable feature that activates a full range of network isolation levels when performing validation on the recovery SDDC.

You can authorize VMware Cloud DR to automatically activate the advanced firewall only for the duration of ransomware recovery or testing. You can pay for the firewall service on-demand, or you can subscribe to NSX Advanced Firewall Advanced Firewall or explicitly enable it in the VMC Console. When you enable NSX Advanced Firewall Advanced and run a ransomware recovery plan, VMs in validation are started in the Quarantined+Analysis network isolation level.

If you activate integrated analysis but do not enable NSX Advanced Firewall Advanced Firewall, and then run a recovery plan, VMs start on the recovery SDDC with full outbound connectivity. To create your own custom network isolation level, see [Create a Custom Network Isolation Level](#).

---

**Note** Applying or changing a network isolation level for VMs overwrites any firewall configurations that were previously set for those VMs.

---

For more information, see [NSX-T Advanced Firewall for VMware Cloud on AWS](#).

Activating ransomware recovery services requires the following user roles: Organization Owner, Global Console Admin, and Orchestrator Admin.

## Procedure

- 1 From the left navigation, select **Settings**.
- 2 Under Integration, click the **Ransomware Recovery Services** button.
- 3 In the Ransomware services integration dialog box, click the **Activate Integrated Analysis** button.
- 4 Select the country where diagnostic data will be analyzed.
- 5 Read and then confirm each of the items described in the dialog box, and then click **Activate**.

If you have a recovery SDDC deployed, then a security workload VM is installed in the SDDC when you activate security and vulnerability scanning. If you have not yet deployed a recovery SDDC, then the workload VM is installed at the time you deploy the SDDC.

After activating security and vulnerability scanning, when you run a recovery plan for ransomware and start VMs in validation, security sensors are installed on Windows VMs. For Linux VMs, you must manually install the security sensors on VMs. For more information, see [Manual Sensor Installation](#).

Integrated analysis might not be compatible with preinstalled security software on VMs. You can [Configure Plan for Ransomware Recovery](#) to pause before VMs start in validation, so you can uninstall the security software when you run the recovery plan and start VMs in validation. To uninstall existing sensors, see [Uninstalling Sensors](#).

- 6 After you activate scanning, you can click **Allow Activation of Advanced Firewall**. If you already have a subscription to the advanced firewall active in your SDDC, the option is already enabled. If you do not have a subscription to NSX Advanced Firewall Advanced Firewall, you can buy one here: [NSX-T Advanced Firewall for VMware Cloud on AWS](#).
- 7 Confirm that you acknowledge the statements in the dialog box, and then click **Activate**.

## What to do next

Once you have activated ransomware recovery services, you can [Create a Protection Group](#) and a [Chapter 11 Set Up Recovery Plans in VMware Cloud DR](#). You can then recover VMs if you experience a ransomware attack.

## Access to Carbon Black Cloud

During ransomware recovery, VMs must be able to reach Carbon Black Cloud servers to send security analysis data.

Typically, ransomware recovery automatically ensures connectivity for VMs to all required Carbon Black Cloud Fully Qualified Domain Names (FQDN)s for the recovery SDDC by programming NSX Advanced Firewall, so no further actions are necessary.

However, if the recovery SDDC has no outbound connectivity to the internet and all outbound traffic is routed through some other external network or a firewall, you might need to perform more configuration steps.

For example, your recovery SDDC might be connected to the on-premises protected site over AWS Direct Connect (DX), or by VPN advertising the default route (0.0.0.0/0) that sends all outbound traffic through the on-premises corporate firewall. In that case, the corporate firewall needs to be configured to enable outbound access to Carbon Black Cloud FQDNs, as described here: [Configure a Firewall](#).

Similarly, the outbound traffic from the recovery SDDC can be routed to a Security Virtual Private Cloud (VPC) in AWS for analysis using a default route advertised by VMware Managed Transit Gateway (vTGW). In that case, the Security VPC must allow outbound connectivity to Carbon Black Cloud FQDNs.

In most cases, you can rely on ransomware recovery to automatically ensure all needed external connectivity for all [Network Isolation Levels](#). However, if the recovery SDDC has no connectivity to the internet, you might need to ensure that Carbon Black Cloud FQDNs are reachable from your VMs in the recovery SDDC by reconfiguring your external firewall.

## Change Country for Ransomware Data Analysis

You can change the country where integrated security and vulnerability analysis data is analyzed.

### Prerequisites

To change the country used to analyze ransomware data, first deactivate ransomware recovery services, then re-enable the services and select a new country. Deactivation can take up to two days to complete.

### Procedure

- 1 From the left navigation, select **Settings**.
- 2 Under Integration, click the **Ransomware Recovery Services** button.
- 3 In the **Ransomware services integration** dialog box, click **Deactivate**.
- 4 In the **Deactivate integrated services** dialog box, confirm that you want to deactivate the services and then click the **Deactivate** button.
- 5 After deactivation has completed, you can [Activate Ransomware Recovery Services](#) and select a new country to process and analyze ransomware data.

## Network Isolation Levels

When you activate ransomware recovery services and start VMs in validation, they are put into a 'Quarantined+Analysis' network, which means those VMs can only connect over the internet to integrated security and vulnerability servers on Carbon Black Cloud and to basic network services like DNS and NTP.

Once a VM has been moved into the validation stage, you can choose different network isolation levels for one or more VMs on the recovery SDDC, depending on your method of analysis and how much isolation you require.

You can select multiple VMs in a running plan, and then from the **Other Actions** menu, select **Change Isolation**.

VM list

VM DETAILS

TRY DIFFERENT SNAPSHOT

POWER OFF AND STAGE

OTHER ACTIONS ▼

☐ VM name ▾

Recovery state

Protection group

<input checked="" type="checkbox"/> IRR-Vm0-centos75_irr_small-0	In validation	PG0
<input type="checkbox"/> IRR-Vm1-centos75_irr_small-0	In validation	PG1

Discard VM in recovery SDDC

Detach VM to recovery SDDC

Open vCenter

Guest file restore

Change isolation

Set base snapshot badge

Or, you can change the network when a VM is started and in validation, by clicking **Change Isolation** from the Toolkit panel.

▼ Toolkit

Inspect, patch, and validate

OPEN VCENTER

COPY IP ADDRESS

Recover data from other snapshots

GUEST FILE RESTORE

Quarantined + Analysis

CHANGE ISOLATION

**Note** You can also create your own [Create a Custom Network Isolation Level](#).

The following table describes the ransomware network isolation levels and the type of network access allowed for each. A check mark means that the type of connectivity is allowed.

**Note** Applying or changing a network isolation level for a VM overwrites any previous firewall configurations that were previously set for the VM.

Some isolation levels require NSX Advanced Firewall. When NSX Advanced Firewall is enabled, VMware Ransomware Recovery creates firewall rules for the various network isolation levels during recovery operations. When all recovery plans for ransomware are ended, all firewall groups and rules are deleted on the recovery SDDC.

If you already have a subscription to the advanced firewall active in your SDDC, VMware Ransomware Recovery leverages the already deployed NSX Advanced Firewall and does not activate or deactivate any NSX Advanced Firewall services, and there are no additional on-demand NSX charges incurred. For more information, see [VMware NSX Advanced Firewall for VMware Cloud on AWS](#).


















If your SDDC does not have NSX Advanced Firewall enabled, VMware Cloud DR enables it each time you run a recovery plan for ransomware (which incurs a cost). When the last concurrent plan is disabled, NSX Advanced Firewall is also deactivated.

## Network Isolation Levels Allowed Connectivity

The following table describes the ransomware network isolation levels and the type of network access allowed for each. A check mark means that the type of connectivity is allowed.

**Note** Applying or changing a network isolation level for a VM overwrites any previous firewall configurations that were previously set for the VM.

Network Access Level	NSX-T Advanced Firewall required for security scanning	DHCP, DNS, NTP	Integrated Security analysis and scanning	Outbound External	Outbound East-West	Inbound External	Inbound East-West
Isolated	No						
Quarantined	No	✓	✓				
Quarantined +Analysis	Yes	✓	✓				
<b>Note</b> Default isolation level when NSX Advanced Firewall is enabled.							
External Outbound	No.	✓	✓	✓			
<b>Note</b> Default isolation level used when NSX Advanced Firewall is not enabled.							


Network Access Level	NSX-T Advanced Firewall required for security scanning	DHCP, DNS, NTP	Integrated Security analysis and scanning	Outbound External	Outbound East-West	Inbound External	Inbound East-West
Internal Inbound	Yes						
Internal	Yes						
Internal + External Outbound	No						
Open	No						


## Network Isolation with NSX Advanced Firewall On


When you activate integrated security and vulnerability analysis and turn **on** NSX Advanced Firewall in the Settings dialog box, and your [Configure Plan for Ransomware Recovery](#) has ransomware recovery and integrated analysis on, VMs in the validation state are placed into the Quarantine+Analysis network isolation level.


Change VM network isolation - RWR-Demo-App01
×


Select the network isolation rule


☐

Isolated
Fully isolated.  
No network access.


☒

Quarantined + Analysis
Only access network and integrated analysis services.

☐

External outbound
Allow outbound access to the internet.  
Use to expose ransomware behavior.

☐

Internal inbound
Allow inbound access from internal network.  
No internet access.

☐

Internal
Allow full access in the internal network.  
No internet access.

☐

Internal + External  
outbound
Allow full access in the internal network,

☐

Open
Full internal and internet access.

[How to create a custom isolation in NSX-T](#)

CANCEL

CHANGE ISOLATION

## Network Isolation with NSX Advanced Firewall Off

If you activate integrated security and vulnerability analysis but turn **off** NSX Advanced Firewall in the Settings dialog box, and your [Configure Plan for Ransomware Recovery](#) has ransomware recovery and integrated analysis on, VMs in the validation state are placed into the External Outbound network isolation level. This isolation level allows the sensors installed on the VMs to connect to Carbon Black Cloud.




Change VM network isolation - Ransomware Test Win2 - Eicar malware

×


Select the network isolation rule

☐

 Isolated


Fully isolated.  
No network access.

☐

 Quarantined


Only access network services (DHCP, DNS, and NTP).  
No access to the integrated analysis service.

☒

 External outbound


Allow outbound access to the internet.  
Use to expose ransomware behavior.

☐

 Internal + External  
outbound

Allow full access in the internal network,

☐

 Open

Full internal and internet access.

CANCEL

CHANGE ISOLATION

## Network Isolation with Integrated Security and Vulnerability Analysis Off

If you deactivate integrated security and vulnerability analysis from the Settings dialog box, VMs in the validation state are placed into the Quarantined network isolation level.


VMware, Inc.

253


Change VM network isolation - c1-w2-hs3-q0615\_cscalescale\_26

Select the network isolation rule.


Integrated analysis is not active; fewer isolation levels are available.

☐

**Isolated**


Fully isolated.  
No network access.

☒

**Quarantined**


Only access network services (DHCP, DNS, and NTP).  
No access to the integrated analysis service.

☐

**External outbound**

Allow outbound access to the internet.  
Use to expose ransomware behavior.

☐

**Internal + External outbound**

Allow full access in the internal network,

☐

**Open**

Full internal and internet access.

[How to create a custom isolation in NSX-T](#)

CANCEL

CHANGE NETWORK ISOLATION

## Video: Network Isolation for Ransomware Recovery

Ransomware recovery provides predefined (and customizable) network isolation levels that allow you to safely analyze and recovery VMs on the recovery SDDC.

VMware Cloud DR leverages NSX Advanced Firewall to provide varying degrees of [Network Isolation Levels](#) you can use to safely recover a VM without infecting other VMs in your production environments.

For a quick understanding of the network isolation levels that you can use for ransomware recovery, watch this video:



([Network isolation levels you can use for ransomware recovery](#))

## Configure Plan for Ransomware Recovery

You can configure a recovery plan for ransomware recovery with integrated vulnerability and security analysis.

In a recovery plan, select the 'Activate ransomware recovery' option to use the plan for ransomware recovery or ransomware recovery tests. Running a plan for a ransomware recovery test is the same as running a ransomware recovery plan, except that a ransomware recovery test has no option to recover VMs to a protected site.

When you activate ransomware recovery in a plan, VMs in the plan are charged for ransomware recovery for VMware Cloud DR. For more information on costs, see <https://www.vmware.com/products/cloud-disaster-recovery.html> and click the **Pricing** tab.

---

**Note** For more details on creating and setting up a recovery plan, see [Configure Recovery Plans](#).

---

When you click either the **Ransomware Recovery** or **Ransomware Test** buttons in the recovery plans list, you make VMs in the plan available for ransomware recovery.

---

**Note** If you configure test network mappings for a recovery plan, and the plan is activated for ransomware recovery, the plan will use the test network mapping by default when you run the plan.

---

Configuring ransomware for a recovery plan requires choosing from the following options:

The screenshot shows a web interface for editing a recovery plan titled "Edit plan - DB-backup-nolan". On the left is a sidebar with a list of configuration categories: General, Sites, Groups, vCenters (checked), vCenter folders (checked), Compute resources (checked), Virtual networks (checked), IP addresses (checked), Script VM (checked), Recovery steps (checked), **Ransomware** (highlighted with a green arrow), and Alerts (checked). The main panel is titled "Ransomware" and contains the following options:

- ☒ **Activate ransomware recovery**  
Allow starting the ransomware test and recovery workflows with this plan.
- A yellow warning box states: "VMs in this plan will be charged the ransomware add-on. [See pricing page.](#)"
- Text: "Confirm you understand the following:"
- ☒ **VMs in this plan will generate the additional ransomware-add-on charge.**
- Security and vulnerability analysis during ransomware recovery**
  - ☒ **Use integrated analysis**  
Install sensors as VMs are restored in the recovery SDDC to perform security and vulnerability analysis.
  - ☒ **Pause when starting a VM to manually remove production security sensors**  
If your VMs have sensors from a security solution, such as Carbon Black, you should uninstall them when starting validation. This is to avoid polluting your production security solution with alerts occurring in the isolated recovery SDDC.
  - ☐ **Do not use integrated analysis**  
Use other tools to test for ransomware. VMs will start in fully isolated mode.

At the bottom right of the window are four buttons: CANCEL, < BACK, NEXT >, and FINISH.

- **Activate ransomware recovery.** Enable ransomware recovery for the plan. When you save the plan, you start being charged for ransomware recovery for all VMs protected by a recovery plan.
- **Use integrated analysis.** Enable integrated security and vulnerability analysis for VMs in the plan. When you run the plan and start VMs for recovery, VMware Cloud DR installs a security sensor on VMs, which enables ransomware analysis. For Linux VMs, you must install the security sensor manually, and uninstall any existing security sensors from the VM. For more information, see [Run Plan and Install Linux Launcher and Sensor](#). Integrated analysis does not trigger any additional charges.
- **Pause when starting a VM to manually remove production security sensors.** Pause when starting VMs during ransomware recovery so you can remove any production sensors or security software, which might interfere with VMware Cloud DR integrated analysis and impact the isolated recovery environment of the recovery SDDC. For more information, see [Uninstalling Sensors](#).

- **Do not use integrated analysis.** If you want to use your own security tools for ransomware recovery on your recovery SDDC, select this option. When selected, no VMware Cloud DR security sensors are installed when you start VMs during ransomware recovery.

## Run a Ransomware Recovery Plan

If you enable ransomware services in a recovery plan, you can run it to recover from a ransomware attack, or run it as a ransomware recovery test.

When you run a plan for ransomware recovery, you enable all VMs included in the plan to be analyzed and validated in a network-isolated recovery SDDC with restricted firewall rules, disconnected from the internet. You can select VMs from the list included in the plan, choose a snapshot from a protection group, and start validating VMs.

When you start VMs on the recovery SDDC, a security sensor is automatically installed to enable [Integrated Security and Vulnerability Analysis](#). The sensor helps you detect malware, repair bad files, and patch software for the VM from the snapshot history. For information on manual sensor installation, see [Manual Sensor Installation](#).

---

**Note** If you do not have ransomware recovery services enabled, VMs are disconnected from internal networks but have external outbound access to the internet.

---

### Uninstalling Existing Sensors

Before the security sensor can be installed when you run a plan, you must uninstall any existing security sensors or software. You also must confirm that you have uninstalled existing security sensors before you can begin ransomware analysis. If any non-VMware Cloud DR security software is left on VMs, it can potentially generate alarms and events for the production environment. For more information, see [Uninstalling Sensors](#).

### Connectivity to Carbon Black Cloud URLs

The network segment the VM is connected to must have internet access, so the VM can reach the security services location within a specific Carbon Black Cloud points of presence (PoP). Make sure that your network and in-guest firewalls do not block access to any Carbon Black Cloud [POP URLs](#).

### Snapshot Expiration Paused While Running the Plan

When you start a ransomware recovery plan, VMware Cloud DR pauses all snapshot expiration for snapshots taken prior to starting the plan. Existing snapshots are deleted upon expiration, regardless of protection group retention policy, until the plan is ended. Any subsequent snapshots taken since the starting of the plan will expire and be deleted according to the configured retention policy.

When the plan is ended, snapshots expiration will resume according to the defined protection group retention policy.

### Ransomware Recovery Test

A ransomware test is the same as ransomware recovery, but a ransomware test has no option to restore VMs on a protected site. Ransomware recovery testing does not require powering off production VMs and pausing of snapshots.

### Prerequisites

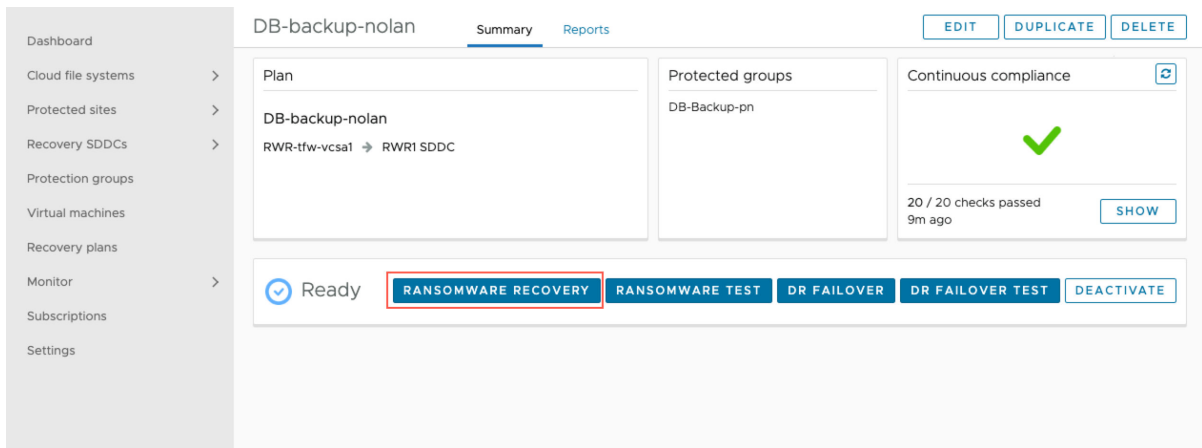
Before you can run a recovery plan for ransomware, you must:

- [Activate Ransomware Recovery Services](#) with integrated security and vulnerability analysis.
- [Create a Protection Group](#) with a recurring snapshot replication schedules for VMs you want to recover.
- [Create a Recovery Plan and Configure Plan for Ransomware Recovery](#) for ransomware.

**Note** If you configure test network mappings for a recovery plan, and the plan is activated for ransomware recovery, the plan will use the test network mapping by default when you run the plan.

### Procedure

- 1 From the left navigation, select **Recovery plans**.
- 2 Select a recovery plan from the list.
- 3 On the recovery plan page, click the **Ransomware Recovery** button. Or, click the **Ransomware Test** button if you only want to perform a test.



- 4 In the **Ransomware recovery** dialog box, click the **Start Ransomware Recovery** button.

### What to do next

Next, you can [Start VMs in Recovery SDDC](#) on the recovery SDDC so you can begin the validation and recovery process.

## Start VMs in Recovery SDDC

Once you run a recovery plan for ransomware, you can start one or more VMs in the plan to begin the validation process on the recovery SDDC.

Once a recovery plan starts, you see a list of all VMs included in the plan in the backup state. In backup, you can start VMs on the recovery SDDC from any of snapshot in the protection groups that the VMs belong to.

Only VMs that belong to protection groups configured in the recovery plan can be started. For example, if a VM belongs to protection group1, protection group2, and protection group3, but the recovery plan only contains protection group1 and protection group2, then snapshots from protection group3 will not be included when you start the VM on the recovery SDDC.

During ransomware recovery, VMware Cloud DR can process a maximum of 50 VMs at a time. This limit applies to starting VMs, staging VMs, recovering VMs, powering off VMs, changing network isolations, selecting new snapshots, and so on. For example, if 25 VMs are currently being started on the recovery SDDC, you can only start 25 more VMs until the other 25 VMs have finished starting.

If you have configured a recovery plan to 'Pause when starting a VM to manually remove production security sensors', you must remove each sensor one VM at a time. You cannot remove a sensor from multiple VMs at a time. For more information, see [Configure Plan for Ransomware Recovery](#) and [Uninstalling Sensors](#).


### Procedure


- 1 To start VMs in validation on the recovery SDDC, after you start the plan select one or more VMs from the list and then click the **Start VMs In Recovery SDDC** button.
- 2 In the **Validate VM in recovery SDDC** dialog box, select a snapshot of the VMs to use for ransomware recovery. If you do not explicitly select a snapshot, then the most recent snapshot is used.

If you are starting one VM, then select a snapshot from the snapshot drop-down menu. For help deciding which snapshot to choose, see [Try a Different Snapshot](#).

If you are starting multiple VMs, then click the **Select Snapshot** button and select a snapshot. If the VMs you are starting belong to more than one protection group, then you must select a snapshot from each group.

Validate VM in recovery SDDC - 3 VMs

  
In backup

  
In validation

Select a snapshot for each protection group. The VMs will be restored in the recovery SDDC, so you can test and validate them.

If a VM is in multiple protection groups, the selected snapshot with the most recent time stamp will be used.

pno-recover

pno-recover - Daily Every 12 hours - 2023-02-07T20:00 UTC

Feb-07-2023 12:00 pm (24m ago)

SELECT SNAPSHOT

pno-test

pno-test - Daily-3 - 2023-02-06T17:15 UTC

Feb-06-2023 09:15 am (1d ago)

SELECT SNAPSHOT

CANCEL

START 3 VMS IN RECOVERY SDDC


- After you have selected snapshots for the VMs, to start analyzing them on the recovery SDDC click any of the VMs in the list. (You can also select a VM and click the **VM Details** button.)



- At this point in the plan execution, if any existing security sensors are on the VMs, uninstall them. (For more information, see [Uninstalling Sensors](#).) After the security sensors are uninstalled, under the Start VM panel, click the **I Uninstalled the Security Sensors** button.


▼ Start VM

Waiting for user to uninstall security sensors...



Starting the VM can take several minutes.  
Current step: Remove security sensors

You should uninstall any sensors to avoid polluting your production security solution with alerts occurring in the isolated recovery SDDC. You can access the VM through the vCenter console.

OPEN VCENTER 

Once you uninstall any security sensors, click the button below.

**I UNINSTALLED THE SECURITY SENSORS**

- Click the **Start VM in Recovery** button to start the VM on the recovery SDDC.

## Search for VMs in a Plan

When you run a plan for ransomware recovery, you can search the VMs in the plan to easily find VMs to recover.

When a recovery plan for ransomware is running, you can search for VMs using [VM Name Pattern](#) (for inclusion and exclusion), you can restrict the search to specific protection groups, and by recovery state (in backup, starting, in validation, staging, staged, recovering, recovered). The search evaluates combined filters as AND operations.

## Procedure

- 1 Start a plan for ransomware recovery.
- 2 From the list of VMs, enter a name pattern to search for a VM.

VM list

VM DETAILS

pno Search tips pno-tester In backup

VM name	Recovery state	Protection group	Network isolation	Snapshot timestamp	Security threat severity	Vulnerability risk
<input type="checkbox"/> pno-centos-app1	In backup	pno-tester	--	--	--	--
<input type="checkbox"/> pno-database-test	In backup	pno-tester	--	--	--	--
<input type="checkbox"/> pno-db-back2	In backup	pno-tester	--	--	--	--
<input type="checkbox"/> pno-web-server	In backup	pno-tester	--	--	--	--
<input type="checkbox"/> pno-webserver	In backup	pno-tester	--	--	--	--
<input type="checkbox"/> pno-win2012-app2	In backup	pno-tester	--	--	--	--

- 3 If the VM belongs to one than one protection group, you can select one of the groups so the search only looks in the selected group.
- 4 To further restrict the search, you can select a VM ransomware recovery state: in backup, starting, in validation, staging, staged, recovering, or recovered.

## Try a Different Snapshot

During validation, you might want to try different snapshots of VMs to find one that can be a candidate for recovery.

From the VM validation page, you can easily try different snapshots during ransomware recovery. When you select a new snapshot, the current iteration of VMs on the recovery SDDC is discarded and a new iteration starts based on the snapshot you select.

## Procedure

- 1 In the VM validation page, **Summary** tab, from the End validation iteration panel click the **Try Different Snapshot** button.

End validating iteration

TRY DIFFERENT SNAPSHOT POWER OFF AND STAGE

OTHER

- 2 In the **Try different snapshot dialog** box, select a new snapshot, either from the timeline or from the **Snapshot** drop-down menu.

---

**Tip** It is a good idea to badge the current snapshot, for future reference

---

- 3 Click the **Try Different Snapshot** button.

When selecting a VM snapshot, consider two factors on the snapshot timeline: change rate and entropy rate. These two factors show any behavioral and structural changes to the VM over time across snapshots, which can indicate possible infection:

- **Change rate.** The amount of bytes changed / (time difference between the current snapshot and the previous snapshot). A high value indicates that too many changes happened during this time. During a ransomware attack, many files are encrypted, so this value is higher. So, if you have a snapshot with a high entropy and high change rate, it might indicate a ransomware attack.

For example, if the change rate for a VM in a snapshot is typically about around 100 KB/s, if the change rate became 500 KB/s, then the snapshot is suspicious.

---

**Note** For the first snapshot of a VM after a product upgrade, change rate is not reported.

---

- **Entropy rate.** 1/compression ratio. Entropy rate is a number between 0 and 1, and the closer it is to 1, the higher the likelihood that the snapshot is encrypted. Sudden jumps in entropy can indicate possible encryption.

For example, if the entropy rate (1/compression ratio) for a VM is .5 or .6, and then it jumps to almost 1, then the snapshot is suspicious.

VMware Cloud DR uses the inverse of the data compression ratio to approximate an entropy rate. The entropy rate of the data is 1 when the data is incompressible: such incompressible data usually means the data is either encrypted or already compressed. The entropy rate of data is smaller because the data is more compressible.

---

**Note** For more information, see [Snapshot Timeline: Change Rate and Entropy Rate](#).

---

## Snapshot Timeline: Change Rate and Entropy Rate

When validating VMs in ransomware recovery, you select historical snapshots of each VM so you can analyze the change rate and entropy rate across all snapshots.

The snapshot timeline appears when you first start VMs in validation and select a snapshot from the **Timeline** tab in validation, and when you [Try a Different Snapshot](#) during validation.

&lt; PLAN - APP-SERV

centos7\_pnolan

Summary

Timeline

Analysis

Events

## Protected site snapshot timeline

This VM is already restored in the recovery SDDC

Current base snapshot

Time stamp

AppServer-Linux - Daily-3 - 2022-09-23T21:40 UTC


Sep-23 02:40 pm (3d ago)

Current snapshot badge



Not badged

⚠ Any changes made to the VM in the recovery SDDC will be lost

⚠ This VM is **powered on** in the recovery SDDC[Learn more](#) about change and entropy rate.

Protected site base snapshot

AppServer-Linux - Daily-3 - 2022-09-23T21:40 UTC

Time stamp

Sep-23-2022 02:40 pm (3d ago)

## Change Rate and Entropy Rate

As you start recovering VMs from a ransomware attack, you become familiar with change rate and entropy rate:

- **Change rate.** The amount of bytes changed / (time difference between the current snapshot and the previous snapshot). A high value indicates that too many changes happened during this time. During a ransomware attack, many files are encrypted, so this value is higher. So, if you have a snapshot with a high entropy and high change rate, it might indicate a ransomware attack.

For example, if the change rate for a VM in a snapshot is typically about around 100 KB/s, if the change rate became 500 KB/s, then the snapshot is suspicious.

**Note** For the first snapshot of a VM after a product upgrade, change rate is not reported.

- **Entropy rate.** 1/compression ratio. Entropy rate is a number between 0 and 1, and the closer it is to 1, the higher the likelihood than the snapshot is encrypted. Sudden jumps in entropy can indicate possible encryption.

For example, if the entropy rate (1/compression ratio) for a VM is .5 or .6, and then it jumps to almost 1, then the snapshot is suspicious.

VMware Cloud DR uses the inverse of the data compression ratio to approximate an entropy rate. The entropy rate of the data is 1 when the data is incompressible: such incompressible data usually means the data is either encrypted or already compressed. The entropy rate of data is smaller because the data is more compressible.

When VMware Cloud DR detects an unusually high change rate and entropy rate, it can indicate unusual activity, such as ransomware attack encrypting the data. The snapshot before the onset of such activity might be a snapshot containing unencrypted data.

For example, a common type of ransomware attack involves encrypting the user files and removing other files from the guest VM. During a malicious encryption operation, the incremental snapshot includes the encrypted data in addition to regular modified VM data.

Because the VMware Cloud DR snapshot is always incremental, only the modified or new data transfers to the cloud backup. When compared with normal snapshot where no ransomware attack is occurring, a problem snapshot has more data transferred, and out of all transferred data, it has a higher percentage of data being encrypted, thus showing a high entropy rate.

The presence of more data on a snapshot is suspicious and shows a higher change rate than normal compared to other snapshots (for example, compared to other snapshot with same time of day, or same time of the week), or compared to other similar VMs that are not under attack.

## Expired Snapshots on the Timeline

On the snapshot timeline, you might see changes in entropy rate and change rate, even where no snapshots show.

When you see entropy rate and change rate metrics where no snapshots exist, you are looking at data from expired snapshots. VMware Cloud DR retains metrics associated with expired snapshots to provide fine-grained data points that can help you discover anomalies on the snapshot history.

For example, you have snapshots A and B, yet on the snapshot timeline you see entropy rate and change rate data between the two snapshots. In this scenario, it indicates that there were snapshots between snapshots A and B that have expired.



If you see variations in entropy rate and change rate in the time interval between the two snapshots, it might indicate suspicious or malicious behavior during that time, so you can decide if you want to select a snapshot prior to A, or a later snapshot after B.

## Integrated Security and Vulnerability Analysis

When you start a recovery plan for ransomware and begin validating VMs on the recovery SDDC, the system analyzes and scans VMs for vulnerabilities, malware signatures, and behavioral anomalies.

VMware Cloud DR ransomware services perform the following analyses and scans on VMs in ransomware recovery:

- **Vulnerability analysis.** Searches for and displays OS and application vulnerabilities discovered on VMs. Results are displayed on the **Vulnerabilities** tab in a VM validation page: **Analysis > Vulnerabilities**
- **Behavior analysis.** Continuous analysis of running software and processes on the guest OS, looking for suspicious behavior. Results are displayed on the **Analysis > Alerts** tab.
- **Malware signature scan.** Scans VMs for known malware and viruses. Results are displayed on the **Analysis > Alerts** tab.

**Note** Currently, malware signature scans for Linux VMs do not report progress in the UI. The scan still occurs, but the progress indicator remains "in progress" even after the scan is finished.

< PLAN - DB-BACKUP-NOLAN RW

pno-web-server      Summary      Timeline      Analysis      Events

TOOLS ▾

Device ID 335665  
Sensor status Installed  
Sensor version 3.9.0.1810

Host name win10b-02  
OS windows9\_64Guest

Vulnerability analysis Finished  
Malware signature scan Finished  
Behavior analysis Started Sep-22 09:37 am (2h ago) - 8h recommended

Alerts      Vulnerabilities

Vulnerability	Severity	Type	App	Description
CVE-2018-8...	8.7	OS	--	A remote code execution vulnerability exists when the Windows Shell does ...
CVE-2019-0...	7.4	OS	--	An elevation of privilege vulnerability exists when Windows AppX Deploye...
CVE-2015-4...	6.3	App	Google Chro...	The TLS protocol 1.2 and earlier, when a DHE_EXPORT ciphersuite is enable...
CVE-2018-8...	5.3	OS	--	A remote code execution vulnerability exists when Microsoft Windows PDF ...
CVE-2019-12...	5.3	OS	--	An elevation of privilege vulnerability exists when the Windows AppX Deplo...
CVE-2009-1...	3.3	App	Google Chro...	Google Chrome executes DOM calls in response to a javascript: URI in the tar...
CVE-2013-6...	3.3	App	Google Chro...	A use-after-free in AnimationController::endAnimationUpdate in Google Chro...
CVE-2019-11...	3.3	OS	--	An elevation of privilege vulnerability exists when Windows AppX Deploye...
CVE-2020-1...	3.2	OS	--	A security feature bypass vulnerability exists in Microsoft Word software wh...
CVE-2019-0...	3	OS	--	A denial of service vulnerability exists when SymCrypt improperly handles a ...
CVE-2020-1...	3	OS	--	A denial of service vulnerability exists in Microsoft Outlook software when th...
CVE-2016-71...	2.7	App	Google Chro...	The HTTPS protocol does not consider the role of the TCP congestion windo...
CVE-2016-71...	2.7	App	Google Chro...	The HTTP/2 protocol does not consider the role of the TCP congestion wind...
CVE-2018-10...	2.7	App	Google Chro...	A hardware vulnerability in GPU memory modules allows attackers to accele...
CVE-2012-4...	2.4	App	Google Chro...	The TLS protocol 1.2 and earlier, as used in Mozilla Firefox, Google Chrome, ...
CVE-2017-8...	2.4	OS	--	Microsoft Edge in Microsoft Windows 10 Version 1703 allows an attacker to o...
CVE-2017-8...	2.4	OS	--	Microsoft Edge in Microsoft Windows Version 1703 allows an attacker to obt...
CVE-2017-8...	2.4	OS	--	The Remote Desktop Protocol (RDP) implementation in Microsoft Windows 1...
CVE-2019-12...	2.4	OS	--	An information disclosure vulnerability exists when the Windows Hyper-V Ne...

30 vulnerabilities loaded      [Load more](#)

## Vulnerability Analysis

VMware Cloud DR scans VMs guest OS's for known vulnerabilities (security flaws which can be exploited) to be patched, to close potential vectors of attack and prevent reinfection. You can view the full context of a vulnerability discovered on VMs, including risk score details and how it impacts your environment, so you can fix the issue.

All vulnerabilities found during a scan receive a Risk Score, between 0.0 (no risk) and 10.0 (maximum risk), which accurately represents the risk of a given vulnerability on VMs in your environment. The analysis is performed by combining Common Vulnerability Scoring System (CVSS) information with proprietary threat data and advanced modeling from [Kenna Security](#).

The vulnerability risk score range and severity are defined as follows.

Score Range	Severity
0.0 - 3.9	Low
4.0 - 6.9	Moderate
7.0 - 8.9	Important
9.0 - 10.0	Critical

The **Vulnerability** tab shows a list of all found vulnerabilities with their [CVE](#) number and a link to the vulnerability article in the National Institute of Standards and Technology ([NIST](#)) database. Each vulnerability is identified by its Common Vulnerabilities and Exposures (CVE) name, which consists of the string CVE plus the four digit year and then a number, such as CVE-2022-1234.

You can also view the vulnerability details in the security console (Carbon Black CloudCarbon Black Cloud console), if you have integrated analysis activated. When you select a vulnerability in the list, details of the vulnerability are displayed.

## Alerts for Malware Scan and Behavior Analysis

When VMs are started on the recovery SDDC, the system analyzes VMs and their guest files for anomalous behaviors, such as running software that has a bad reputation, running processes that repeatedly make outbound connections, reaching out to known suspicious IP addresses, malicious interference with the Windows Registry, or other system files or processes on VMs.

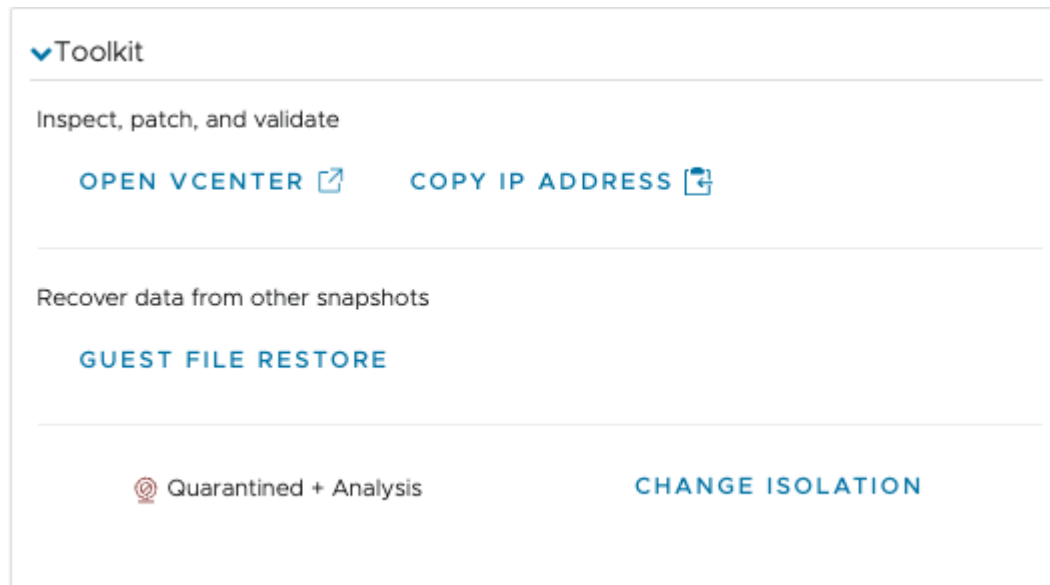
The system also scans VM guest files for malware and viruses, calculating hashes of binaries, parsing binary components, looking for attributes known as a malware 'signature' so you can remove them.

The results of the malware scan and behavior analysis displayed in the **Alerts** tab and ranked according to severity, with a higher score being worse than a lower score:



Score Range	Severity
1 - 2	Activities such as port scans, malware drops (installation of malware), changes to system configuration files, persistence of malware, and more.
3 - 5	Activities such as running malware, generic virus-like behavior, monitoring user input, potential memory scraping, password theft, and more.
6 - 10	Activities such as reverse command shells, process hollowing, destructive malware, hidden processes and toolsets, applications talking on the network that should not be, and more.

If you find vulnerabilities or malware, or observe suspicious behavior, you can use some of the tools provided on the VM ransomware page, such as [Copy IP Address for VM Access](#) of a VM, open the VM in vCenter, or perform [Guest File Recovery](#):



You can also view VMs in the Carbon Black Cloud console for further analysis and remediation, from the **Tools** drop-down menu when you select the **Analysis > Vulnerabilities** tab. The, from the Tools menu, select **Open in security console**.

## Video: Next Generation Antivirus and Behavioral Analysis with Ransomware Recovery

When you start a VM in ransomware recovery on the recovery SDDC, a next generation antivirus scan looks for known vulnerabilities and malware signatures, and a deep behavior analysis begins of all running software and processes on the VM guest OS, looking for suspicious behavior.

To learn about how VMware Cloud DR analyzes VMs during ransomware recovery, watch this short video:



(Next Generation Antivirus and Behavioral Analysis with Ransomware Recovery)

## Discard or Detach VMs

During validation, you have the option to either discard or detach staged VMs.

When you **discard** VMs in the recovery SDDC during validation, VMs are deleted from the recovery SDDC and then moved to the backup state. All staging snapshots, badges, and notes remain. You can also **discard** VMs that are in the staged state, which moves the VMs back into the backup state.

You can discard multiple VMs at a time.

When you **detach** a VM to the recovery SDDC during validation, the currently running instance of the VM is renamed, suspended, and no longer tracked as part of the plan. The original VM from the plan is moved to the in backup state. Detaching a VM is useful if you find a VM in an interesting state that you want to save for further analysis after the plan ends.

You can only detach one VM at a time.

**Note** When you detach a VM, you must delete the detached VM before starting a new validation iteration of the same VM.

### Procedure

- 1 When VMs are in validation, from the VMs list select one or more VMs and from the **Other Actions** menu, select **Discard VM in recovery SDDC**.

VM list

VM DETAILS

POWER OFF AND STAGE 2 VMS

OTHER ACTIONS ▾

☐ VM name ▾

Recovery state

Prot

<input checked="" type="checkbox"/>	pno-app-serv	In validation	pno-recover	Quarantin...
<input checked="" type="checkbox"/>	pno-db-1	In validation	pno-recover	Quarantin...
<input type="checkbox"/>	pno-db-2	In backup	pno-recover	--

Discard 2 VMs in recovery SDDC

Change network isolation

Set base snapshot badge

- 2 In the **Discard VM** dialog box, badge the snapshots (optional) and then type the phrase to confirm the deletion. Click the **Discard VMs** button to discard the VMs.
- 3 To detach a VM on the recovery SDDC, select a VM from the VMs list and then from the **Other Actions** menu select **Detach VM in recovery SDDC**.

- 4 From the **Detach VM** dialog box, enter a name for the new VM, badge the snapshot (optional), and then type **DETACH VM** to confirm.
- 5 Click the **Detach VMs** button.

## Open Security Console

If you want to do further analysis of a VM in validation, you can open the VM in the security console ('Carbon Black Cloud console').

Depending on the deployment and your user role, you can use the Carbon Black Cloud console to view information about a VM, investigate threats, manage security policies, view and query audit events, manage workload VMs, and more. For more information, see [VMware Carbon Black Cloud on VMware Cloud services Platform User Guide](#).

---

**Note** Opening the Carbon Black Cloud console requires at least one Carbon Black Cloud user role. For more information, see [User Roles](#).

---

**Note** Currently, if you are using VMware Cloud DR on the VMware Cloud Partner Program, you cannot launch the security console from the VMware Cloud DR UI, or from your organization.

---

### Procedure

- ◆ From the **Analysis** tab of a VM during validation, select **Open in security console** from the **Tools** menu.

The Carbon Black Cloud console opens showing details on the VM in validation.

## Guest File Recovery

VMware Cloud DR snapshots allow you to download guest files of individual VMs to recover those files to a safe site.

Guest files downloaded as ZIP packages, which you can unzip and manually restore them to the destination of your choice.

Guest files are downloaded to the browser host where guest file recovery is being performed. For example, you can use guest file recovery to find a clean guest file from an older but known good snapshot of a VM.

You can run a recovery plan for ransomware and when a VM is in validation, you can open a browser on the VM running in the recovery SDDC and download imported guest files from a more recent snapshot directly into the running VM. After clean guest files have been downloaded to the VM, you can complete security analysis and recover the VM back to production.

Every guest file download is also available as a link from the **Monitor > Tasks** list that you can send to other users. The download link expires after six hours. The user on the local system must have file-level permissions to unzip the package.

---

**Note** If you are using [access lists](#) for VMware Cloud DR, only IP addresses listed in the Management access list can download a guest file for recovery.

---

**Best Practice** For security reasons, guest file download links are valid for six hours after the task completes. Any attempt to access an expired link results in an HTTP 403 forbidden error. Download links for this feature use enterprise grade encryption at the source and only allow SSL-based connections for download. Each download link contains both proof of identity and means of authentication, so anyone with the link can download the file. Share these links only at the discretion of the backup administrator. As a security best practice, use great caution while sharing these links.

---

Guest file recovery supports the following file systems:

- **Windows:** NTFS and FAT32.
- **Linux:** Ext3 and Ext4.

Guest file recovery does not support the following technologies:

- Windows dynamic volumes.
- Linux VMs that use Logical Volume Manager (LVM)
- Linux VMs formatting with the XFS file system.
- Microsoft Storage Spaces.

Current caveats for guest file recovery:

- You can run one guest file download at a time.
- Maximum path length of the download file directory = 255 characters.
- Maximum number of files or folder paths per ZIP package = 25. A folder path that contains multiple files is only counted as one item in the ZIP package out of a maximum of 25.
- Maximum individual file size allowable for download = 40 GB. This means that any given file in a download package cannot be larger than 40 GB.
- Maximum ZIP package export size = 100 GB.
- Windows OS unzip utility. Currently, restoring guest files does not support using the Windows OS default unzip utility in the File Explorer. Use 7ZIP or WinRAR utility on Windows systems for guest file restore operations.

## Recovering Guest Files on a Recovery SDDC

If you are restoring VM guest files or folders directly on a VM in a recovery SDDC on VMware Cloud on AWS, you must first configure access to the cloud file system S3 bucket in AWS. The cloud file system, where protection group snapshots are stored, uses S3 as a repository of recovery points.

You have two options for configuring access for guest file restore from the cloud file system into a VM on a recovery SDDC;

- Use an S3 endpoint in a linked customer account. Create the endpoint in a linked account, and then add VMC firewall rules, described here: [Access an S3 Bucket Using an S3 Endpoint](#). Or
- Use an internet gateway to access your S3 bucket. This method also requires deactivating the S3 option on the connected Amazon VPC for your SDDC. For information, see [Access an S3 Bucket Using the Internet Gateway](#).

## Recover VM Guest Files

You can recover virtual machine guest files from VMware Cloud DR snapshots, which you can then manually restore to any destination of your choice.

You can recover guest files from three different locations:

- The virtual machines list.
- A snapshot inside a protection group.
- The **Other** menu during [Chapter 14 Ransomware Recovery](#), when VMs are in the validation state.

---

**Note** For supported file systems on Windows and Linux, and for a full list of caveats and limitations with guest file recovery, see [Guest File Recovery](#).

---

### Procedure

- 1 From the left navigation, select **Virtual Machines**. Or, when a VM is in validation during ransomware recovery, from the **Other** menu.

If you are accessing guest file recovery during ransomware recovery, select **Recover guest files** from the **Other** menu after you start a VM in validation.

---

**Note** If you have more than one cloud file system, select one from the upper left of the list. If you have deployed only one cloud file system, it is already selected.

---

- 2 In the search box at the top you can search for specific VMs using [VM Name Pattern](#) (except for exclusion patterns). By default, the search box uses the \* wildcard to search for VMs across all snapshots.
- 3 Decide which VM to recover files from, and then click the **Recover guest files** button next to the VM.

- 4 In the **Recover guest files** dialog box, the most recent snapshot of the VM is selected. (If you select a VM snapshot from a protection group list, then that snapshot is the one selected.) To select a different snapshot, click the left or right arrow, or click the **Use different snapshot button**.
- 5 In the **Recover guest files** dialog box, you can select files from the list of Available files and folders. Click the down arrow to download individual files or folders. When you click the down arrow after selecting a file or folder, the ZIP package downloads immediately.

Recover guest files for win2012-vm3 ×

Snapshot

win-server - Daily - 2021-09-24T19:30 UTC ← → Use different snapshot

Sep-24 12:30 pm (2h ago)

Guest file system path

/ SCSI 0:0-Partition-1 (C:) / System Volume Information /

Available files and folders

File name	Type	Size	Modified	
MountPointManagerRemoteDatabase	File	--	Sep-24 12:26 pm	
tracking.log	Log file	20.0 KiB	Jun-29 04:47 am	

Selected files and folders ×

File name

Click the arrow on any file or folder to add it to the selected file and folder list for ZIP archive creation.

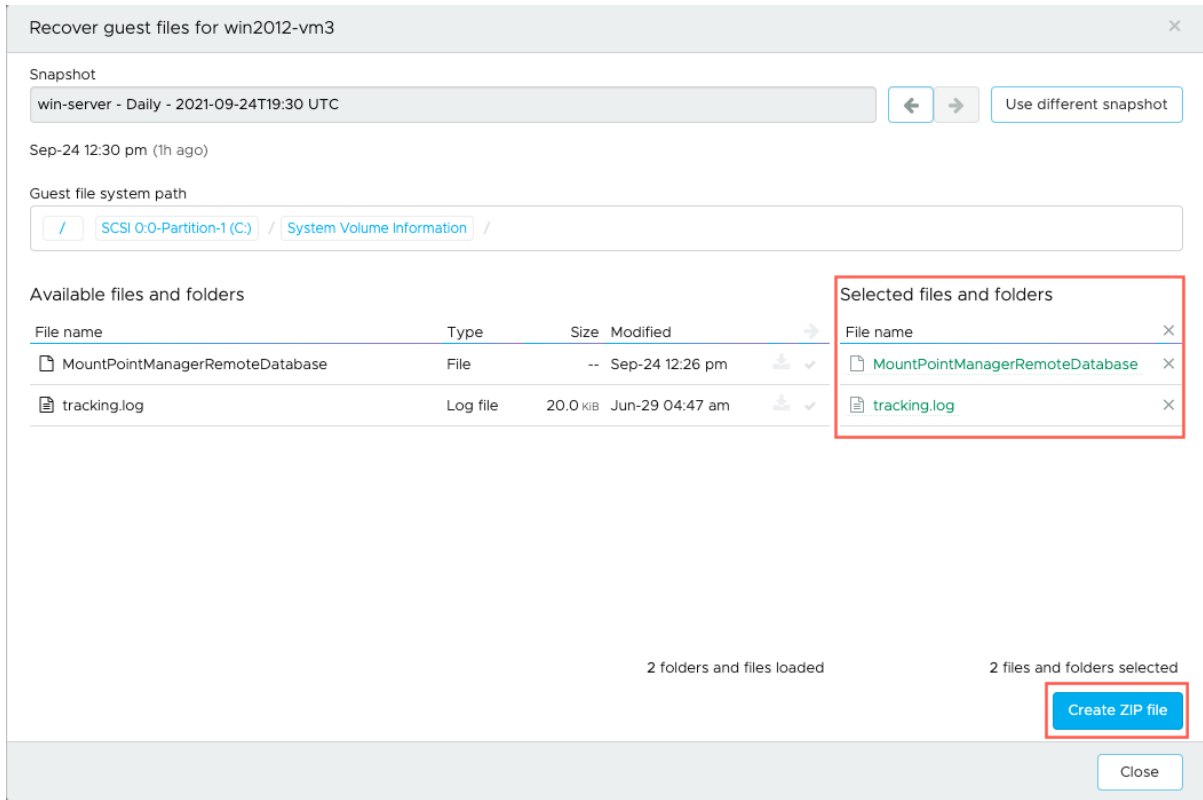
2 folders and files loaded

0 files and folders selected

Create ZIP file

Close

- 6 To download multiple files, click the right arrow to select multiple files. After you have selected files for download, click the **Create ZIP file** button to start the download.



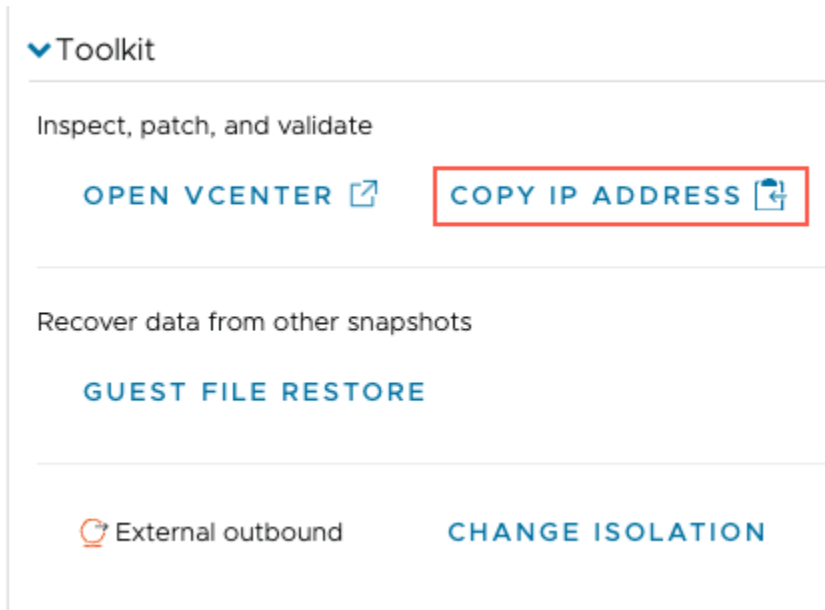
- 7 After the ZIP file downloads, click the **Close** button. Your user must have file-level permissions to unzip the package.
- 8 To access a URL for the guest file package, select **Monitor > Tasks**.
- 9 Find the guest file download task in the list. To filter the list, select the Protection filter.
- 10 From the menu at the far right of the task entry, right-click the download icon and select **Copy Link Address**.

## Copy IP Address for VM Access

When a VM is in validation, you can get its IP address to log in to the VM for further investigation. From the VM page during validation, you can copy the IP address of the VM.

## Procedure

- ◆ From the **Summary** tab of a VM in validation, in the Toolkit panel click the Copy IP address link to copy the address.



## Badging Snapshots

As you analyze VMs from snapshots, badge the snapshots to distinguish between good and bad ones.

### Prerequisites

Badging snapshots helps you and others on your team know which ones are good recovery candidates and those which are not. Badge bad snapshots to help others to avoid infected data. Badge clean snapshots to mark those VMs that are clean and ready to be recovered on a protected site.

Snapshot badges include:

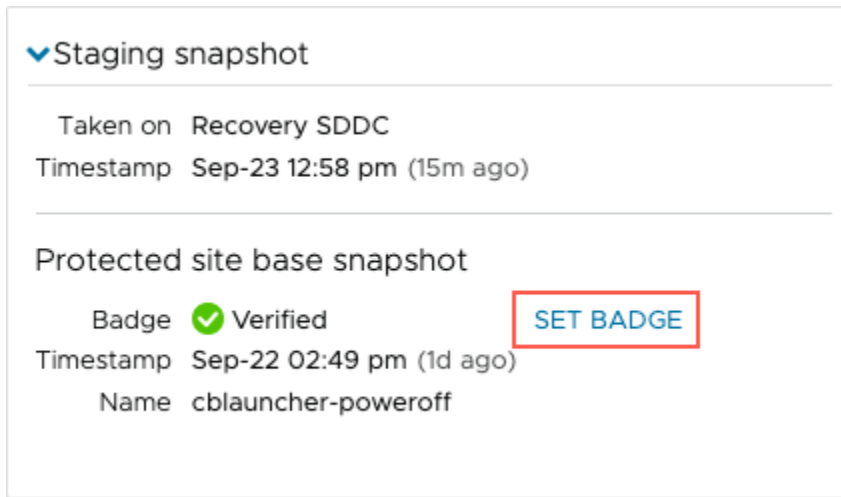
**Table 14-1.**

<b>Not badged</b>	No information on the security status of this snapshot.
<b>Verified</b>	Snapshot is clean and safe
<b>Warning</b>	Some of the data in this snapshot might be compromised, but overall infection is uncertain.
<b>Compromised</b>	Vulnerabilities and malware infections were found on the VM snapshot.
<b>Encrypted</b>	Data in this snapshot is encrypted by ransomware.



## Procedure

- ◆ You can badge a snapshot from several different locations:
  - From the VMs list, when you are running a recovery plan. You can select VMs from the list and from the Other **Actions** menu, select **Set base snapshot badge**.
  - From the **Summary** tab of an individual VM in validation.



You can also badge snapshots in these scenarios:

- When you are trying a different snapshot during validation.
- When you [Power Off and Stage VMs](#).
- When you [Discard or Detach VMs](#).
- When you [Restart Validation from Protected Site Snapshots](#).

## User Notes

It is a good idea to write notes during validation of a VM in ransomware recovery, to document your findings for yourself and for others on your team.

When you start a VM in validation, it is considered an 'iteration'. Every time you change snapshots of the VM in validation, it is considered a new iteration.

For each iteration of VM validation, you can write notes to document which VM snapshots are good and which are bad. You can also describe what methods you used to recover a VM or files, such as removing malicious files or patching malware during scans.

You can add iteration notes in the VM info panel for any of the iterations:

▼ VM info	
Protection group	AppServer-Linux
IP address	192.168.1.3 fe80::250:56ff:feb4:5705
Operating system	CentOS 7 (64-bit)
VMware Tools	10282
Protected site folder	DKRIEGER-STG/vm SDDC-
Recovery SDDC folder	Datacenter/vm/Workloads
<b>User notes</b>	
Iteration 1	
First snapshot deemed infected so trying a later date.	
Iteration 2	
Second snapshot also infected.	
Iteration 3	
Third snapshot has many known vulnerabilities, but could be a candidate for recovery.	
Iteration 4	
Fourth snapshot is clean.	

## Restart Validation from the Staging Snapshots on recovery SDDC

During the process of validating VMs, you can restart validation and revert to previously staged snapshots on the recovery SDDC.

## Procedure

- 1 From the VMs list, select one or more VMs.
- 2 From the **Other Actions** menu, select **Restart validation in recovery SDDC from the staging snapshot**.

- 3 In the **Restart validation from staging snapshot** dialog box, click the **Restart Validation** button.

## Restart Validation from Protected Site Snapshots

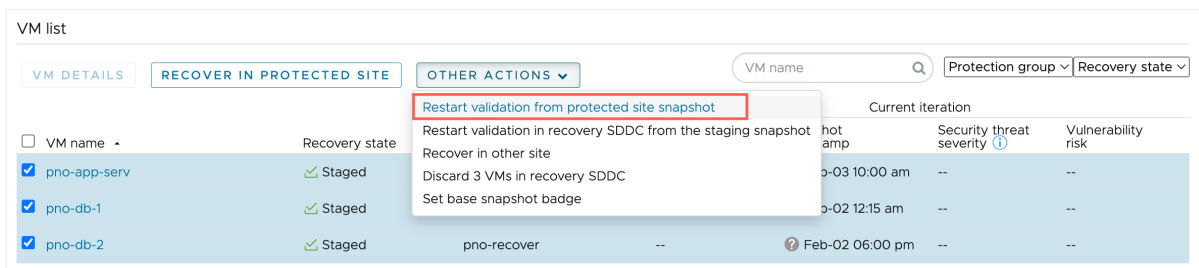
You have the option to restart VM validation using protected site snapshots.

During ransomware recovery, you can restart VM validation from different snapshots as many times as you wish. You can restart the validation process based on any snapshot in the recovery plan from the protected site.

You can also restart validation from the snapshot currently staged on the recovery SDDC. For more information, see [Restart Validation from the Staging Snapshots on recovery SDDC](#).

### Procedure

- 1 From the VMs list, select one or more VMs.
- 2 From the **Other Actions** menu, select **Restart validation from protected site snapshot**.



- 3 In the **Restart validation from protected site snapshot** dialog box, you can select a snapshot for the VMs. If you are working one VM, then select a snapshot from the Snapshot drop-down menu. If working with multiple VMs, or if the VMs belong to more than one protection group, click the **Select Snapshot** button for each group.
- 4 Click the **Try Different Snapshot** button.
- 5 Or, from the End staging panel select **Restart validation from staging snapshot**.
- 6 In the **Restart validation from the staging snapshot** dialog box, click the **Restart Validation** button.

## Power Off and Stage VMs

When you find good VM candidates to recover to a protected site, you are ready to power off and stage the VMs for recovery.

### Prerequisites

When powering off and staging VMs, VMware Cloud DR takes a snapshot of the VMs to prepare them for recovery to a protected site. The staging snapshot reflects the current state of the VMs in validation.

You can also choose to:

- Discard all changes made to the VMs in the recovery SDDC and stage with the protected site base snapshot (the first snapshot you started using when you began validating the VMs).
- Discard all changes made to the VMs in the recovery SDDC and start over with the last staged snapshots you used to validate the VMs.

For Windows VMs, when you power off and stage, the security sensor is uninstalled.

For Linux VMs, however, you must uninstall the sensor before clicking the **Power Off and Stage** button. For more information, see [Uninstalling Sensors](#).

Ransomware recovery supports recovering VMs to the original protected site from either of the following:

- A snapshot that originated from the on-premises protected site.
- A snapshot that was created on the recovery SDDC after cleansing the VM of ransomware and is staged for recovery.

A snapshot can be used to recover a VM to either [Recover VMs in Protected Site](#) or [Recover VMs in Other Protected Site](#). However, ransomware recovery does not support restoring a snapshot from an on-premises protected site to a protected VMware Cloud on AWS SDDC.

#### Procedure

- 1 From the VMs list of the recovery plan, select one or more VMs and click the **Power Off and Stage** button.

Figure 14-1.

The screenshot shows the 'VM list' interface. At the top, there is a search bar labeled 'VM name'. Below it, there are three buttons: 'VM DETAILS', 'POWER OFF AND STAGE 2 VMS' (which is highlighted with a red box), and 'OTHER ACTIONS' with a dropdown arrow. Below the buttons is a table with four columns: 'VM name', 'Recovery state', 'Protection group', and 'Network isolation'. The table contains three rows of VMs. The first two rows are selected with checkboxes and have a blue background. The third row is not selected and has a white background.

VM name	Recovery state	Protection group	Network isolation
<input checked="" type="checkbox"/> pno-app-serv	In validation	pno-recover	Quarantine...
<input checked="" type="checkbox"/> pno-db-1	In validation	pno-recover	Quarantine...
<input type="checkbox"/> pno-db-2	Starting...	pno-recover	--

- 2 In the **Power off and stage** dialog box, badge the base snapshot you were working from. If you are ready to recover the VMs, select the **Take a new staging snapshot** option. Taking a new staging snapshot allows you to restart validation using this current snapshot, instead of having to use the original snapshot.

You can also choose to discard changes and stage either the first snapshot you started with, or the most recently staged snapshot. Previously staged snapshots are overwritten.

- 3 Click the **Power Off and Stage** button.

#### What to do next

After you powered off and staged the VMs, you can now [Recover VMs in Protected Site](#), or [Recover VMs in Other Protected Site](#).

## Recover VMs in Protected Site

After you find good snapshots of VMs, badge them, power off and stage them, you can recover those VMs in the protected site.

When you recover VMs in a protected site where the VMs originated from, you replace the original VMs with the snapshots of the VMs you staged for recovery. The system preserves all the data that is common between the protected site and staged VMs and transfers only differences during recovery.

You can also recover VMs in a different protected site you have configured with VMware Cloud DR. If you recover in another protected site, the other site inventory must match exactly as defined in the recovery plan where the VMs live, or VM recovery will not succeed.

#### Procedure

- 1 From the VMs list, select one or more VMs.
- 2 From the **Other Actions** menu, click the **Recover in Protected Site** button.

VM list

VM DETAILS

RECOVER IN PROTECTED SITE

OTHER ACTIONS ▾

<input type="checkbox"/> VM name ▾	Recovery state	Protection group
<input checked="" type="checkbox"/> pno-app-serv	✓ Staged	pno-recover
<input checked="" type="checkbox"/> pno-db-1	✓ Staged	pno-recover
<input checked="" type="checkbox"/> pno-db-2	✓ Staged	pno-recover

- 3 In the **Recover in protected site** dialog box, confirm that you understand that the VM on the protected site will be replaced (overwritten) by VMs from the staging snapshots.
- 4 When the recovery operation finishes, the VMs are listed as Recovered.

#### What to do next

You can start validating other VMs, or you can [End Ransomware Recovery](#) for the plan.

## Recover VMs in Other Protected Site

After you badge VMs, power them off and stage them, you can recover those VMs in any protected site you have configured with VMware Cloud DR.

#### Prerequisites

If you recover VMs to a protected site different than the site where the VMs originated from, the other site inventory must match exactly as defined in the recovery plan where the VMs originated from, or the VMs will fail to recover.

---

**Note** You must have more than one protected site configured to perform this task. Only protected sites that share the same cloud file system can be used for this operation. For example, if you want to recover VMs from protected site A to protected site B, both protected sites must be registered to the same cloud file system.


---


#### Procedure

- 1 From the VMs list of a running recovery plan, select one or more staged VMs.
- 2 From the **Other Actions** menu, select Recover in other site.
- 3 In the **Recover in other site** dialog box, select a different protected site and vCenter to recover the VMs. Under

Required inventory, all required valid mappings are listed. If the check marks are green, then you can recover the VMs. If any of the mappings are invalid, cancel this dialog box and ensure that the destination inventory mappings match that of the new VMs. Once the mappings are valid, then you can resume this task.

Recover in other site - centos7\_pnolan


  
Staged



  
Recovered

The VM will be restored back into the protected site based on the current staging snapshot.

Staging snapshot

Taken on Recovery SDDC  
Timestamp Sep-22 11:37 am (38m ago)


Protected site base snapshot


Badge  Verified  
Timestamp Sep-21 02:40 pm (1d ago)  
Name AppServer-Linux - Daily-3 - 2022-09-21T21:40 UTC

Destination


Protected site


vCenter


 dk-PS-23Aug2033


 192.168.247.8

Required inventory

Compute  Resources

VM folder  vm

Datastore  Datastore

Networks  VM Network

CANCEL

RECOVER

4 Click the **Recover** button.

#### What to do next

When you have finished recovering VMs, you can [End Ransomware Recovery](#).

## End Ransomware Recovery

After you finish recovering VMs, you can end ransomware recovery by stopping the plan.

After you end a ransomware recovery plan, it reverts to the ready state. Ending the recovery plan deletes all staging snapshots, removes running VM instances in the recovery SDDC, and resumes snapshot retention schedules for VMs in the plan.

**Note** You cannot end a recovery plan for ransomware recovery if any VMs are in validation or staged. Cancel or recover VMs in those states before ending the plan.

## Procedure

- 1 From the Recovery plan page, click the **End Recovery** button.

The screenshot shows the 'App-Serv' recovery plan page. At the top, there are tabs for 'Summary' and 'Reports', and buttons for 'VIEW PLAN', 'DUPLICATE', and 'DELETE'. The 'Summary' tab is active. It contains three main sections: 'Plan' (showing 'App-Serv' and 'dk-PS-23Aug2033'), 'Protected groups' (showing 'AppServer-Linux'), and 'Continuous compliance' (showing '23 / 24 checks passed 18m ago' and a 'SHOW' button). Below these is a banner that says 'Recovering from ransomware...' with a red 'END RECOVERY' button highlighted by a red box. At the bottom, there is a 'Ransomware recovery summary' section with a progress bar showing four stages: 'In backup' (0 VMs), 'In validation' (0 VMs), 'Staged' (0 VMs), and 'Recovered' (1 VM). The 'Recovered' stage is highlighted with a green checkmark.

- 2 The **End ransomware recovery** dialog box launches and states how many VMs you recovered. Click **End Ransomware Recovery** to stop the plan.

The screenshot shows the 'End ransomware recovery - App-Serv' dialog box. It has a title bar with a close button (X). The main text reads: 'Once all the VMs of interest are recovered, end ransomware recovery to bring the plan back to the ready state'. Below this, it says: 'Staging snapshots in the recovery SDDC and expired snapshots in the protected site will be deleted.' and '1 VM has been recovered.' At the bottom, there are two buttons: 'CANCEL' and 'END RANSOMWARE RECOVERY'. The 'END RANSOMWARE RECOVERY' button is highlighted with a red box.



## Ransomware Events

When you run a recovery plan for ransomware and start a VM in validation, all events related to this workflow display in the **Events** tab.

View the **Events** tab during ransomware recovery of a VM to see all events related to the ransomware recovery process. For example, you can see events when you start a recovery plan for ransomware, start a VM in validation, stage a VM, recover it, and more. Event severity levels highlight important and critical events. For example, if a sensor does not install, or a VM did not successfully start, or if a snapshot fails to stage, you see events about them.

The same events are also displayed on the global events list found at **Monitor > Events**.

< PLAN - APP-SERV

centos7\_pnolan      Summary    Timeline    Analysis    **Events**

Severity	Description	Target	Timestamp	User name
Info	Waiting for user action	centos7_pnolan	Sep-22 01:40 pm	
Info	VM centos7_pnolan powered on	centos7_pnolan	Sep-22 01:40 pm	
Info	Reconfigured Ethernet adapters for VM centos7_pnolan	centos7_pnolan	Sep-22 01:39 pm	
Info	Network isolation for VM 'centos7_pnolan' set to NETWORK_ISOLATION_LEVEL_ISOLATED	centos7_pnolan	Sep-22 01:39 pm	
Info	VM centos7_pnolan registered with vCenter 10.2.224.4	centos7_pnolan	Sep-22 01:38 pm	
Info	VM centos7_pnolan restored from snapshot AppServer-Linux - Daily-3 - 2022-09-21T21:40 UTC	centos7_pnolan	Sep-22 01:38 pm	
Info	VM centos7_pnolan powered on	centos7_pnolan	Sep-22 12:26 pm	
Info	Reconfigured Ethernet adapters for VM centos7_pnolan	centos7_pnolan	Sep-22 12:26 pm	
Info	VM centos7_pnolan restored from a snapshot	centos7_pnolan	Sep-22 12:25 pm	
Info	VM centos7_pnolan restored from a snapshot	centos7_pnolan	Sep-22 12:24 pm	
Info	VM centos7_pnolan powered off	centos7_pnolan	Sep-22 12:24 pm	
Info	VM centos7_pnolan deleted	centos7_pnolan	Sep-22 11:37 am	
Info	Staging finished for VM 'centos7_pnolan'	centos7_pnolan	Sep-22 11:37 am	
Info	Snapshot successfully taken for VM 'centos7_pnolan'	centos7_pnolan	Sep-22 11:37 am	

37 events loaded

## Manual Sensor Installation

When you start a VM in ransomware recovery, VMware Cloud DR installs a security sensor on the VM that analyzes its behavior and scans its files for malware and known vulnerabilities.

VMware Tools 11.2 or higher is required to install the sensor. VMware Tools must also have the Carbon Black launcher.

For a Windows VM, if VMware Tools 11.2 or later is installed on the VM, then the VMware Cloud DR automatically installs the security sensor when you start the VM when running a recovery plan for ransomware recovery. For Linux VMs, VMware Tools 11.2 or later is required, but the sensor does not install automatically.

You have two methods to install the sensor manually: from vCenter on the recovery SDDC or from the Carbon Black Cloud console.

The network segment the VM is connected to must have internet access, so the VM can reach the security services location within a specific Carbon Black Cloud point of presence (PoP). Make sure that your network and in-guest firewalls do not block access to the URLs listed below.

Common URLs for all Carbon Black Cloud PoPs (UK, US, EU):

```
https://content.carbonblack.io
https://updates.cdc.carbonblack.io
https://packages.vmware.com
```

URLs for specific Carbon Black Cloud PoPs:

#### UK

```
https://ew2-device.carbonblackcloud.vmware.com
```

#### US

```
https://dev-prod05.conferdeploy.net
```

#### EU

```
https://dev-prod06.conferdeploy.net
```

For more information about setting up firewalls to allow access to Carbon Black Cloud, see [Configure a Firewall](#).

## Install Sensor from vCenter on Recovery SDDC

You can [Enable Carbon Black on Virtual Machines](#) if you do not have access to the Carbon Black Cloud service tile, or if you are a partner using the CPN (Cloud Provider Network) console. When you enable Carbon Black Cloud for a VM, you also install the security sensor.

When you enable integrated security and vulnerability analysis, VMware Cloud DR deploys a Carbon Black Cloud plug-in that allows you to install the sensors from vCenter on the recovery SDDC. VMs must be powered on in the recovery SDDC prior to beginning sensor installation, which is done when you run the plan and install the sensors on [Run Plan and Install Windows Sensor](#) or [Run Plan and Install Linux Launcher and Sensor](#) VMs.

## Install Sensor from Carbon Black Cloud Console

To install the sensor from the Carbon Black Cloud console, your user account requires at least one Carbon Black Cloud user role.

To install the sensor on a Windows VM, see: [Run Plan and Install Windows Sensor](#).

To install the sensor on a Linux VM, see: [Run Plan and Install Linux Launcher and Sensor](#).

## Run Plan and Install Windows Sensor

To install the Carbon Black Cloud launcher and sensor on a Windows VM, run the recovery plan, start the VM, and then install the sensor.

When you run a recovery plan for ransomware recovery and configure the plan to pause to manually install the sensor on a VM, you have two methods to install the sensor:

- Log into the recovery SDDC vCenter and [Enable Carbon Black on Virtual Machines](#) (steps 1 through 7 in this task). Use this option if you do not have access to the Carbon Black Cloud service tile, or if you are a partner using the CPN (Cloud Provider Network) console.
- Use the Carbon Black Cloud security console, which requires setting up a configuration file before you install and having at least one Carbon Black Cloud role associated with your user.

### Carbon Black Cloud Console Sensor Installation Prerequisite

If you are installing the sensor from the Carbon Black Cloud, before you perform this task you need to create a configuration file that you will upload to Carbon Black Cloud when you install the Windows sensor.

---

**Note** If you do not have access to the Carbon Black Cloud console, then you do not need a configuration file.

---

Create the configuration file using either the INI, TXT, CONF, or CFG file extension, and then add the following to the contents of the file, exactly as is:

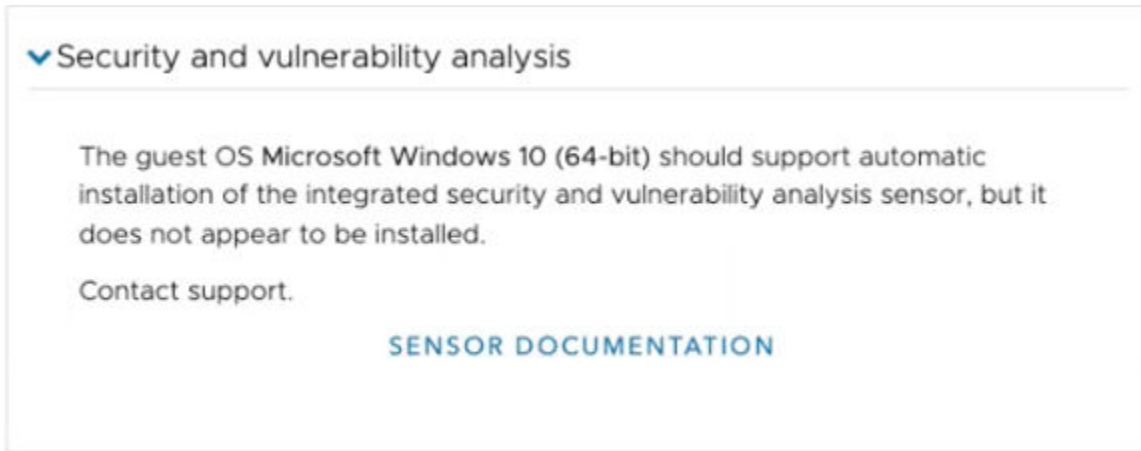
```
[customer]
CurlCrlCheck=false
GroupName=VMware Cloud DR
PolicyName=VMware Cloud DR
DelaySigDownload=0
```

Have this file available before you start the recovery plan plan.

### Procedure

- 1 Start the plan by clicking the **Ransomware Recovery** button.
- 2 In the VMs list of the plan, click the Windows VM.
- 3 In the Ransomware recovery page, click **Start Ransomware Recovery**.
- 4 On the ransomware recovery iteration page, click the **Start VM in Recovery SDDC** button.
- 5 In the **Validate VM in recovery SDDC** dialog box, select a snapshot of the VM you want to analyze. Then, click **Start VM In Recovery SDDC**.

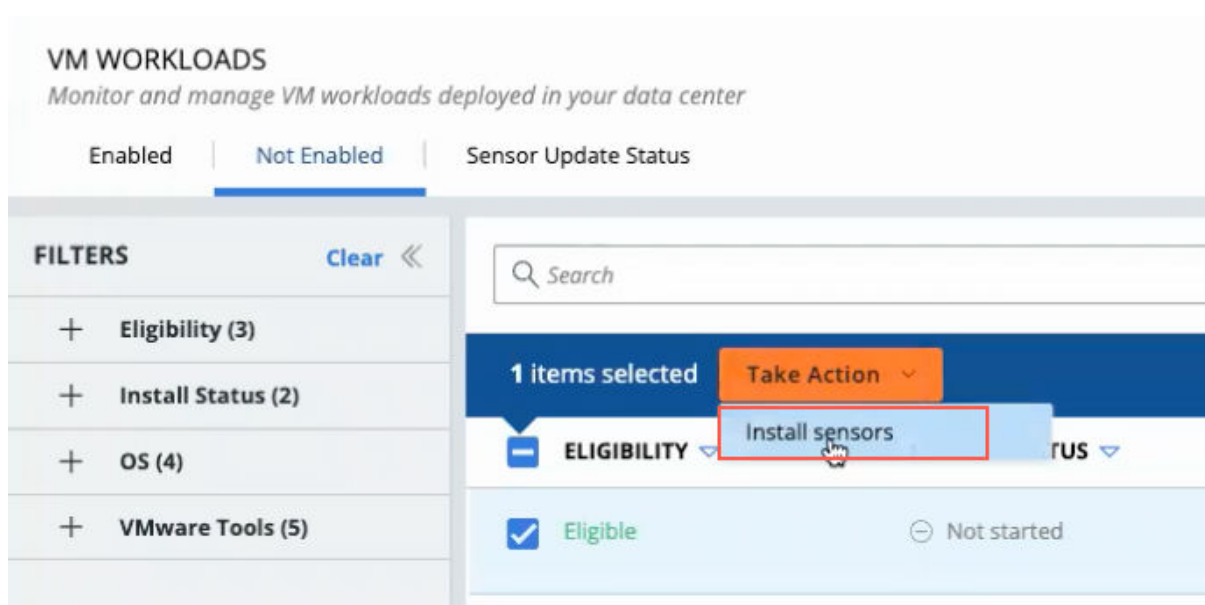
- 6 As the VM starts, the system pauses the VM recovery so you can install the Windows sensor. Under Start VM, you see the following:



- 7 If you are installing the sensor manually using the recovery SDDC vCenter, then follow the steps to [Enable Carbon Black on Virtual Machines](#). You do not need to follow the remaining steps in this task.
- 8 If you are using the Carbon Black Cloud console to install the sensor, go to your VMware Cloud organization home page and click the Carbon Black Cloud tile to launch the console.

**Note** Opening the Carbon Black Cloud console requires at least one Carbon Black Cloud user role.

- 9 In the Carbon Black Cloud console, browse to **Inventory > VM Workloads**.
- 10 Select the **Not Enabled** tab.
- 11 Next, select the Eligible check box next to the VM, and then from the **Take Action** menu, select **Install Sensors**.



- 12 In the **Install Sensor** dialog box, click the Upload File dialog box.
- 13 Select the file you created before starting this task. You see the file uploaded.

Install Sensors

Install sensors on 1 selected workload(s)

Sensors will not be installed on 1 ineligible workload(s)

SENSOR VERSION

Learn more in [Sensor Release Notes](#) and [Sensor Install Guide](#)

OS	SENSOR VERSION
Windows 64-bit	3.8.0.722

SENSOR CONFIGURATION FILE

[Download a template](#) | Learn more in [Sensor Install Guide](#)

Authenticated proxy is not supported as part of configuration setting

config.txt [ [Remove](#) ]

```

1 [customer]
2 CurlCrICheck=false
3 GroupName=VMware Cloud DR
4 PolicyName=VMware Cloud DR
5 DelaySigDownload=0

```

Install

Cancel

- 14 click the **Install** button. After the sensor installs, it appears on the **Enabled** tab.
- 15 When the sensor successfully installs, restart the VM.

#### What to do next

Now that you have installed the sensor, return to the VMware Cloud DR UI so [Integrated Security and Vulnerability Analysis](#) can begin on the VM.

## Run Plan and Install Linux Launcher and Sensor

To install the launcher and sensor on a Linux VM, you run a recovery plan, start the VM, install the Carbon Black Cloud launcher, and then install the sensor.

When you run a recovery plan and start the VM on the recovery SDDC, you are prompted to install the sensor. At this point in the plan, you can install the launcher and then install the sensor on the Linux VM.

You have two methods of installing the sensor on a Linux VM:

- Log into the recovery SDDC vCenter and [Enable Carbon Black on Virtual Machines](#) (steps 1 through 7 in this task). Use this option if you do not have access to the Carbon Black Cloud service tile, or if you are a partner using the Cloud Provider Network (CPN) console
- Use the Carbon Black Cloud security console. This method requires having access to the Carbon Black Cloud console, which requires at least one Carbon Black Cloud user role. For more information, see [Predefined User Roles](#).

#### Procedure

- 1 Start the plan by clicking the **Ransomware Recovery** button.
- 2 In the VMs list of the plan, click the Linux VM.
- 3 In the Ransomware recovery page, click **Start Ransomware Recovery**.
- 4 On the ransomware recovery iteration page, click the **Start VM in Recovery SDDC** button.
- 5 In the **Validate VM in recovery SDDC** dialog box, select a snapshot of the VM you want to analyze. Then, click **Start VM In Recovery SDDC**.
- 6 As the VM starts, the system pauses the VM recovery so you can install the Linux sensor. Under Start VM, you see the following:

#### ▼ Security and vulnerability analysis

The guest OS CentOS 7 (64-bit) does not support automatic installation of the integrated security and vulnerability analysis sensor.

Install it manually to access security.

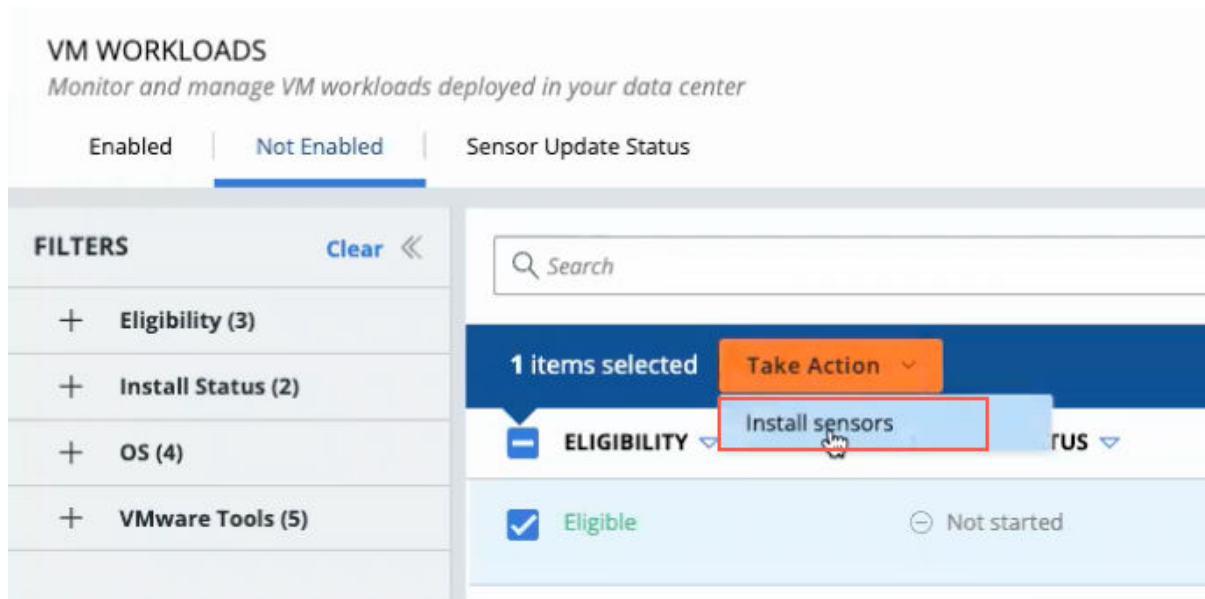
[SENSOR DOCUMENTATION](#)

- 7 If you are installing the sensor manually using the recovery SDDC vCenter, follow the steps to [Enable Carbon Black on Virtual Machines](#). You do not need to follow the remaining steps in this task.
- 8 If you are using the Carbon Black Cloud console to install the sensor, first follow these instructions to install the [Carbon Black Launcher for Linux VMs](#).

- 9 Then, open the Carbon Black Cloud console and navigate to your VMware Cloud organization home page and click the Carbon Black Cloud tile.

**Note** Opening the Carbon Black Cloud console requires at least one Carbon Black Cloud user role.

- 10 In the Carbon Black Cloud console, browse to **Inventory > VM Workloads**.
- 11 Select the **Not Enabled** tab.
- 12 Next, select the Eligible check box next to the VM, and then from the **Take Action** menu, select **Install Sensors**.



- 13 In the **Install Sensor** dialog box, click the **Install** button.
- 14 When the sensor finishes installing, return to the VMware Cloud DR UI. You do not need to restart the VM.

#### What to do next

Now that you have installed the sensor, [Integrated Security and Vulnerability Analysis](#) starts on the VM.

## Uninstalling Sensors

Before you can start VMs in ransomware recovery, you must uninstall any existing Carbon Black Cloud sensors from the VM endpoints.

When you start VMs in ransomware recovery, VMware Cloud DR installs a Carbon Black Cloud sensor on them to enable scanning and VM behavior analysis.

If VMs in your recovery plan already have Carbon Black Cloud sensors installed, uninstall the existing sensors to avoid generating security events for the production organization. Once the existing sensors are removed, when you start a recovery plan, VMware Cloud DR installs new sensors that are associated with a recovery organization.

For Windows instructions, see [How to Uninstall Windows Sensors via Command Prompt](#).

For Linux instructions, see [Uninstall a Linux sensor from an Endpoint](#).

If your Carbon Black Cloud administrator requires code to uninstall a sensor, see here: [Require Codes to Uninstall Sensors at an Endpoint](#) for information about finding that code.

## Pausing a Plan to Uninstall Pre-existing Sensors

To keep Carbon Black Cloud sensors running on VMs until the moment you start a ransomware recovery plan, on the Ransomware Recovery page of the plan select the option named 'Pause when starting VMs to manually remove production security sensors':

Edit plan - DB-backup-nolan

General
Sites
Groups
vCenters
vCenter folders
Compute resources
Virtual networks
IP addresses
Script VM
Recovery steps
Ransomware
Alerts

### Ransomware

☒ **Activate ransomware recovery**  
Allow starting the ransomware test and recovery workflows with this plan.

VMs in this plan will be charged the ransomware add-on.  
[See pricing page.](#)

Confirm you understand the following:

☒ VMs in this plan will generate the additional ransomware-add-on charge.

**Security and vulnerability analysis during ransomware recovery**

☒ **Use integrated analysis**  
Install sensors as VMs are restored in the recovery SDDC to perform security and vulnerability analysis.

☒ **Pause when starting a VM to manually remove production security sensors**  
If your VMs have sensors from a security solution, such as Carbon Black, you should uninstall them when starting validation. This is to avoid polluting your production security solution with alerts occurring in the isolated recovery SDDC.

☐ **Do not use integrated analysis**  
Use other tools to test for ransomware. VMs will start in fully isolated mode.

CANCEL
< BACK
NEXT >
FINISH



## Create a Custom Network Isolation Level

You can create your own custom network isolation level to gain more control over the network environment on the recovery SDDC.

When you start a ransomware recovery plan, VMware Cloud DR creates a predefined set of default firewall rules for your recovery SDDC called network isolation levels.

Each network isolation level (Isolated, Quarantined, External Outbound, and more) exists as a networking and security compute SDDC group on VMware Cloud on AWS. The name of the SDDC group roughly corresponds to the name of corresponding network isolation level in the VMware Cloud DR UI.

You can create your own custom network isolation level in the VMC Console console by creating your own networking and security compute SDDC group on the recovery SDDC, set the group membership criteria, and then configure networking and security for the group. You can create or apply existing firewall rules and other network configuration to the group, which then serves as a network isolation level in VMware Cloud DR.

Do not edit the VMware Cloud DR created firewall rules/SDDC groups because they are deleted when plan gets committed. Custom groups and rules are not deleted by VMware Cloud DR. Custom isolation levels are also visible to multiple plans.

To create your own custom network isolation level, perform these two tasks:

- Create a networking and security compute SDDC group on the recovery SDDC. Do not use the SDDC groups that VMware Cloud DR creates for the pre-defined network isolation levels, as they are deleted when a recovery plan is deactivated.
- Create firewall rules and apply to the SDDC group.

---

**Warning** Using custom network isolation levels is fully supported, but be aware that if you move a VM from a pre-defined network isolation level to a custom isolation level, it removes all network isolation controls provided by VMware Cloud DR. Ensure that you have appropriately configured your custom isolation levels to match your recovery SDDC network isolation requirements.

---

## Create a Group

The first step required to create a custom network isolation level is to create a networking and security compute group for the recovery SDDC you are using for ransomware recovery.

### Procedure

- 1 Log in to VMware Cloud Services at <https://console.cloud.vmware.com>.
- 2 Click **Inventory > SDDC**, then on the recovery SDDC card click **View Details**.
- 3 Click the **Open NSX Manager** button on the upper-right.
- 4 In the **Open NSX Manager** dialog box, click the **Access Via the Internet** button.
- 5 Select **Inventory > Groups**.

- 6 Under Groups, select **Compute Groups**.
- 7 Click **Add Group**.
- 8 Under name, enter a name for the group. The name you give here is the name displayed in the VMware Cloud DR UI when you view network isolation levels.
- 9 Under Compute Members, click the Set Members link.
- 10 In the **Select Members** dialog box, click **Add Criteria**.
- 11 Set the criteria to the following parameters:

Virtual Machine | Tag | Equals | <CustomTagName> | CloudDR-Custom-Scope.

The Tag scope here must be exactly as shown above for VMware Cloud DR to detect the custom isolation level. You must type this scope manually.

---

**Note** Do not use the CloudDR-System-Scope, which is reserved for predefined isolation levels.

---

- 12 Close the **Set Members** dialog box, and then click the **Save** button to save the group.
- 13 Next, click the refresh icon next to Uninitialized to initialize the group.

#### What to do next

At this point, the group has no firewall rules associated with it. To set firewall rules for the custom isolation level, you can create new distributed firewall rules and associate those firewall rules with the SDDC group. (You can also create any other networking configuration available for the group, such as a gateway firewall, distributes IDS/IPS, and so on.) For instruction, see [Create Firewall Rules and Apply to the SDDC Group](#).

## Create Firewall Rules and Apply to the SDDC Group

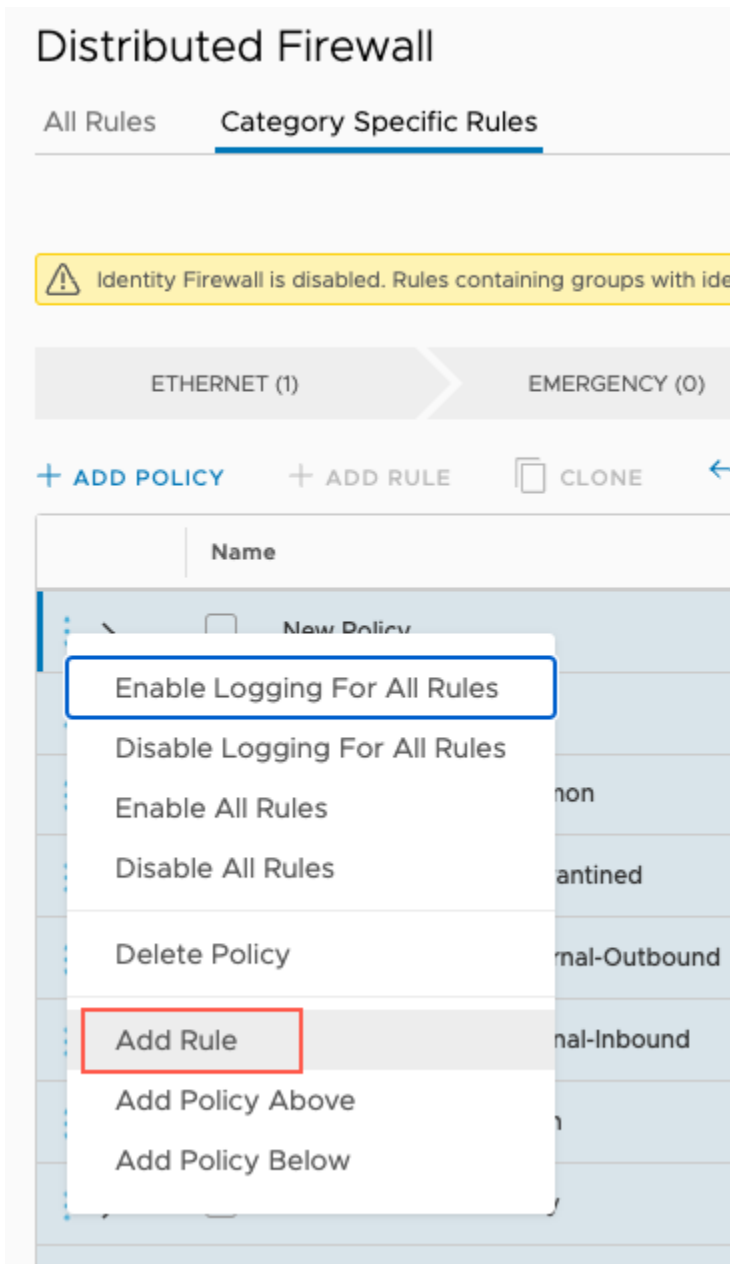
When creating a custom network isolation level, after you have created the SDDC group, you can now apply firewall rules and any other network configuration you want.

These instructions show you how to create a firewall rule to allow a VM in ransomware recovery communicate with the NTP service.

#### Procedure

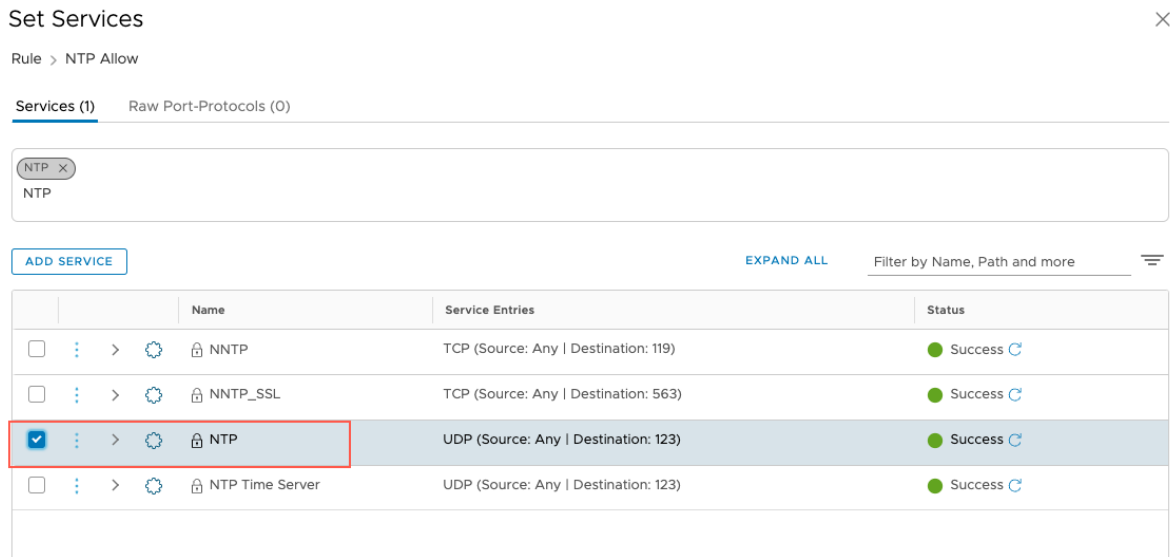
- 1 From the SDDC group select **Networking and Security > Security > Distributed Firewall**.
- 2 Under Distributed Firewall, click **Add Policy**. For the name of the policy, type **CloudDR-Custom-Policy**.

- 3 Click the drop-down menu to the left of the new policy and select **Add Rule**.



- 4 The new rule appears under the policy. Type a name for the policy, for example: **VCDR-Custom-Isolation-NTP-Allow**.
- 5 Next, under the Sources column click in the box which shows the Any selector.
- 6 In the **Set Source** dialog box, select the SDDC group you created for your custom network isolation level.
- 7 Scroll down and then click **Apply**.
- 8 Next, click the Destination to set the destination for the firewall rule.

- 9 Click Services, and in the **Set Services** dialog box, find and select the NTP service.



**Set Services** ×

Rule > NTP Allow

**Services (1)** Raw Port-Protocols (0)

NTP ×

ADD SERVICE EXPAND ALL Filter by Name, Path and more ≡










	Name	Service Entries	Status
<input type="checkbox"/>	NNTP	TCP (Source: Any   Destination: 119)	Success <span>↻</span>
<input type="checkbox"/>	NNTP_SSL	TCP (Source: Any   Destination: 563)	Success <span>↻</span>
<input checked="" type="checkbox"/>	NTP	UDP (Source: Any   Destination: 123)	Success <span>↻</span>
<input type="checkbox"/>	NTP Time Server	UDP (Source: Any   Destination: 123)	Success <span>↻</span>

- 10 Under the Applied To column, click in the field and then in the **Set Applied To** dialog box, select the **Groups** option.
- 11 Select your custom SDDC group, and then scroll down and click **Apply**.
- 12 Under Applied To, click in the field.
- 13 In the **Set Applied To** dialog box, make sure that the Select Applied is configured to Groups.
- 14 Find your custom SDDC group and select it, then scroll down and click **Apply**. When you select the SDDC group here, the firewall rule is associated with the new firewall rule.
- 15 Under the Action column, select **Allow**.

- 16 Last, when you are finished configuring the firewall rule to allow NTP traffic, click the **Publish** button. Publishing the rule might take a moment. Now if you open the Change VM network isolation dialog box in VMware Cloud DR, you can see the new isolation level.

Change VM network isolation - RWR-A-WinSvr3 ×

Select the network isolation rule

<input type="radio"/>		Isolated	Fully isolated. No network access.
<input type="radio"/>		Quarantined + Analysis	Only access network and integrated analysis services.
<input type="radio"/>		External outbound	Allow outbound access to the internet. Use to expose ransomware behavior.
<input type="radio"/>		Internal inbound	Allow inbound access from internal network. No internet access.
<input type="radio"/>		Internal	Allow full access in the internal network. No internet access.
<input type="radio"/>		Internal + External outbound	Allow full access in the internal network,
<input type="radio"/>		Open	Full internal and internet access.
<input checked="" type="radio"/>		VCDR-Custom-Isolation- NTP-Allow	(Custom network access)
<input type="radio"/>		VCDR-Custom-Isolation- DropAll	(Custom network access)

[How to create a custom isolation in NSX-T](#)

CANCEL
CHANGE ISOLATION

### What to do next

Now, you can begin adding other firewall rules and network configurations as needed for the group. Log into VMware Cloud DR and when you run a recovery plan for ransomware recovery and start the VM in validation, you can click the **Change Isolation Level** button and see the new isolation level in the dialog box.

You can create PDF reports for failover and test failover operations, recovery plan configuration changes, and plan compliance checks.

You can create PDF reports for the following:

## Failover and Test Failover Reports

Failover and test failover reports provide information about a completed recovery plan operation.

After a failover or test failover plan has completed (and you have committed or acknowledged the plan), you can generate a PDF report of the plan operation by clicking the **Reports** tab on a plan's details page. Select the Runtime toggle to display the list of finished plan operations. Select a completed plan to see a summary of plan operations. Click **Create PDF report** to generate a report for download.

The generated report contains summaries of the plan configuration at the time of recovery, and a summary of the recovery. The report also includes details for the recovery mappings, the plan's recovery steps, and a detailed report on each action taken during the recovery operation, and any errors that occurred.

---

**Note** Currently, you cannot download VMware Cloud DRPDF reports using Microsoft Edge Browser.

---

---

**Note** Reports for ransomware recovery operations are not currently available. However, you can generate reports for ransomware recovery plan compliance and configuration.

---

## Configuration Reports

From the **Plan Details > Reports** page, click the **Configuration** tab to display a history of plan configuration changes. Each time a plan gets changed and saved, a new version of the plan configuration is created with a time stamp.

Select an item in the list to see a brief summary of the configuration in the area under the list. Click the **Create PDF report** button to generate a PDF download of the report.

The generated report contains a summary of the plan configuration, failover mapping details, and the configured failover steps.

## Compliance Reports

You can generate PDF reports for the automatic compliance checks that run regularly on the system. The manually generated compliance report covers the details for the last completed compliance check, including any configuration and mapping errors that you want to fix.

You can generate a compliance report PDF by first clicking the **Show** button in the Continuous compliance section on a plan's detail page. And then from the Continuous compliance dialog box, click **Create PDF report**.

You can download a PDF of the report, which provide information about each compliance check and also shows detailed information for each compliance check that fails, so you can fix the errors before the plan is tested or run.

## Automated PDF Report and Email

You can configure some events in a plan that triggers report alert emails and to specific email addresses configured with VMware Cloud DR. The system sends automatic PDF reports to the email addressed configured for all failover and test failover operations, and for compliance check events.

# Monitor Events, Tasks, and Alarms

# 16

You can monitor a global collection of events, tasks, and alarms from all the recovery processes for all plans and other system events.

For those events and alarms more important than others and require attention, VMware Cloud DR sends email alerts for users whose email addresses were added to the dialog box.

- **Events.** An event is an indication of activity on a protected site. Events represent the full set of observations and issues raised by VMware Cloud DR. Events can provide information or they can describe a situation that requires attention or action.

When a plan is started and recovery begins, events are populated on the plan details page for the current run of the plan. Each event signifies an action within the substeps in the recovery process.

- **Tasks.** Tasks are operations either in progress or completed, and include things such as the running of a recovery plan, a snapshot, deploying a new recovery SDDC, and more.
- **Alarms.** An alarm indicates an outstanding issue that requires attention. Alarms are triggered by events. VMware Cloud DR displays these notices in the Alarms section which remains visible until you cancel them. The system retains a maximum of 100 active alarms at any given time. If the maximum is exceeded, the system automatically clears the oldest alarms.

Read the following topics next:

- [Forward Events to vRealize Log Insight Cloud](#)
- [SLA Status](#)
- [Viewing Events](#)
- [Viewing Tasks](#)
- [Alarms](#)
- [Running Tasks and Recent Alarms](#)
- [Replication Progress Statistics](#)
- [Configure Email Alerts](#)



## Forward Events to vRealize Log Insight Cloud

You can forward VMware Cloud DR events to VMware vRealize Log Insight Cloud.

You can also forward VMware Cloud DR events to other VMware Cloud Services organizations (in the same region) that have vRealize Log Insight Cloud enabled. You can forward events to any organization you have access to, but you can only forward events to one organization at a time.

For information about data ingestions limits, see [Getting Started with vRealize Log Insight Cloud](#).

---

**Note** You must add the vRealize Log Insight Cloud service to your organization before you can enable event forwarding.

---

### Prerequisites

You can forward all events, including those related to protected sites, snapshot replication, cloud file systems, recovery plans, recovery SDDCs, and user interactions in the following ways:

- From a time in the past to the present and going forward.
- From the present going forward.

You set up event forwarding to vRealize Log Insight Cloud by creating an API key and generating a URL associated with that API key. You can then use the API key and URL to forward events to vRealize Log Insight Cloud.

---

**Note** VMware vRealize Log Insight Cloud has been renamed to VMware Aria Operations for Logs SaaS. This change will be made in this documentation in the next release.

---

### Procedure

- 1 From the left navigation, click **Settings > Log Insight Cloud**.
- 2 In the vRealize Log Insight Cloud integration dialog box, click **Open Log Insight**. vRealize Log Insight Cloud launches.
- 3 In the vRealize Log Insight Cloud UI, from the left navigation click **Configuration > API Keys**.
- 4 Click **New API Key** (upper-right).
- 5 In the New API Key window, enter a name for the key. This name cannot contain spaces and must be unique.
- 6 Click the **Create** button.
- 7 In the generated API Key window, click the **Copy URL** and paste the text, and then click **Copy Key** and paste the text. Keep both values so you can enter them in the VMware Cloud DR UI.
- 8 Switch back to the VMware Cloud DR UI and to the vRealize Log Insight Cloud integration dialog box. If the dialog box has closed, navigate to **Settings > Log Insight Cloud**.
- 9 In the vRealize Log Insight Cloud integration dialog box, enter both the URL and Key.
- 10 Click **Validate**.

11 In the Event forwarding section, select an event forwarding option:

- Forward events starting now.
- Forward events starting from a past date.

---

**Note** Selecting to "Forward events starting from a past date" might result in duplicate events in vRealize Log Insight Cloud. For example, if you configure event forwarding to vRealize Log Insight Cloud, stop event forwarding, and then later restart event forwarding using the "Forward events starting from a past date," you might see duplicate events in vRealize Log Insight Cloud.

---

12 If you choose 'From a past date,' click one of the pre-set time buttons. Or you can use the calendar picker to select a date in the past that you want to start forwarding events from.

13 If forwarding from a past date, you can optionally use the Calendar picker to select a stop date and time for event forwarding.

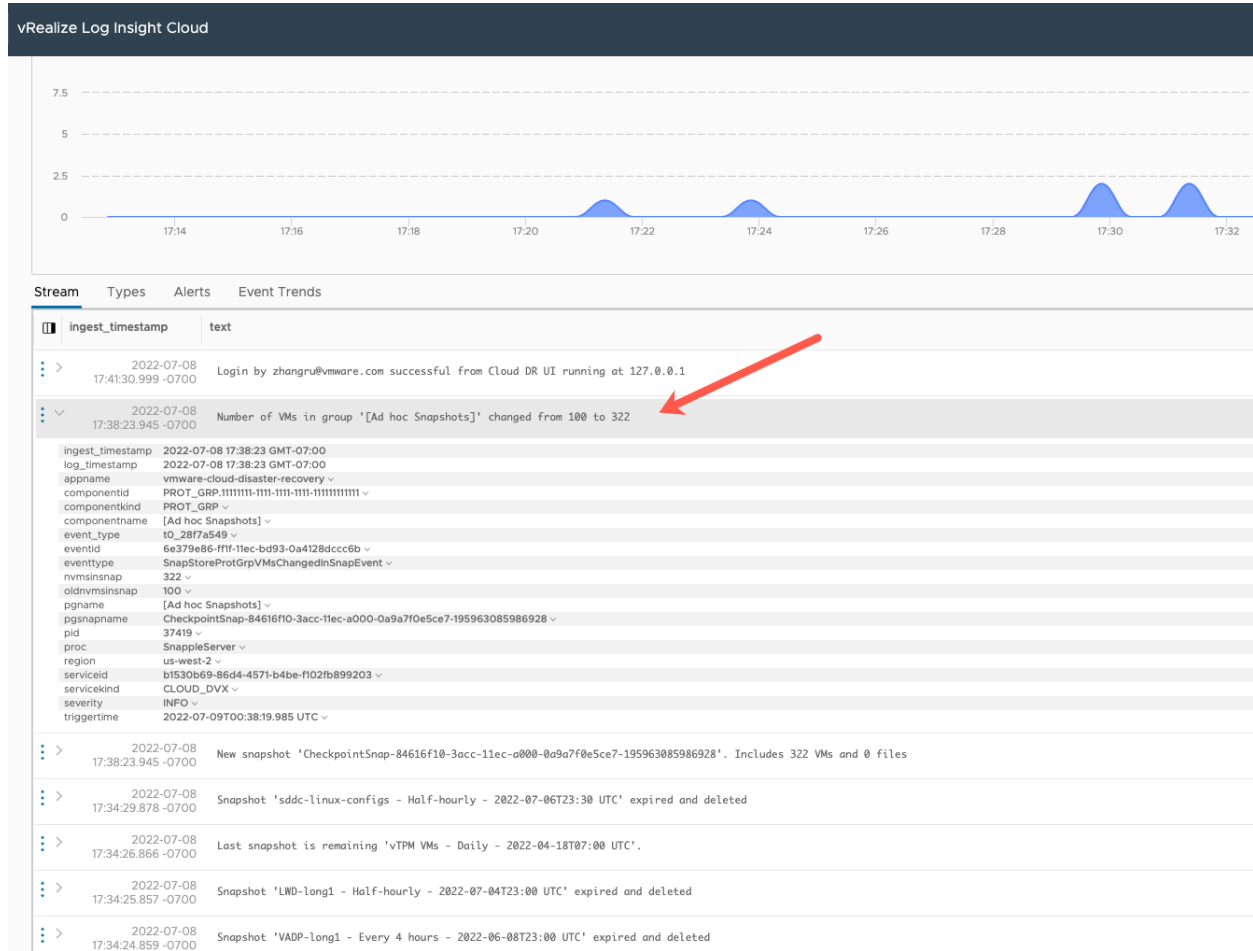
If there are no past events, then only events starting from the present time are forwarded. The first time you connect, it can take up to a minute or more to begin forwarding events.

14 Click **OK**. You see a green checkmark icon if the operation was successful.

#### What to do next

Once you have configured event forwarding, you can see VMware Cloud DR events as they appear in vRealize Log Insight Cloud. For example, this image shows an information event from

VMware Cloud DR indicating that the number of VMs in a protection group has increased:



## Update Event Forwarding API Key

You can update the vRealize Log Insight Cloud API key you use to forward events.

### Prerequisites

If you want to rotate your event forwarding API key, or someone has deactivated an API key, you can regenerate a new one in vRealize Log Insight Cloud. When you regenerate a new API key, you also see the API key and the paired ingestion URL. You can then replace the API key and URL in VMware Cloud DR.

Make sure you have the new or regenerated API key and URL before you start this task.

### Procedure

- 1 From the left navigation, click **Settings > Log Insight Cloud**.
- 2 In the vRealize Log Insight Cloud integration dialog box, click **Update API Key**.
- 3 In the vRealize Log Insight Cloud integration dialog box, enter both the API key and ingestion URL.
- 4 Click **Validate**.

- 5 In the Event forwarding section, choose an event forwarding option:
  - Forward events starting now.
  - Forward events starting from a past date.
- 6 If you choose 'From a past date,' click one of the quick-select buttons. Or in 'Since date' you can use the calendar picker to select a date in the past to start forwarding events from.
- 7 In the Until box, use the Calendar picker to select an event forwarding stopping date and time.
- 8 Click **OK**.

## Stop Event Forwarding

You can stop event forwarding to vRealize Log Insight Cloud any time.

### Prerequisites

If you stop forwarding event, you must reconfigure event forwarding to send events to vRealize Log Insight Cloud. Event forwarding stops automatically if the vRealize Log Insight Cloud ingestion limit is reached. If you reach an ingestion limit, VMware Cloud DRreports the error and stops forwarding events. When you set up a new URL and API key and restart forwarding events, you can choose the 'Forward from last stop' option in the vRealize Log Insight Cloud integration dialog box.

### Procedure

- 1 From the left navigation, click **Settings > Log Insight Cloud**.
- 2 In the vRealize Log Insight Cloud integration dialog box, click **Stop Forwarding**.
- 3 In the Stop forwarding dialog box, click **Stop Forwarding**.

## SLA Status

The SLA status page gives you a high level operational status of your most important configurations related to VM protection and recoverability.

### Service Level Agreement (SLA) Status for Protection and Recoverability

SLA Status is the service level agreement between DR administrators and their users. From the left navigation select **Monitor > SLA Status** to see SLA status for important configurations, grouped as follows:

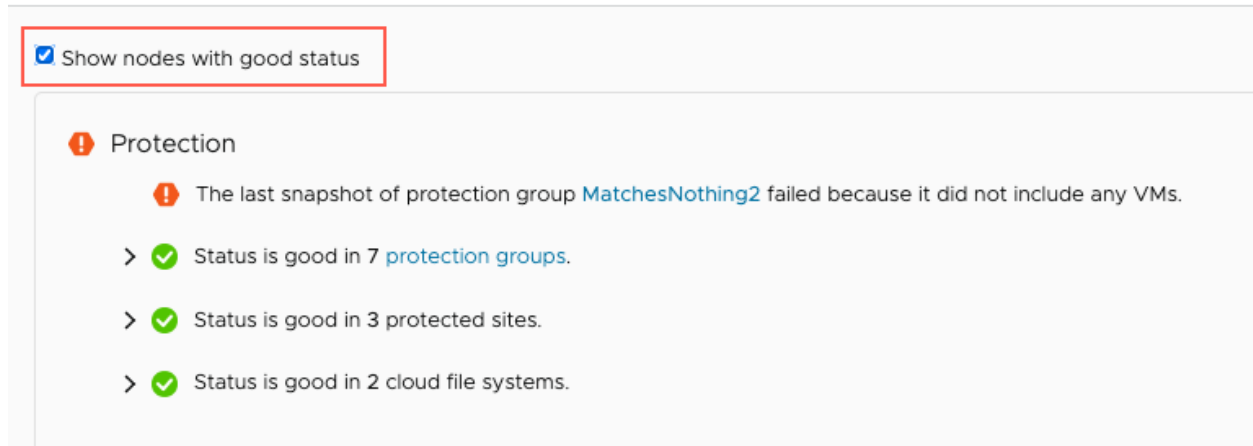
- **Protection.** Protected sites, protection groups, snapshots, and cloud file systems.

- **Recoverability.** Recovery plans, recovery SDDCs, and cloud file systems.

**Note** SLA Status reporting is not real time. Typically, there is a short delay of ~1 minute for reported data to reflect in the SLA Status page. Delays can be slightly longer (from ~1 through 6 minutes) in cases where an SLA status is delayed, for example, if an entity is not available or was removed.

You can toggle the 'Show nodes with good status' button to show or hide all nodes with a good status:

## Monitor



When an SLA node is not in good health (warning or critical), an SLA status banner displays on both the dashboard and individual configuration pages, providing links to nodes that require attention.

**Note** SLA banners display continuously until you resolve the issue.

For example, the following image shows critical SLA status for two snapshots in a protection group. You can click the links for each node to investigate:

## Dashboard for US West (Oregon)

▼ There are 2 critical issues in [protection groups](#).

- The last snapshot of protection group [max\\_test](#) includes only 2 out of 3 VMs.
- The last snapshot of protection group [MatchNothing](#) failed because it did not include any VMs.

Recovery region summary [Show setup guide](#)

<b>Status</b> Protection Recoverability	<b>Cloud file systems</b> 1 <b>299.6 MiB</b> <small>Calculated every 12h</small>	<b>Protected sites</b> 0 VMware Cloud 2 On-premises
<b>Recovery SDDCs</b> 1	<b>Protection</b> 3 Protection groups 4 VMs <small>VMs in groups</small>	<b>Disaster recovery</b> 0 DR plans 0 VMs <small>VMs in plans</small>

SLA banners and links also display at the top of each configuration page. For example, the following image shows SLA status banners at the top of the Protection groups page:

## Protection groups

▼ There are 2 critical issues in [protection groups](#).

- The last snapshot of protection group [max\\_test](#) includes only 2 out of 3 VMs.
- The last snapshot of protection group [MatchNothing](#) failed because it did not include any VMs.

Protection groups [CREATE PROTECTION GROUP](#)

Protection group	Site	Cloud file system	Health	Schedule	Frequency	Quiesce	Last snapshot	VMs
<a href="#">MatchNothing</a>	Denver	cloud-backup-1	Critical	Active	Standard	No	--	--
<a href="#">max_test</a>	Denver	cloud-backup-1	Critical	Active	Standard	No	Jul-13 11:01 pm (9h ago)	2 VMs
<a href="#">NoProblems</a>	HFSCapable	cloud-backup-1	OK	Active	Standard	No	Jul-14 07:53 am (27m ago)	1 VM

## SLA Status Severity Levels

SLA status severities include:

Status	Meaning
Good	SLA health is normal. Your protection groups are successfully replicating snapshots to one or more cloud file systems. Recovery plans are in compliance and ready for failover.
Warning	SLA health has some issue and needs attention. Possible issues: an expired snapshot schedule, replication errors such as missing VM, connectivity issues with protected sites, and more.

Status	Meaning
Critical	<p>An SLA node is either not functioning normally or might stop functioning soon. Critical issues can include loss of connectivity to a protected site, a snapshot replication failure, or a recovery plan failover malfunction.</p> <p>A critical status can also display if VMware Cloud DR cannot determine the status of an entity. For example, if a connector cannot connect to the cloud file system, it cannot report its status, and so after some time it is reported as critical.</p>
NA/Unmonitored	<p>VMware Cloud DR cannot determine the SLA node status. An NA or unmonitored status can occur when you create a node and the SLA Status check has not been run yet, or if you stop a protection group snapshot schedule.</p>

## Viewing Events

The Events list shows a running list of all VMware Cloud DR events that occur in system, both system-generated events and user actions.

You can filter the Events list by:

- Event category, such as user operations ("Audit"), protection group and snapshot operations ("Protection"), protected site events ("Site"), recovery plan failovers, recovery SDDC events, and more.
- Event severity such as Info, Warning, Error, Emergency.
- Event duration, such as Today, Last 24 hours, Last 7 days, or any specific time frame you select.

---

**Note** The event category filters and severity filters are mutually exclusive. When you filter by event category, you cannot also filter by severity. Conversely, if you filter by event severity, you cannot also filter by event category. You can however, combine the timestamp filter with either category filters or severity filters.

---

You can also filter the events list page results further by text strings. This search filters the currently displayed list of events in the UI based on the entered text:

## Monitor

Filter events...

Audit
Site
Recovery SDDC
Failover
Protection
Other

*i*
⚠️
❗
❌

Anytime

Severity	Description	Target	Timestamp
<i>i</i> Info	Snapshot 'GPT-PG - Daily - 2022-02-10T18:30 UTC' expired and deleted	GPT-PG	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'cust-script-grp - Daily - 2022-02-10T18:30 UTC' expired and deleted	cust-script-grp	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'win-notools-wintool-mix - Daily - 2022-02-10T18:30 UTC' expired an...	win-notools-wintool-mix	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'win2012-vm3-pg - Daily - 2022-02-10T18:30 UTC' expired and deleted	win2012-vm3-pg	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'win2016-sql2016-pg - Daily - 2022-02-10T18:30 UTC' expired and del...	win2016-sql2016-pg	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'windows-vm11-pg - Daily - 2022-02-10T18:30 UTC' expired and delet...	windows-vm11-pg	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'hammerDB-PG - Every 4 hours - 2022-02-10T18:30 UTC' expired an...	hammerDB-PG	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'bitlocker-vm - Daily - 2022-02-10T18:30 UTC' expired and deleted	bitlocker-vm	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'Build_187 - Every 4 hours - 2022-02-10T18:30 UTC' expired and dele...	Build_187	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'pg1-centos-win-manish - Daily - 2022-02-10T18:30 UTC' expired and...	pg1-centos-win-manish	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'centos-vm1 - Daily - 2022-02-10T18:30 UTC' expired and deleted	centos-vm1	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'centos-vm13-pg - Daily - 2022-02-10T18:30 UTC' expired and deleted	centos-vm13-pg	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'quiesce-centos-50VM - Daily - 2022-02-10T18:30 UTC' expired and ...	quiesce-centos-50VM	Feb-17 10:51 am
<i>i</i> Info	Snapshot 'open-vm-tools-pg - Daily - 2022-02-10T18:30 UTC' expired and dele...	open-vm-tools-pg	Feb-17 10:51 am
<i>i</i> Info	Task drc-a0334d8c-901f-11ec-a6c8-0a379a9c601d for PG centos-40vdisk-vss-...	centos-40vdisk-vss-pg1	Feb-17 10:51 am
<i>i</i> Info	New snapshot 'centos-40vdisk-vss-pg1 - Daily - 2022-02-17T18:30 UTC'. Includ...	centos-40vdisk-vss-pg1	Feb-17 10:51 am
<i>i</i> Info	Task drc-a0330ef8-901f-11ec-a6c7-0a379a9c601d for PG quiesce-centos-50V...	quiesce-centos-50VM	Feb-17 10:45 am
<i>i</i> Info	New snapshot 'quiesce-centos-50VM - Daily - 2022-02-17T18:30 UTC'. Includes...	quiesce-centos-50VM	Feb-17 10:45 am
<i>⚠️</i> Warr	Failing to use app-consistent snapshot 49 out of 50 VMs in group 'quiesce-cen...	quiesce-centos-50VM	Feb-17 10:45 am
<i>⚠️</i> Warr	Failed to use app-consistent snapshot 49 out of 50 VMs in snapshot 'quiesce-c...	quiesce-centos-50VM	Feb-17 10:45 am

50 events loaded
Load more

The Events list shows a maximum of 100 events at first viewing. To view more events, click the **Load more** button at the bottom of the page.



## Event Details

When you select an event, a details panel opens to show more information about the event, such as event severity, timestamp, target system (replication or failover tasks), and description. For a snapshot task, the panel provides details about the protection group name, snapshot name, and VMs included in the snapshot. Some events, such as snapshot tasks, generate a log when completed, which you can access by clicking the **Snapshot log** button in the panel below the events.

The screenshot displays the VMware Cloud Disaster Recovery event details interface. At the top, there is a search bar labeled 'Filter events...' and a set of tabs: Audit, Site, Recovery SDDC, Failover, Protection, and Other. Below these tabs is a list of events. Each event row includes a radio button, a severity icon (Info), a description, a target, a timestamp, and a user name. One event is selected, and its details are shown in a panel below the list. The details panel includes fields for Severity (Info), Timestamp (Feb-22 09:06 am (1h ago)), Target (tfw hw v17 linux and windows), and Description (Task drc-7d00fd50-93f8-11ec-ad90-060988610bad for PG tfw hw v17 linux and windows completed successfully and protection group snapshot tfw hw v17 linux and ... - Every 8 hours - 2022-02-22T16:00 UTC is created). On the right side of the details panel, there are fields for Protection group (tfw hw v17 linux and windows), Snapshot (tfw hw v17 linux and ... - Every 8 hours - 2022-02-22T16:00 UTC), Start time (Feb-22 09:02 am), End time (Feb-22 09:06 am), VMs in snapshot (2 Quiesced VMs 0), Duration (4 minutes, 35 seconds), Data transferred (543.6 gib), and Snapshot log (a button highlighted with a red box).

## Event Types

Events are grouped by the following types:

Event Type	Description
Audit	Events that capture user actions, such as configuration changes and recovery plan executions. Includes the event name, the user who performed it, and the source IP address where the event was initiated.
Site	Events related to sites, such as adding or removing a protected site, downloading a DRaaS Connector VM, registering vCenter.
Recovery SDDC	Events related to a recovery SDDC, such as deploying or deleting a recovery SDDC, adding or removing a network, and more.

Event Type	Description
Failover	Events related to recovery plan operations, such as a recovery or failback operation, acknowledging and committing a plan, deleting a plan, and more.
Protection	Events related to protection groups, such as creating or deleting protection groups, snapshots, and more.
Other	Miscellaneous events, such as space reclamation tasks, user notifications for alert recipients, and more.

## Event Severities

You can filter by the event severities listed in the following table:

Severity	Description
Info	Descriptive information about an operation or state of the system, including information that might be useful.
Warning	Indicates a situation that requires attention. A warning condition does not affect system operation.
Error	Indicates that an operation failed or that there was a hardware error. An error condition does not affect continued system operation.
Emergency	Indicates that a fatal event occurred. A fatal event affects continued system operation.

## Event Timestamps

You can restrict the events displayed by timestamp, which allows you to set a specific time or time duration by which to filter events.

You can use the event timestamp filter when you want to see recent events, old events, or events that occurred during a specific time period.


You can select 'quick select' times, or use the Since and Up to calendar selectors to specify a time frame by which to filter the events.

Search by timestamp ×

Select the start and end time to search for events.


Quick select: Today Last 24 hours Last 7 days Anytime

Since

July 15, 2021 12:00 AM 

×

Up to

July 22, 2021 12:00 AM 

×


Cancel

OK

## Viewing Tasks

The Tasks list shows all running and completed tasks in your system.

Tasks can include snapshot replication, system upgrade, guest file download, failover operations and more. You can filter the task list by category, such as disaster recovery or ransomware operations, protection (snapshot replication, guest file ZIP packaging for file download), and infrastructure (system upgrades, cloud file system deployments).

Monitor							
Events Alarms <b>Tasks</b>							
0 running 12 finished		0 disaster recovery <b>3 protection</b> 9 infrastructure					
Status	Type	Description	Start time	End time / progress	Data transferred	Logical throughput	
✓ Finished	Protection	Snapshot PG DB-Test-SQL - Daily - 2021-09-22T22:00 UTC	Sep-22 03:00 pm	Sep-22 03:01 pm	51200 MiB	--	
✓ Finished	Protection	Preparing guest file archive for recovery	Sep-22 01:41 pm	Sep-22 01:41 pm	--	--	
✓ Finished	Protection	Snapshot PG DB-Test-SQL - Daily - 2021-09-22T20:25 UTC	Sep-22 01:25 pm	Sep-22 01:30 pm	51200 MiB	--	

## Alarms

The Alarms list shows a history of all notifications that require your attention.

The Alarms list shows all issues that require your attention. Events trigger alarms when something did not occur as planned in the event. You can search the alarms list and also filter it by severity.

Alarms can be cleared by clicking the small X at the far right of the alarm entry. If you click the small X at the top of the column, it clears all alarms. To see cleared alarms, select the Show cleared alarms option.

Monitor
Events
Alarms
Tasks

Filter alarms...
☐ Show cleared alarms





Severity	Description	Target	Timestamp	
Warning	PLANNED_FAILOVER failover started from site tfw-...	tfw-vcsa - dvx23	Oct-07 09:15 pm	x
Warning	UNPLANNED_FAILOVER failover started from site t...	tfw-vcsa1 - Cloud2	Oct-07 09:02 pm	x
Warning	UNPLANNED_FAILOVER failover started from site t...	tfw-vcsa1 - Cloud2	Oct-07 08:55 pm	x
Warning	PLANNED_FAILOVER failover started from site tfw-...	tfw-vcsa - dvx23	Oct-07 02:46 pm	x
Warning	PLANNED_FAILOVER failover started from site tfw-...	tfw-vcsa - dvx23	Oct-07 02:15 pm	x
Warning	UNPLANNED_FAILOVER failover started from site t...	tfw-vcsa1 - Cloud2	Oct-07 01:52 pm	x
Warning	PLANNED_FAILOVER failover started from site tfw-...	tfw-vcsa - dvx23	Oct-07 12:50 pm	x

## Running Tasks and Recent Alarms



The right side of VMware Cloud DR UI displays status information about running tasks, finished tasks, and recent alarms.

Here you can view currently running tasks, cancel some tasks, view completed tasks, and any alarms that might be important to you.





## Recent alarms

-  SDDC VCHA SDDC will expire in 1 days  
17h ago [x](#)
-  SDDC VCHA SDDC is provisioned and running for 59 days  
17h ago [x](#)
-  SDDC VCHA SDDC is provisioned and running for 58 days  
2d ago [x](#)
-  SDDC VCHA SDDC will expire in 2 days  
2d ago [x](#)

## Running tasks

-  Recovering from ransomware... [R WR windows](#)  
4d ago
-  Recovering from ransomware... [R WR1](#)  
4d ago

## Recently finished tasks

-  Snapshot PG w2-hs3-q0607\_cj\_100 - Half-hourly - 2022-11-14T17:00 UTC  
13m ago
-  Snapshot PG w2-hs3-q0607\_cj\_100 - Half-hourly - 2022-11-14T16:30 UTC  
43m ago
-  Snapshot PG w2-hs3-q0607\_cj\_100 - Half-hourly - 2022-11-14T16:00 UTC  
1h ago
-  Snapshot PG live - dvx38 - Every 8 hours - 2022-11-14T16:00 UTC  
1h ago

For a more detailed list of all tasks and events in the VMware Cloud DRUI, see [Chapter 16 Monitor Events, Tasks, and Alarms](#).

## Replication Progress Statistics

You can monitor the real time progress of snapshot replication tasks from the dashboard.

Replication progress shows you the percentage of data being transferred and the rate of data transferred ("logical throughput") during and after the task.

### Replication Progress in the Task List

You can monitor the progress of snapshot replication tasks from the Tasks list and view the following information:

**Note** Replication statistics only measure replication from a protected site to a cloud file system. These statistics do not represent failover or failback progress.


Metric	Description
Start time	Time when the task started.
End time/Progress	During the task, you can see the percentage of the total task completed. When the task is finished, it shows the time of completion.
Logical throughput	The progress during snapshot replication from a protected site to a cloud file system, since the start of the task. Measured in Megabits (Mbps) and Gigabits per second (Gbps). This measurement only displays while the task is in progress.

**Note** This view does not show progress for single VM restore tasks

### Replication Statistics on Protected Site Page

The protected sites page shows replication statistics for both replication throughput and throttle (if configured), and also shows throughput for restore operations (during the operation).

Cloud backup target



prasanna-scfs-1

Backup 778 Mbps

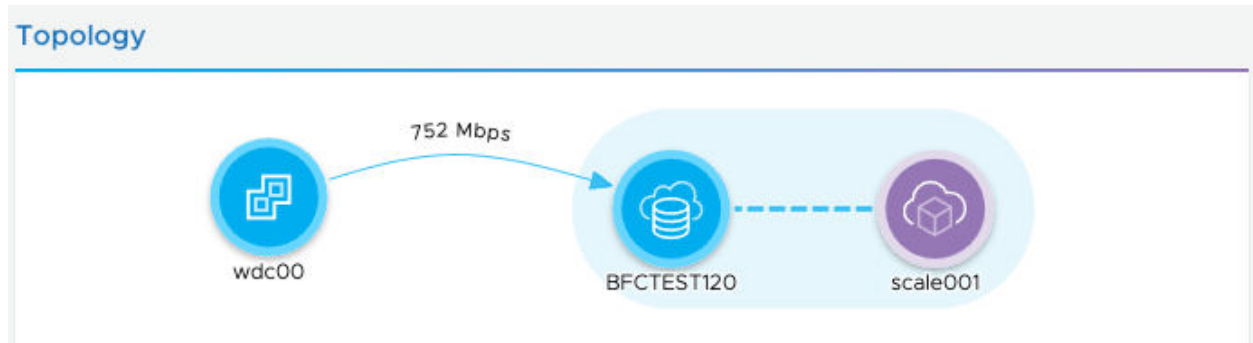
Throttle None

The protected site page also shows the current [Throttle Replication](#) maximum, if configured. You can also view replication transfer rate in the Topology pane.

If no snapshots are replicating to a cloud file system, then the throughput value is empty.

## Replication Throughput in the Dashboard Topology Pane

The Topology map shows The rate of data being transferred to and from the protected site to a cloud file system, measured in Mbps and Gbps.



## Configure Email Alerts

You can configure VMware Cloud DR to send an email when SLA statuses change and when a recovery plan finishes running.

Email addresses added here also receive emails for VMware Cloud DR system alerts. For more information, see [Chapter 16 Monitor Events, Tasks, and Alarms](#). To send an email to users when a recovery plan finishes running, you first must add their email addresses in the [Alerts](#) page of a plan.

---

**Note** VMware Cloud DR uses the AWS mail service for email notifications. Recipients of these emails must respond to the AWS Email Address Verification Request before getting email from VMware Cloud DR.

---

### Procedure

- 1 From the left navigation, select **Settings**.
- 2 Click the **Email alerts** button.
- 3 Select the SLA status alerts option if you want to send an email to all listed recipients when any SLA status changes.
- 4 Under Email alert recipients, click the **Add** button to add user email addresses. Each user is sent an email verification link from AWS that they must click to validate their email address. After verifying their email address with AWS, they can start receiving alert emails from VMware Cloud DR.
- 5 Under the Email alert sender section, select a sender email address.

- 6 You can optionally add a Sender name for the emails.

### Configure email alerts ✕

---

#### SLA status alerts

---

On changes of SLA status, send email alerts to the email alert recipient list.

☒ Send SLA change email alerts

#### Email alert recipients

---

Global VMware Cloud DR alerts will be sent to this list. Recipients in this list can also be selected on each plan to get plan-specific alerts.

VMware Cloud DR uses the AWS mail service. Recipients must respond to the AWS Email Address Verification Request before getting email from VMware Cloud DR.

ADD

#### Email alert sender

---

Choose one of the email recipients to use as the "From" for all VMware Cloud DR emails.

Sender email address

< select > ▼

CANCEL OK

- 7 Click **OK**.



# Upgrade Process

# 17

VMware Cloud DR provides frequent software upgrades each year.

## Before the Upgrade

Before a VMware Cloud DR upgrade, you receive a series of email notifications from VMware support, starting two weeks before the upgrade. You then receive notices at one week and then 24 hours before the upgrade, so you have plenty of time to prepare for the upgrade.

The email notifications specify the maintenance timeframe of your upgrade, which can be a roughly eight hours. The software upgrade takes about 30 minutes to complete during the maintenance timeframe.

Before the upgrade window starts, make sure you commit all recovery plans and stop any running plans. VMware Cloud DR software cannot upgrade if any recovery plans are uncommitted or running. You also cannot start any failover operations during the upgrade, or the upgrade fails. Software upgrades cannot run if a recovery plan is in the middle of a failover operation or if a plan is running.

## During the Upgrade

All of your users must log out of VMware Cloud DR during the upgrade. If you are logged in and attempt a failover operation, the VMware Cloud DR UI displays an error message that an upgrade is in progress, and the system prompts you to try again later. The VMware Cloud DR UI logs out users until the upgrade is complete.

During the upgrade, it might be possible to log in to the VMware Cloud DR UI. If you log in before the upgrade is over, the overall system health status shows as Unknown and the Topology map icons temporarily disappear then reappear, which is normal and expected behavior.

## After the Upgrade

After the maintenance timeframe is over, you can log in and begin using the service.

# Deactivate VMware Cloud DR

# 18

If you want to stop using VMware Cloud DR in a recovery region, you must deactivate the service.

You can also perform these steps if you want to stop using VMware Cloud DR in one region and want to activate a new region (although deactivation is not required to activate a new region).

Deactivating VMware Cloud DR does not cancel any previously purchased protected capacity, VM, or ransomware recovery subscriptions.

---

**Note** When you deactivate a region, all snapshots in the region are deleted.

---

To deactivate VMware Cloud DR, perform the following tasks in this order:

- 1 Remove all DRaaS Connectors from all protected sites. See [Remove a DRaaS Connector from a Protected Site](#).
- 2 Delete all recovery SDDCs. See [Delete a recovery SDDC](#).
- 3 Deactivate the recovery region from the Global DR Console. (Do this step last.) See [Deactivate Recovery Region](#). Usage charges for VMware Cloud DR are not stopped until this step is completed.