RecoverPoint for VMs

Cloud Solutions Guide

5.2

August 2020
Rev. 3

Contents

Figures

Tables

Preface

As part of an effort to improve product lines, we periodically release revisions of software. Therefore, some functions described in this document might not be supported by all versions of the software currently in use. The product release notes provide the most up-to-date information on product features.

Contact your technical support professional if a product does not function properly or does not function as described in this document.

NOTE: This document was accurate at publication time. Go to Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

Purpose

This document includes conceptual information on managing a RecoverPoint for Virtual Machines system.

Audience

This document is intended for use by vSphere administrators who are responsible for managing the RecoverPoint for Virtual Machines system.

Related documentation

The following publications provide additional information:

- *RecoverPoint for Virtual Machines Release Notes*
- *RecoverPoint for Virtual Machines Installation and Deployment Guide*
- *RecoverPoint for Virtual Machines FLEX Plugin Administrator's Guide*
- *RecoverPoint for Virtual Machines Deployment REST API Programming Guide*
- *RecoverPoint for Virtual Machines REST API Programmer's Guide*
- *RecoverPoint for Virtual Machines Security Configuration Guide*
- *RecoverPoint for Virtual Machines Scale and Performance Guide*
- *RecoverPoint for Virtual Machines FAQ*
- *Recoverpoint for Virtual Machines Simple Support Matrix*

In addition to the core documents, we also provide White papers and Technical Notes on applications, arrays, and splitters.

Typographical conventions

This document uses the following style conventions:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
`Monospace italic`	Used for variables
`Monospace bold`	Used for user input

[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

Technical support, product, and licensing information can be obtained as follows:

Product information For documentation, release notes, software updates, or information about products, go to Online Support at https://support.emc.com.

Technical support Go to Online Support and click Service Center. You will see several options for contacting Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

1

Introduction

This section contains the following topics:

Topics:

Cloud solution overview and architecture

RecoverPoint for VMs 5.2.1 introduces the RecoverPoint for VMs cloud solution.

The RecoverPoint for VMs cloud solution uses RecoverPoint for VMs (5.2.1 or later) to protect your production VMs, and Cloud DR (18.4 or later) to orchestrate recovery, and recover your protected VMs on the Amazon cloud. RecoverPoint for VMs works in conjunction with Amazon Web Services (AWS) to protect your production VMs, by replicating them to the AWS cloud. Copy data is compressed, encrypted, and stored as incremental snapshots in an Amazon S3 bucket. You can manually retrieve each copy's snapshots from the cloud using the Cloud DR Server (CDRS) user interface, and the data is decompressed upon retrieval.

In the RecoverPoint for VMs cloud solution, the **RecoverPoint for VMs vCenter plugin** is used to register an existing AWS cloud account and S3 bucket, deploy a CDRS, and protect your VMs. During VM protection, you can modify the default copy retention policy, and specify the period of time during which copies should be available for recovery. Once every 24 hours, a temporary **Retention Service** EC2 instance is automatically launched in AWS to consolidate the snapshots of all copies whose retention policy has expired. You can also modify the default copy **Snap Replication** policy, and specify the interval between replicated snapshots. By default, copy snapshots are replicated once every hour, and they are retained on AWS for 5 days.

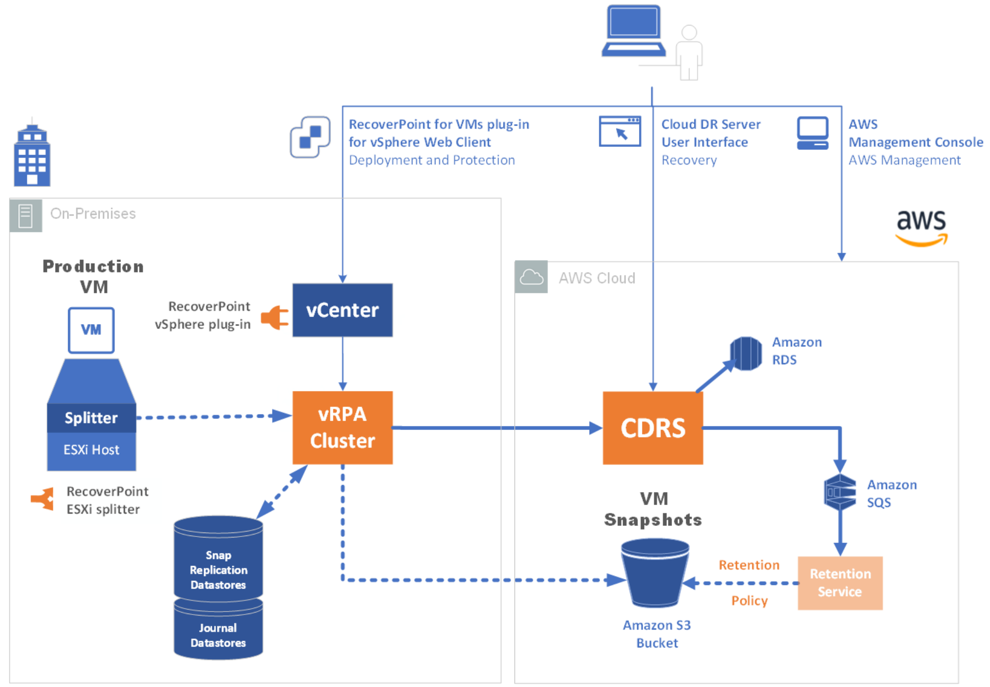

Figure 1. Protection architecture

After your VMs are protected, the **Cloud DR Server** user interface is used to create, manage, test and fail over DR plans in AWS. Cloud DR provides crash-consistent, image-level, VM recovery to native AWS EC2 instances or VMware Cloud on AWS. DR plans are used to recover multiple VMs and pre-configure recovery orchestration, including network and security groups association, VM boot order definition, and VM type selection.

Using Cloud DR, you can also accelerate the recovery process by creating rapid recovery images for protected VMs. Creating a rapid recovery image starts a rehydration process and converts the copy VMDK files to the format required by Amazon before recovery. When a recovery operation is requested by the user through the CDRS, a temporary **Restore Service** EC2 instance is launched, on-demand, to handle the recovery process.

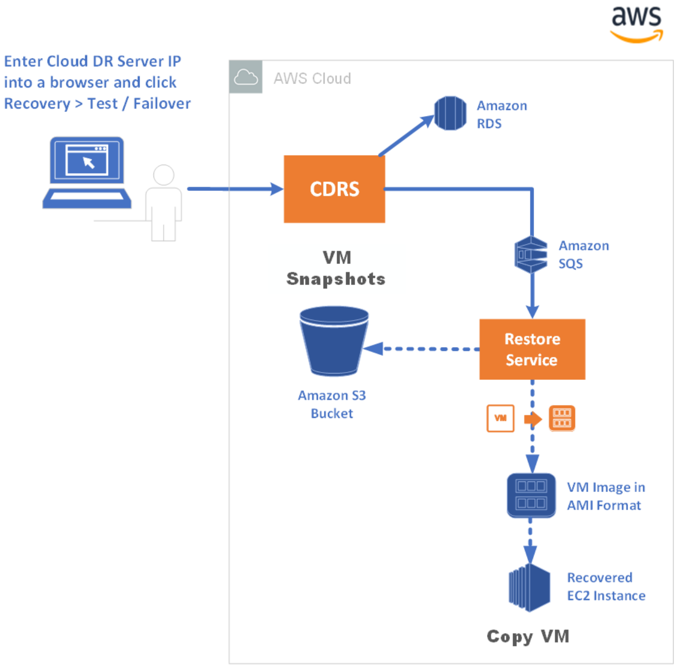

Figure 2. Test and failover to AWS architecture

After failover, you can fail back from AWS to an on-premises vCenter (requires an on-premises CDRA).

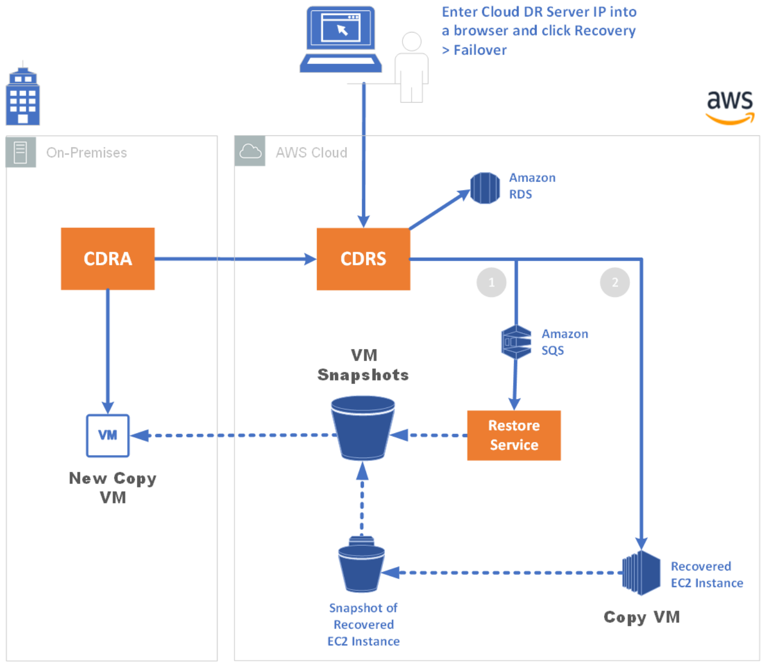

Figure 3. Failback from AWS architecture

You can also recover the VM snapshots in your Amazon S3 bucket to VMware Cloud on AWS (requires a CDRA on the VMware Cloud on AWS). This architecture can also be used to recover VM snapshots to an on-premises vCenter.

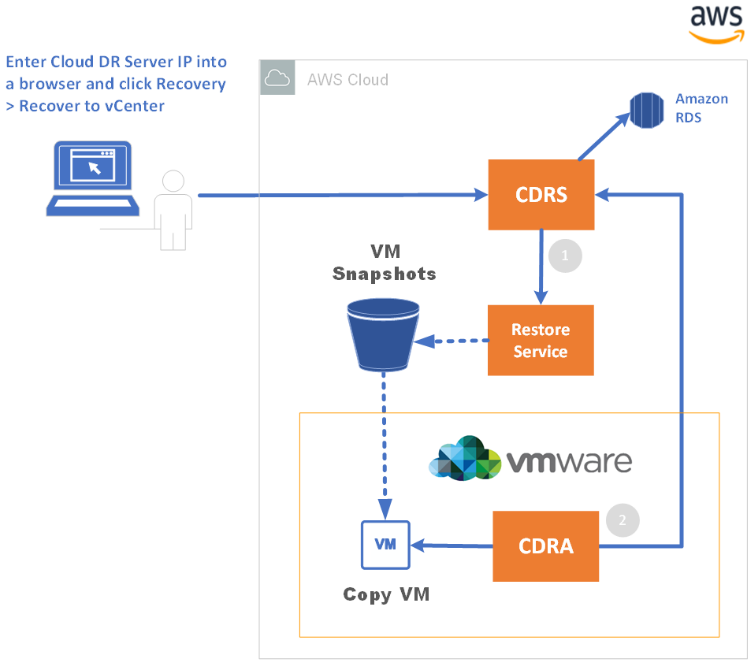
aws
Enter Cloud DR Server IP into a browser and click Recovery > Recover to vCenter
AWS Cloud
Amazon RDS
CDRS
1
VM Snapshots
Restore Service
vmware
VM
Copy VM
CDRA
2

Figure 4. Recover to vCenter or VMware Cloud on AWS architecture

The RecoverPoint for VMs cloud solution can be deployed as a standalone solution, or alongside the on-premises copies of an existing RecoverPoint for VMs vSphere solution. See the *RecoverPoint for Virtual Machines Scale and Performance Guide* for the maximum number of supported components, and the *RecoverPoint for Virtual Machines Release Notes* for limitations.

NOTE:

If you have already protected your VMs with on-premises copies, and you now want to add a cloud copy alongside your existing on-premises copies, refer to this *RecoverPoint for Virtual Machines Cloud Solutions Guide* for all concepts and procedures relevant only to your cloud copy. Refer to the *RecoverPoint for Virtual Machines Administrator's Guide* for all concepts and procedures related to your on-premises copies.

Cloud solution network architecture

The only network requirement for the RecoverPoint for VMs cloud solution is to ensure that TCP/IP port 443 open for communication between every vRPA cluster that protects a production VM, and AWS, and CDRS, as described in the *RecoverPoint for Virtual Machines Security Configuration Guide*. You should also ensure that all servers (ESXi, vRPA clusters, vCenters and CDRAs) have their clocks synchronized with NTP servers. The rest of the network architecture is created in AWS automatically, and on-demand.

NOTE:

This guide is specific to the RecoverPoint for VMs cloud solution. If you are using RecoverPoint for VMs to protect your VMs with both an on-premises copy and a cloud copy, see the *RecoverPoint for Virtual Machines Installation and Deployment Guide* for the networking requirements of the on-premises architecture.

When you Install and register a Cloud DR Server on page 25, the **Cloud DR Server** is installed in a **Public subnet**. For high-availability, a database cluster of 2 **Amazon RDS** database instances is created. Each **Amazon RDS** instance is created in a separate **Private subnet** and **Availability zone**.

When you Recovering VMs on page 43, one or more **Restore Services** are created (on-demand) in a separate **VPC** in the same **AWS Region** as the **Amazon S3 Bucket**. The **S3 Bucket** and **Restore Service VPC** can reside in a different **AWS Region** than, or the same **AWS Region** as the **VPC** with the CDRS and RDSs.

Once every 24 hours, a **Retention Service** is created in a separate **VPC** in the same **AWS Region** as the **Amazon S3 Bucket**. The **S3 Bucket** and **Restore/Retention Service VPC** can reside in a different **AWS Region** than, or the same **AWS Region** as the **VPC** with the CDRS and RDSs. See Managing the cloud copy retention policy on page 42 for more details.

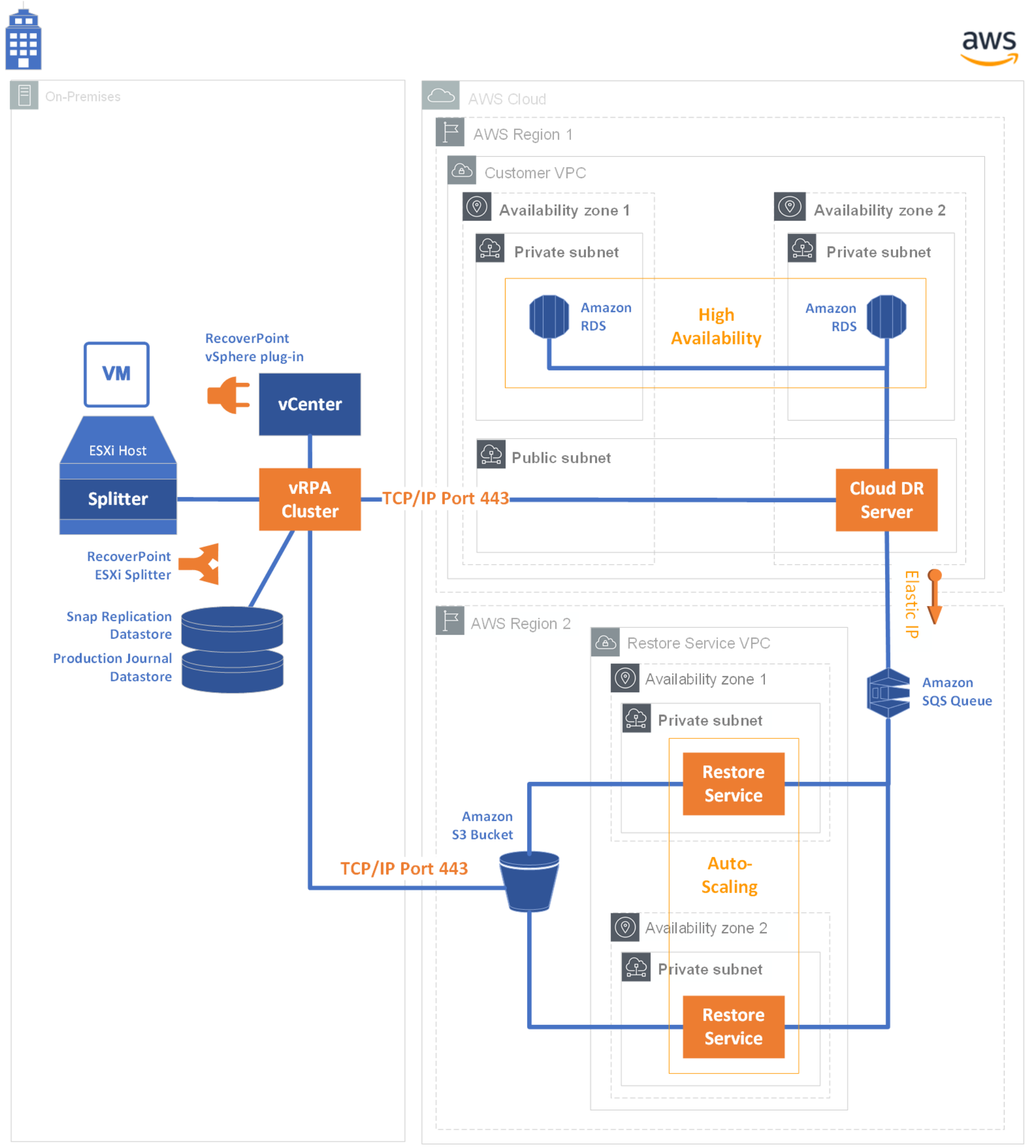

Figure 5. Cloud solution network architecture

Cloud solution with VMware Cloud on AWS

Copies are protected in AWS S3, and they are recovered in VMware Cloud on AWS.

VMware Cloud on AWS can be used on demand, when DR is needed. Since VMware Cloud on AWS is not needed for protection, the user can deploy a software-defined data center (SDDC) only when failover is required. The user connects the VMC to the Cloud DR solution by deploying a CDRA in VMC and connecting it to the CDRS. Then failover of VMs can begin.

Since the production site and DR site are both using VMware, failover to the VMware Cloud on AWS does not require launching an EC2 instance or converting VMDKs to AMIs.

For more information about VMware Cloud on AWS, read the VMware Cloud on AWS Technical Overview.

In this solution, the general recovery workflow is:

1. When recovery is needed, deploy an SDDC.
2. Deploy a CDRA in the SDDC and connect it to the CDRS.
3. From the VMC CDRA, when you define the recovery staging area, ensure that you enable direct failover to the VMC vCenter.
4. During recovery operations, select the VM that you want to recover, and then click **FAILOVER TO VCENTER**.
5. The recovery process fails over the AWS S3 copy to the VMware Cloud on AWS.

If failback is required, use vMotion to move the recovered VM from the VMC vCenter to the vCenter at the production site on premises.

2

Solution deployment

Deploy the RecoverPoint for VMs cloud solution in the provided sequence.

Before deploying the cloud solution, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

Topics:

Before you begin

Before deploying the RecoverPoint for VMs Cloud Solution, ensure you have complied with all of the requirements, and performed all of the pre-requisites in this section.

Credentials for cloud solution deployment

Before you begin cloud solution deployment, ensure that you have access to the usernames and passwords for the RecoverPoint for VMs cloud solution components.

Security best practices recommend that you change default passwords to something unique.

Table 1. Cloud solution component credentials

Cloud DR component	Notes
Cloud DR Server	Credentials are set during CDRS deployment, and can be changed through the CDRS interface, using the procedure to Change the CDRS admin user account password on page 95. In the RecoverPoint for VMs cloud solution, a CDRS **admin** user is created and the password for the CDRS **admin** user is defined during CDRS deployment. The CDRS can be deployed using the **RecoverPoint for VMs vSphere plug-in**, as described in Install and register a Cloud DR Server on page 25, or you can connect the RecoverPoint for VMs cloud solution to an existing CDRS (for example, a CDRS that is already being used to protect Avamar/Data Domain systems).
Amazon Web Services	Credentials are needed to establish a connection to the AWS account with the S3 bucket with the snapshots of your protected VMs. • AWS IAM user credentials are managed through the **AWS Management Console** > **IAM Console**. • AWS root user credentials are managed through the **AWS Management Console** > **Security Credentials Page**. In order to deploy a CDRS, you must have an IAM user with the minimum permissions described in Define the AWS IAM policy on page 17.
RecoverPoint vRPA Cluster	Credentials are defined during vRPA cluster installation, as described in the *RecoverPoint for Virtual Machines Installation and Deployment Guide*.
vCenter Server	Credentials are needed to establish a connection to the vCenter server that supports the production environment.
Cloud DR Add-on	Created during CDRA OVA deployment, the initial username/password is admin/admin.

Cloud solution requirements checklist

Use the following checklist to ensure that you have complied with all of the pre-requisites to deploying your RecoverPoint for VMs cloud solution. The RecoverPoint for VMs cloud solution requires RecoverPoint for VMs version 5.2.1 or later and Cloud DR Server version 18.4 or later.

> **NOTE: Installing the Cloud DR Server is the last step in a RecoverPoint for VMs cloud solution deployment. Ensure that you have complied with all the requirements in the following checklist, before you Install and register a Cloud DR Server on page 25.**

Table 2. Prerequisite checklist

Prerequisite	Requirement
Operational training	Familiarity with Cloud DR, RecoverPoint for VMs, AWS, and VMware concepts and terminology. See the *Glossary* at the end of this publication, for general definitions of the terms and concepts used throughout.
RecoverPoint for VMs	· Familiarity with the support and limitation statements for each RecoverPoint for VMs release. ○ See the *RecoverPoint for Virtual Machines Simple Support Matrix* (ESSM) for detailed support statements for third-party platforms and operating systems. ○ See the *RecoverPoint for Virtual Machines Release Notes* for the supported component versions and limitations. ○ See the *RecoverPoint for Virtual Machines Scale and Performance Guide* for the maximum number of supported components in a RecoverPoint for VMs system. · An on-premises installation of RecoverPoint for VMs 5.2.1 or later, with a network architecture and installed vRPA clusters as described in the *RecoverPoint for Virtual Machines Installation and Deployment Guide*. **NOTE: All vRPAs must be able to resolve `amazonaws.com` addresses, so all vRPA clusters will require an appropriate DNS.** · TCP/IP port 443 open for communication between every vRPA cluster that protects a production VM and AWS, and vRPA cluster that protects a production VM and CDRS, as described in the *RecoverPoint for Virtual Machines Security Configuration Guide*. · An on-cloud installation of Cloud DR Server 18.4 or later. · One public Amazon cloud account, S3 bucket, Cloud DR Server, and on-premises datastore (for snap replication), that are registered with every vRPA cluster that protects a production VM.
Clock synchronization via NTP	All servers (ESXi, vRPA clusters, vCenters and CDRAs) should have their clocks synchronized with NTP servers.
vSphere environment	· An on-premises vSphere environment, release 6.0U2 and later. · Network connectivity between on-premises environment and AWS. · Virtual machines compatible with Cloud DR. Comply with: ○ Virtual machine specifications for Cloud DR with AWS on page 16. ○ Supported operating systems for Cloud DR and AWS on page 17 contains information about supported VM operating systems. ○ http://docs.aws.amazon.com/vm-import/latest/userguide/vmie_prereqs.html for information about VM compatibility with AWS.
Amazon Web Services	· An AWS account. · AWS Marketplace terms for CentOS must be accepted before deploying the Cloud DR Server. Accept Amazon Web Services Marketplace terms on page 16 contains information about accepting AWS Marketplace terms for CentOS. (only for AWS users) · Network connectivity between on-premises environment and AWS. · An S3 bucket to be used as a Cloud DR target in one of the supported regions. See AWS/AWS GovCloud regions for CDRS deployment on page 16. · An AWS IAM policy, as described in Define the AWS IAM policy on page 17. · Enable downloads of Cloud DR logs from AWS on page 18.

Virtual machine specifications for Cloud DR with AWS

The following tables list the required specifications for the VMs used for Cloud Disaster Recovery components.

NOTE: To support recovery operations for production VMs, ensure that each VM has a unique identifier (UID).

Table 3. Cloud DR AWS components specifications

Component	Specification
CDRS instance type **NOTE: If m5.large is not available, m4.large will be deployed.**	m5.large
Temporary **Restore Service** instance type	c4.8xlarge
Temporary **Retention Service** instance type	m4.xlarge
RDS **NOTE: If db.t3.small is not available, db.t2.small will be deployed.**	db.t3.small

NOTE: For auto-scale handling, up to 100 Restore Service instances can be created for recovery, and up to 20 can be created for failback.

Table 4. Cloud DR Add-on VM specifications

Component	Specification
vCPU	4 (2x2)
RAM	4 GB
HDD	16 GB

In the RecoverPoint for VMs cloud solution, a CDRA is required only if you want to fail back from AWS to an on-premises vCenter or recover to vCenter or VMware Cloud on AWS.

Accept Amazon Web Services Marketplace terms

Before you deploy Cloud DR, you must accept the AWS Marketplace terms.

Steps

1. To connect to https://aws.amazon.com/marketplace/pp/B00O7WM7QW/, open a browser.
   The **CentOS 7 (x86_64) - with Updates HVM** page displays.
2. Click **Continue**.
   The **Sign In or Create an AWS Account** page appears.
3. Sign in using the AWS account.
   The **Launch on EC2** page appears.
4. Click **Manual Launch with EC2 Console, API, or CLI**.
5. Click **Accept Software Terms**.

   Clicking **Accept Software Terms** subscribes you to the CentOS software and indicates that you agree to the End User's License Agreement (EULA).

   The **Thank you for subscribing...** page appears. Verify that the subscription has been completed.

AWS/AWS GovCloud regions for CDRS deployment

The list of regions for CDRS deployment is subject to change. The most up-to-date list of supported regions where you can deploy the CDRS is maintained in the *Cloud DR Simple Support Matrix*, which is available at Dell EMC Online Support

The AWS Service Endpoints web page contains further information about AWS regions.

For AWS GovCloud regions, see Endpoints for the AWS GovCloud (US) Regions

Supported operating systems for Cloud DR and AWS

The list of operating systems for Cloud DR and AWS is subject to change. The most up-to-date list of supported operating systems is maintained in the *Cloud DR Simple Support Matrix*, which is available here:

https://www.dell.com/support/

Supported browsers and resolutions

The following browsers and resolutions are supported with Cloud DR.

Supported browsers

- Chrome - The latest version at the time of the release of Cloud DR.
- Firefox - The latest version at the time of the release of Cloud DR.

Supported desktop resolutions

- 1280 x 800
- 1366 x 768
- 1920 x 1080

Define the AWS IAM policy

To deploy CDRS, you must have an AWS Identity and Access Management (IAM) user with the following minimum permissions:

Create group policy in AWS

NOTE: You can also do this after you Install RecoverPoint for VMs on page 20, as you Register an AWS account on page 24.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "iam:GenerateCredentialReport",
                "iam:GenerateServiceLastAccessedDetails",
                "iam:Get*",
                "iam:List*",
                "iam:CreateRole",
                "iam:DeleteRole",
                "iam:AttachRolePolicy",
                "iam:DetachRolePolicy",
                "iam:DeleteRolePolicy",
                "iam:CreatePolicy",
                "iam:DeletePolicy",
                "iam:PutRolePolicy",
                "iam:CreateInstanceProfile",
                "iam:DeleteInstanceProfile",
                "iam:AddRoleToInstanceProfile",
                "iam:RemoveRoleFromInstanceProfile",
                "iam:PassRole",
                "iam:SimulateCustomPolicy",
                "iam:SimulatePrincipalPolicy"
            ],
            "Effect": "Allow",
            "Resource": "*"
        },
        {
            "Action": "ec2:*",
            "Effect": "Allow",
            "Resource": "*"
        },
        {
        "Effect": "Allow",
        "Action": "cloudwatch:*",
        "Resource": "*"
        },
        {
        "Effect": "Allow",
        "Action": "s3:*",
        "Resource": "*"
        },
        {
        "Effect": "Allow",
        "Action": [
            "ses:SendEmail",
            "ses:SendRawEmail",
            "ses:Verify*",
            "ses:ListVerifiedEmailAddresses"
            ],
        "Resource": "*"
        },
        {
        "Effect": "Allow",
        "Action": "cloudformation:*",
        "Resource": "*"
        },
        {
        "Effect": "Allow",
        "Action": "rds:*",
        "Resource": "*"
        },
        {
        "Effect": "Allow",
        "Action": "sqs:*",
        "Resource": "*"
        }
    ]
}
```

To create a policy in AWS using this IAM policy:

1. At the AWS **Identity and Access Management Console** (https://console.aws.amazon.com/iam/home?#/home), click **Policies**.
2. Click **Create policy**.
3. Click **Select** for **Create Your Own Policy**.
4. Enter a name and description for the policy.
5. In **Policy Document**, paste the above IAM policy.
6. Click **Create Policy**.

Enable downloads of Cloud DR logs from AWS

To enable downloads of Cloud DR logs from the AWS, log in to the AWS console.

About this task

This procedure enables public access to the logs folder for downloading the log files. When public access is enabled, the CDRS dashboard displays a recommendation to remove public access. After retrieving Cloud DR logs, ensure that you remove public access.

Steps

1. Log into the **S3 Dashboard** of the AWS Console (https://console.aws.amazon.com/s3/).
2. Select the S3 bucket that contains the logs.
3. Click the **Permissions** tab.
4. Click **Bucket Policy**.
5. Enter the following text in the **Bucket policy editor**:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AddPerm",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::bucket-name/logs/*"
    }
```

```
        ]
    }
```

where *bucket-name* is the name of the bucket that contains the logs.

6. Click **Save**.

Results

The log files are now accessible via the link provided in the RecoverPoint for VMs vSphere plugin > Administration > Log Collection tab.

Create your license files

When a RecoverPoint for Virtual Machines sales order is approved, a License Authorization Code is automatically sent to the email addresses provided during order entry. The License Authorization Code contains your entitlements. You must activate each entitlement and then save it as a license file before it can be entered into the RecoverPoint for VMs vSphere plugin.

About this task

> **NOTE: If you have an already licensed RecoverPoint for VMs system, you can skip this section. The RecoverPoint for VMs cloud solution uses your existing license.**

For a detailed description of how RecoverPoint for VMs is licensed, see RecoverPoint for VMs licensing on page 104.

Steps

1. Access the entitlements on support.emc.com:
   - If you have the **License Authorization Code** email, open it and click the **Click here** link. Clicking the link automatically accesses **Powerlink Licensing** on the support site, and searches for all entitlements associated with the License Authorization Code.
   - If you do not have the **License Authorization Code** email but you do have the LACs or sales order numbers, log into support.emc.com, and:

   a. Select **Support** > **Service Center** from the main menu.
   b. Select **Get and Manage Licenses**.
   c. Select **RecoverPoint for Virtual Machines**.
   d. Type the customer's License Authorization Code and click **Activate** to search for all inactive entitlements that are associated with a customer's profile, or access all of the features of the Licensing site by clicking **Manage Entitlements**. Whichever option you chose, the **Search Entitlements to Activate** screen is displayed.
2. Activate the entitlements and download the license files:
   a. In the **Search Entitlements to Activate** screen, select an entitlement to activate. Each entitlement must be selected and activated separately.
   b. Click **Start Activation Process**.
   c. In the **Search Machines** dialog box, click **Add a Machine**.
   d. In the **Add Machine** dialog box, type a new machine name, and click **Save**. A unique machine name must be specified for each entitlement.

      A machine name is like a folder. It is used to group items together logically.
   e. In the **Register** screen, verify the machine name, and click **Next**.
   f. In the **Activate** screen, type the **Locking ID**, and click **Next.**

      The Locking ID is the field that is displayed in the Machine Information column. Its value is the entity that the license is enforced for, namely, the vCenter Server ID. To find the vCenter Server ID, type **`https://<vCenterServerIP>/mob`** into the browser address bar or SSH client, and type the credentials to log in to the vCenter Server. Select **Content** > **About**. The instanceUuid is the vCenter Server (Locking) ID that the license is enforced for.
   g. In the **Confirm** screen, type the email addresses of the recipients of the license file in the **Email to** field of the **Additional Email Options** section, and click **Finish**. Separate multiple email addresses with commas.
   h. In the **Complete** screen, click **Save to File** to download the license file and save the file locally. The resulting license file has a `*.lic` extension and is in plain text format (can be opened in any text editor).
   i. Repeat this procedure for all inactive entitlements in each License Authorization Code email.

Results

The entitlements are converted to license files.

Next steps

Transfer the license files to the computer from which you will be running RecoverPoint for VMs.

Install RecoverPoint for VMs

To start protecting your VMs, you must first install and deploy RecoverPoint for VMs.

> **NOTE:**
>
> **The rest of this *RecoverPoint for Virtual Machines Cloud Solution Guide* assumes that you have RecoverPoint for VMs 5.2.1 or later installed and configured as described in the *RecoverPoint for Virtual Machines Installation and Deployment Guide*.**

Access the RecoverPoint for VMs vSphere plugin

There are two ways in which you can access the RecoverPoint for VMs plugin in the vSphere Web Client.

Prerequisites

Connect to the vSphere Web Client of your production site.

Steps

1. Click the **RecoverPoint for VMs** menu item in your **vSphere Web Client** > **Navigator**.
2. Click the **RecoverPoint for VMs** icon in your **vSphere Web Client** > **Navigator** > **Inventories**.

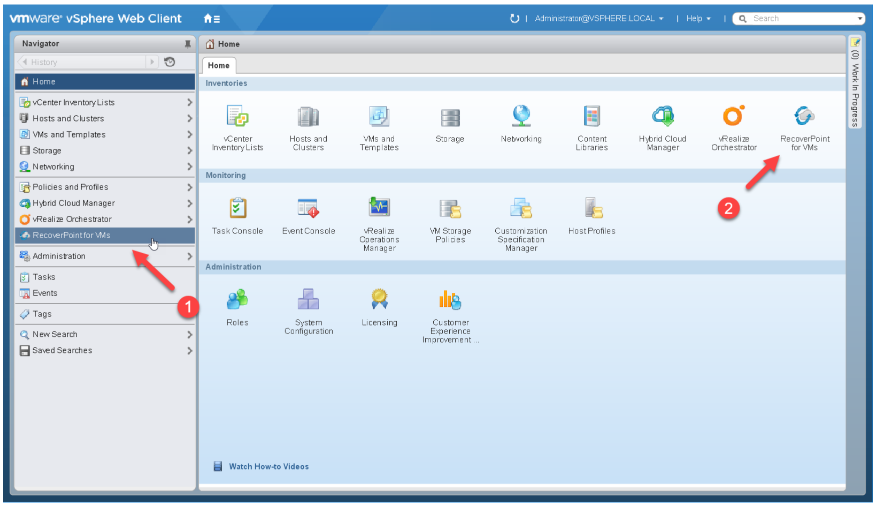

Results

The RecoverPoint for Virtual Machines **Dashboard** is displayed.

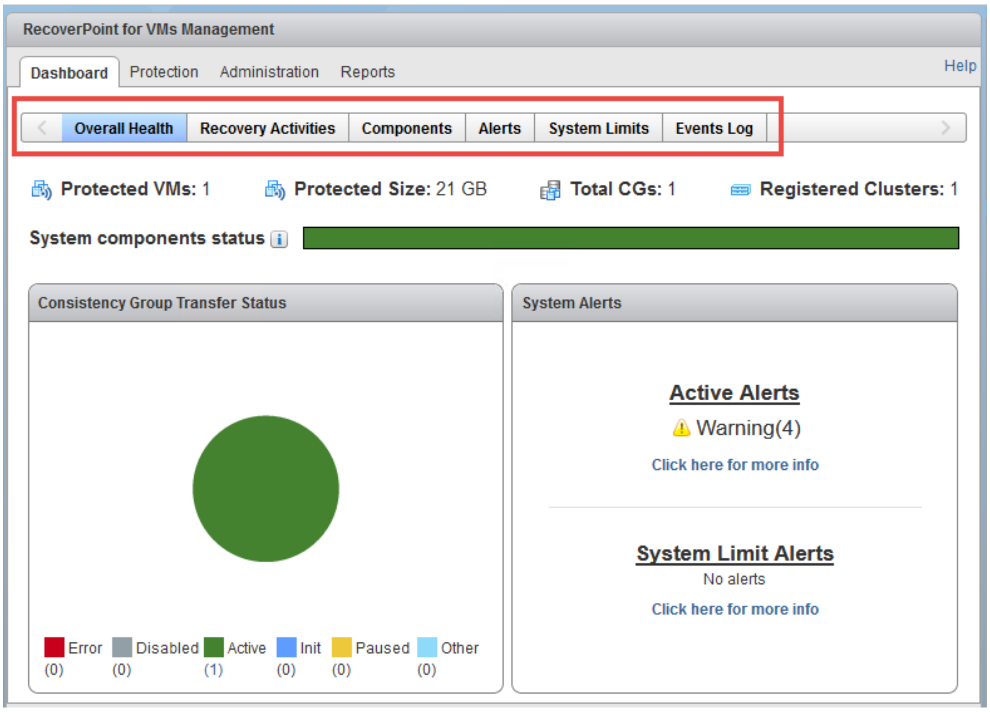

Protected VMs: 1 **Protected Size:** 21 GB **Total CGs:** 1 **Registered Clusters:** 1

System components status

Active Alerts

⚠ Warning(4)

Click here for more info

System Limit Alerts

No alerts

Click here for more info

To monitor the status of your vRPA clusters, click the **Dashboard** > **Components** tab and ensure that a green checkbox appears next to each vRPA cluster and that the vRPA cluster **Status** column contains an `OK`.

To monitor the environment, click the **Overall Health**, **Components**, **Alerts**,**System Limits**, and **Events Log** sub-tabs of the system **Dashboard**.

License and register RecoverPoint for VMs

The **Getting Started Wizard** will guide you through the process of entering a license file, registering the product, and enabling system support.

Prerequisites

- To transfer system reports and alerts using SMTP or Secure Remote Services, ensure that port 25 is open and available for SMTP traffic.
- To transfer system reports and alerts using FTPS, ensure that ports 990 and 989 are open and available for FTPS traffic.

About this task

> **NOTE: If you have an existing RecoverPoint for VMs system that is already licensed and registered, you can skip this section. The RecoverPoint for VMs cloud solution uses your existing license files and support settings.**

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, click **Administration** > **vCenter Servers** > **Licensing**.
2. Click **Add...** under the **Registered Licenses** table.
   The **Getting Started Wizard** is displayed.
3. In the **Welcome** screen, click **Next**.
4. In the **Licensing** screen, click **Browse...** to locate and select the license file (*.lic extension). Click **Next**.
5. In the **Support** screen, to provide communication between the RecoverPoint for VMs system and the System Reports database, select **Enable pre-emptive support for RecoverPoint for VMs**.
   a. Define the transfer method:
      - To transfer system notifications through an SMTP server, in the **Transfer Method** section, select **SMTP**. In the **SMTP server address** field, specify the IP address or DNS name of the dedicated SMTP server, in IPv4 format. In the **Sender address** field, specify the email address to send the system notifications from.
      - To transfer system notifications through the FTPS server, in the **Transfer Method** section, select **FTPS**.

- To transfer system notifications through the Secure Remote Services gateway, in the **Transfer Method** section, select **ESRS**. In the **ESRS gateway IP address** field, specify the IP address of the Secure Remote Services gateway in IPv4 format..

b. Click **Test Connectivity**. Wait 10 minutes. Then, click **Dashboard** > **Events Log** and look for event 1020: "`Failed to send system report`".

- If this event does not appear in the **Events Log**, the system notifications mechanism is correctly configured.
- If you do receive event 1020: "`Failed to send system report`", check whether there is an issue with the selected method of transfer. If a problem exists, fix it, re-configure support, and click **Test Connectivity** again. If the problem persists, contact Customer Support.

c. Click **Next**.

6. In the **Registration** screen:

a. Register your RecoverPoint for VMs system, at the current vRPA cluster:

- **Company name**: The name of your company, as it appears on your sales order.
- **Connect home method**: The method that is used to send configuration reports and alerts to Dell EMC. The connect home method allows Dell EMC to pro-actively address issues within the RecoverPoint for VMs environment, should they arise.
- **Connect in method** The method that is used to allow remote connectivity to the RecoverPoint environment. Enabling this feature is recommended as it enables secure access to the RecoverPoint for VMs environment to gather logs and resolve issues as quickly as possible. If you already have a Secure Remote Services Gateway servicing other products, use the Secure Remote Services Config Tool to add the RecoverPoint devices to the list of Secure Remote Services monitored environments. When the device is added, click the request **update** button to send the new device information to EMC and contact the local Customer Engineer to approve the update. Refer to the *Secure Remote Services Gateway Operation Guide* for further instructions on Config Tool usage. If you do not have a Gateway at the site, contact the Account Manager to find out more about the benefits of Secure Remote Services.
- **License type**: Displays the type of the license that has been registered in RecoverPoint for VMs. Ensure the displayed license type is **RecoverPoint for VMs**.
- **Location**: The city, state, and country where your company is located.
- **Sales order number** If you don't have your sales order, your Customer Engineer can provide it.
- **Site (party) ID**: The unique ID of the customer site. This value is automatically retrieved from the registered license file and can only be modified by Customer Service.

b. If your company does not have outside connectivity, click **Export to CSV** to export the registration information to a CSV file.

c. Enter the email address to which a verification email should be sent when the registration information is updated in the Install Base in the **Send verfiication email to** field.

ⓘ **NOTE: Skip this step if your company does not have outside connectivity.**

d. Click **Next**.

7. In the **Ready to complete** screen, verify that the information is correct, and click **Finish**.

Results

If your company has outside connectivity, a service request is opened and sends an email to the specified verification email address from Customer Support to verify that the registration details were updated successfully in the Install Base for every vRPA cluster in the RecoverPoint for VMs system.

Next steps

If your company does not have outside connectivity, use the exported CSV file to register by email or phone, as described in Register RecoverPoint by email or phone on page 105.

Register cloud services and install CDRS

Before you start protecting VMs, use the **Cloud Services Tab** to register your AWS account and bucket, and install and register the Cloud DR Server.

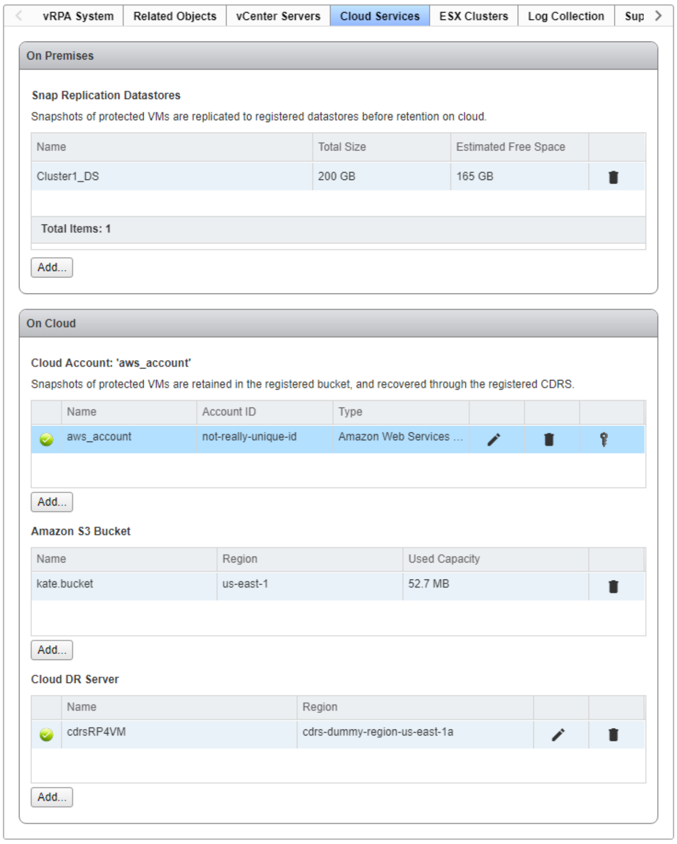

On Premises

Snap Replication Datastores

Snapshots of protected VMs are replicated to registered datastores before retention on cloud.

Name	Total Size	Estimated Free Space	
Cluster1_DS	200 GB	165 GB	🗑

Total Items: 1

Add...

On Cloud

Cloud Account: 'aws_account'

Snapshots of protected VMs are retained in the registered bucket, and recovered through the registered CDRS.

	Name	Account ID	Type			
●	aws_account	not-really-unique-id	Amazon Web Services ...	✎	🗑	🔑

Add...

Amazon S3 Bucket

Name	Region	Used Capacity	
kate.bucket	us-east-1	52.7 MB	🗑

Add...

Cloud DR Server

	Name	Region		
●	cdrsRP4VM	cdrs-dummy-region-us-east-1a	✎	🗑

Add...

Figure 6. Cloud Services tab

Register on-premises datastores for snap replication

Register one or more on-premises datastores for snap replication. Your production VM snapshots will be replicated to the datastore(s), before they are replicated to your S3 bucket. Register on-premises datastore(s) with every vRPA cluster that will protect a production VM.

Prerequisites

- Registered datastores should be shared (exposed to all ESXi servers hosting vRPA clusters).
- See the *RecoverPoint for Virtual Machines Scale and Performance Guide* for the amount of free space that should be available in the registered datastore(s).
- Registered datastore(s) must have performance capabilities equal to, or greater than, the datastore with the highest performance that is used for your production VMs.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select a vRPA cluster.
3. Select the **Cloud Services** tab.
4. Click the **Add...** button under the **Snap Replication Datastores** table.
   The **Register Snap Replication Datastores** dialog box is displayed.
5. In the **Register Snap Replication Datastores** dialog box:
   a. Ensure the correct vCenter Server is selected
   b. Select one or more datastores in which to store the snapshots of your production VMs.

6. Click **Register**.
7. Repeat this procedure for every vRPA cluster that will protect a production VM.

Results

The on-premises datastore(s) for snap replication are registered at the specified vRPA cluster(s).

Register an AWS account

Register the AWS account that you want to use with every vRPA cluster that will protect a production VM.

Prerequisites

- Ensure you have an on-premises vSphere environment, release 6.0U2 or later.
- Ensure you have an existing Amazon Web Services (AWS) public cloud account.
- Ensure you have an AWS access key and secret access key.
- Ensure TCP/IP port 443 is open for communication between the vRPA clusters and AWS.
- Ensure all vRPA clusters have a configured DNS server that enables vRPAs to resolve amazonaws.com addresses. To configure the vRPA cluster DNS server, refer to the *RecoverPoint for Virtual Machines Installation and Deployment Guide*.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster.
3. Select the **Cloud Services** tab.
4. Click **Add...** button under the **Cloud Account** table.
5. In the **Register Cloud Account** dialog box:
    a. Enter the account name, the AWS access key and the AWS secret access key.
    b. Click **Register**.
    c. Ensure the registered AWS account contains an IAM policy with the displayed permissions.

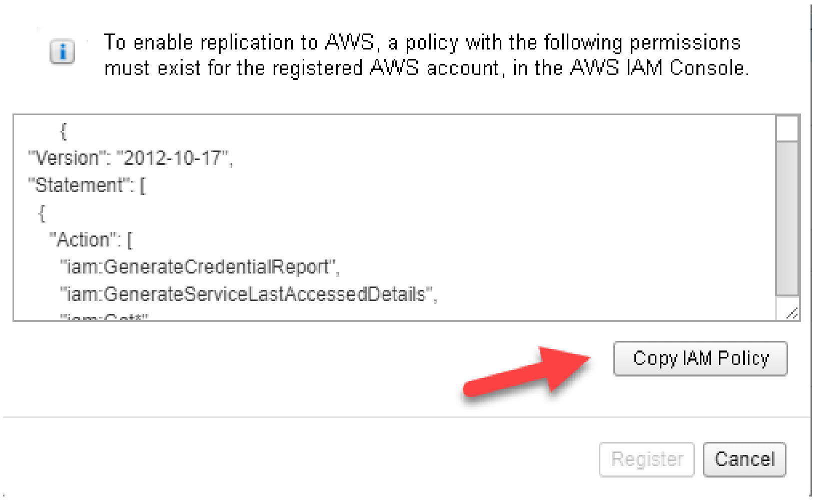

To enable replication to AWS, a policy with the following permissions must exist for the registered AWS account, in the AWS IAM Console.

```
    {
"Version": "2012-10-17",
"Statement": [
 {
   "Action": [
     "iam:GenerateCredentialReport",
     "iam:GenerateServiceLastAccessedDetails",
     "iam:Get*"
```

If required, click **Copy IAM Policy** to copy the required IAM permissions in JSON format and paste them into the AWS IAM Console, as described in Define the AWS IAM policy on page 17.

6. Repeat this procedure for every vRPA cluster that will protect a production VM.

Results

The cloud account is registered at the specified vRPA cluster.

Register an S3 bucket

After registering your cloud account, register the Amazon Simple Storage Service (S3) bucket in which to store the VM snapshots, with every vRPA cluster that will protect a production VM.

Prerequisites

- Ensure TCP/IP port 443 is open for communication between the vRPA clusters and AWS.
- Ensure you have registered your AWS account and granted the required IAM permissions, as described in Register an AWS account on page 24.
- Ensure you have an existing **Amazon S3 bucket** in the registered AWS account.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster.
3. Select the **Cloud Services** tab.
4. Ensure your AWS account is selected in the **Cloud Account**.
5. Click the **Add...** button under the **Amazon S3 Bucket** table.
   The **Register Amazon S3 Bucket** dialog box is displayed.
6. In the **Register Amazon S3 Bucket** dialog box, select the S3 bucket of the AWS account in which you want to store the snapshots of your production VMs.
7. Click **Register**.
8. Repeat this procedure for every vRPA cluster that will protect a production VM.

Results

The bucket is registered at the specified vRPA cluster(s).

Install and register a Cloud DR Server

After registering your AWS account and bucket, install and register the Cloud DR Server (CDRS) that you will use to recover your protected VMs, with every vRPA cluster that will protect a production VM.

Prerequisites

- Ensure TCP/IP port 443 is open for communication between every vRPA cluster that protects a production VM, and AWS, and CDRS.
- Ensure you have completed to Register an AWS account on page 24.
- Ensure you have completed to Register an S3 bucket on page 25.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster.
3. Select the **Cloud Services** tab.
4. Click the **Add...** button under the **Cloud DR Server** table.
   The **Register Cloud DR Server** dialog box is displayed.
5. In the **Register Cloud DR Server** dialog box:
   - If no CDRS has previously been installed in the cloud account:

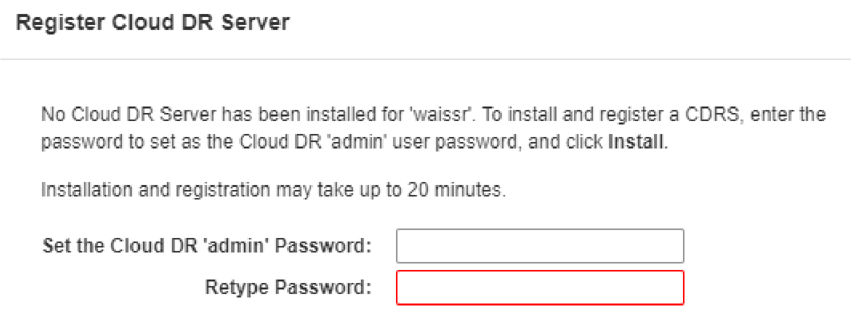

Register Cloud DR Server

No Cloud DR Server has been installed for 'waissr'. To install and register a CDRS, enter the password to set as the Cloud DR 'admin' user password, and click **Install**.

Installation and registration may take up to 20 minutes.

Set the Cloud DR 'admin' Password:

Retype Password:

a. Set the **Password** for the Cloud DR 'admin' user. The Cloud DR 'admin' user password must be at least 8 characters in length, and it must contain at least one uppercase (A-Z) and one lowercase (a-z) character, and at least one numeric (0-9) or special (#@!) character.

   > **NOTE: Save this password. You will have to enter the password that you specify here when you register the CDRS at every subsequent vRPA cluster.**

b. Retype the password for verification.

c. Click **Install**.

- If a CDRS has already been installed in the cloud account:

  a. Enter the password of the Cloud DR **admin** user.

     > **NOTE: This is the password that was set during initial installation of the CDRS.**

  b. Click **Register**.

6. Repeat this procedure for every vRPA cluster that will protect a production VM.

Results

If no CDRS had previously been installed in the cloud account:

> **NOTE: CDRS deployment may take up to 30 minutes.**

- In the AWS **VPC**:
  - The CDRS is deployed in a **Public subnet**. The **m4.large** instance type is used for the CDRS instance. An elastic IP address is automatically assigned to the CDRS instance. You cannot change this IP address. The progress bar in the **Cloud DR Server** area of the **Cloud Services** tab displays the progress of CDRS installation.

    > **NOTE: To reduce deployment costs, you may want to purchase reserved instances from AWS; otherwise an on-demand instance is used.**
  - For high-availability, two **Amazon RDS** are created in their own **Availability zone** and **Private subnet**.
- The CDRS is registered with the specified vRPA cluster.
- The specified password is set as the Cloud DR **admin** user password, and will be required to access the CDRS, when Recovering VMs on page 43.

If a CDRS had previously been installed in the cloud account, the CDRS is registered at the specified vRPA cluster(s).

> **NOTE:**
>
> **If an error occurs during CDRS deployment, Uninstall the cloud solution on page 101 to delete any previously installed cloud resources, and perform this procedure again, from step 5.**

Next steps

Best practice is to Define the email address for CDRS password recovery on page 96 as soon as possible after CDRS deployment. To change the password of the CDRS **admin** user after it has been defined, follow the instructions in Change the CDRS admin user account password on page 95.

3

Protecting VMs

In RecoverPoint for VMs, consistency groups are used to protect virtual machines and replicate virtual machine application data to a consistent point in time. A consistency group is a logical entity that constitutes a container for virtual machines and all of their copies.

Consistency groups can protect many VMs. If this is the first time you are using RecoverPoint for VMs, protect your virtual machines by creating new consistency groups for them, or by adding them to an existing consistency group. If you already have RecoverPoint for VMs consistency groups, you can create a new copy to protect your production VMs, alongside your existing copy.

Before protecting VMs, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

Topics:

Protect a virtual machine in a consistency group

The RecoverPoint for VMs **Protect VMs Wizard** will guide you through the process of protecting your production VMs.

Prerequisites

Ensure you have completed Solution deployment on page 14.

Steps

1. Connect to the vSphere Web Client of your production site.
2. Select **VMs and Templates** view.
3. Power on the virtual machine that you want to protect.
4. Right-click on the virtual machine and select **All RecoverPoint for Virtual Machines Actions** > **Protect**.

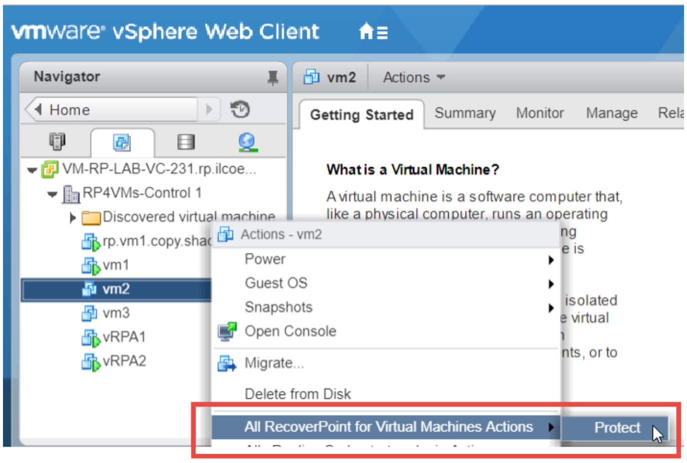
vmware vSphere Web Client
Navigator
Home
VM-RP-LAB-VC-231.rp.ilcoe...
RP4VMs-Control 1
Discovered virtual machine
rp.vm1.copy.shac
vm1
vm2
vm3
vRPA1
vRPA2
vm2
Actions
Getting Started
Summary
Monitor
Manage
Rela
What is a Virtual Machine?
A virtual machine is a software computer that, like a physical computer, runs an operating
Actions - vm2
Power
Guest OS
Snapshots
Open Console
Migrate...
Delete from Disk
All RecoverPoint for Virtual Machines Actions
Protect

ⓘ **NOTE: Protecting a virtual machine with fault tolerance enabled is not supported.**

The **Protect VMs Wizard** is displayed.

The minimum input required in each screen of the **Protect VMs Wizard** is indicated in red, in the following screenshots.

5. In the **Select VM protection method** screen:

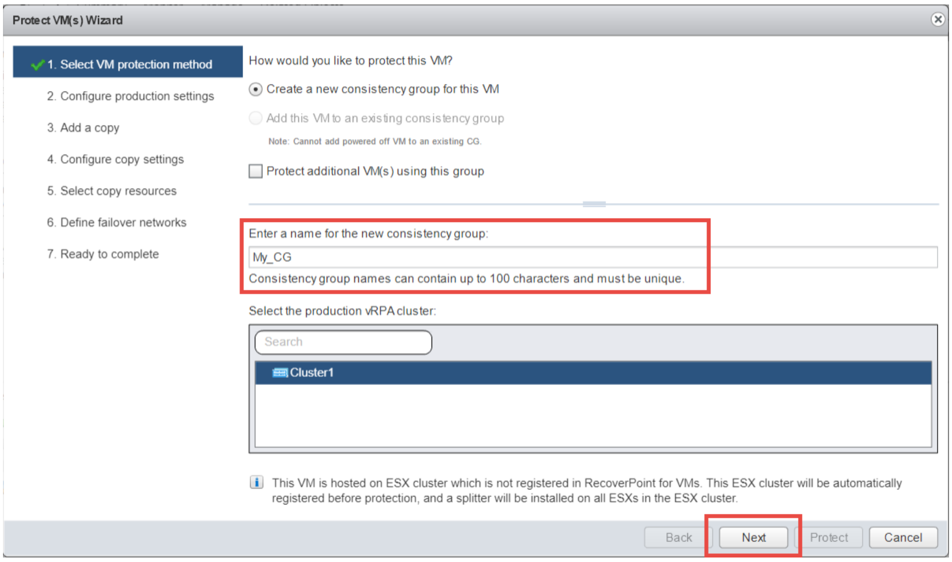
Protect VM(s) Wizard
✓ 1. Select VM protection method
2. Configure production settings
3. Add a copy
4. Configure copy settings
5. Select copy resources
6. Define failover networks
7. Ready to complete
Back
Next
Protect
Cancel

How would you like to protect this VM?

● Create a new consistency group for this VM

○ Add this VM to an existing consistency group

Note: Cannot add powered off VM to an existing CG.

□ Protect additional VM(s) using this group

Enter a name for the new consistency group:

My_CG

Consistency group names can contain up to 100 characters and must be unique.

Select the production vRPA cluster:

Search

Cluster1

This VM is hosted on ESX cluster which is not registered in RecoverPoint for VMs. This ESX cluster will be automatically registered before protection, and a splitter will be installed on all ESXs in the ESX cluster.

- **Create a new consistency group for this VM**. Type a descriptive name for the new consistency group. Best practice is to use the VM or application name as your consistency group name. Ensure the production vRPA cluster is selected. If you want to add additional virtual machines to protect, mark the **Protect additional VM(s) using this group** checkbox, select the additional virtual machines to protect in the consistency group, and click **Add**. If you do not want to add additional virtual machines, click **Next**.
- **Add this VM to an existing consistency group**. Select an existing consistency group. If you want to add additional virtual machines to protect, mark the **Protect additional VM(s) using this group** checkbox select the additional virtual machines to protect in the consistency group, and click **Add**. If you do not want to add additional virtual machines, click **Next**.

6. In the **Configure production settings** screen:

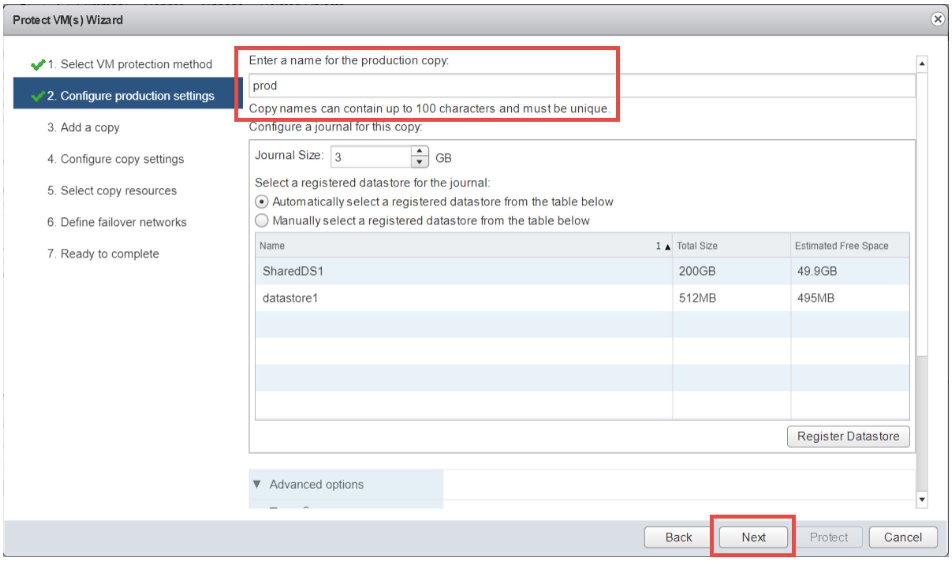

Enter a name for the production copy:

prod

Copy names can contain up to 100 characters and must be unique.

Configure a journal for this copy:

Journal Size: 3 GB

Select a registered datastore for the journal:

● Automatically select a registered datastore from the table below

○ Manually select a registered datastore from the table below

Name	Total Size	Estimated Free Space
SharedDS1	200GB	49.9GB
datastore1	512MB	495MB

a. Enter a name for the production copy. Best practice is to differentiate the production copy name from the replica copy name (for example, use "Production" or the production site location).

b. If you chose to create a new consistency group in the previous step (not relevant if adding a VM to an existing group):

   - Accept or define the minimum **Journal Size** for the production copy. The default minimum journal size (3GB) is sufficient to enable RecoverPoint for VMs cloud services.
   - Optionally, select a specific datastore to use for the production journal. By default, RecoverPoint automatically registers up to 15 datastores for the production journal and automatically selects the datastore with the most free space.

     NOTE: RecoverPoint for VMs will attempt to create the journal on the selected datastore. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.

     - If you want to select a specific datastore from the registered datastores, select **Manually select a registered datastore from the table below** and select it in the table.
     - If you want to select a datastore that hasn't been registered, click **Register Datastore**. Select the datastore and click **Register**. Ensure the required datastore is selected in the table.

c. Expand and configure the **Advanced options** per virtual machine:

- **VMDK(s)**: Displays the number of included VMDKs at the relevant production copy, and their total size. Uncheck a VMDK to exclude it from replication.

  > **NOTE: Excluded VMDKs are not displayed in the list of snapshots for selection, when Recovering VMs on page 43. For Linux-based production VMs: If in the file systems table (in `/etc/fstab`), a VM with an excluded VMDK is configured to be automatically mounted when the operating system boots up , add a *nofail* option in the file systems table or the cloud copy of the VM will fail to boot when Recovering VMs on page 43.**

- **Protection policy**: Default = `Enabled`. Selecting **Automatically protect new VMDKs** ensures all new VMDKs are automatically protected.

> **NOTE:**
>
> **The Disk provisioning, Hardware changes, and MAC address replication to local copy VMs on the same vCenter settings are not relevant in the RecoverPoint for VMs cloud solution. In this solution, every VM snapshot is saved together with its OVF. When Recovering VMs on page 43 from a VM copy snapshot, the recovered VM will have the same hardware settings that the protected VM had, at the point in time that the snapshot was replicated.**

d. Click **Next**.

7. In the **Add a copy** screen, define the copy.

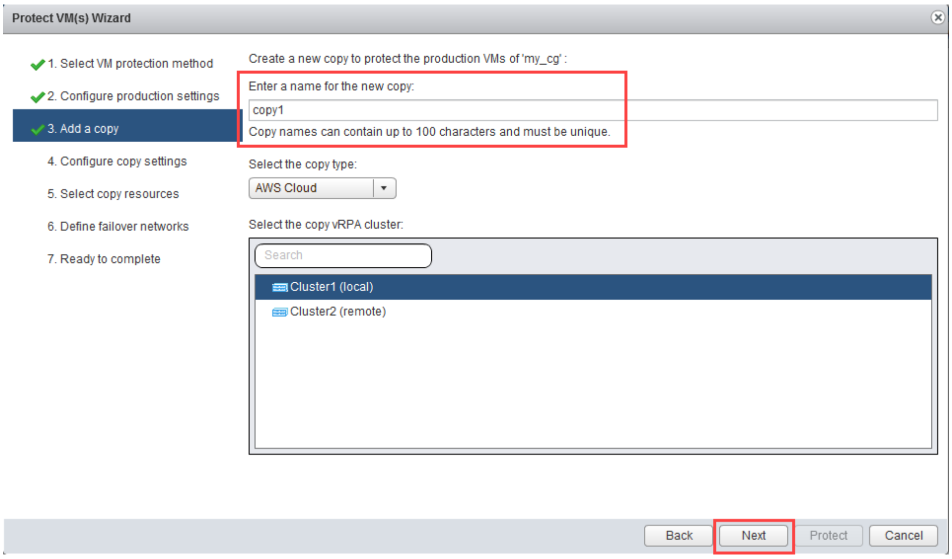

Protect VM(s) Wizard

✓ 1. Select VM protection method
✓ 2. Configure production settings
✓ 3. Add a copy
4. Configure copy settings
5. Select copy resources
6. Define failover networks
7. Ready to complete

Create a new copy to protect the production VMs of 'my_cg' :

Enter a name for the new copy: copy1

Copy names can contain up to 100 characters and must be unique.

Select the copy type: AWS Cloud

Select the copy vRPA cluster:

a. Enter a name for the copy. Best practice is differentiate the replica copy name from the production copy name (for example: "AWS copy").

b. From the **Select copy type** field, select **AWS Cloud**.

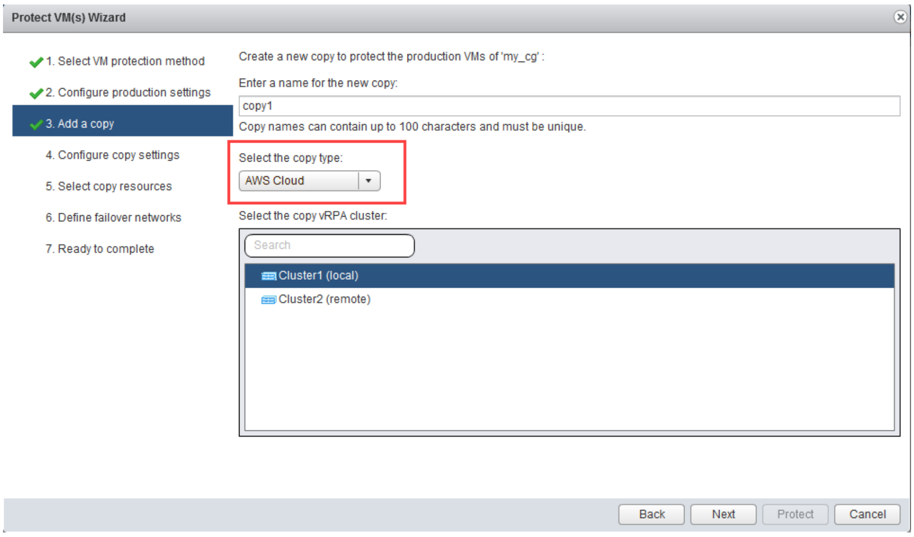

Protect VM(s) Wizard

✓ 1. Select VM protection method
✓ 2. Configure production settings
✓ 3. Add a copy
4. Configure copy settings
5. Select copy resources
6. Define failover networks
7. Ready to complete

Create a new copy to protect the production VMs of 'my_cg' :

Enter a name for the new copy: copy1

Copy names can contain up to 100 characters and must be unique.

Select the copy type: AWS Cloud

Select the copy vRPA cluster:

c. Ensure the vRPA cluster that will manage the group data is selected.

Cloud copies can only be managed by a local vRPA cluster.

d. Click **Next**

8. In the **Configure copy settings** screen:

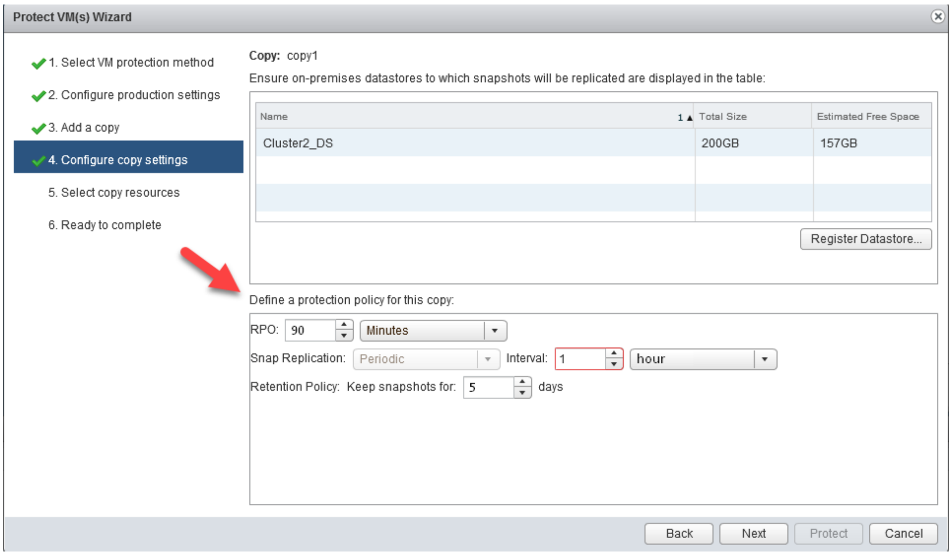

Copy: copy1

Ensure on-premises datastores to which snapshots will be replicated are displayed in the table:

Name	Total Size	Estimated Free Space
Cluster2_DS	200GB	157GB

Define a protection policy for this copy:

RPO: 90 Minutes

Snap Replication: Periodic Interval: 1 hour

Retention Policy: Keep snapshots for: 5 days

a. The datastores you registered during Solution deployment on page 14 are displayed in the table, and will be used for on-premises snap replication. In the rare case that the datastore you want to use for snap replication is not displayed in the table, click **Register Datastore...** and register it now.

b. Define the copy protection policy:

- **RPO**: Default = *90 minutes*

  The maximum data lag that is required between the production copy and the latest snapshot uploaded to the S3 bucket. If the specified RPO is exceeded, a warning is displayed in the The RecoverPoint for VMs Dashboard on page 33. RPO can be defined in `Minutes`, `Hours`, or `Days`.

  NOTE: Specify an RPO value that is higher than the specified snap replication Interval. Best practice is to specify an RPO value that is 1.5 times the specified snap replication Interval. For example, if you require an RPO of *1 hour*, specify a snap replication interval of *90 minutes*.

- **Snap Replication**: Default = *Periodic* at *1 hour* intervals

  Sets the periodic **Interval** between snapshots in `Minutes`, `Hours`, or `Days`. The minimum interval value is *15 minutes* and the maximum interval value is *7 days*. A new snapshot starts after the specified interval has passed since the previous snapshot was started. If the time interval has passed and the previous snapshot is incomplete, the next snapshot will start as soon as the previous one has completed.

  NOTE: Specify a snap replication Interval value that is lower than the specified RPO value. Best practice is to specify an RPO value that is 1.5 times the specified snap replication Interval. For example, if you require an RPO of *1 hour*, specify a snap replication interval of *90 minutes*.

- **Retention Policy**: Default = *5 days*

  Defines the period of time that snapshots will be retained in the cloud. Snapshots can be retained for a minimum of `1` day, and a maximum of `90` days. Once every 24 hours, a temporary **Retention Service** EC2 instance is launched in AWS to consolidate the snapshots of every copy whose retention policy has expired.

c. Click **Next**.

9. The **Select copy resources** screen displays the AWS account and S3 bucket to which the production VMs will be replicated. Click **Next**.

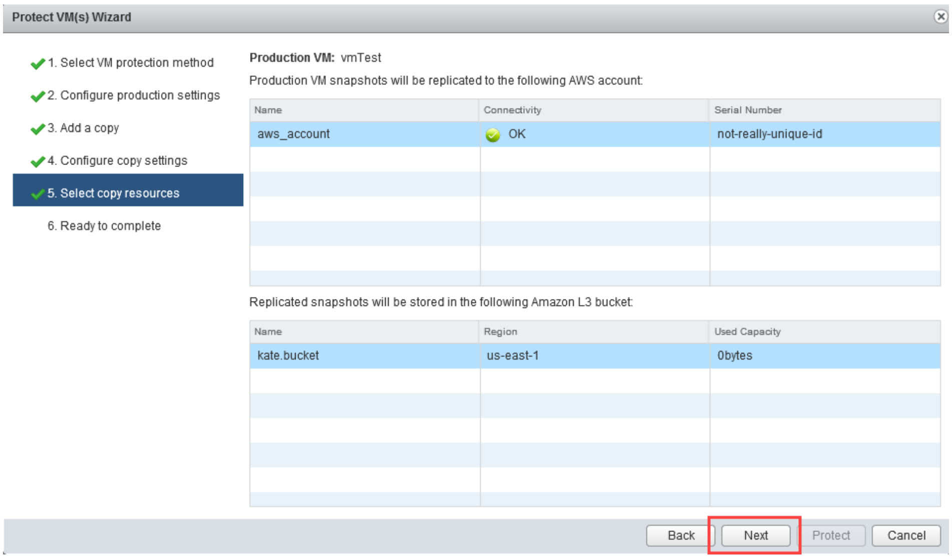

Protect VM(s) Wizard

- ✓ 1. Select VM protection method
- ✓ 2. Configure production settings
- ✓ 3. Add a copy
- ✓ 4. Configure copy settings
- ✓ 5. Select copy resources
- 6. Ready to complete

Production VM: vmTest

Production VM snapshots will be replicated to the following AWS account:

Name	Connectivity	Serial Number
aws_account	● OK	not-really-unique-id

Replicated snapshots will be stored in the following Amazon L3 bucket:

Name	Region	Used Capacity
kate.bucket	us-east-1	0bytes

10. In the **Ready to complete** screen:

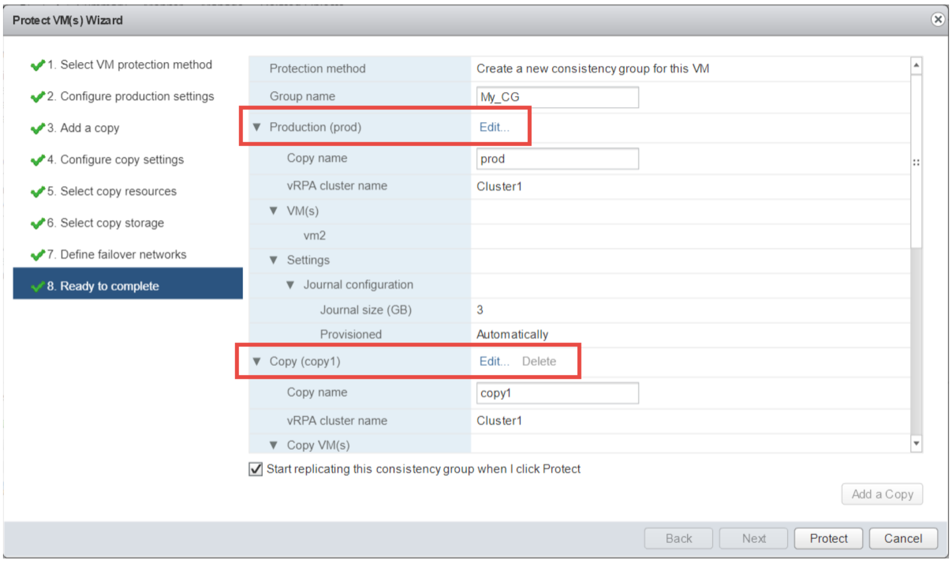

Protect VM(s) Wizard

- ✓ 1. Select VM protection method
- ✓ 2. Configure production settings
- ✓ 3. Add a copy
- ✓ 4. Configure copy settings
- ✓ 5. Select copy resources
- ✓ 6. Select copy storage
- ✓ 7. Define failover networks
- ✓ 8. Ready to complete

Protection method	Create a new consistency group for this VM
Group name	My_CG
▼ Production (prod)	Edit...
Copy name	prod
vRPA cluster name	Cluster1
▼ VM(s)	
vm2	
▼ Settings	
▼ Journal configuration	
Journal size (GB)	3
Provisioned	Automatically
▼ Copy (copy1)	Edit... Delete
Copy name	copy1
vRPA cluster name	Cluster1
▼ Copy VM(s)	

☑ Start replicating this consistency group when I click Protect

a. Ensure your protection settings are as required:
   - Expand the **Production** and **Copy** settings to ensure that they are correct.
   - If need be, click **Edit...** to change a setting before clicking **Protect**.
   - If you do not want to start replicating data from the production VM immediately, uncheck **Start replicating this group when I click Protect**.
   - Click **Add a Copy** to add more copies to the group.

b. Click **Protect** to create the copy(s) and enable VM protection.

Results

The specified virtual machine(s) are protected. If you added a virtual machine to an existing consistency group, a volume sweep occurs on the newly added virtual machine and a short initialization on all other virtual machines in the consistency group.

Next steps

- Use the **Protection** > **Consistency Groups** screen for Monitoring replication on page 34.
- Use The RecoverPoint for VMs Dashboard on page 33 to monitor the system.
- To stop replicating a VM, see Stop protecting a virtual machine on page 84.

Add a copy to a consistency group

The RecoverPoint for VMs **Add a Copy Wizard** will guide you through the process of protecting your virtual machines with a new cloud copy.

Prerequisites

Ensure you have completed Solution deployment on page 14.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Consistency Groups**.
2. Select the consistency group to which you want to add the cloud copy.
3. Click the **Add a copy** icon:
4. Follow the instructions for protecting VMs, starting from **step** 7 on page 29.

4

Monitoring protection

After protecting your VMs, use the RecoverPoint for VMs **Dashboard** to monitor the system. Use the **Protection** > **Consistency Groups** screen to monitor replication.

Topics:

The RecoverPoint for VMs Dashboard

Use the RecoverPoint for VMs **Dashboard** to attain a high-level overview of the RecoverPoint for VMs system. The RecoverPoint for VMs **Dashboard** and its sub-tabs present important system information to help you analyze and monitor your RecoverPoint for VMs system.

About this task

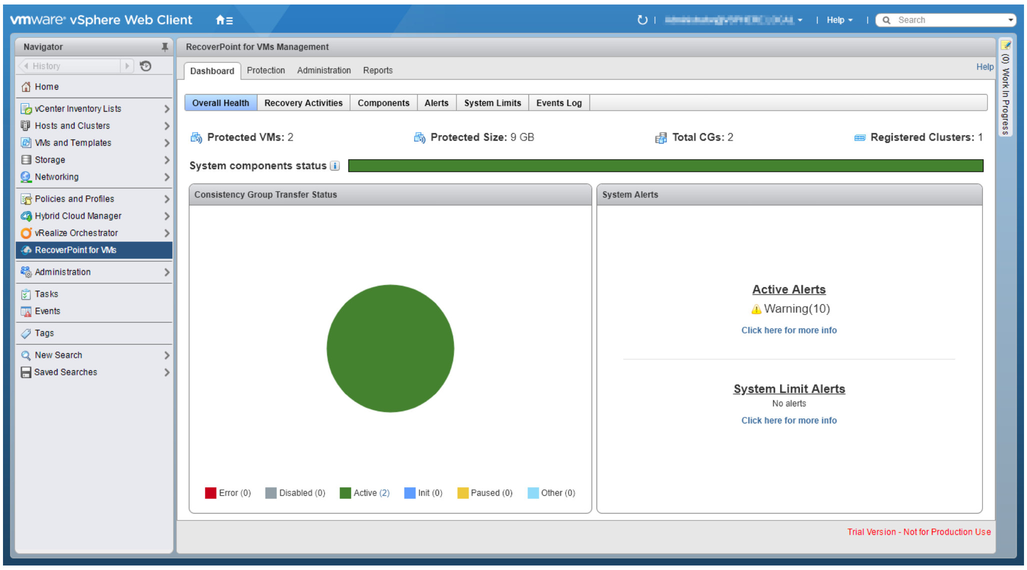

Protected VMs: 2

Protected Size: 9 GB

Total CGs: 2

Registered Clusters: 1

Warning(10)

No alerts

🟥 Error (0)
▨ Disabled (0)
🟩 Active (2)
🟦 Init (0)
🟨 Paused (0)
🟦 Other (0)

Steps

1. In your **vSphere Web Client** > **Navigator**, select **RecoverPoint for VMs** to access the system **Dashboard**.
2. To monitor your RecoverPoint for VMs system, click the **Overall Health**, **Components**, **Alerts**, **System Limits**, and **Events Log** sub-tabs.

> **NOTE: The Recovery Activities tab is not relevant in the RecoverPoint for VMs cloud solution. To monitor recovery, use The CDRS user interface on page 56 and The CDRS Dashboard on page 56.**

Monitoring replication

Use the RecoverPoint for VMs **Protection** tab to monitor all aspects of replication.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Consistency Groups**.

   In the **Transfer** column, note the **Transfer State** of each consistency group.

   - **Replicating Snap (n%)**: Data is being replicated to a copy in snap-based replication mode.
   - **Snap Idle**: The system is not replicating. The system is about to begin or has completed replicating a snap.
   - **Snap Init**: During snap-based replication, the current copy snapshot is being initialized.
   - **Snap Error**: An error occurred during snap-based replication.
   - **Snap High-load**: During snap-based replication, the system enters a temporary high-load state if the thin device used for snap replication cannot handle the load of the protected VM writes. The system will resolve the high-load state without user action, when the write load goes back to normal.
   - **Paused**: Data is not being transferred to a copy, because transfer has been paused by the user.
   - **Paused by System**: Data is not being transferred to a copy, because transfer has been paused by the system. If this state occurs for long periods of time, check the system alerts and events in the **Dashboard** for more information.
   - **N/A**: Data is not being transferred to a copy, because the copy has been disabled by the user.

2. Select a specific consistency group. The state of transfer to each of the copies in the group, and other important information about the replication process, is displayed in the group **Topology** diagram.

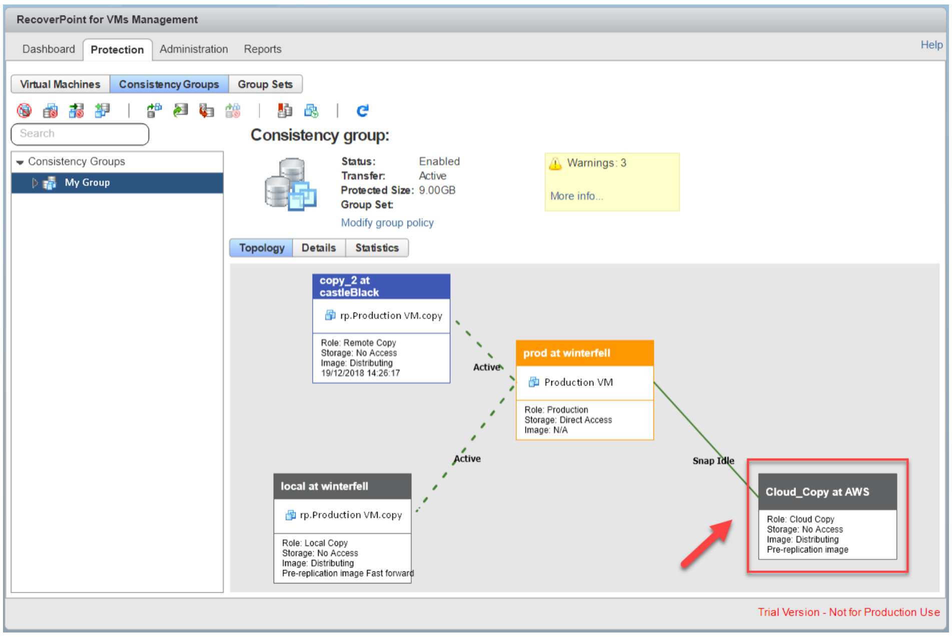

Consistency group:

Status: Enabled
Transfer: Active
Protected Size: 9.00GB
Group Set:

Warnings: 3

copy_2 at castleBlack
rp.Production VM.copy
Role: Remote Copy
Storage: No Access
Image: Distributing
19/12/2018 14:26:17

prod at winterfell
Production VM
Role: Production
Storage: Direct Access
Image: N/A

local at winterfell
rp.Production VM.copy
Role: Local Copy
Storage: No Access
Image: Distributing
Pre-replication image Fast forward

Cloud_Copy at AWS
Role: Cloud Copy
Storage: No Access
Image: Distributing
Pre-replication image

3. Select the **Details** and **Statistics** sub-tabs for more detailed information about your replication environment, configuration, and performance.

Monitor cloud solution components

Use the RecoverPoint for VMs **Cloud Services** tab to monitor the registration, usage, and connectivity of the components of the RecoverPoint for VMs cloud solution.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster of your cloud copy.
3. Click the **Cloud Services** tab:

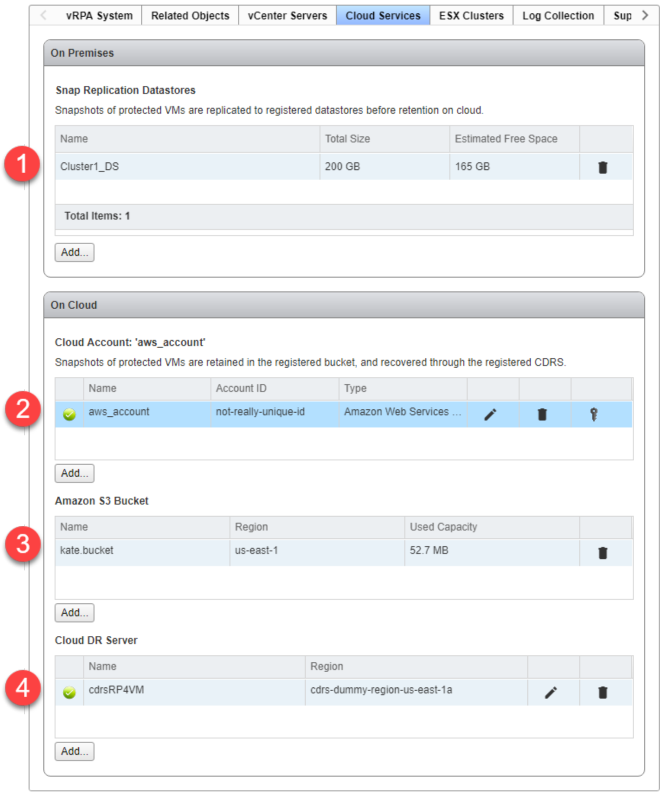

On Premises

Snap Replication Datastores

Snapshots of protected VMs are replicated to registered datastores before retention on cloud.

Name	Total Size	Estimated Free Space	
Cluster1_DS	200 GB	165 GB	🗑

Total Items: 1

Add...

On Cloud

Cloud Account: 'aws_account'

Snapshots of protected VMs are retained in the registered bucket, and recovered through the registered CDRS.

	Name	Account ID	Type			
✓	aws_account	not-really-unique-id	Amazon Web Services ...	✎	🗑	🔑

Add...

Amazon S3 Bucket

Name	Region	Used Capacity	
kate.bucket	us-east-1	52.7 MB	🗑

Add...

Cloud DR Server

	Name	Region		
✓	cdrsRP4VM	cdrs-dummy-region-us-east-1a	✎	🗑

Add...

a. Note which components of the cloud solution are registered, and which are not. Select the **Cloud Account** in the table to display the S3 bucket and CDRS that have been registered in the account (if any).
b. Note the information that is displayed for each component of the cloud solution:

   i. **Snap Replication Datastores**:

      - **Name:**The names of the on-premises datastore(s) that have been registered for snap replication, in RecoverPoint for VMs.
      - **Total Size:**
      - **Estimated Free Space:**

   ii. **Cloud Account**:

      - **Connectivity:**: A green checkmark to the left of the account name means there is connectivity between the vRPA cluster and the registered cloud account.
      - **Name:**The name of the registered cloud account, in RecoverPoint for VMs.
      - **Account ID:** The account ID associated with the registered cloud account.

   iii. **Amazon S3 Bucket**:

      - **Name:**The name of the registered S3 bucket, in RecoverPoint for VMs.
      - **Region:**The AWS Region of the registered S3 bucket.
      - **Used Capacity:**The amount of space that is used to contain snapshots of protected VMs, and other system information, in the registered S3 bucket.

   iv. **Cloud DR Server**:

      - **Connectivity:** A green checkmark to the left of the Cloud DR Server name means there is connectivity between the vRPA cluster and the registered Cloud DR Server.
      - **Name:**The name of the registered CDRS, in RecoverPoint for VMs.
      - **Region:**The AWS Region of the registered CDRS.

Next steps

If a red error icon is displayed to the left of the Cloud Account or Cloud DR Server, see Monitoring system alerts and events on page 36 to troubleshoot.

Monitoring system alerts and events

Use RecoverPoint for VMs events and alerts to understand and troubleshoot events in your RecoverPoint for VMs environment.

About this task

To monitor your system events, use the **RecoverPoint for VMs vSphere plugin** > **Dashboard** > **Events Log** tab.

An event is a notification that a change has occurred in the state of a system component. In some cases, the change indicates an error or warning condition for the component. Multiple events can occur simultaneously on a single component. A single incident can generate events across multiple system components.

In RecoverPoint for VMs, events have a:

- **Level**: **`Info`**, **`Warning`**, or **`Error`**
- **Scope**: **`Normal`**, **`Detailed`**, or **`Advanced`**
- **Topic**: **`All`**, **`vRPA Cluster`**, **`vRPA`**, **`Group`**, **`Splitter`**, or **`Management`**

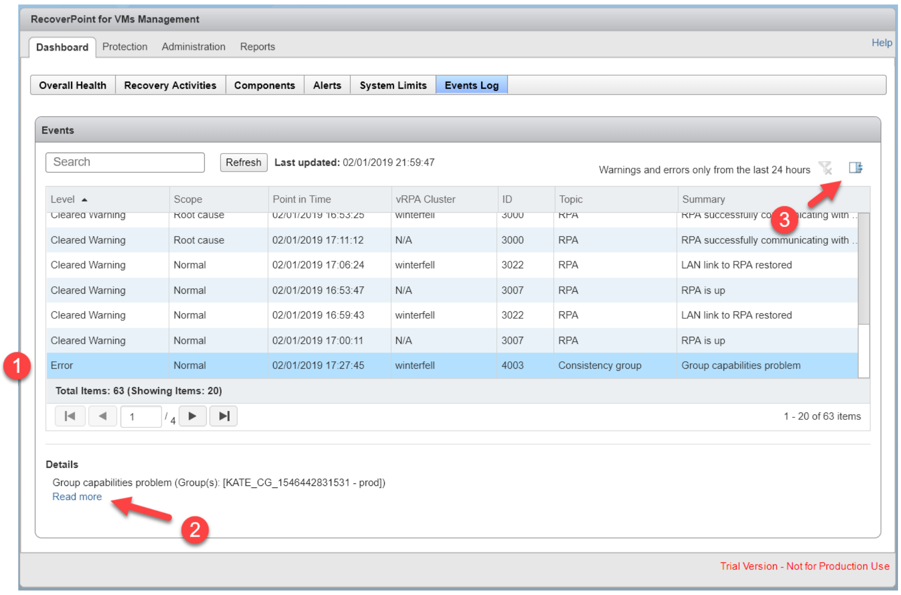

Last updated: 02/01/2019 21:59:47

Warnings and errors only from the last 24 hours

Level	Scope	Point in Time	vRPA Cluster	ID	Topic	Summary
Cleared Warning	Root cause	02/01/2019 16:53:25	winterfell	3000	RPA	RPA successfully co icating with ...
Cleared Warning	Root cause	02/01/2019 17:11:12	N/A	3000	RPA	RPA successfully communicating with ...
Cleared Warning	Normal	02/01/2019 17:06:24	winterfell	3022	RPA	LAN link to RPA restored
Cleared Warning	Normal	02/01/2019 16:53:47	N/A	3007	RPA	RPA is up
Cleared Warning	Normal	02/01/2019 16:59:43	winterfell	3022	RPA	LAN link to RPA restored
Cleared Warning	Normal	02/01/2019 17:00:11	N/A	3007	RPA	RPA is up
Error	Normal	02/01/2019 17:27:45	winterfell	4003	Consistency group	Group capabilities problem

Total Items: 63 (Showing Items: 20)

1 / 4

1 - 20 of 63 items

Details

Group capabilities problem (Group(s): [KATE_CG_1546442831531 - prod])

Read more

In the events log:

1. Select an event to display the event **Details**.
2. If there is detailed event information that can help you to troubleshoot, a **Read more** link is displayed. Click this link to display additional information about the selected event.
3. Click the **Event Filter** icon to configure which events are displayed in the **Events Log**.

> **NOTE: For more event monitoring and troubleshooting options, log into the RecoverPoint for VMs Command Line Interface (CLI) and run the `get_events_log` command. For more information, see the *RecoverPoint for Virtual Machines CLI Command Reference Guide*.**

To monitor your system alerts, click the **RecoverPoint for VMs vSphere plugin** > **Dashboard** > **Alerts** tab. System alerts are a mechanism that allows vRPAs to send events about system components in real-time, to a specified email, or the system reports database, via SMTP.

- To manage your system alerts settings, see Managing cloud solution support on page 89.
- You can also monitor the alerts of specific consistency groups from the **Protection** > **Consistency Groups** screen, when you select a specific consistency group. Click the **More info...** link under the alert box for more details about these alerts.

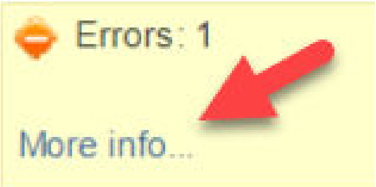

Identifying a RecoverPoint for VMs system

When a vRPA cluster is selected, the GUI displays all other vRPA clusters (besides the one you are connected to) that constitute a RecoverPoint for VMs system.

Steps

1. In the **RecoverPoint for VMs vSphere plugin**, select **Administration** > **vRPA Clusters** > **vRPA System**
2. Select a vRPA cluster.
3. Note the value of **Other vRPA clusters in system**.

5

Protection automation

RecoverPoint for VMs provides the following features that automate the protection of your VM copies.

Topics:

Create a bookmark

Create a snapshot of a virtual machine, a consistency group, or a group set, and label it for easy identification during testing and recovery.

About this task

NOTE:

You can also create bookmarks using the Cloud DR Server user interface. See Create a tag on page 67 for more details.

Steps

1. In the **RecoverPoint for VMs vSphere plugin**, click **Protection**.

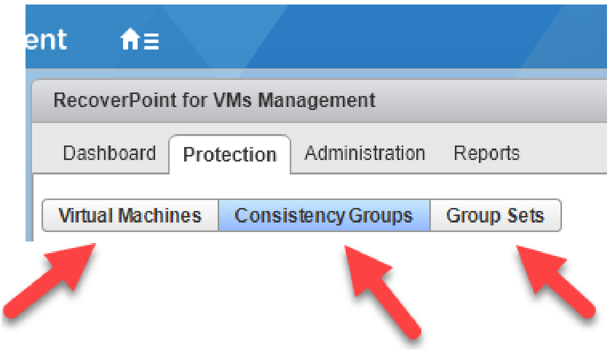

   - To bookmark a virtual machine, select the **Virtual Machines** tab.
   - To bookmark a consistency group, select the **Consistency Groups** tab.
   - To bookmark a group set, select the **Group Sets** tab.
2. Select the consistency group or group set that you want to bookmark.
3. Click the **Create Bookmark** button at the top of the screen:
4. In the **Create Bookmark** dialog box:

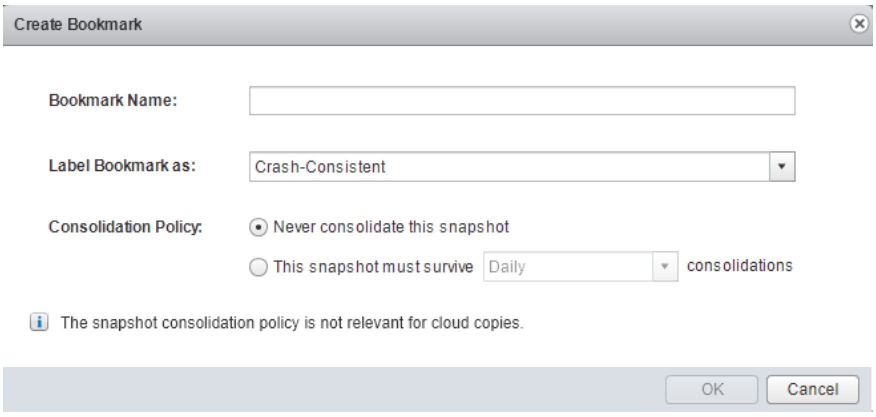

Create Bookmark

Bookmark Name:

Label Bookmark as: Crash-Consistent

Consolidation Policy: ● Never consolidate this snapshot

○ This snapshot must survive Daily consolidations

The snapshot consolidation policy is not relevant for cloud copies.

- **Name** - Type a name for the snapshot. This is the bookmark. The bookmark is the name that will be used to identify the snapshot during testing and recovery.
- **Label Bookmark As** - Select:
  - **Crash-Consistent**: Labels the snapshot as crash-consistent.
  - **Application-Consistent**: Labels the snapshot as application-consistent. Selecting this option does not create an application-consistent snapshot, it only labels the snapshot as application-consistent.

5. Click **OK**.

Results

A crash-consistent snapshot is created with the specified name for the specified virtual machine, consistency group, or group set.

Create a group set

Create a group set to bookmark, enable/disable, modify the start-up sequence, and pause/start replication for multiple groups, simultaneously.

About this task

NOTE: If a group set contains a consistency group with both an on-premises copy and a cloud copy, parallel bookmarks can be enabled even though this feature is not supported for cloud copies. In this case, parallel bookmarks are applied only to the on-premises copy.

Steps

1. Click **Protection** > **Group Sets**.
2. Click the **Add Group Set** icon:

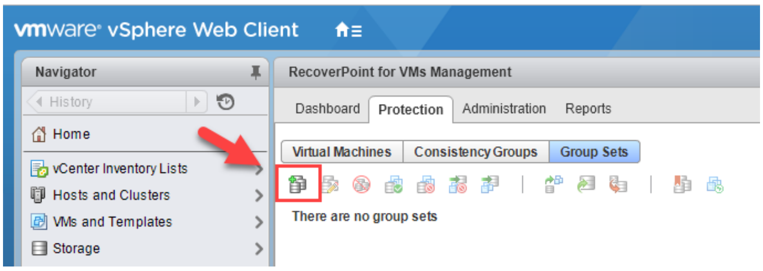

There are no group sets

3. In the **Add Group Set** dialog box:

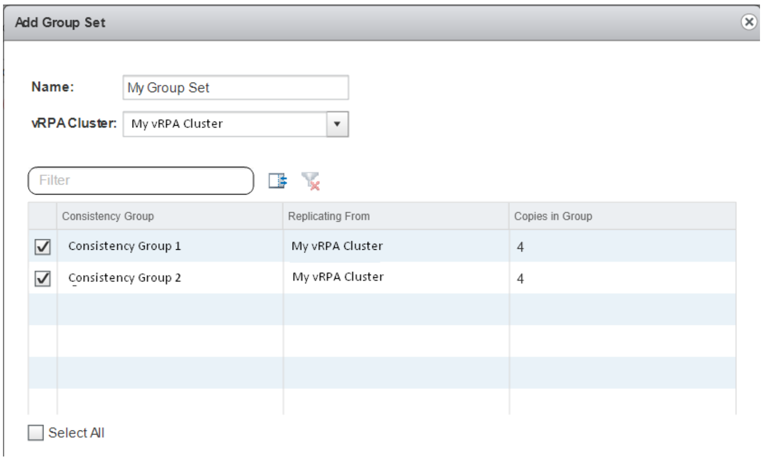

Add Group Set

Name: My Group Set

vRPA Cluster: My vRPA Cluster

Filter

	Consistency Group	Replicating From	Copies in Group
☑	Consistency Group 1	My vRPA Cluster	4
☑	Consistency Group 2	My vRPA Cluster	4

□ Select All

a. Enter a descriptive name for the group set.
b. Choose the vRPA cluster from which to select consistency groups.
c. Select the consistency groups to add to the group set.

4. Click **OK**.

Results

The group set is created.

Next steps

Select a group set and use the buttons at the top of the screen to bookmark, enable/disable, modify the start-up sequence, and pause/start replication for multiple groups, simultaneously.

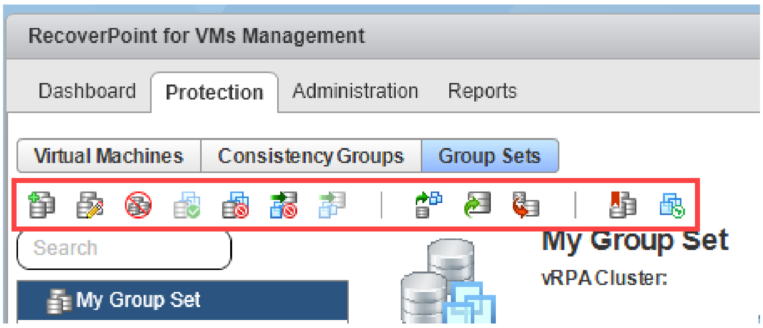

Disabling automatic protection of new VMDKs

By default, all newly added VMDKs are automatically protected. Use this procedure to change the default behavior.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Virtual Machines**.
2. Select the virtual machine for which you want to disable the automatic protection of any newly added VMDKs, in the future.
3. Under the **Protected VMDKs** widget, click **Edit...**.
4. In the **Edit VMDK Protection Policy** dialog, clear the **Automatically protect new VMDKs** checkbox to disable automatic protection of any newly added VMDKs to this VM.
5. Click **OK**.

Results

Any VMDKs added to this VM in the future will not be automatically protected.

Excluding a VMDK from replication

If required, you can mark individual VMDKs for exclusion from replication. For example, virtual machines containing shared or non-persistent VMDKs cannot be replicated. You can, however, change the VMDK type in vSphere, or mark these VMDKs to be excluded from replication in the **RecoverPoint plugin for vSphere**, and replicate the virtual machines without them. You can also use the following procedure to include an excluded VMDK.

About this task

- **For Linux-based production VMs:** If in the **file systems table** (in `/etc/fstab`), a VM with an excluded VMDK is configured to be automatically mounted when the operating system boots up , add a *nofail* option in the file systems table or the cloud copy of the VM will fail to boot when Recovering VMs on page 43.
- Changing the disk type of an excluded shared or non-persistent VMDK to a supported type (such as non-shared or persistent) does not automatically include the VMDK, regardless of the value of the **Automatically protect new VMDKs** setting.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Virtual Machines**.
2. Under the **Protected VMDKs** widget, click **Edit...**.
3. Clear the checkbox next to the VMDKs that you want to exclude from replication.
4. Click **OK**.

Results

In the future, the excluded VMDKs are not displayed in the list of snapshots that you can select when Recovering VMs on page 43.

Adding a new VMDK

RecoverPoint for VMs automatically detects when a hard disk is added to a protected virtual machine using the **vSphere client** > **Virtual Machine Properties**, and orchestrates VMDK addition.

About this task

When adding VMDKs to a protected VM, RecoverPoint for VMs automatically starts protecting each added VMDK by creating copies of it, as specified by the VM's consistency group link and copy settings.

- To disable automatic protection for all VMDKs added to a protected VM in the future, see Disabling automatic protection of new VMDKs on page 40.
- To exclude specific VMDKs of a protected VM from protection, see Excluding a VMDK from replication on page 41.

RecoverPoint for VMs does not automatically protect VMDKs of type **`shared`** when they are added to a protected VM.

Results

A volume sweep occurs on the added VMDK(s) and a short initialization occurs on all other VMDKs in the consistency group, but no history is lost.

Removing a VMDK

RecoverPoint for VMs automatically detects when a hard disk is removed from a protected virtual machine using the **vSphere client** > **Virtual Machine Properties**, and orchestrates VMDK removal.

About this task

If you don't want to replicate a specific VMDK of a protected VM to the copy, you can remove the VMDK from the production VM through the vSphere client **Virtual Machine Properties**, or exclude it from replication, as described in Excluding a VMDK from replication on page 41.

Results

Future VM snapshots will not include the removed VMDK.

Next steps

The removed VMDKs can be recovered by Recovering VMs on page 43, and selecting a snapshot that precedes the removal of the VMDK.

Managing the cloud copy retention policy

The RecoverPoint for VMs cloud solution automates copy retention on AWS.

Each copy has a **Retention Policy** that dictates the period of time during which the copy snapshots should be available for recovery. Once every 24 hours, a temporary **Retention Service** EC2 instance is launched in AWS to consolidate the snapshots of every copy whose retention policy has expired.

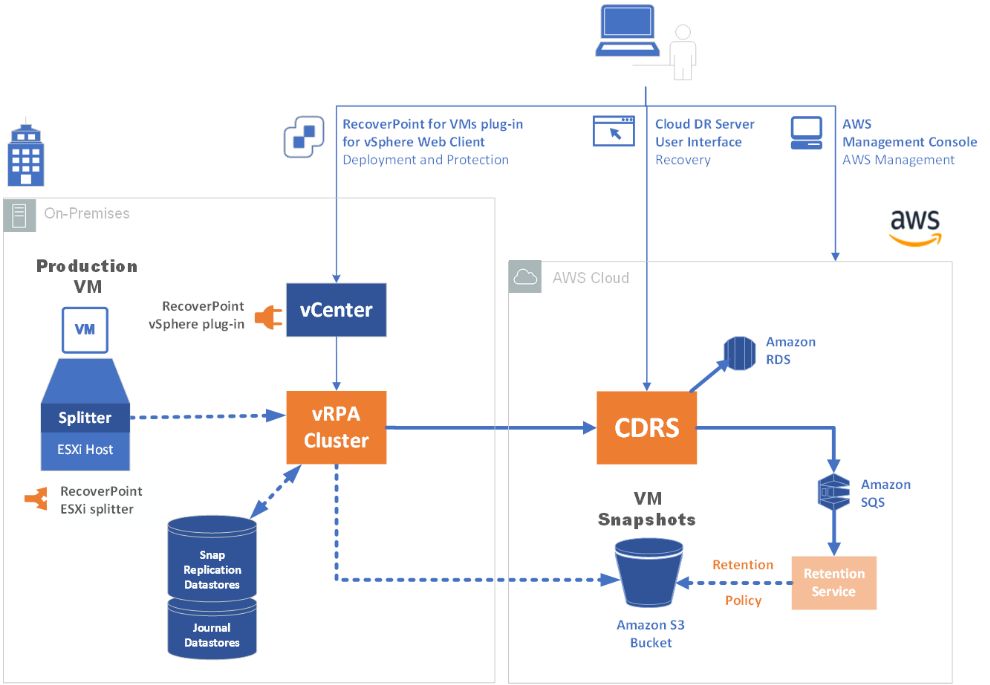

The copy retention policy is set while Protecting VMs on page 27. After protecting your VMs, you can change the copy retention policy by Managing the protection policies of cloud copies on page 86.

NOTE:

Snapshots whose Retention Policy has expired are no longer available when Recovering VMs on page 43.

6

Recovering VMs

Use Cloud DR Server user interface to periodically test copy images and fail over to a secondary cloud copy in the case of a disaster or during system maintenance. You can also fail back from AWS to an on-premises vCenter, or recover to a vCenter or VMware Cloud on AWS.

Before recovering VMs, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

Topics:

Log into the CDRS interface

To log into the cloud-based CDRS component of the Cloud Disaster Recovery solution, you need a username and password.

Steps

1. From a host that has network access to the CDRS virtual appliance, use a browser to connect to the appliance:

   ```
   https://CDRS_hostname
   ```

   Where *CDRS_hostname* is the hostname or IP address of the address that was created when the CDRS was deployed from the CDRA. You can find the CDRS hostname on the **Cloud DR Server** page in the Cloud DR Add-on window by selecting **Configuration** > **Cloud DR Server**.
2. For **Username**, enter **`admin`**.
3. For **Password**, enter the password for the admin user.

   If you have forgotten the password:

   a. Click **Forgot Password?**.
   b. Enter 'admin' as the username, and click **Send**.

   > **NOTE: CDRS checks whether the User email address (see Define the email address for CDRS password recovery on page 96) exists in (and has been verified by) the AWS root user account. If a valid User email address has been defined, an email is sent to the specified email address, with instructions for resetting the password.**

Results

On logging in, the Cloud DR Server window opens and the **Welcome** page appears.

The menu bar on the Cloud DR Server window shows the current location in the user interface. To log out of the Cloud DR Server user interface, click the icon on the right side of the menu bar and select **Sign out**. To leave feedback, click **Tell us what you think** at the bottom of the window, enter the comments, and click **Send Feedback**.

Recovery workflows

Use following guides as references during testing, failover and failback, to understand what to expect and how to proceed, at every step in the recovery process.

Test

A DR test enables temporary access to a cloud instance to verify that a recovered asset works before you perform a failover. Testing DR scenarios before a real disaster occurs is a recommended best practice that saves time and ensures that production assets on premises can be quickly recovered in the cloud.

Figure 1 shows the basic test workflow. Table 1 lists the user actions that are available for each workflow state.

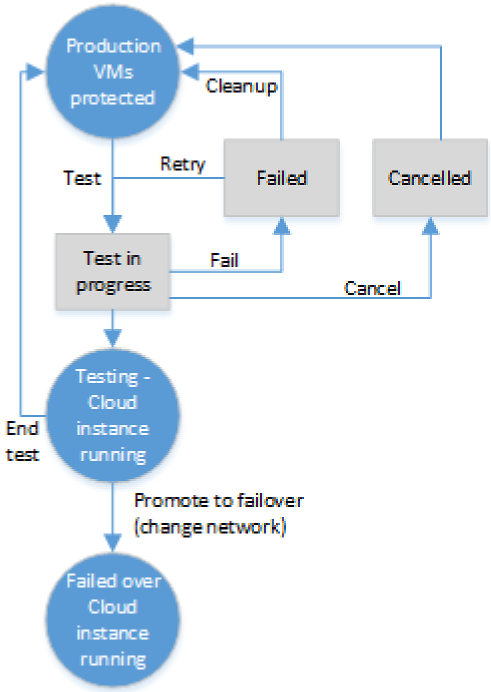
Production VMs protected
Cleanup
Retry
Test
Failed
Cancelled
Test in progress
Fail
Cancel
Testing - Cloud instance running
End test
Promote to failover (change network)
Failed over Cloud instance running

Figure 7. DR test workflow

To understand the workflow and available user actions for each state, read Table 1 from left to right and from top to bottom.

Table 5. Test workflow states and related user actions

Workflow state	User Actions	Next state
Starting state: Production VMs are protected in cloud and remain protected during the test	Select VM/DR Plan Select test network Select cloud instance, security group Start test	Test in progress
Test in progress	Cancel	Canceled
Canceled	--	Starting state
Failed	Retry	Test in progress
	Clean up	Starting state
Succeeded: Testing - cloud instance running	Promote to failover (can change network)	Failed over - cloud instance running
	End test (removes cloud instance)	Starting state

Failover

You perform a failover to the cloud when an on-premises disaster occurs and the production VMs are not running.

During a failover, shut down the on-premises production VMs to prevent users from writing new data to them.

Figure 1 shows the basic failover workflow. Table 1 lists the user actions that are available for each workflow state.

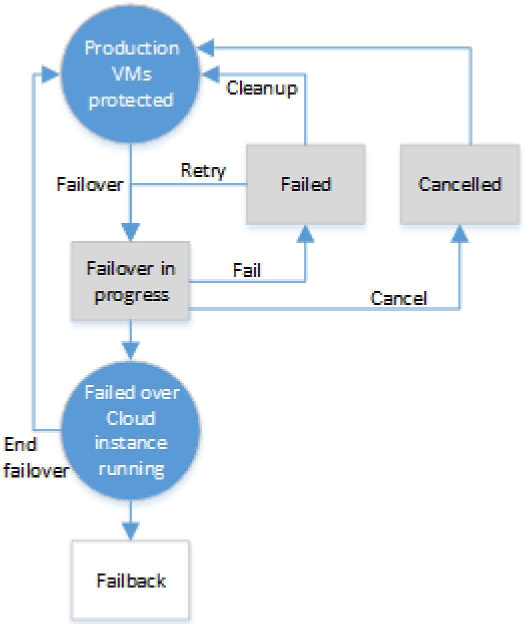

Figure 8. Failover workflow

To understand the workflow and available user actions for each state, read Table 1 from left to right and from top to bottom.

Table 6. Failover workflow states and related user actions

Workflow state	User Actions	Next state
Starting state: Production VMs are protected in cloud and remain protected during failover	Select VM/DR Plan Select failover network Select cloud instance, security group Start failover	Failover in progress
Failover in progress	Cancel	Canceled
Canceled	--	Starting state
Failed	Retry	Failover in progress
	Clean up	Starting state
Succeeded: Failed over - cloud instance running	Fail back	Failed back
	End failover (removes cloud instance)	Starting state

Failback

A failback transfers a failed-over VM (cloud instance) back to the on-premises vSphere environment. A failback is only crash-consistent, not application-consistent.

Before starting failback, it is a best practice to shut down services on the cloud instance.

Figure 9. Failback workflow on page 46 shows the basic failback workflow. Table 7. Failback workflow states and related user actions on page 46 lists the user actions that are available for each workflow state.

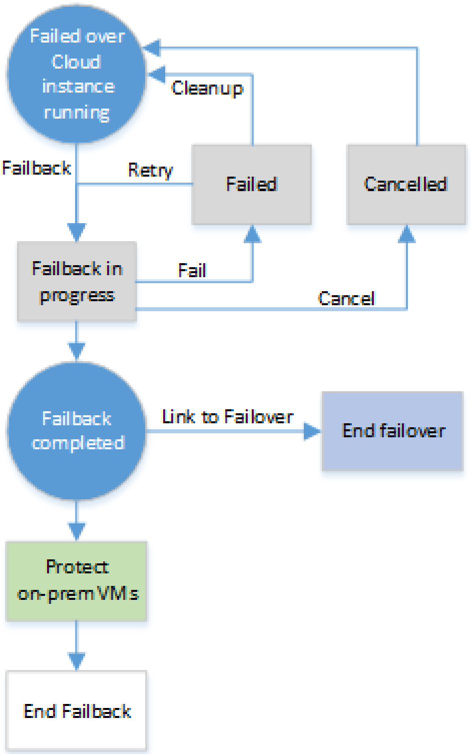

Figure 9. Failback workflow

To understand the workflow and available user actions for each state, read Table 7. Failback workflow states and related user actions on page 46 from left to right and from top to bottom.

Table 7. Failback workflow states and related user actions

Workflow state	User Actions	Next state
Starting state: Failed over - cloud instance running	Select VM/DR plan. Start failback.	Failback in progress
Failback in progress	Cancel.	Canceled
Canceled	--	Starting state
Failed	Retry	Failback in progress
	Clean up.	Starting state
Succeeded: Failback completed, new VM copies restored on premises	Link to failover activity card. End failover to terminate recovered cloud instances.	--

Test a protected VM and fail over to AWS

Testing and failover takes the VM snapshots in your S3 bucket, converts them to AMI format and places a copy of the VMs in AMI format on another storage and network on AWS.

NOTE:

Conversion to AMI format takes time. To enable speedy disaster recovery, create rapid recovery images before you fail over, and create tags and DR plans to manage recovery, as described in Recovery orchestration on page 67.

The following diagram illustrates the testing and failover workflows:

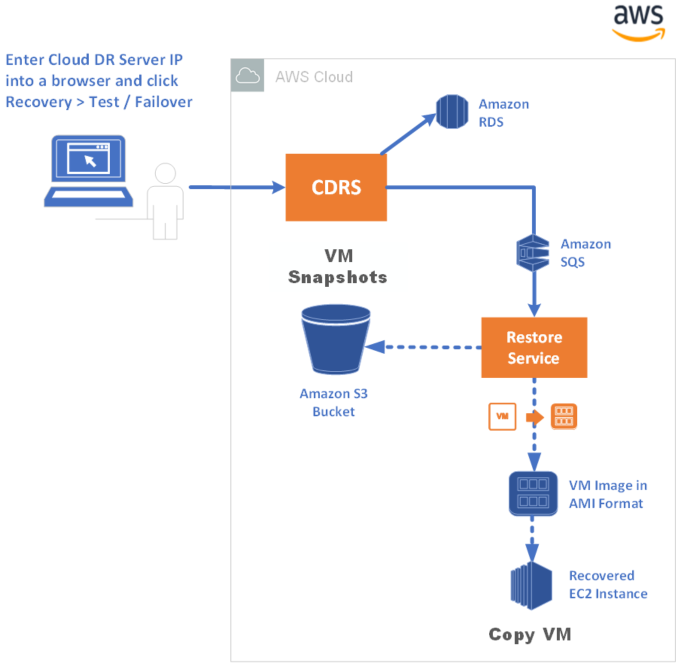

After you test a VM, using Test or fail over a protected VM on AWS cloud on page 47, you can Promote a DR test to failover.

Test or fail over a protected VM on AWS cloud

This procedure describes how to test or fail over a VM protected by RecoverPoint for VMs, on the AWS cloud, when an operational error or disaster occurs on premises.

Prerequisites

- To ensure a successful failover, and better prepare for a disaster, best practices recommend testing various disaster recovery scenarios. After performing a test, you can promote the test to a failover.
- To perform a DR test or failover of an asset, you must have VMs that are protected and copied to the cloud.
- To fail over to a vCenter or VMware Cloud environment, see Failover to vCenter or VMware Cloud on AWS on page 53.
- If you intend to use tags, you must first create the tags. See Create a tag on page 67.

Steps

1. In the Cloud DR Server user interface, select **Recovery > Asset Recovery**

   You can also open the **Asset Recovery** page from the dashboard by clicking **See All** in the **Recovery** pane.

   The **Asset Recovery** page is displayed.
2. Select the asset that you want to recover and click **Test** or **Failover**.

   If you click **Failover** and the asset has never been tested, a dialog box opens and reminds you that running a DR test is recommended before implementing a failover. The message also recommends that you shut down the production VM to avoid a possible data loss that is caused by accidental user access. Click **Select Copy** to continue.

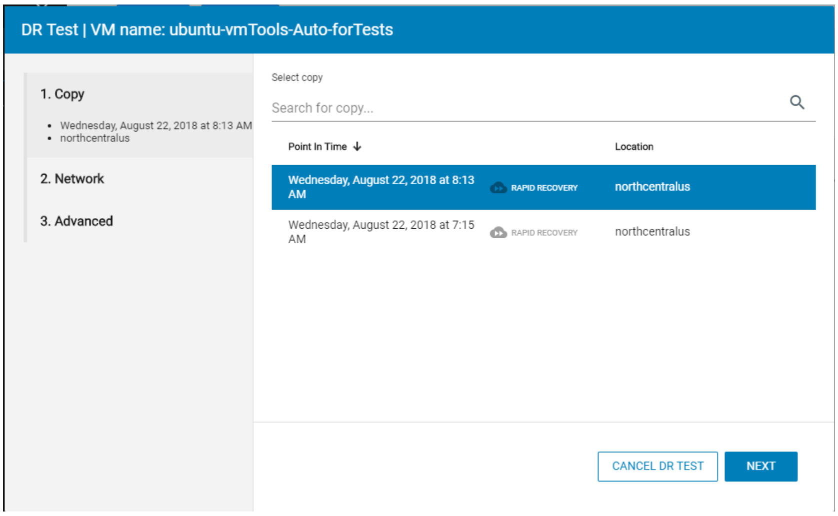

DR Test | VM name: ubuntu-vmTools-Auto-forTests

1. Copy
- Wednesday, August 22, 2018 at 8:13 AM
- northcentralus

2. Network

3. Advanced

Select copy

Search for copy...

Point In Time ↓		Location
Wednesday, August 22, 2018 at 8:13 AM	RAPID RECOVERY	northcentralus
Wednesday, August 22, 2018 at 7:15 AM	RAPID RECOVERY	northcentralus

3. In the wizard that opens, in the **Copy** step, select a point-in-time copy of a protected VM that you want to test or fail over, and then click **Next**.

   All bookmarks created using RecoverPoint for VMs are displayed.
4. In the **Network** step, select the network where you want to launch the EC2 instance, and then click **Next**.
5. (Optional) In the **Advanced** step:
   a. In the **Security Groups** tab, select a security group.
   b. In the **EC2 Instance Type & Tags** tab, select an EC2 instance type and a tag.
   c. In the **IP settings** tab, to enter a private IP address for the recovered instance, select the checkbox for this setting and enter the address. The system prevents you from selecting an IP address that is already in use.
6. Click **Start DR Test** or **Start Failover**.

Results

The recovery process begins and you can monitor progress on the **DR Activities** page. During recovery:

1. A temporary **Restore Service** instance is launched in each region where recovery is needed (unless the VM is enabled for rapid recovery). This instance performs hydration during recovery, and is automatically terminated after 10 minutes of idle time.
2. The Cloud DR Server converts the VMDK to an AMI and launches an EC2 instance that is based on the AMI.
3. When the EC2 instance is running, the Cloud DR Server deletes the VMDK and AMI.

Network communications

After failover to the cloud, the customer is responsible for ensuring proper networking communications from restored VM instances on the cloud to their local network, such as using a VPN or similar networking solution, load balancing, and other networking-related issues.

Promote a DR test to failover

From the **DR Activities** page, you can promote a test of a single asset to failover.

Prerequisites

Before promoting a test to failover, shut down the on-premises production VM. This action ensures that users do not accidentally write new data to the on-premises VM when they should be accessing the cloud-based VM instead.

Steps

1. To view status and other information about recovery activities, select **Recovery** > **DR Activities**.
   The **DR Activities** page displays.
2. For a DR test that is in the running state, click **Promote to Failover**.
   The **Promote to Failover** dialog box is displayed. It reminds you shut down the production VM to avoid possible data loss. To continue, click **Select Network**.
3. In the **Promote to Failover** dialog box, select the network for the failover operation:

Option	Description
Keep current network	Retains the network that was used during the test.
Select a network/security group	Enables selecting a different network for the failover.

4. If you select a different network for the failover, you can also select the default security group or a different security group.
5. To select a private IP address for the recovered instance, select the checkbox for this setting, and enter the address. The system prevents you from selecting an IP address that is already in use.
6. Click **Failover**.

End a DR test

When a DR test on a single VM or a DR plan has completed and is in the running state, you can end the test from the **DR Activities** page.

Steps

1. To view status and other information about recovery activities, select **Recovery** > **DR Activities**.
   The **DR Activities** page is displayed.
2. For a test that is in the running state, click **End DR Test**.
3. In the **End this DR Test** dialog box, click **End Test**.

Results

When you end a DR test, CDRS clears all used resources from the cloud, and the recovered instances are terminated.

ⓘ **NOTE:**

You can also terminate a recovery instance from the cloud provider console. When you terminate the recovery instance, the CDRS DR Activities page indicates an Instance Terminated status.

End a failover

You can end a failover at any time after a failback transfers a VM from the cloud to the on-premises vSphere environment.

Steps

1. Select **Recovery** > **DR Activities**.
2. If available, click **Open Failover Activities** for the VM.

   ⓘ **NOTE: The Open Failover Activities option is displayed only if there are VMs in a successful failback state.**

   The **Failover Details** dialog box opens.
3. Click **End Failover**.

Results

When a failover ends, CDRS clears all used resources from the cloud, and the recovered instances are terminated.

ⓘ **NOTE:**

You can also terminate a recovery instance from the cloud provider console. When you terminate the recovery instance, the CDRS DR Activities page indicates an Instance Terminated status.

Fail back from AWS to an on-premises vCenter

After failing over to AWS, use this procedure to fail back to an on-premises vCenter.

ⓘ **NOTE:**

You must first deploy and configure a CDRA as described in Deploying the CDRA on page 76.

Failback workflow

A failback operation allows a failover instance to be copied back to an on-premises vCenter.

1. Failback is initiated from a failover instance by using the CDRS user interface.
2. CDRS powers off the instance and creates snapshots of its disks.
3. A Restore Service:
   a. Creates disks from the snapshots.
   b. Attaches the new disks to itself.
   c. Reads the data and creates segments of data, compressing and encrypting the data stored in the cloud target for that specific region.
4. When the CDRA receives a new failback request, it creates a Restore VM, including a boot disk, at the on-premises vCenter in the failback staging area. The failback staging area is defined during Cloud DR deployment at the **Connect to vCenter Server** page.
5. The Restore VM copies the data from the cloud storage. Disks (VDMKs) are directly attached to the Restore VM and allocated as thick lazy-zeroed.
6. When the restore process completes, the CDRA powers off the Restore VM, deletes the boot disk, configures the failed-back VM as necessary, and relaunches the VM.

   At this point, you can vMotion the VMs from the failback staging area to their original locations or new locations. The IP addresses used for Restore VMs are not used for failed back VMs, so assign appropriate IP addresses to failed back VMs and ensure that DHCP can resolve them.
7. The CDRS performs any required clean-up of temporary resources in the cloud provider environment. However, the user must use the cloud provider console or the CDRS user interface to manually terminate the original failover instance in the cloud. This instance was used to launch the failback process.

Failback from the cloud

When an operational error or a disaster occurs in the on-premises environment, you can fail over a VM or DR plan to the cloud. After a failover to the cloud, the failed-over workloads run on cloud instances (VMs) with data that is stored in cloud storage. When the on-premises issue is resolved, you may want to fail the cloud instance back to the on-premises environment to continue running the workloads locally, instead of in the cloud. This procedure provides steps to fail back workloads that were failed over to the cloud.

Prerequisites

- Ensure your cloud instances are in a failed-over state.
- Ensure that you have deployed an on-premises CDRA.

NOTE: To support failback operation, you must deploy a CDRA on premises, connect it to the existing CDRS in the cloud, enter the on-premises vCenter details, and define the recovery staging area.

Steps

1. To perform a failback, select **Recovery** > **DR Activities**.
   The **DR Activities** page displays.
2. Click **Failback** for the VM or DR plan that you want to recover from the failover state.

   The **Failback** option is available only for VMs or DR plans in a successful failover state.

   The **Failback** dialog opens.

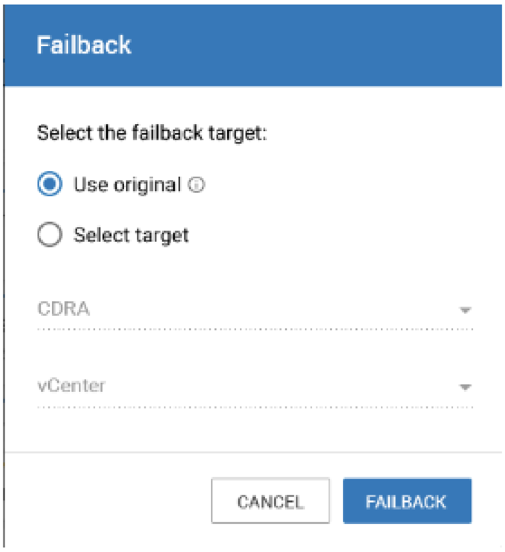

Failback

Select the failback target:

● Use original ⓘ

○ Select target

CDRA

vCenter

3. In the **Failback** dialog, select one of these options:

Option	Description
Use original	Enables you to fail back to the original VM location on premises.
Select target	Enables you to select the target CDRA and vCenter for the failback.

4. Click the **FAILBACK** button.
   The failback activity begins. The VM or DR plan is restored to the recovery staging area that you specified.
5. To verify that the VM is being restored, open vCenter. To display the **Summary** tab for the VM, click the VM in the list.

   The VM that you failed back does not have an assigned IP address.
6. Open the console for the VM or DR plan that you failed back, and assign IP addresses for the failback VMs.

   You can either assign an IP address or obtain an IP address from a DHCP server.
7. Manually install VMware tools on the failed back VM. (AWS removes VMware tools during AMI conversion.)

Results

After the failback has completed successfully, you can vMotion the VMs from the failback staging area to their original locations or new locations. The IP addresses used for Restore VMs are not used for failed back VMs, so assign appropriate IP addresses to failed back VMs and ensure that DHCP can resolve them.

The CDRS performs any required clean-up of temporary resources in the cloud provider environment. However, the user must use the cloud provider console or the CDRS user interface to manually terminate the original failover instance in the cloud. This instance was used to launch the failback process.

ⓘ **NOTE: The maximum number of failback activities is limited by the range of pool IP addresses that you configured for failback. If all IPs in the IP range pool already have failback operations in progress, a message informs you that the operation cannot be started until one or more of the running activities ends.**

Fail over to a vCenter or VMware Cloud on AWS

Recover the data of the VM copy in your Amazon S3 bucket to VMware Cloud on AWS. This architecture can also be used to recover your data to an on-premises vCenter.

ⓘ **NOTE:**

You must first deploy and configure a CDRA as described in Deploying the CDRA on page 76.

Requirements and limitations for VMware Cloud on AWS

Observe the requirements and limitations of Cloud DR with VMware Cloud on AWS (VMC).

Requirements

Recovery to VMware Cloud requires:

- AWS cloud account
- VMware Cloud deployed in AWS cloud environment (used on demand)
- On-premises RecoverPoint for VMs version 5.2.1 or later
- CDRA that is deployed in VMware Cloud (requires the same version level as the CDRS).

NOTE: Direct recovery to vCenter/VM Cloud is supported for UEFI enabled virtual machines.

Limitations

Failover to the VMware Cloud on AWS has these limitations:

- You cannot test a copy or promote a DR test to failover (only direct failover is supported).
- You cannot fail over from rapid recovery copies to VMC.
- You cannot use DR plans to fail over to VMC.
- You cannot use automated failback from VMC to the on-premises production site. Instead, use vCenter vMotion.

Prerequisites to enable failover to VMC

When you enable failover to VMware Cloud on AWS (VMC), ensure that you observe the detailed prerequisites in this section.

Provide a Cloud DR environment

Provide a Cloud DR environment, including a Cloud DR Server (CDRS). Procedures for deploying a CDRA and CDRS are described in this chapter.

About this task

You can deploy CDRS in any AWS region. To avoid the high costs of cross-region recovery, Dell EMC recommends to deploy CDRS within the VMC supported regions.

In a typical scenario, VMC is used on demand. When recovery operations are needed, you deploy a VMC SDDC, deploy a CDRA in the VMC, connect the CDRA to the CDRS, enable failover to the VMC vCenter, and then fail over the protected VM.

Create VMware Cloud on AWS

Prerequisites

Review VMware documentation about VMware Cloud on AWS: Getting Started with VMware Cloud.

Steps

1. Obtain a VMware Cloud on AWS (VMC) account.
2. Select an AWS region for VMC from the VMC supported regions list.
3. Connect VMC to the AWS account that is running Cloud DR.
4. Connect the VPC and subnet from the same region that you selected for VMC (in step 2 on page 52).
5. Configure networking for the VMC software defined data center (SDDC).

Configure SDDC networking

About this task

Details about configuring the SDDC networking are described in this white paper:

Creating a VMware Software-Defined Data Center.

High-level steps include:

Steps

1. To connect between the Management Gateway (MGW) and the Compute Gateway (CGW), create VPN gateway details:
   a. Add a VPN from MGW to CGW.
   b. Add VPN from CGW to MGW.
2. Create network connection firewall rules for MGW:
   a. To enable network connection from web to VMC, add rule: vCenter access from Web with `HTTPS(TCP 443)` service.
   b. To enable provisioning from inbound to ESXi, add rule: inbound to ESXi provisioning with `Provisioning (TCP 902)` service.
   c. To enable TCP connection from inbound to ESXi, add rule: inbound to ESXi 443 with `HTTPS(TCP 443)` service.
3. Create network connection firewall rules for CGW:
   a. To enable network connection from VMC to outside network, add rule: outbound any with `All traffic` service.
   b. To enable network connection from VPC (AWS) to VMC, add rule: any from VPC with `All traffic` service.

Deploy the CDRA on the vCenter in the SDDC

Steps

1. Ensure that the CDRA that is deployed in VMC is accessible from the customer network. Use one of these methods:
   - Create a jump host, a machine on the same VPC that you selected for the VMC (in step 4 on page 52).
   - Assign a public IP address to the CDRA. See Assign a Public IP Address to a VM.
   - Configure a VPN that connects the on-premises network to the SDDC. See the Network section in the VMware FAQs: https://aws.amazon.com/vmware/faqs/.
2. Deploy the CDRA.ova file (using the same version as the CDRS) on the vCenter in the SDDC. Use the VMC internal IP address.

Connect CDRA in VMC to the CDRS

Steps

1. Ensure that the cloud account credentials are the same as the configuration for the Cloud DR environment.
2. Connect the CDRA in the VMC to the existing CDRS.
3. Add the vCenter in the SDDC as the vCenter server. Define the recovery staging area and enable direct failover to this vCenter.

   > ⓘ **NOTE: The number of IP addresses that you allocate to direct failover defines the number of simultaneous recoveries that you can run.**

Failover to vCenter or VMware Cloud on AWS

This procedure describes how to fail over a VM to a recovery-enabled vCenter (for example, the vCenter where the VMware Cloud on AWS is deployed).

Prerequisites

Deploy the CDRA, and enable direct failover to the target vCenter, as described in Define a recovery staging area on page 82.

Steps

1. In the Cloud DR Server user interface, select **Recovery > Asset Recovery**
   The **Asset Recovery** page displays.
2. Select a VM and click **FAILOVER TO VCENTER**.

   The **Failover to vCenter** dialog box opens.
3. In the **Failover to vCenter** dialog box, in the **Copy** step, select a **Point in Time** copy and click **NEXT** to go to the **Failover Target** step.

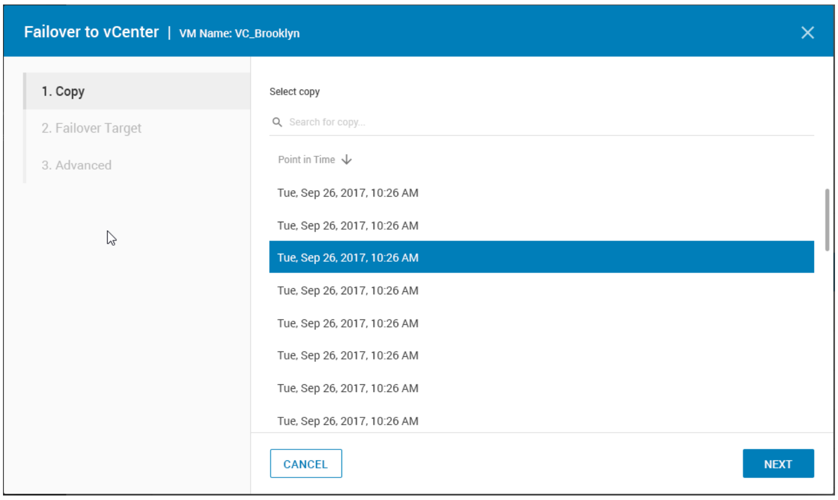

1. Copy

2. Failover Target

3. Advanced

Select copy

Search for copy...

Point in Time ↓
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM
Tue, Sep 26, 2017, 10:26 AM

All bookmarks created using RecoverPoint for VMs are displayed. Every copy snapshot (**Point in Time**) is replicated together with its OVF, so the failed over VM will have the same hardware settings that the protected VM had, at the selected **Point in Time**.

4. In the **Failover Target** step, select a CDRA/vCenter failover target.

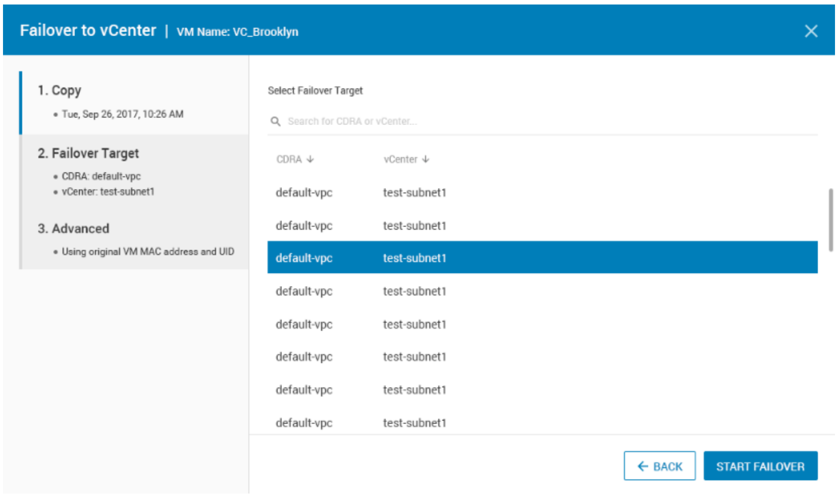

1. Copy
- Tue, Sep 26, 2017, 10:26 AM

2. Failover Target
- CDRA: default-vpc
- vCenter: test-subnet1

3. Advanced
- Using original VM MAC address and UID

Select Failover Target

Search for CDRA or vCenter...

CDRA ↓	vCenter ↓
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1

5. Optionally, in the **Advanced** section, update the **Keep original VM MAC address and UID** checkbox setting.

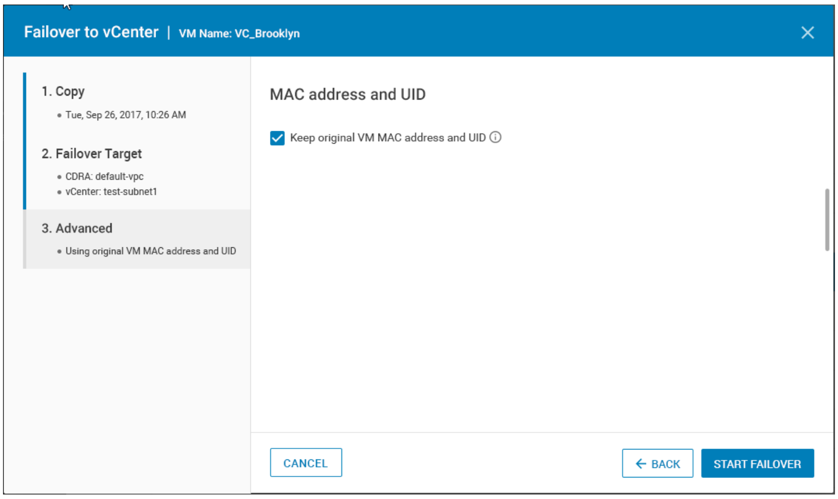

1. Copy
- Tue, Sep 26, 2017, 10:26 AM

2. Failover Target
- CDRA: default-vpc
- vCenter: test-subnet1

3. Advanced
- Using original VM MAC address and UID

MAC address and UID

☑ Keep original VM MAC address and UID ⓘ

If you are failing over to the same network as the production VM, to avoid IP conflicts, clear this checkbox to ensure that the failed over VM has a different MAC address and UID than that of the production VM.

> ⓘ **NOTE: When a production VM is protected, the hardware settings of the production VM (including the MAC address) are also replicated, with these exceptions:**

- **RAW disk is not supported. In the failed-over VM, it becomes a VMDK.**
- **Single-root I/O virtualization (SR-IOV) pass-through is not supported. In the failed-over VM, it becomes an e1000 virtual NIC.**

6. Click **START FAILOVER**.

Results

The failover process begins and you can monitor progress on the **DR Activities** page.

7

Monitoring recovery

The Cloud DR Server user interface contains tools for monitoring and managing all aspects of recovery.

Topics:

The CDRS user interface

The Cloud DR Server user interface provides a dashboard representation of the CDRS environment and the capability to perform and monitor recoveries of protected virtual machines and configuration tasks that are related to the CDRS.

The CDRS Dashboard

The CDRS dashboard provides insight into key product information and operational behavior. The dashboard is divided into panes that display unique information.

To open the dashboard, click **Overview** in the navigation pane of the **Cloud DR Server** window.

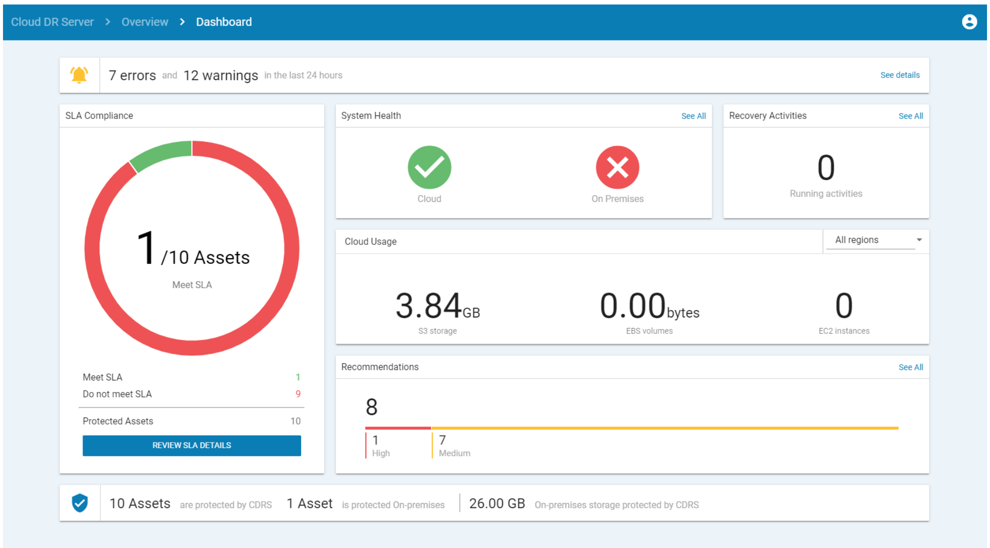

System Health pane

The **System Health** pane of the CDRS dashboard provides general system health status.

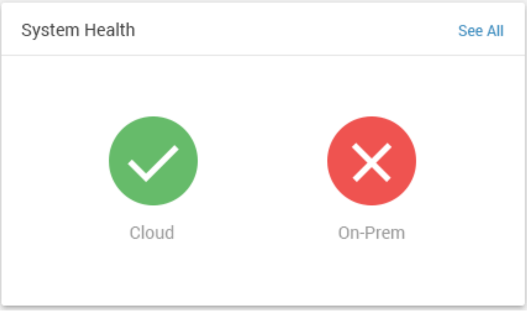

To view system health details, click **See All** in the **System Health** pane, or select **System** > **Health** from the navigation pane of the CDRS Dashboard. System Health on page 66 provides information about system health details.

Cloud Usage pane

The **Cloud Usage** pane of the CDRS dashboard provides a summary of the amount of storage being used in the cloud.

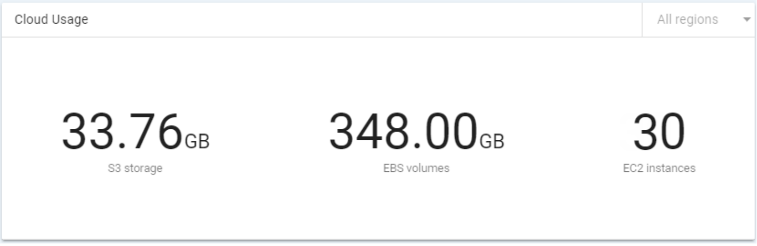

33.76GB S3 storage

348.00GB EBS volumes

30 EC2 instances

NOTE: The information displayed varies depending on the cloud provider environment and the operating mode.

To filter the cloud usage based on region, click the down-arrow in the upper right of the **Cloud Usage** pane and select a specific region.

On-premises assets and storage information pane

The **On-premises assets and storage information** pane of the CDRS dashboard identifies the number of on-premises VMs that are protected by the RecoverPoint for VMs cloud solution. It also identifies the amount of on-premises storage that Cloud DR is protecting.

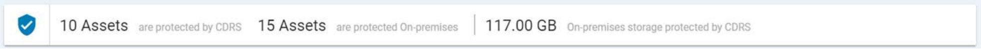

10 Assets are protected by CDRS 15 Assets are protected On-premises | 117.00 GB On-premises storage protected by CDRS

Events pane

The **Events** pane of the CDRS dashboard provides a summary of system events.

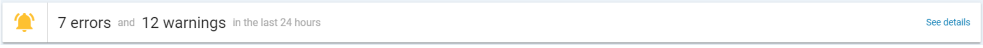

7 errors and 12 warnings in the last 24 hours

To view event details, click **See Details** in the **Events** pane, or select **System** > **Events** from the navigation pane of the CDRS dashboard. Events on page 65 provides information about Events details.

Navigation pane

The Cloud DR Server navigation pane provides links to the various pages of the interface.

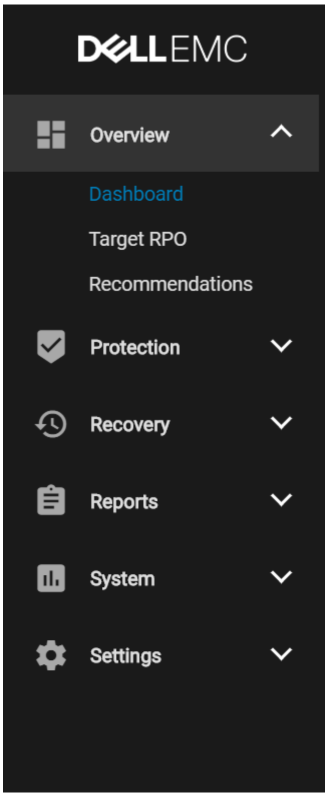

The following sections describe the pages that you access through the navigation pane. You can also access many of these pages through the dashboard.

SLA Compliance pane

The **SLA Compliance** pane of the CDRS dashboard provides a summary of the compliance of protected assets with the service level agreements (SLAs) that were established in the backup software when backup policies were configured for protection.

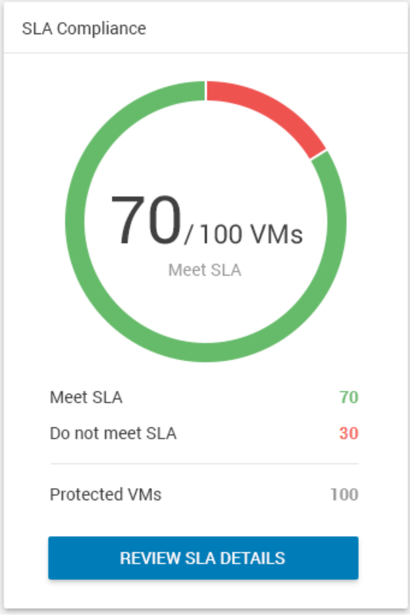

Meet SLA	70
Do not meet SLA	30
Protected VMs	100

To view SLA compliance details, click **Review SLA Details** in the **SLA Compliance** pane, or select **Overview** > **Target RPO** from the navigation pane of the CDRS dashboard. SLA Compliance page on page 59 provides information about SLA compliance of the protected assets.

Recommendations pane

The recommendations pane provides a summary of recommendations that are based on their severity: high, medium, or low.

To view greater details, click **See All**. The resulting list provides a description of each recommendation.

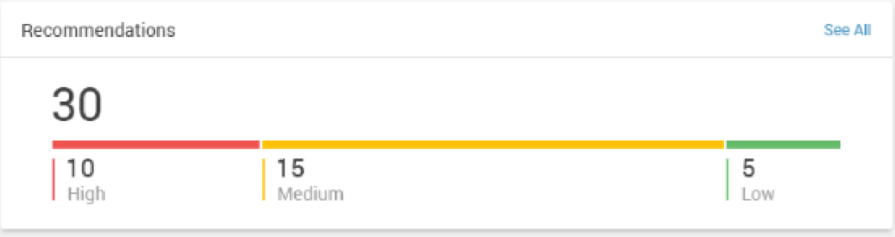

Figure 10. Recommendations pane

Recovery Activities pane

The **Recovery Activities** pane of the CDRS dashboard displays information about current running recovery activities, which include DR test and failover.

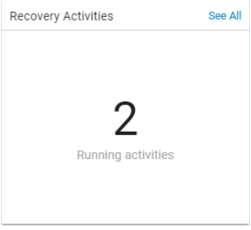

For more information about recovery activities, click **See All** in the **Recovery Activities** pane, or select **Recovery** > **DR Activities** from the navigation pane of the CDRS dashboard to open the **DR Activities** page. DR Activities page on page 61 contains information about the **DR Activities** page.

SLA Compliance page

The **SLA Compliance** page provides details about the compliance of protected assets with the service level agreements (SLAs) that were established when policies were configured for protection.

This compliance represents the Recovery Point Objective (RPO) that is defined in the RecoverPoint for VMs **RPO** setting of each cloud copy.

Access the **SLA Compliance** page from the dashboard by clicking **Review SLA Details** in the **SLA Compliance** pane. or from the navigation pane by selecting **Overview** > **Target RPO**

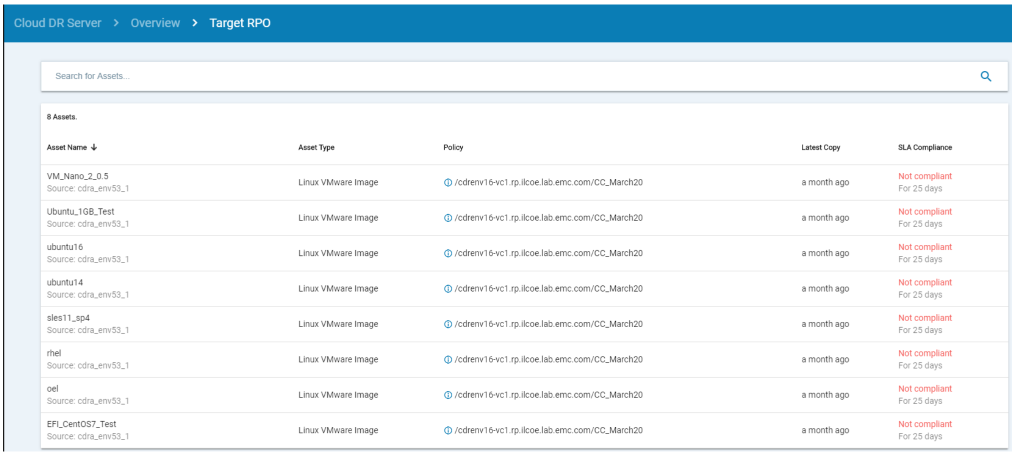

8 Assets.

Asset Name ↓	Asset Type	Policy	Latest Copy	SLA Compliance
VM_Nano_2_0.5 Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
Ubuntu_1GB_Test Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
ubuntu16 Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
ubuntu14 Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
sles11_sp4 Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
rhel Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
oel Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days
EFI_CentOS7_Test Source: cdra_env53_1	Linux VMware Image	ⓘ /cdrenv16-vc1.rp.ilcoe.lab.emc.com/CC_March20	a month ago	Not compliant For 25 days

Noncompliant protected assets are at the top of the list, with the most severe type listed first. You can search the list by asset name by using the search bar at the top of the page.

The SLA Compliance page provides information about compliance. Changes cannot be made in this page.

Asset Recovery page

The **Asset Recovery** page provides a list of protected VMs that you can test or fail over.

You access the **Asset Recovery** page from the navigation pane by selecting **Recovery** > **Asset Recovery**.

From the **Asset Recovery** page, you can search for VMs to recover, select a VM to test or failover, or recover to a specific vCenter (if previously enabled).

When you select a VM from the list, buttons appear at the top of the dialog box to enable DR actions for you to perform.

Recover assets to a vCenter

If you enabled recovery to at least one vCenter for at least one on-premises source, the **Asset Recovery** page is displayed. When you select the asset, an additional action button is displayed: **RECOVER TO VCENTER**.

The **RECOVER TO VCENTER** button displays only when:

- The selected VM contains a copy in the cloud
- At least one recovery-enabled vCenter is available

When you click **RECOVER TO VCENTER**, you are prompted to select a copy, a failover target, and configure other settings before starting the failover.

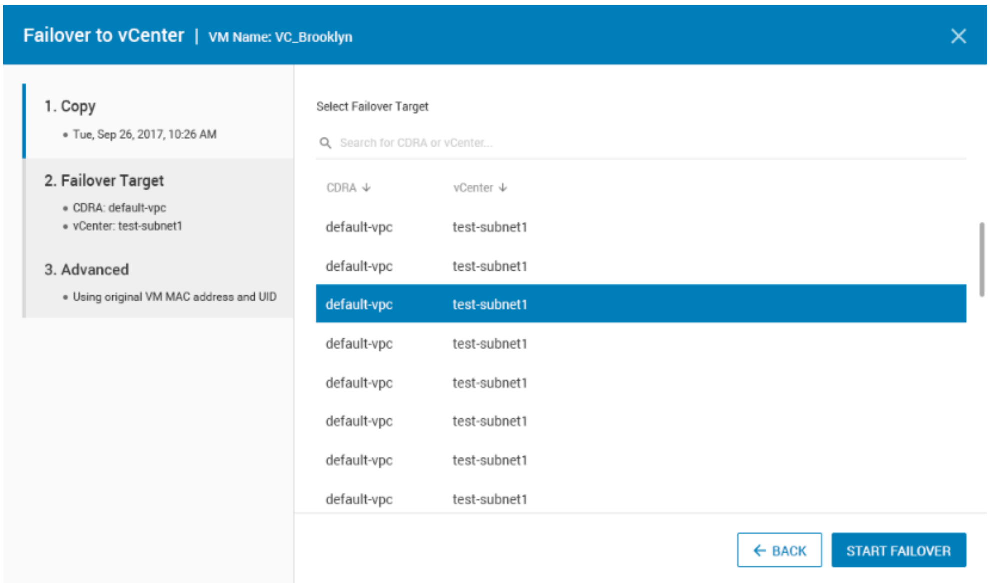

Failover to vCenter | VM Name: VC_Brooklyn

1. Copy
- Tue, Sep 26, 2017, 10:26 AM

2. Failover Target
- CDRA: default-vpc
- vCenter: test-subnet1

3. Advanced
- Using original VM MAC address and UID

Select Failover Target

Search for CDRA or vCenter...

CDRA ↓	vCenter ↓
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1
default-vpc	test-subnet1

Figure 11. Failover to vCenter

DR Activities page

The **DR Activities** page displays recovery activities for DR test and failover and enables you to promote DR tests to failover and end DR tests.

Access the **DR Activities** page from the dashboard by clicking **See All** in the **Recovery Activities** pane, or from the navigation pane by selecting **Recovery** > **DR Activities**.

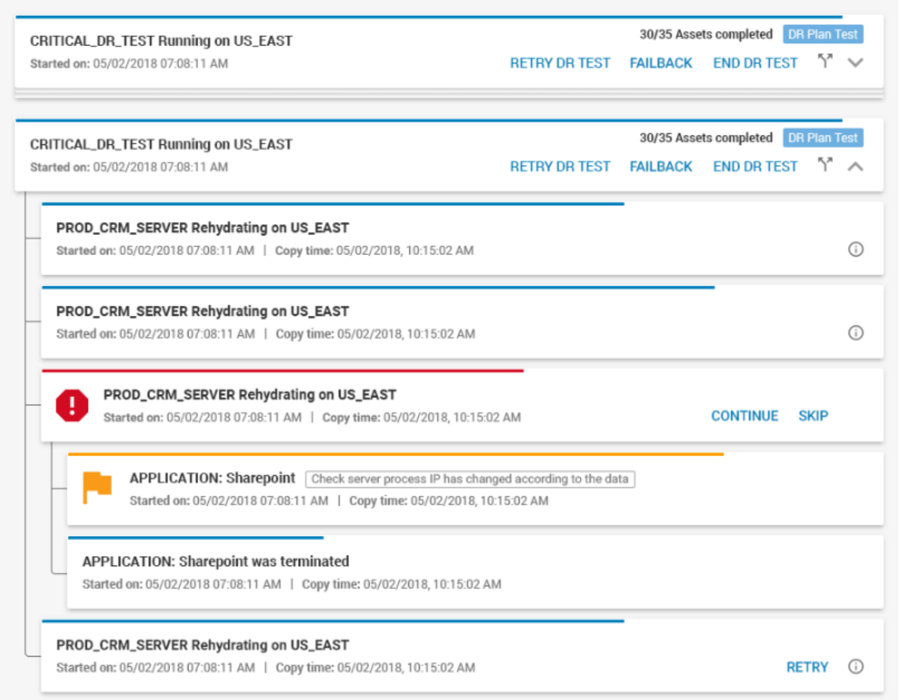

CRITICAL_DR_TEST Running on US_EAST
Started on: 05/02/2018 07:08:11 AM
30/35 Assets completed DR Plan Test

CRITICAL_DR_TEST Running on US_EAST
Started on: 05/02/2018 07:08:11 AM
30/35 Assets completed DR Plan Test

PROD_CRM_SERVER Rehydrating on US_EAST
Started on: 05/02/2018 07:08:11 AM | Copy time: 05/02/2018, 10:15:02 AM

PROD_CRM_SERVER Rehydrating on US_EAST
Started on: 05/02/2018 07:08:11 AM | Copy time: 05/02/2018, 10:15:02 AM

PROD_CRM_SERVER Rehydrating on US_EAST
Started on: 05/02/2018 07:08:11 AM | Copy time: 05/02/2018, 10:15:02 AM

APPLICATION: Sharepoint Check server process IP has changed according to the data
Started on: 05/02/2018 07:08:11 AM | Copy time: 05/02/2018, 10:15:02 AM

APPLICATION: Sharepoint was terminated
Started on: 05/02/2018 07:08:11 AM | Copy time: 05/02/2018, 10:15:02 AM

PROD_CRM_SERVER Rehydrating on US_EAST
Started on: 05/02/2018 07:08:11 AM | Copy time: 05/02/2018, 10:15:02 AM

Searching for DR activities

To search the list of DR activities by name, enter the asset name in the search bar at the top of the page and click the magnifying glass icon. You can also click the filter (≡) icon to select filters to include in the search parameters, including the activity status, activity type, region, and creation time of the DR activity. When you identify the search filters, they are displayed below the search pane. To clear the filters from the search, click **Clear Filters**.

Promoting a failover to failback

From the **DR Activities** page, you can also select a VM or DR plan in a failover state and fail it back to an on-premises vCenter server. When promoting a single VM to failback, you can change the network and security group. However, this action is not possible when promoting a DR plan.

Ending recovery instances from the cloud provider console

You can also terminate a recovery instance from the cloud provider console. When you terminate the recovery instance, the CDRS **DR Activities** page indicates an **Instance Terminated** status.

DR activity statuses

Each DR activity (test, failover, or failback) can have one of several statuses that indicates the progress of the activity.

Table 1 provides a definition and example of each DR activity status.

Table 8. DR activity statuses

DR activity status	Definition
Successfully running	The operation is complete. Disaster recovery is now active. The recovered cloud instance is now available.
Failed	The DR activity failed. The recovered cloud instance is not available. The user may retry the operation.
In progress	DR activity was started and is underway. This status is displayed from the time the DR activity was activated until the operation is complete.
Ending	The "End" operation has been activated. For the test or failover activity, the recovered cloud instance is being terminated.
Successfully completed	DR activity has ended.
Partially successful	The DR plan activity includes successful and failed VMs. This status is relevant only for DR plans.

DR activity states for AWS environments

The **DR Activities** page enables you to monitor the progress of ongoing activity states for DR tests and failovers.

You may notice system messages that indicate the current state of an activity while it is in progress. Table 9. Ongoing activity states for AWS environments on page 63 describes the activity states.

Table 9. Ongoing activity states for AWS environments

State	Description
Rehydrating	When you start a recovery, a temporary **Restore Service** instance is created for each region on which the CDRS must perform recovery. In this state, the **Restore Service** instance constructs the VMDK file from raw data chunks that are stored in Cloud DR target. The **Restore Service** instances are created in a private subnet, in a separate VPC. The **Restore Service** instances automatically terminate after 10 minutes of idle time.
Converting	When the **Restore Service** instance completes rehydration of the VMDK file, CDRS converts the file into an AMI.
Launching	When conversion is complete, CDRS launches a cloud instance that is based on the AMI.
Running	When the launch completes successfully, the restored VM is running. This state is the final step of the recovery.

Each step in this process can take several minutes to complete.

View recovery details

The **DR Activities** page enables you to view detailed information about the assets that are listed.

Steps

1. For any asset listed in the **DR Activities** page, click the information icon ⓘ.

   ⓘ **NOTE: For DR plans, you must first click the down-arrow icon ⌄ to access the individual assets.**

   A detailed list of information about the asset is displayed. For example:

   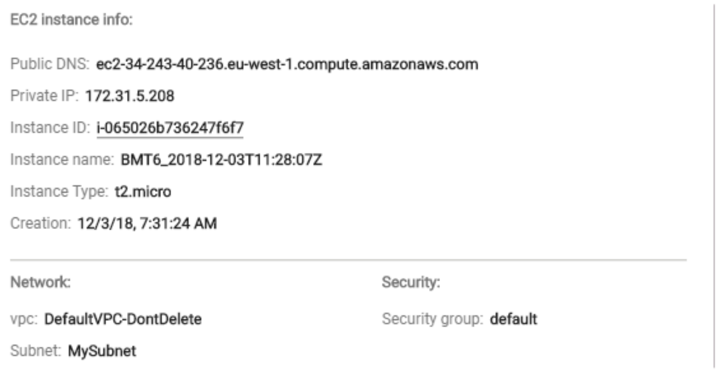

   EC2 instance info:

   Public DNS: ec2-34-243-40-236.eu-west-1.compute.amazonaws.com

   Private IP: 172.31.5.208

   Instance ID: i-065026b736247f6f7

   Instance name: BMT6_2018-12-03T11:28:07Z

   Instance Type: t2.micro

   Creation: 12/3/18, 7:31:24 AM

   Network:

   vpc: DefaultVPC-DontDelete

   Subnet: MySubnet

   Security:

   Security group: default

   Original machine info:

   CPU: 1 CPUs

   RAM: 256.00 MB

   DISKS: 6.00 GB

2. To collapse the detailed information view, click the information icon ⓘ again.

Reports

CDRS enables you to generate reports that help you to monitor resources in the Cloud DR solution.

Protected Copies Cloud Consumption

CDRS enables you to define the reporting parameters for cloud consumption. You select **Reports** > **Generate Report**, and then define parameters.

- Region
- Consumption for copies:
  - Only asset copies
  - Only rapid recovery copies
  - Asset copies and rapid recovery copies
- Time interval:
  - Last week
  - Last month
  - Last year

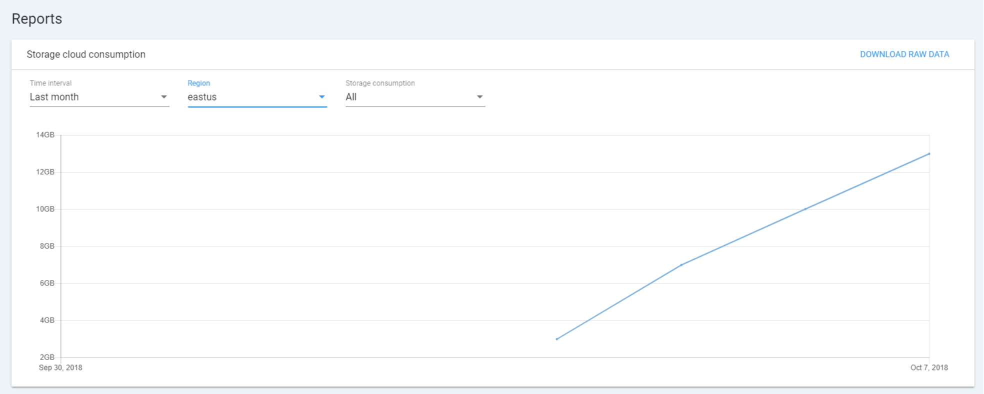

Reports

Storage cloud consumption

DOWNLOAD RAW DATA

Time interval: Last month

Region: eastus

Storage consumption: All

NOTE: **The storage consumption is calculated once a day.**

Click the **DOWNLOAD RAW DATA** link (upper right) to retrieve the report in CSV format.

DR Activities

You can also generate a report to show the DR activities based on status, type, region, and selected date range:

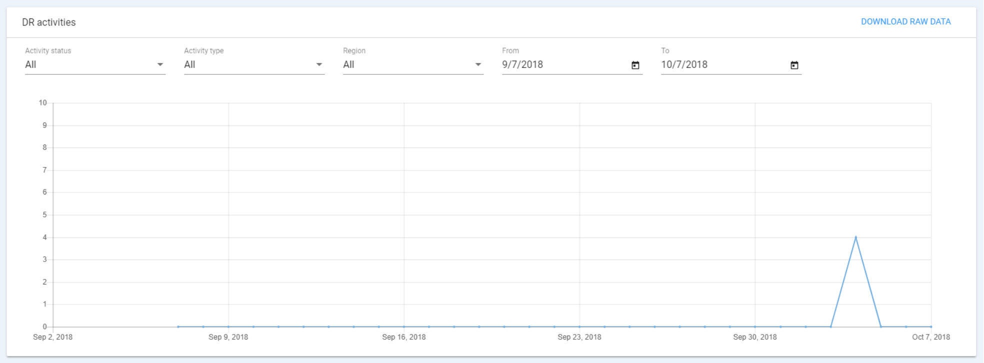

DR activities

DOWNLOAD RAW DATA

Activity status: All

Activity type: All

Region: All

From: 9/7/2018

To: 10/7/2018

Events

Use the **Events** page, which is accessible by selecting **System** > **Events** from the navigation pane, to review system events by date, severity, title, and category.

To view details about an event, click the down-arrow icon (**v**) to the right of the event. A details pane provides the event ID and detailed information about the event.

To search the list of events for various event types, type a search string in the search bar at the top of the page and click the magnifying glass icon. For example, to limit the event list to only those events that contain the word "Failover," type `Failover` in the search bar and click the magnifying glass icon.

You can also click the filter (⏷) icon to select filters to include in the search parameters, including security level, category, the Cloud DR Add-on, and event creation time. The search filters you identify appear below the search pane. To clear the filters from the search, click **Clear Filters**.

Export events to Syslog

You can configure CDRS to export events to a syslog server where they can be viewed using external monitoring systems.

Steps

1. In the CDRS user interface navigation tree, select **System** > **Syslog**.
2. Select **Add Syslog Server** and provide the following information about the Syslog server:
   - IP or hostname.
   - Transfer protocol.

- Port number.
- Facility name.

3. To return to the **Syslog** page, click **Save & Connect**.
4. To verify connection to the syslog server, click **Test Syslog**.
5. Click **Add Event Filter** and specify the following information on the **Defined Events** window:
   - Filter name.
   - One or more categories to send to syslog.
   - One or more severity levels to send to syslog.
6. To return to the **Syslog** page, click **Add**.

   The **Enable Syslog log transfer** switch is automatically toggled to on.

Results

The events data is exported to the syslog server. You can disable these exports by toggling off **Enable Syslog log transfer**.

System Health

The **Health** page, which is accessed by clicking **See All** in the **System Health** pane of the dashboard, or selecting **System** > **Health** from the navigation pane, provides information about the health of the Cloud DR implementation. Cloud-based and on-premises components are listed.

To view details about a component that is listed in this screen, click the down-arrow icon (**v**) to the right of the component. A details pane provides information about the status of the component. Component issues are identified so corrective action can be taken.

Identifying your CDRA

The **Registered Components** page, which is accessible by selecting **System** > **Registered Components** from the navigation pane, enables you to view registered components and unregister them.

A registered component includes name, IP address, version, and on-premises source (CDRA or vRPA). Click **UNREGISTER** next to the component that you want to unregister.

NOTE: In the RecoverPoint for VMs cloud solution, the vRPA is also registered as a CDRA.

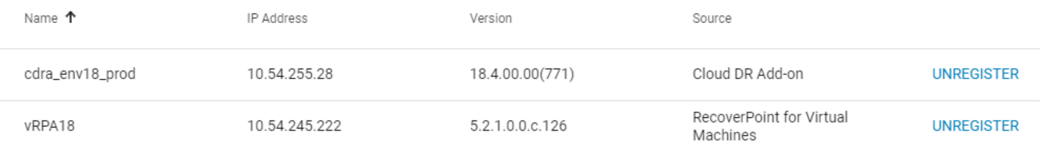

Name ↑	IP Address	Version	Source	
cdra_env18_prod	10.54.255.28	18.4.00.00(771)	Cloud DR Add-on	UNREGISTER
vRPA18	10.54.245.222	5.2.1.0.0.c.126	RecoverPoint for Virtual Machines	UNREGISTER

8

Recovery orchestration

The RecoverPoint for VMs cloud solution offers the following features for the orchestration of VM recovery:

Topics:

Create a tag

You can create one or more tags to enable tagged-based resources management. Examples of use cases include Cloud Snapshot Manager (CSM) tag-based policy protection, applying bulk updates or security patches, upgrading applications, opening or closing ports to network traffic, collecting specific logs, or monitoring data from recovered instances.

About this task

An important use case for tag-based management is protection during failover operation. You can create tags in CDRS and leverage CSM to protect tagged workloads that are being failed over to the cloud. Read more about tag-based management with CSM here:

https://support.emc.com/docu86938_Cloud-Snapshot-Manager:-Manage-Copy-Sprawl-in-Amazon-Web-Services-.pdf?language=en_US

Steps

1. From the Cloud DR Server UI, select **Settings** > **Tags**, and then click the **Create Tag** button.

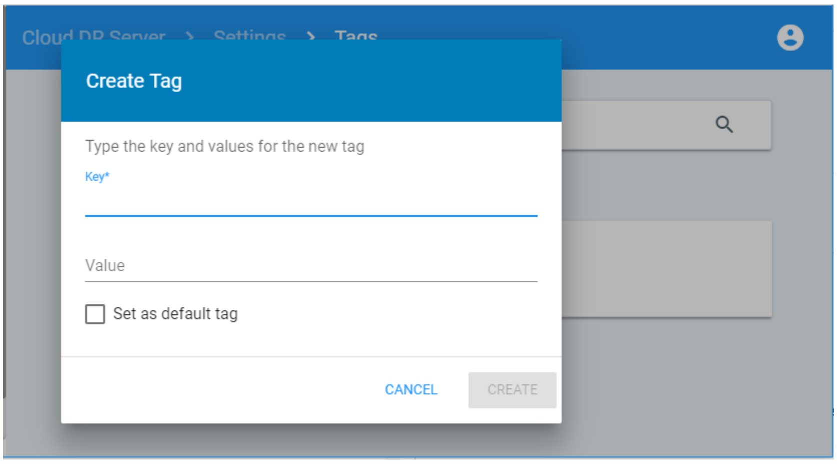

Cloud DR Server > Settings > Tags
Create Tag
Type the key and values for the new tag
Key*
Value
☐ Set as default tag
CANCEL
CREATE

2. Enter the key and values for the new tag. Optionally, set the tag as a default. Then click **Create**.
3. Repeat steps above to create additional tags.

ⓘ **NOTE: Once you create the tags, you can apply them whenever you run a test or failover.**

Create rapid recovery copies for protected assets

You can accelerate the recovery process ahead of time by creating rapid recovery copies for protected assets. Creating a rapid recovery copy reduces the RTO for a protected asset but consumes additional cloud resources and incurs additional costs.

About this task

Creating a rapid recovery copy starts the rehydration process and converts the VMDK files to an Amazon Machine Image (AMI). The recovery process (test or failover) then launches the recovered instance from the AMI.

Perform this procedure when a copy is available in the cloud storage.

Rapid recovery is supported for the VM and its associated applications. To enable rapid recovery for an application, apply rapid recovery to its associated VM.

NOTE: Failover of rapid recovery copies to a vCenter or VMware Cloud is not supported.

Steps

1. In the CDRS user interface, select **Protection > Asset Protection** in the navigation pane.

   The existing protected assets are displayed in the right pane. The **Rapid Recovery Image** column indicates whether the asset is enabled for rapid recovery.
2. Select one or more VMs and click **Set Rapid Recovery Image**.
3. In the **Set Rapid Recovery Image** dialog box, select the number of rapid recovery copies that you want to keep (from 1 to 5), and then click **Set**.

   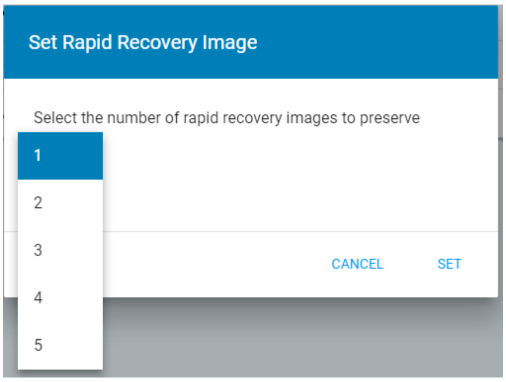

   Set Rapid Recovery Image

   Select the number of rapid recovery images to preserve

   NOTE: Configuring more than one rapid recovery copy for selected VMs enables you to quickly recover to an older point in time in case the latest point-in-time copy cannot be used because of inconsistent or corrupt data.

Results

- The CDRS creates the rapid recovery copy and removes the oldest machine image to maintain the number of copies that you configured.
- You can verify the results by reviewing the **Rapid Recovery Image** column where the number of copies is indicated. The ⏩ icon is displayed in some CDRS windows and designates a copy that is enabled for rapid recovery.

Next steps

- You can disable rapid recovery for an asset by selecting it and clicking **Disable Rapid Recovery Images**.
- You can set the minimal time interval during which rapid recovery copies are not created. See Set rapid recovery interval on page 69.

Set rapid recovery interval

You can set the minimal time interval during which rapid recovery copies are not created. The minimum time interval is 6 hours, and the maximum is 24 hours. The default setting is 12 hours.

Prerequisites

The rapid recovery process can be performed only when a copy is available in the cloud storage after rapid recovery is enabled. Rapid recovery does not occur for copies that are uploaded before rapid recovery is enabled.

Steps

1. In the CDRS user interface, select **Settings** > **General** in the navigation pane.
2. In the **Set Rapid Recovery Interval** section, move the slider to select the time interval during which rapid recovery copies are not created.
   The change takes effect immediately.

Results

The CDRS runs the rapid recovery process that is based on the time interval that you set.

Example

For example, if the time interval is set to 10 hours, then no rapid recovery copy is created within 10 hours of the previous rapid recovery copy.

DR plan activities

A disaster recovery (DR) plan is a collection of assets that enables you to define run book recovery plans, including batch operations on multiple assets, network and security group association, VM boot order definition, and selection of cloud instance type. You can manage, recover, and fail back DR plans through the CDRS. If you want to manage each asset separately, you can split the DR plan into its individual assets.

This section provides the basic procedures for DR plan activities.

DR plans

A disaster recovery (DR) plan is a collection of assets (VMs and their applications) that enables you to define run book recovery plans, including batch operations on multiple assets, network and security group association, VM boot order definition, and selection of cloud instance type.

A DR plan is associated with a single region and on-premises source (CDRA or vRPA). You can add to the plan only those assets that are protected by the designated source (CDRA or vRPA) and are in the designated region.

The assets that you add to the DR plan are called DR plan members. If required, you can add the same asset to multiple DR plans. For example, you might want to create several DR plans to test various DR scenarios. You can also create a master DR plan that contains all the assets on premises.

For each VM in the DR plan, you can specify a startup priority, called a boot order, from 1 to 5, where a lower number represents a higher priority. For example, a VM with a boot order of 1 begins recovery before a VM with a boot order of 2 to 5. All VMs with the same boot order begin recovery at approximately the same time (actual start times may vary depending on when each VM recovery operation ends).

ⓘ **NOTE: Boot order, network, security group, and cloud instance type apply to VMs, not to individual applications.**

You can test, fail over, or fail back a DR plan in the same way that you might perform those operations on a single asset. There are minor differences in the workflows.

When you test or fail over a DR plan, that operation is applied to all the assets contained in the plan. If one asset in the plan fails, the operation continues on the other assets in the plan (the default behavior). You may choose to retry the operation for the failed asset while the DR plan operation continues. A partially successful DR test means that the batch operation continues even when one or more assets in the DR plan encounter a test failure. Optionally, you may configure the DR plan to fail when any asset in the plan fails by enabling the **Fail on error** option.

When a DR plan is partially successful (that is, recovery of some assets has succeeded while others have failed), the user has three options:

- Retry - This action retries the operation only for the failed assets. Cloud instances that are already recovered remain available.

- End test or failover - This action terminates the cloud instances of successfully recovered VMs.

  ⓘ **NOTE: Ending a failback operation for a DR plan only closes the failback card.**

- Split - This action splits a partially successful DR plan into its individual members so you can manage each asset separately.

Depending on the number of members in a DR plan, it may take some time for the plan operation to complete. One convenient feature of a DR plan is that when you run a DR plan, you can immediately begin editing the plan or even delete it without affecting the completion of the original plan.

Create a DR plan

You can create a DR plan for a specific region/location and CDRA. Then you can add assets to the DR plan.

Prerequisites

You can add to the DR plan only those assets that are protected by the selected on-premises source in the designated region.

Steps

1. From the CDRS user interface, select **Protection** > **DR Plans**.

   The **Select or Create a New Plan** window is displayed.
2. To create a DR plan, click **Create Plan**.
3. In the **Plan Details** tab, enter a unique name for the DR plan and select an on-premises source, and location.

   ⓘ **NOTE: You cannot edit the on-premises source name or region after you select members for the plan.**

4. If you want the DR plan to fail when any asset in the plan fails, select the **Fail plan on error** checkbox. If you want the DR plan to continue running when one or more assets fail, clear the checkbox.
5. Select a default network, default security group, and, if you are using tags, a tag.
6. In the **Plan Members** tab, click **Add Members**.

   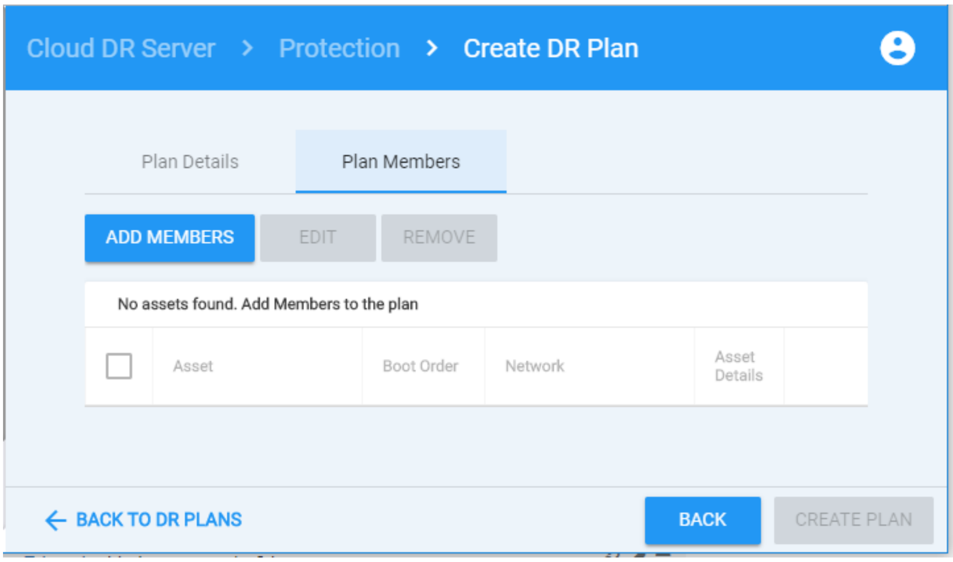

   No assets found. Add Members to the plan

   The **Add Members** dialog box displays a list of assets.
7. In the **Add Members** dialog box, select the checkbox for each asset that you want to add to the DR plan, and then click **Add**.
8. To change the asset boot order, default network, default security group, virtual machine type, tags, or private IP address selection, click the **Edit** button for the asset. Make the change, then click **Apply**.
9. Review the list of assets that you added to the new DR plan. If you require additional changes, select one or more of the assets to edit (by using the **Edit** button) or remove (by using the **Remove** button).
10. When you are satisfied with the DR plan, its assets, and properties, click **Create Plan**.

Results

The DR plan is created and may be used for testing or failover.

Test or fail over a DR plan to AWS cloud

To verify that the operations of a DR plan work as expected, you test the DR plan. To start a failover of the assets in the DR plan, you fail over the DR plan. This procedure describes how to test or fail over a DR plan by using the Cloud DR Server interface.

Prerequisites

To perform a test or failover of a DR plan, you must have instances of virtual machines that are protected in the cloud.

About this task

To ensure a successful failover and prepare for a disaster, best practices entail testing various disaster recovery scenarios.

When an operational error or disaster occurs on premises, you can fail over a DR plan to the cloud. When the on-premise issue is resolved, you may fail back the DR plan to the on-premises environment.

ⓘ **NOTE: When you fail over a DR plan, CDRS fails over the assets in the DR plan according to the VM boot order.**

Steps

1. In the CDRS user interface, select **Recovery > Plan Recovery**

   The **Plan Recovery** page displays a list of DR plans on which recovery activities can be performed.
2. Select the DR plan that you want to recover, and click **DR Test** to test the plan or **Failover** to fail it over to the cloud.
   A dialog box is displayed and prompts you to select copies. Any bookmarks that are applied in RecoverPoint for VMs are displayed. Corrupted copies are clearly identified, and you are prevented from selecting them.
3. Select one of the copy options:

Option	Description
Latest available copies	Recovery uses the latest copies of the asset in the recovery operation.
Select a point in time	Recovery uses asset copies that are based on the time, date, and selection that you specify.

   If you configured the DR plan to fail on error, the plan fails if the VM copy is not available.
4. Click **Next**.

   A dialog box is displayed and prompts you to review the list of copies and their status.

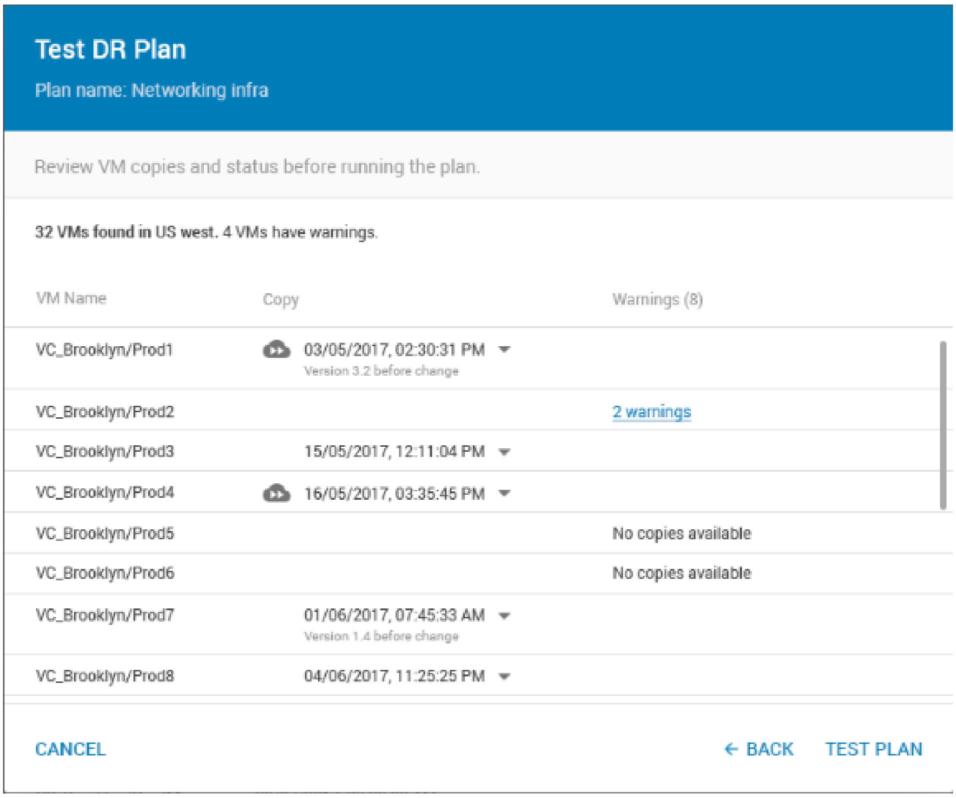

Test DR Plan

Plan name: Networking infra

Review VM copies and status before running the plan.

32 VMs found in US west. 4 VMs have warnings.

VM Name	Copy	Warnings (8)
VC_Brooklyn/Prod1	☁ 03/05/2017, 02:30:31 PM ▾ Version 3.2 before change	
VC_Brooklyn/Prod2		2 warnings
VC_Brooklyn/Prod3	15/05/2017, 12:11:04 PM ▾	
VC_Brooklyn/Prod4	☁ 16/05/2017, 03:35:45 PM ▾	
VC_Brooklyn/Prod5		No copies available
VC_Brooklyn/Prod6		No copies available
VC_Brooklyn/Prod7	01/06/2017, 07:45:33 AM ▾ Version 1.4 before change	
VC_Brooklyn/Prod8	04/06/2017, 11:25:25 PM ▾	

5. If you are:
    - Unsatisfied with the copy selections, make the necessary changes before continuing.
    - Satisfied with the copy selections, continue with a test or failover of the DR plan.

Results

Depending on the selection, the Cloud DR Server starts the test or failover of the DR plan.

Edit a DR plan

You can edit the properties of a DR plan except for the region and the on-premises source.

About this task

If the plan is active (running or in failover or test), editing the plan does not affect the active DR plan.

Steps

1. From the CDRS user interface, select **Protection** > **DR Plans**.

   The **Select or Create a New Plan** window is displayed (not shown).

2. Click the edit icon ✎ for the plan that you want to edit.

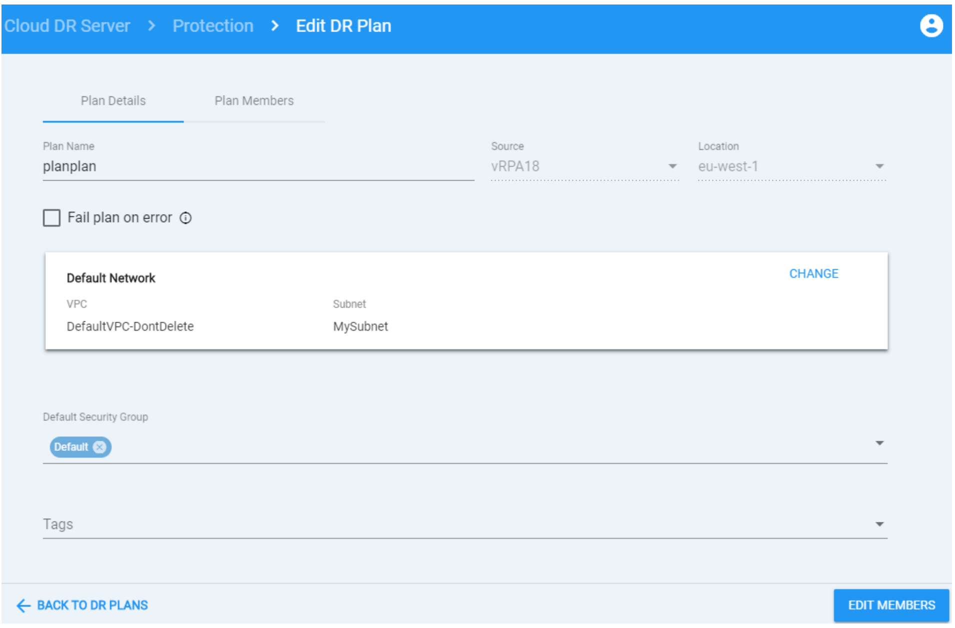

Cloud DR Server > Protection > Edit DR Plan

Plan Details | Plan Members

Plan Name: planplan

Source: vRPA18

Location: eu-west-1

□ Fail plan on error

Default Network

VPC: DefaultVPC-DontDelete

Subnet: MySubnet

Default Security Group: Default

Tags

The **Edit DR Plan** window is displayed.

3. If required, change the **Fail plan on error** setting.
4. If you want to change the default network, click **CHANGE** and pick a different network.
5. If required, pick a different security group.
6. If required, select a different tag.
7. If you want to change the members that belong to the DR plan or edit the settings for any selected member:
   a. Click the **EDIT MEMBERS** button.

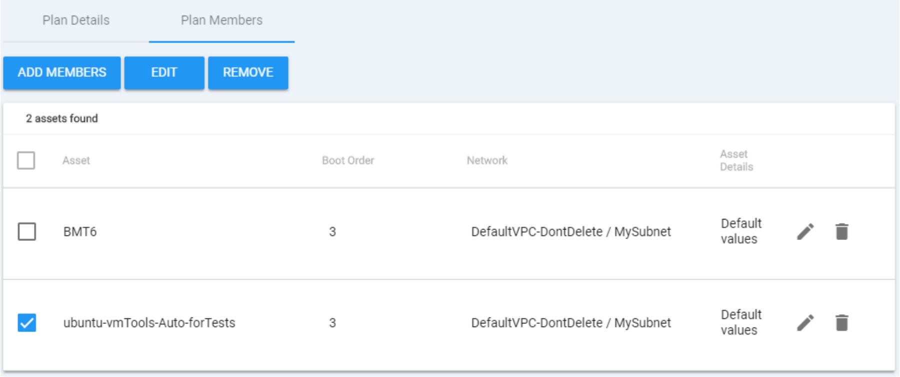

Plan Details | Plan Members

2 assets found

□	Asset	Boot Order	Network	Asset Details
□	BMT6	3	DefaultVPC-DontDelete / MySubnet	Default values
☑	ubuntu-vmTools-Auto-forTests	3	DefaultVPC-DontDelete / MySubnet	Default values

The **Plan Members** window is displayed.

   b. Select one or more members of the plan.
   c. If you want to remove one or more selected members, click the **REMOVE** button.
   d. If you want to edit settings for one or more selected plan members, click the **EDIT** button.
      The **Edit Member** dialog box is displayed.
   e. In the **Network** tab of the **Edit Member** dialog, if required, change the boot order, default network, and default security group of the member.
   f. In the **Advanced** tab, if required, change the virtual machine type, tags, or the private IP address checkbox.

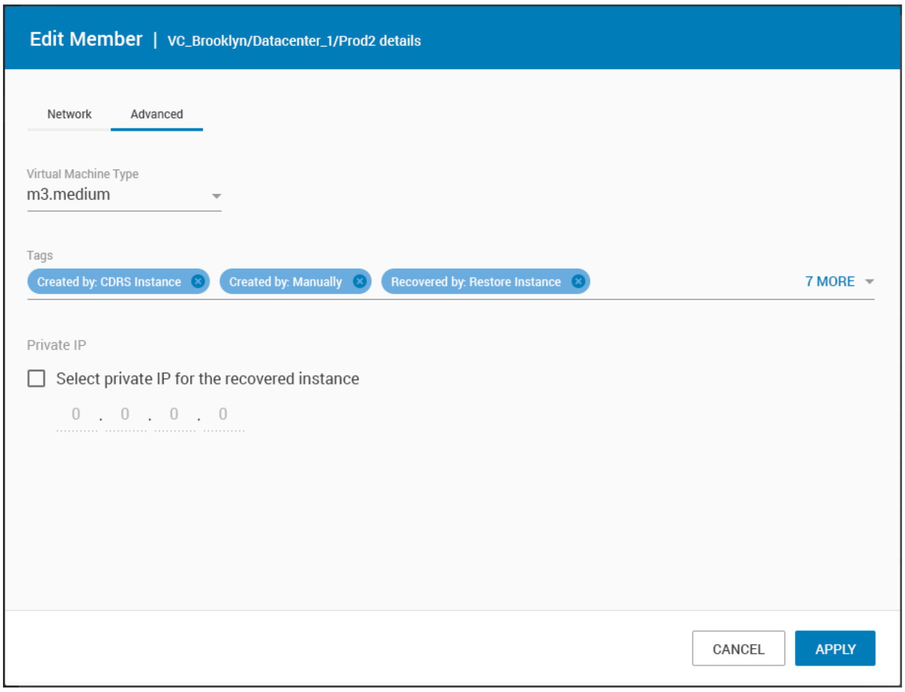

Edit Member | VC_Brooklyn/Datacenter_1/Prod2 details

Network | Advanced

Virtual Machine Type: m3.medium

Tags: Created by: CDRS Instance, Created by: Manually, Recovered by: Restore Instance, 7 MORE

Private IP

□ Select private IP for the recovered instance

0 . 0 . 0 . 0

CANCEL APPLY

g. Click **APPLY** to apply changes to the edited member.
h. In the **Edit DR Plan** window, click **APPLY** to apply changes to the edited DR plan.

Results

The DR plan is updated and may be used for testing or failover.

Split a DR plan activity

If you want to manage each asset separately, you can split the DR plan.

About this task

In the **DR Activities** window, DR plan activities are organized by card types: DR test cards, DR failover cards, and DR failback cards. If you have a DR plan in test and you split it, the DR test cards are split apart and you can individually end them or promote them to failover. The assets in the DR plan are separated, and the DR plan is removed. When you split apart a DR plan activity, the action is irreversible.

Steps

1. From the CDRS user interface, select **Recovery** > **DR Activities**.
2. Locate the DR plan activity that you want to split.

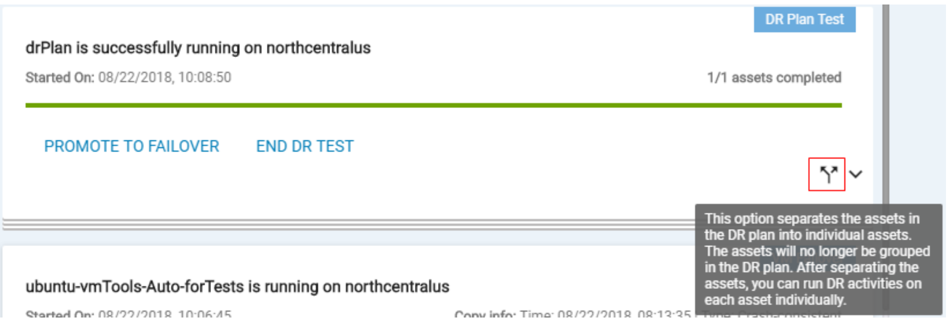

DR Plan Test

drPlan is successfully running on northcentralus

Started On: 08/22/2018, 10:08:50

1/1 assets completed

PROMOTE TO FAILOVER END DR TEST

This option separates the assets in the DR plan into individual assets. The assets will no longer be grouped in the DR plan. After separating the assets, you can run DR activities on each asset individually.

ubuntu-vmTools-Auto-forTests is running on northcentralus

Started On: 08/22/2018, 10:06:45

Copy info: Time: 08/22/2018 08:13:35

3. To split the DR plan into its individual assets, click the ⑂ icon.

Results

The DR plan is split into its individual assets, and the cards in the DR plan activity are split into individual activities.

Delete a DR plan

When you no longer require a DR plan and the VMs it contains, you can delete the plan.

About this task

If the plan is active (running or in failover or test), deleting the plan does not affect the active DR plan.

Steps

1. From the CDRS user interface, select **Protection** > **DR Plans**.

   The **Select or Create a New Plan** window appears.
2. Select a DR plan to delete.
3. To delete the plan, click the delete (trash can) icon for the plan, and confirm the action.

Results

The DR plan is deleted.

A

Deploying the CDRA

Use the procedures in this section to deploy CDRA for fail back from AWS or to recover to a vCenter or VMware cloud on AWS.

Topics:

Deployment architectures

You can deploy a CDRA on-premises or in a VMware Cloud on AWS.

To fail back from AWS to an on-premises vCenter, you will want to deploy an on-premises CDRA.

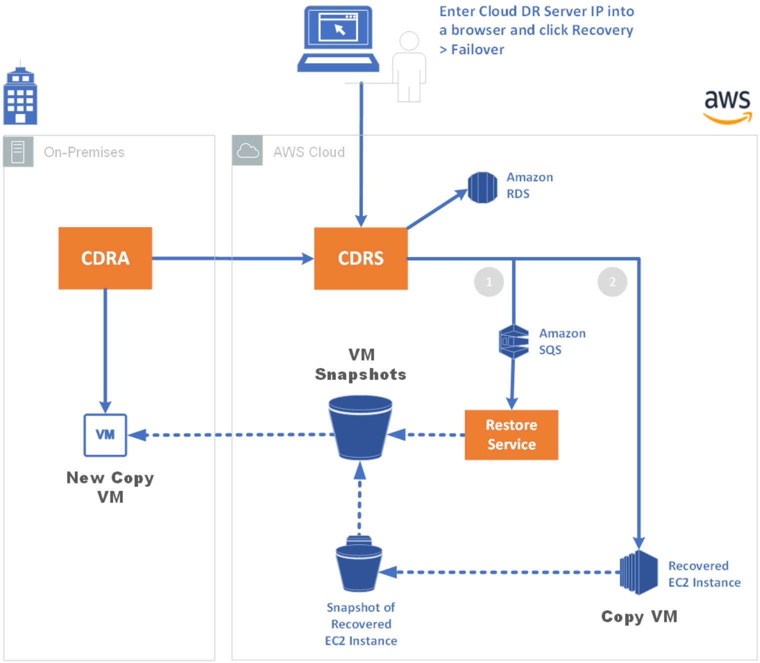

To recover VMs from AWS S3 to a VMware Cloud on AWS, you will want to deploy a CDRA in a VMware Cloud on AWS. This architecture can also be used with an on-premises vCenter.

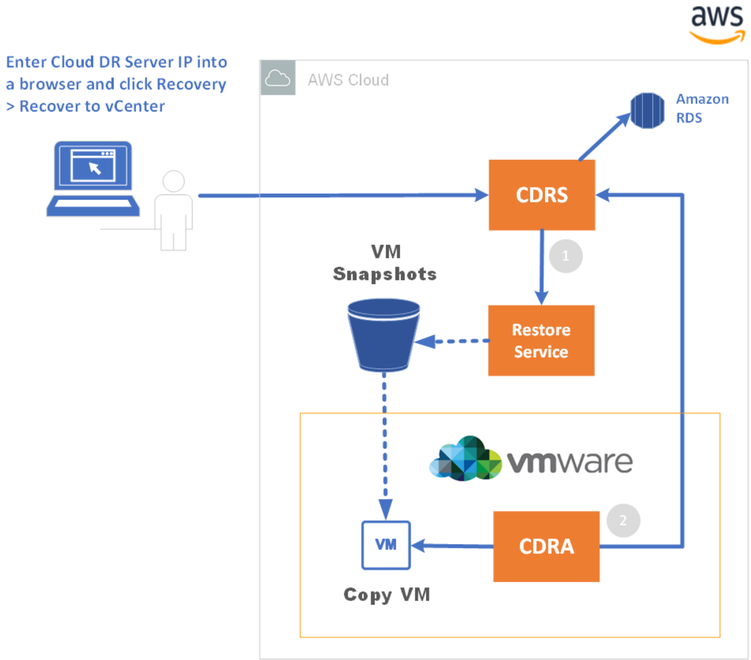
aws
Enter Cloud DR Server IP into a browser and click Recovery > Recover to vCenter
AWS Cloud
Amazon RDS
CDRS
1
VM Snapshots
Restore Service
vmware
2
VM
CDRA
Copy VM

Deploy the CDRA OVA

The Cloud DR Add-on (CDRA) is a Cloud DR component, and it is provided as an OVA deployed on a VMware vCenter Server environment.

The following table lists the required specifications for Cloud DR Add-on VMs.

Table 10. Cloud DR Add-on VM specifications

Component	Required specification
vCPU	4 (2x2)
RAM	4 GB
HDD	16 GB

ⓘ **NOTE: To support recovery operations for production VMs, ensure that each VM has a unique identifier (UID).**

Download the OVA from the link that was provided when you purchased the Cloud DR solution. Use the vSphere client to deploy the OVA in the vSphere environment.

In the network-mapping step, one network interface is required for the CDRA VM. Map the CDRA network interface to a VLAN that provides network access to the cloud.

CDRA supports dual NIC configurations for CDRS deployment. See Configuring the CDRA and deploying the CDRS for more information.

ⓘ **NOTE: After the CDRA is deployed, changing its IP address is not supported.**

Log into the CDRA

You can log in to the CDRA with the username and password.

Steps

1. From a host that has network access to the CDRA virtual appliance, use a browser to connect to the appliance:

   ```
   https://CDRA_hostname
   ```

   Where *CDRA_hostname* is the hostname or IP address of the address that you created when the CDRA was deployed to the vCenter server.

2. In the **Admin username** and **Admin password** fields, enter the username and password that were provided when you purchased the product.

   NOTE:

   - **The default admin password is *admin*.**
   - **Passwords expire based on the specified expiration period. By default, the expiration period is 90 days.**

   If this login is the first login or the password has expired, the **Cloud DR Add-on Change Admin Password** window opens for you to change the password. Passwords must be at least eight characters in length and contain a minimum of three of the following character types:

   - English uppercase: A-Z
   - English lowercase: a-z
   - Numeric character: 0–9
   - Special (non-alphanumeric) characters

   NOTE: If you forget the password, click Forgot password?. Then enter the username and click Send.

   When the admin user account's email address is initially provided or changed, AWS sends a verification email to the email address. This email address must be verified before receiving the password reset email. You can request a new verification email through the AWS console by signing into the console and selecting the US East (N. Virginia) region. Then, open https://console.aws.amazon.com, select Email Addresses, select the email address, and click resend.

Results

The **Cloud DR Add-on** window opens and the **Welcome** page is displayed.

Configuring the CDRA

To begin, click **Configuration** in the navigation pane.

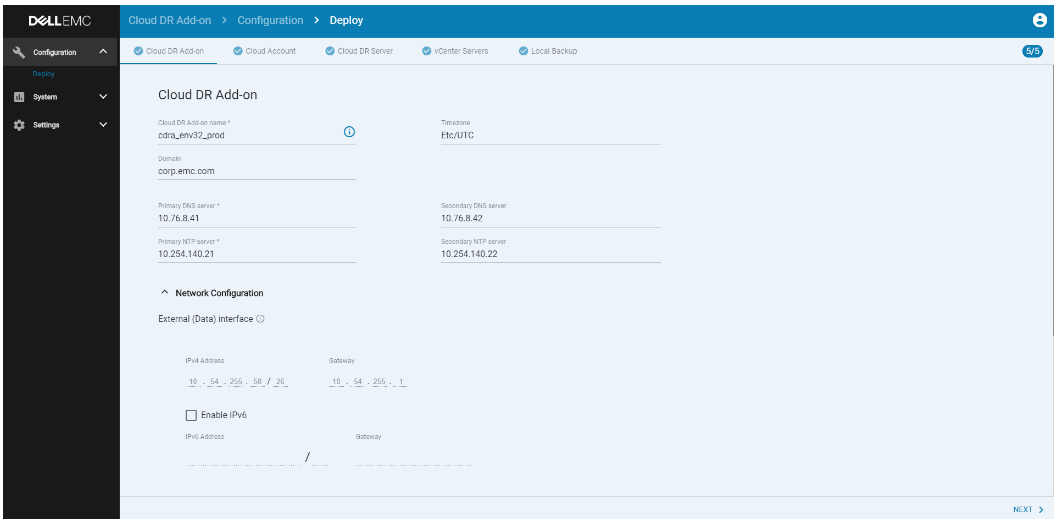

Cloud DR Add-on

Cloud DR Add-on name * cdra_env32_prod

Timezone Etc/UTC

Domain corp.emc.com

Primary DNS server * 10.76.8.41

Secondary DNS server 10.76.8.42

Primary NTP server * 10.254.140.21

Secondary NTP server 10.254.140.22

Network Configuration

External (Data) interface

IPv4 Address 10 . 54 . 255 . 58 / 26

Gateway 10 . 54 . 255 . 1

☐ Enable IPv6

IPv6 Address /

Gateway

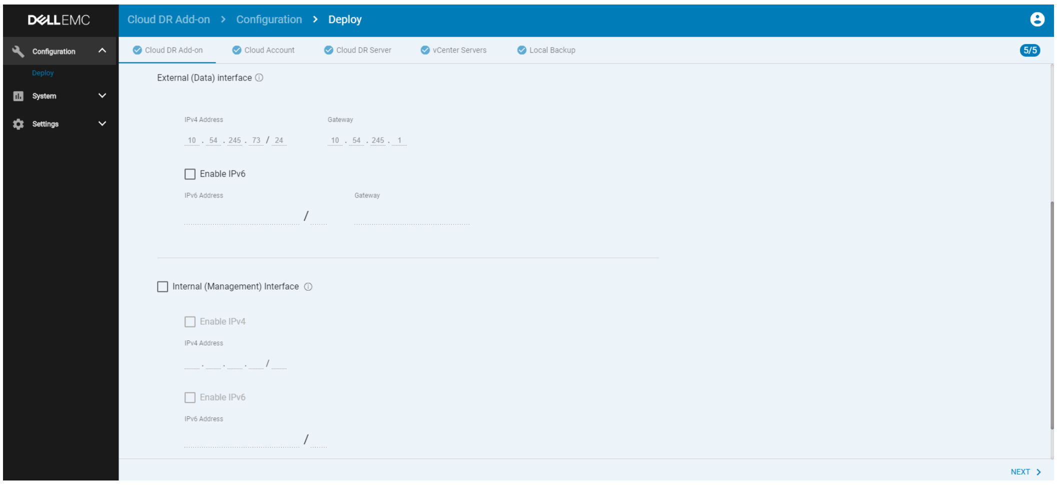

External (Data) interface

IPv4 Address 10 . 54 . 245 . 73 / 24

Gateway 10 . 54 . 245 . 1

☐ Enable IPv6

IPv6 Address /

Gateway

☐ Internal (Management) Interface

☐ Enable IPv4

IPv4 Address . . . /

☐ Enable IPv6

IPv6 Address /

The menu bar (across the top) displays the steps that are required to complete the configuration and deployment process. The Cloud DR solution is fully deployed when you complete these tasks.

Generally, you complete the steps working from left to right. For example, you must connect to the Cloud Account and create Cloud DR targets before you deploy the Cloud DR Server.

CDRA supports single NIC (NIC-0) and dual NIC (NIC-0 and NIC-1) configuration. Dual NIC configuration is the default configuration. The single NIC configuration (NIC-0) is used for both internal and external networks. In single NIC configuration, IPv4 is mandatory whereas IPv6 is optional. There are two types of Dual NIC configurations.

- External (Cloud and Data) NIC-0
  - NIC-0 is used for external network.
  - IPv4 is mandatory.
  - IPv6 is optional.
- Internal - NIC-1
  - NIC-1 is used for internal network.
  - NIC-1 supports either IPv4 and IPv6.
  - NIC-1 supports dual stack configuration where both IPv4 and IPv6 are defined.

Set up the CDRA

To configure networking and other settings for the CDRA, use the **Setup CDRA** page of the Cloud DR Add-on window.

Steps

1. For **Cloud DR Add-on name**, enter a name for the CDRA.
2. Enter the hostname or IP address for the primary and secondary DNS servers.

3. Enter the hostname or IP address for the primary and secondary NTP servers.
4. Select a time zone that is the same as the on-premises RecoverPoint for VMs time zone.
5. Expand the **Network Configuration** section. If CDRA is configured with dual NIC, **External (Data) interface** and **Internal (Management) Interface** are displayed . Only **External (Data) interface** is displayed in single NIC configuration.

   The **External (Data) interface** connects cloud provider to the Data Domain path. The **Internal (Management) Interface** is used for internal components.
6. Enter the IPv4 address and the gateway id in the **External (Data) interface** section.
7. Select the **Enable IPv6** checkbox to configure the IPv6 address and gateway.
8. Select the **Internal (Management) Interface** to enable the **Enable IPv4** and **Enable IPv6** sections.
9. Select the **Enable IPv4** checkbox to configure the IPv4 address.
10. Select the **Enable IPv6** checkbox to configure the IPv6 address.
11. Click **Save**.

Add AWS cloud account

Add the AWS cloud account and connect the CDRA to the account.

Prerequisites

Ensure that you have an AWS account that is already configured before connecting to the cloud account.

Steps

1. Click **Cloud Account** on the menu bar.
   The **Connect to Cloud Account** page is displayed.
2. Click **Add Cloud Account**.
3. In the **Connect to Cloud Provider Account** dialog box, select AWS.

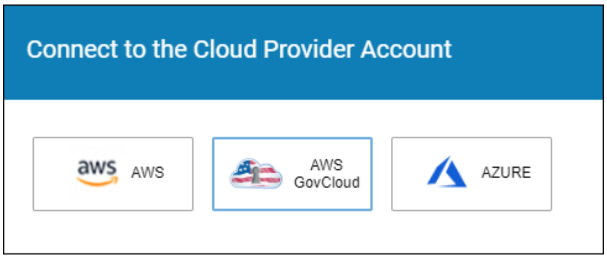

4. In the **Connect to Cloud Provider Account** dialog box, enter the **Access Key ID** and the **Secret Access Key** for the AWS account. http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html provides information about obtaining the access and secret keys.
5. To copy the IAM policy, click **Copy IAM Policy**.

   This action copies to the buffer a JSON version of the minimum AWS user account permissions that are required for Cloud DR implementation. This implementation is then applied to AWS and to set the permissions policy for the appropriate user. Define the AWS IAM policy on page 17 also provides the IAM policy and instructions for creating an AWS policy that uses this IAM policy.
6. To view the Identity and Access Management (IAM) policy that represents the minimum user account permissions that are required for Cloud DR implementation, click **Show IAM Policy**.
7. To save the AWS / cloud account, click **Verify & Save**.

   The CDRA verifies that the account exists before saving the cloud account information and closing the **Connect to Cloud Provider Account** dialog box.

   NOTE: The user cannot change to a different AWS account, after the account is linked to Cloud DR.

Add AWS cloud targets

You can add one or more AWS cloud targets to the cloud account by selecting an Amazon S3 bucket and an encryption method.

Steps

1. Click **Cloud Account** on the menu bar.
   The **Cloud Account** page is displayed.
2. Click **ADD CLOUD TARGET** to set up one or more Cloud DR targets on the cloud account.

   The Cloud DR target is the S3 bucket on AWS where data is written when VMs are backed up to the cloud. The Cloud DR Server is deployed on one of the targets.

   The **Add Cloud DR Target** dialog box opens.
3. Enter a name for the Cloud DR target.

   Enter the same name that is displayed in the RecoverPoint for VMs plug-in for vSphere, when creating a cloud copy.
4. Select an Amazon S3 bucket and region for the Cloud DR target.
5. Click **Advanced security option** and select an encryption method.

   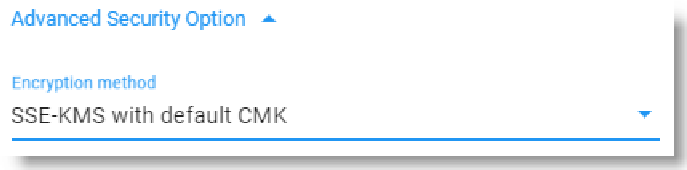

   Advanced Security Option ▲

   Encryption method

   SSE-KMS with default CMK

Option	Description
SSE-S3	Default encryption (no cost)
SSE-KMS	Key management service encryption (incurs a cost)

   NOTE: If you select the SSE-KMS encryption method, only the default customer-managed key is supported. Changing the encryption key might cause errors with the files in the Amazon S3 bucket.

   For more information about these encryption methods, see:

   - SSE-S3 - https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingServerSideEncryption.html
   - SSE-KMS - https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingKMSEncryption.html
6. Click **ADD**.
7. For each Cloud DR target that you want to add, repeat the steps in this procedure.

Connect the CDRA to your CDRS

Connect the CDRA to the CDRS that you deployed using the RecoverPoint for VMs vSphere plugin.

Prerequisites

To connect to the CDRS you will need the password of the CDRS **admin** user. The CDRS **admin** user password is defined when you Install and register a Cloud DR Server on page 25, and can be updated through the CDRS user interface as described in Change the CDRS admin user account password on page 95.

Steps

1. On the **Cloud DR Server** page of the CDRA UI, click the link for the CDRS hostname.

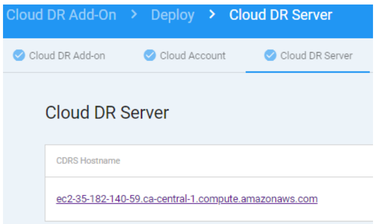

CDRS Hostname

ec2-35-182-140-59.ca-central-1.compute.amazonaws.com

2. When the Cloud DR Server log-in is displayed, enter the username and password for the CDRS.

Results

The CDRA is connected to the CDRS that you deployed using the **RecoverPoint for VMs vSphere plugin**.

Connect to vCenter servers

You can connect the CDRA to vCenter servers that manage VMs in the Cloud DR solution. You can also define recovery settings.

Steps

1. Click **vCenter Servers** on the menu bar.
   The **Connect to vCenter Servers** page appears.
2. Click **Add vCenter Server**.
   The **Connect to vCenter Server** dialog box appears.
3. Enter the hostname or IP address of the vCenter server.
4. Enter the port number for the vCenter server.
5. Enter the Admin username and password.
6. Click **Save**.
7. In the **Confirm vCenter's SSL Certificate** dialog box, click **Confirm**.

   A dialog box prompts you to define a recovery staging area.
8. Define the recovery settings as described in "Define a recovery staging area." To define recovery settings later, click **Define Later**.
9. To add additional vCenter servers, repeat steps in this procedure for each vCenter server.

Results

The **vCenter Servers** page lists vCenter servers that you add to the CDRA.

Define a recovery staging area

Recovery is the process of transferring protected VMs from the cloud to the designated vCenter environment. The **Define Recovery Staging Area** dialog box enables you to configure settings for the operation.

Prerequisites

If you are defining a recovery staging area for the VMware Cloud on AWS (VMC), follow these guidelines when performing this procedure:

- When prompted to select a network, select the network for the VMC software-defined data center (SDDC).
- When enabling direct failover to a vCenter, select the VMC vCenter.

About this task

ⓘ **NOTE: If you do not define a recovery staging area during initial Cloud DR configuration, you can define it later. However, recovery operations do not work unless these settings are configured.**

Steps

1. In the **vCenter Servers** tab, select a vCenter, and click the edit icon ✏. To update information about the vCenter, select **Edit vCenter Details**. To update the failback settings, select **Edit Failback Setting**.

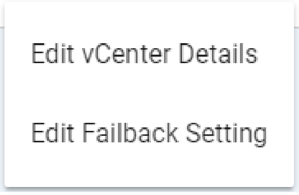

When you click **Edit vCenter Details**, the **Define Recovery Staging Area** dialog box is displayed.

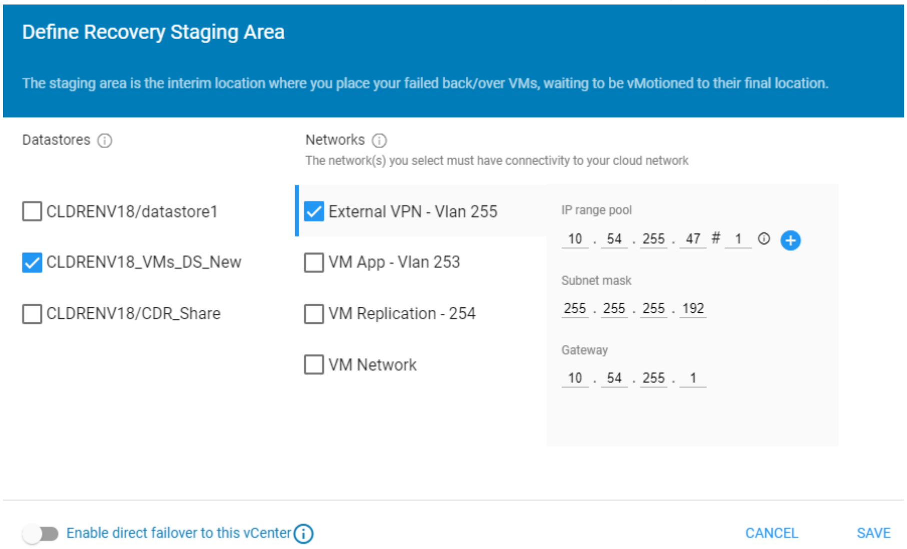
Define Recovery Staging Area

The staging area is the interim location where you place your failed back/over VMs, waiting to be vMotioned to their final location.

Datastores ⓘ

☐ CLDRENV18/datastore1

☑ CLDRENV18_VMs_DS_New

☐ CLDRENV18/CDR_Share

Networks ⓘ

The network(s) you select must have connectivity to your cloud network

☑ External VPN - Vlan 255

☐ VM App - Vlan 253

☐ VM Replication - 254

☐ VM Network

IP range pool: 10 . 54 . 255 . 47 # 1

Subnet mask: 255 . 255 . 255 . 192

Gateway: 10 . 54 . 255 . 1

2. Select one or more datastores or datastore clusters on the vCenter server.
3. Select one or more networks for the recovery staging area.

   Selected networks must connect to the cloud.
4. For each selected network:
   a. Highlight the network.
   b. Configure the **IP range pool** by typing the first IP address in the pool and the number of IP addresses in the subnet to be included in the pool. To enter additional IP range pools, click the plus ⊕ button.
   c. Enter the network **Subnet mask**.
   d. Enter the network default gateway for the **Gateway**.
5. To enable a direct failover to the selected vCenter, click the toggle button at the bottom of the dialog box:

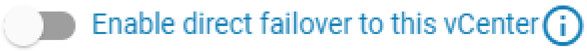

ⓘ **NOTE: You may define multiple vCenters as recovery targets.**

6. Click **Save**.

B

Managing the cloud solution

Use the instructions in this manage the components of your cloud solution.

Before changing the configuration of the RecoverPoint for VMs cloud solution, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

Topics:

Managing virtual machines

This section describes how to manage the protection of virtual machines, after they are initially protected.

About this task

After initial protection, virtual machines are managed through the **RecoverPoint for VMs vSphere plugin** > **Protection** > **Virtual Machines** tab. Select a virtual machine to display the management options for that machine. For a detailed description of how to protect virtual machines, see Protecting VMs on page 27.

Stop protecting a virtual machine

Unprotect a VM to stop replication and remove it from its consistency group.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Virtual Machines**.
2. Select the production VM that you want to stop protecting.
3. Click the **Unprotect** button to the top of the screen:

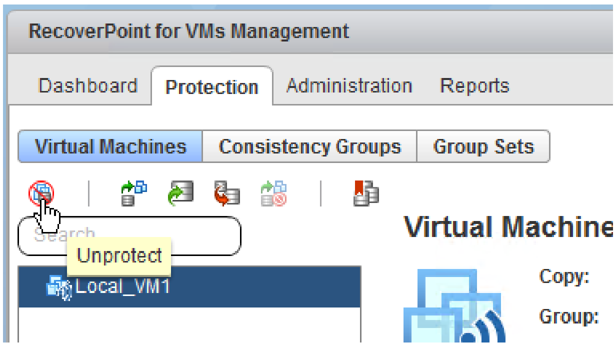

Results

Replication stops and the virtual machine is removed from its consistency group. The copy VM is not automatically deleted. If there are no other virtual machines in the consistency group, the consistency group is removed.

Managing consistency groups

This section describes how to manage consistency groups, after they are created.

After initial creation, consistency groups are managed through the **RecoverPoint for VMs vSphere plugin** > **Protection** > **Consistency Groups** tab. Select a consistency group to display the management options for that group. For a detailed description of consistency groups and how to create them, see Protecting VMs on page 27.

Disabling or enabling a consistency group

About this task

Disabling a consistency group stops all replication, and deletes journals. Enabling a consistency group starts replication and causes a full sweep.

Steps

1. In the **vSphere Web Client** home page, click **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Consistency Groups**.
2. Select the consistency group that you want to enable or disable. Click the **Enable Group** icon or the **Disable Group** icon:

or

Managing the protection policies of groups

About this task

Change the protection policies of groups.

Steps

1. In the **RecoverPoint for vSphere plugin**, select **Protection** > **Consistency Groups**.
2. Expand the list of consistency groups, and select the consistency group whose policies you want to change.
3. Click the **Modify group policy** link:

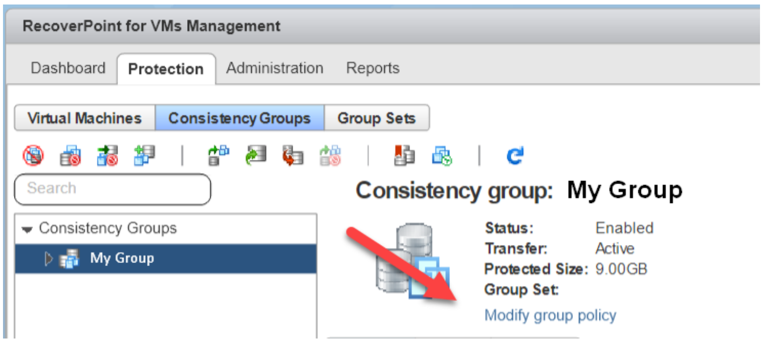

Consistency group: My Group

Status: Enabled
Transfer: Active
Protected Size: 9.00GB
Group Set:
Modify group policy

Modify the policy settings as required:

- **Name**: The name of the consistency group.
- **Primary vRPA**: The vRPA that you prefer to replicate the consistency group. When the primary vRPA is not available, the consistency group will switch to another vRPA in the vRPA cluster. When the primary vRPA becomes available, the consistency group will switch back to it.

> ⓘ **NOTE: If your vRPA cluster is the only vRPA cluster in the system, it is a single point of failure in cases of disaster. Consider adding additional vRPAs to this cluster to ensure high availability.**

4. Click **OK**.

Results

The group protection policies are updated.

Managing copies

This section describes how to manage copies, after they are initially created.

After initial creation, copies are managed through the **RecoverPoint for VMs vSphere plugin** > **Protection** > **Consistency Groups** tab. Select a copy to display the management options for that copy. For a detailed description of copies and how to create them, see Protecting VMs on page 27.

Managing the protection policies of cloud copies

Modify the protection policies of cloud copies.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Consistency Groups**.
2. Expand the list of consistency groups.
3. Expand the consistency group whose copy protection policies you want to edit.
4. Select the cloud copy.
5. Click the **Modify copy policy** link.

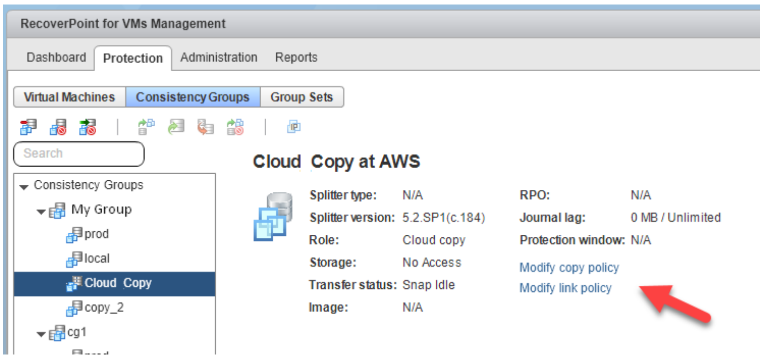

Cloud Copy at AWS

Splitter type: N/A
Splitter version: 5.2.SP1(c.184)
Role: Cloud copy
Storage: No Access
Transfer status: Snap Idle
Image: N/A

RPO: N/A
Journal lag: 0 MB / Unlimited
Protection window: N/A

Modify copy policy
Modify link policy

Edit the copy policy settings, as required:

- **Copy Name**: Default = *copy<num>*

  The name of the cloud copy. Best practice is to differentiate the cloud copy name from the production copy name.
- **Retention Policy**: Default = *5 days*

  Defines the period of time that snapshots will be retained in the cloud. Snapshots can be retained for a minimum of `1` day, and a maximum of `90` days. Once every 24 hours, a temporary **Retention Service** EC2 instance is launched in AWS to consolidate the snapshots of every copy whose retention policy has expired.

6. Click the **Modify link policy** link, and edit the settings, as required:

   - **RPO**: Default = *90 minutes*

     The maximum data lag that is required between the production copy and the latest snapshot uploaded to the S3 bucket. If the specified RPO is exceeded, a warning is displayed in the The RecoverPoint for VMs Dashboard on page 33. RPO can be defined in `Minutes`, `Hours`, or `Days`.

     > ⓘ **NOTE: Specify an RPO value that is higher than the specified snap replication Interval. Best practice is to specify an RPO value that is 1.5 times the specified snap replication Interval. For example, if you require an RPO of *1 hour*, specify a snap replication interval of *90 minutes*.**

   - **Snap Replication**: Default = *Periodic* at *1 hour* intervals

     Sets the periodic **Interval** between snapshots in `Minutes`, `Hours`, or `Days`. The minimum interval value is *15 minutes* and the maximum interval value is *7 days*. A new snapshot starts after the specified interval has passed since the previous snapshot was started. If the time interval has passed and the previous snapshot is incomplete, the next snapshot will start as soon as the previous one has completed.

     > ⓘ **NOTE: Specify a snap replication Interval value that is lower than the specified RPO value. Best practice is to specify an RPO value that is 1.5 times the specified snap replication Interval. For example, if you require an RPO of *1 hour*, specify a snap replication interval of *90 minutes*.**

Next steps

See Managing the protection policies of groups on page 85 for protection policies specific to groups.

Managing group sets

This section describes how to manage group sets, after they are created.

After initial creation, group sets are managed through the **RecoverPoint for VMs vSphere plugin** > **Protection** > **Group Sets** tab. Select a group set to display the management options for that group set. For a detailed description of group sets and how to create them, see Create a group set on page 39.

Modifying a group set

Steps

1. Click **Protection** > **Group Sets**.
2. Select the group set that you want to modify.
3. Click the **Edit Group Set** icon:
4. In the **Edit Group Set** dialog box.
   a. If required, modify the name of the group set.
   b. Select the consistency groups to add or remove from the group set.
5. Click **OK**.

Removing a group set

Steps

1. Click **Protection** > **Group Sets**.
2. Select the group set to remove.
3. Click the **Remove Group Set** icon:

Managing cloud solution licenses

You can remove a RecoverPoint for VMs license from the system or add a new license.

Prerequisites

- For a detailed description of RecoverPoint for VMs licensing, see RecoverPoint for VMs licensing on page 104.
- To add a new license, you must first Create your license files on page 19.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in,** select **Administration** > **vCenter Servers** > **Licensing**.
2. Modify your system license configuration:
    - To remove an existing license, select the license file and click **Remove**.
    - To add a new license file, click **Add**. The **Getting Started Wizard** is displayed to guide you through the process. Follow the instructions in License and register RecoverPoint for VMs on page 21 to add the licence to the system, and register the cloud solution.

Managing cloud solution registration

Register your RecoverPoint for VMs cloud solution whenever you complete a RecoverPoint system installation, connect vRPA clusters in a RecoverPoint cloud solution, or upgrade a RecoverPoint cloud solution.

Prerequisites

A permanent RecoverPoint for VMs license must exist in the system, see Managing cloud solution licenses on page 88. System registration does not work with a temporary license.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters** > **Support**.
2. Select a vRPA cluster.
3. In the **Registration** widget, enter the new registration settings for your vRPA cluster:
    - **Connect in method:** The method that is used to allow remote connectivity to the RecoverPoint environment. Enabling this feature is recommended as it enables secure access to the RecoverPoint environment to gather logs and resolve issues as quickly as possible. If you already have a Secure Remote Services Gateway servicing other products, use the Secure Remote Services Config Tool to add the RecoverPoint devices to the list of Secure Remote Services monitored environments. When the device is added, click the request **update** button to send the new device information to EMC and contact the local Customer Engineer to approve the update. Refer to the *Secure Remote Services Gateway Operation Guide* for further instructions on Config Tool usage. If you do not have a Gateway at the site, contact the Account Manager to find out more about the benefits of Secure Remote Services.
    - **Location:** The city, state, and country where the customer is located.
    - **Sales order number:** The customer or Customer Engineer should provide this information.
    - **Site (party) ID:** The unique ID of the customer site. This value is automatically inserted and taken from the license file and can only be modified by contacting Customer Support.
    - **Activity type:** The kind of activity you are performing (upgrade, installation).

  - **Resource performing this upgrade/installation:** The role of the person performing this upgrade or installation activity.
  - **Connect home method:** The method that is used to send configuration reports and alerts to Dell EMC. Enabling this feature is recommended as it allows Dell EMC to pro-actively address issues within the RecoverPoint environment, should they arise.

4. Click **Register**.

Results

A service request is opened and sends an email to the specified verification email address from Customer Support to verify that the registration details were updated successfully in the Install Base.

Next steps

If your company does not have outside connectivity, export the registration information to a CSV file and register by email or phone, as described in Register RecoverPoint by email or phone on page 105.

Managing cloud solution support

Prerequisites

- A permanent RecoverPoint for VMs license must exist in the system, see Managing cloud solution licenses on page 88. System reports and alerts do not work with a temporary license. Best practice is to keep both system reports and alerts, and compression and encryption enabled.
- To transfer system reports and alerts using SMTP or Secure Remote Services, ensure that port 25 is open and available for SMTP traffic.
- To transfer system reports and alerts using FTPS, ensure that ports 990 and 989 are open and available for FTPS traffic.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters** > **Support**.
2. Select a vRPA cluster.
3. Select **Enable pre-emptive support for RecoverPoint for VMs** to provide communication between the RecoverPoint for VMs system and the System Reports database.
4. Define the transfer method:
   - To transfer system notifications through an SMTP server, in the **Transfer Method** section, select **SMTP**. In the **SMTP server address** field, specify the IP address or DNS name of the dedicated SMTP server, in IPv4 format. In the **Sender address** field, specify the email address to send the system notifications from.
   - To transfer system notifications through the FTPS server, in the **Transfer Method** section, select **FTPS**.
   - To transfer system notifications through the Secure Remote Services gateway, in the **Transfer Method** section, select **ESRS**. In the **ESRS gateway IP address** field, specify the IP address of the Secure Remote Services gateway in IPv4 format..
5. Click **Test Connectivity**.

Next steps

Wait 10 minutes. Then, click **Dashboard** > **Events Log** and look for event 1020: "`Failed to send system report`".

- If this event does not appear in the **Events Log**, the system notifications mechanism is correctly configured.
- If you do receive event 1020: `Failed to send system report`, check whether there is an issue with the selected method of transfer. If a problem exists, fix it, configure support, and click **Test Connectivity** again. If the problem persists, contact Customer Support.

Managing cloud solution component registration

This section describes how to manage the registration of the components of your RecoverPoint for VMs cloud solution, after the system has already been deployed and configured.

About this task

.

After initial system deployment and configuration, manage the cloud solution configuration through the **RecoverPoint for VMs vSphere plugin** > **Administration** tab. Select a vRPA cluster to display the management options for that vRPA cluster.

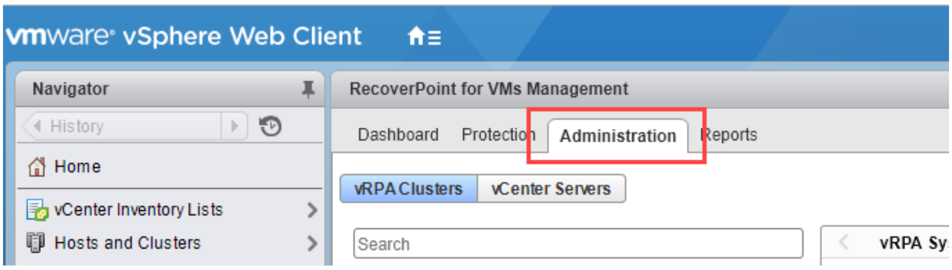

For a detailed description of how to deploy and configure the RecoverPoint for VMs cloud solution, see Solution deployment on page 14.

Managing snap replication datastore registration

Re-define which on-premises datastores are used to store the snapshots of your production VMs before replication to your S3 bucket. Ensure snap replication datastores are registered at every vRPA cluster that protects a production VM. To create a cloud copy, you must have at least one datastore that is registered for snap replication.

Prerequisites

- Registered datastores should be shared (exposed to all ESXi servers hosting vRPA clusters).
- See the *RecoverPoint for Virtual Machines Scale and Performance Guide* for the amount of free space that should be available in the registered datastore(s).
- Registered datastore(s) must have performance capabilities equal to, or greater than, the datastore with the highest performance that is used for your production VMs.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select a vRPA cluster.
3. Select the **Cloud Services** tab.
4. Under the **Snap Replication Datastores** heading:
   - To register a datastore, click the **Add...** button, and in the **Register Snap Replication Datastores** dialog box:
     a. Ensure the correct vCenter Server is selected
     b. Select one or more datastores in which to store the snapshots of your production VMs.
     c. Click **Register**
   - To unregister a datastore, click the **Delete** icon.
5. Repeat for every vRPA cluster that protects a production VM.

Results

Datastore registration is updated at the specified vRPA cluster(s).

Managing cloud account registration

Modify the registration of an AWS account or delete the account. There can only be one cloud account per RecoverPoint for VMs system. Ensure your AWS account is registered at every vRPA cluster that protects a production VM.

Prerequisites

- Ensure you have an on-premises vSphere environment, release 6.0U2 or later.
- Ensure you have an existing Amazon Web Services (AWS) public cloud account.
- Ensure you have an AWS access key and secret access key.
- Ensure TCP/IP port 443 is open for communication between the vRPA clusters and AWS.
- Ensure the vRPA clusters have a configured DNS server that enables vRPAs to resolve amazonaws.com addresses. To configure the vRPA cluster DNS server, refer to the *RecoverPoint for Virtual Machines Installation and Deployment Guide*.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select a vRPA cluster.
3. Select the **Cloud Services** tab.
4. Under the **Cloud Account** header:
   - To edit the registration information of a cloud account, select the registered account and click the **Edit** icon.
     a. Modify the account name, the AWS access key, and/or the AWS secret access key.
     b. Click **Register**.
     c. Ensure the registered AWS account contains an IAM policy with the displayed permissions.

     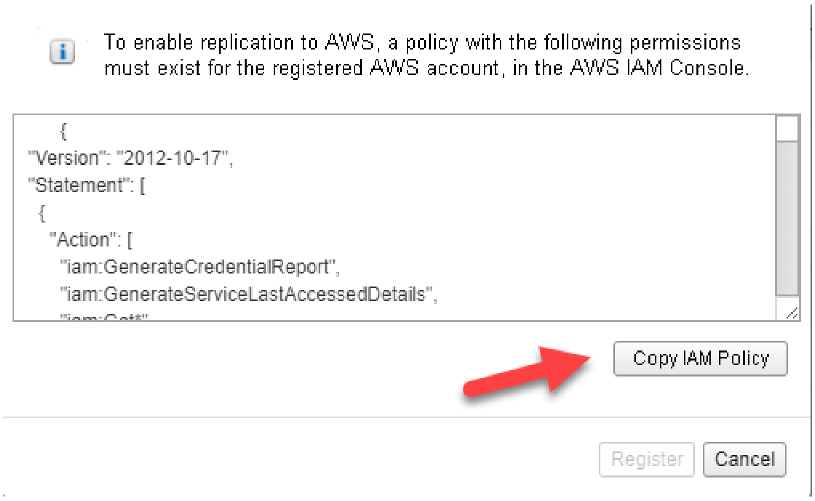

     To enable replication to AWS, a policy with the following permissions must exist for the registered AWS account, in the AWS IAM Console.

     ```
        {
     "Version": "2012-10-17",
     "Statement": [
      {
        "Action": [
          "iam:GenerateCredentialReport",
          "iam:GenerateServiceLastAccessedDetails",
          "iam:Get*"
     ```

     If required, click **Copy IAM Policy** to copy the required IAM permissions in JSON format and paste them into the AWS IAM Console, as described in Define the AWS IAM policy on page 17.
   - To unregister a cloud account, select the registered account and click the **Delete** icon.
5. Repeat for every vRPA cluster that protects a production VM.

Results

Cloud account registration is updated at the specified vRPA cluster(s).

Managing S3 bucket registration

Change or unregister the Amazon S3 bucket in which to store the VM snapshots. There can only be one registered bucket per cloud account. Ensure your S3 bucket is registered at every vRPA cluster that protects a production VM.

Prerequisites

- You cannot unregister a bucket that contains protected VMs.
- To register an Amazon S3 bucket, you will need:
  - an AWS account that has already been registered at the vRPA cluster.
  - an existing Amazon S3 bucket, created in the registered AWS account, with the required IAM permissions, as described in Register an AWS account on page 24.

Steps

1. In the vSphere Web Client home page, select **RecoverPoint for VMs Management** > **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster.
3. Select the **Cloud Services** tab.
4. Under the **Amazon S3 Bucket** header:
   - To unregister an S3 bucket, first ensure there are no protected VMs in the bucket. Select the **Cloud Account** of the S3 bucket, select the registered S3 bucket and click the **Delete** icon.

- To register an S3 bucket, click **Add...**, select an S3 bucket and click **Register**.

5. Repeat for every vRPA cluster that protects a production VM.

Results

S3 bucket registration is updated at the specified vRPA cluster(s).

Managing Cloud DR Server registration

After installing a Cloud DR Server (CDRS), you can unregister and register the CDRS that will be used to recover your cloud copy. You also may want to register the CDRS if there is a loss in communication or if the CDRS 'admin' user password changes (for example, for security purposes). There can only be one registered CDRS. Ensure your CDRS is registered at every vRPA cluster that protects a production VM.

Prerequisites

To register the CDRS, you will need:

- an installed CDRS, as described in Install and register a Cloud DR Server on page 25.
- the password of the CDRS 'admin' user. This password is set when you Install and register a Cloud DR Server on page 25. and can be updated through the CDRS user interface, as described in Change the CDRS admin user account password on page 95.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster.
3. Select the **Cloud Services** tab.
4. Under the **Cloud DR Server** heading:
   - To unregister a CDRS, select the Cloud DR Server in the table, and click the **Delete** icon.
   - To register a CDRS, click **Add...**. Enter the password of the Cloud DR 'admin' user. Click **Register**.
5. Repeat for every vRPA cluster that protects a production VM.

Results

Cloud DR Server registration is updated at the specified vRPA cluster(s).

Managing vCenter Server registration

Registers the vCenter Servers used to manage your production VMs, at a vRPA cluster.

Prerequisites

- All vCenters that manage production VMs must be registered at the relevant vRPA cluster before you protect VMs.
- When a vCenter is registered, all ESX clusters hosted by the vCenter are automatically registered, and a splitter is installed on all ESXs in the cluster.
- Best practice is to configure the vCenter Server to require a certificate, because once RecoverPoint has read the certificate, it does not need further access to the location.
- The default certificate locations are:
  - Windows 2003 Server: `C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt`.
  - Windows 2008 Server: `C:\Users\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt`.

For more information about the location of the security certificate, refer to "Replacing vCenter Server Certificates in VMware vSphere 5.0, 5.5 and 6.0," available at www.vmware.com.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select the **Administration** tab.

2. Access the vCenter registration information:
   - To manage the registration of all vCenter servers in a RecoverPoint for VMs system select **vCenter Servers** > **Registration** and use the **Edit** icon to edit the vCenter settings. Use this option to:
     - Edit the vCenter server information, upload a new vCenter certificate, or delete an existing certificate.
     - Propagate the changes to the specified vCenter server at the specified vRPA cluster using the **Apply** button.
     - Propagate the changes to all vRPA clusters in the system using the **Apply changes to all clusters** button.
   - To manage the registration of a vCenter server at a specific vRPA cluster select **vRPA Clusters** > **vCenter Servers**, select a vRPA cluster, and:
     - To edit the registration details of an existing vCenter server at the selected vRPA cluster, click the **Edit** icon.
     - To register a new vCenter server at the selected vRPA cluster, click **Add...**.
3. Click **OK**.

Results

The specified vCenter Server is registered at the specified vRPA cluster. All ESX clusters hosted by the vCenter are automatically registered with the specified vRPA cluster, a splitter is installed on all ESXs in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Managing ESX cluster registration

Registers the ESX cluster of a production VM or copy VM, at a vRPA cluster.

About this task

By default, ESX clusters are automatically registered in RecoverPoint for VMs during VM protection and copy addition. Use this procedure to register ESX clusters in the rare case that the system cannot automatically register an ESX cluster.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration** > **vRPA Clusters**.
2. Select the vRPA cluster at which you want to register ESX clusters.
3. Select the **ESX Clusters** tab.
4. Click **Add**.
5. In the **Register ESX Clusters** dialog box:
   a. Select the ESX cluster that you want to register.
   b. Click **OK**.

Results

The specified ESX cluster is registered at the specified vRPA cluster.

NOTE: When an ESX cluster of an unregistered vCenter Server is registered with a vRPA cluster, a splitter is installed on all ESXs in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Managing the CDRA

The following options are available for managing the CDRA.

Change the password for the CDRA admin account.

You can maintain security by changing the password for the CDRA admin account.

Steps

1. From the **Settings** menu option, select **Users**.
   The **User Management** page appears.
2. Click the edit (pencil) icon.
   The **Edit User Details** dialog box opens.

3. Click **Change Password**.
4. Enter the new password.

   The password must:

   - Be at least eight characters in length
   - Contain characters of a minimum of three of the following types:
     - English uppercase: A-Z
     - English lowercase: a-z
     - Numeric character: 0–9
     - Special (non-alphanumeric) characters
5. Confirm the new password by entering it again.
6. Click **Save**.

Change the CDRA password expiration period

You can change the password expiration period for the CDRA **admin** account.

Prerequisites

Log in as the **admin** user.

Steps

1. From the **Settings** menu option, select **Users**.
   The **User Management** page appears.
2. Click the edit (pencil) icon.
   The **Edit User Details** dialog box opens.
3. Click **Change Password**.
4. Select a different expiration period. To set the password to never expire, select **Never**.
5. Click **Save**.

Results

The expiration period of the CDRA **admin** user password is updated.

Upgrading the CDRA

The CDRA is upgraded as part of the RecoverPoint for VMs cloud solution.

See Upgrading the cloud solution on page 98 for detailed instructions on how to upgrade the RecoverPoint for VMs cloud solution.

Managing the CDRS admin user account

The Cloud DR Server **User Management** page (**Settings** > **Users**) displays the user accounts that are associated with the CDRS. Use this page to view warning messages for a CDRS user, or to change the password, password expiration period, and password recovery email address, of a CDRS user.

NOTE: You cannot create or delete a user account.

CDRS user accounts are comprised of a **Username** (admin) and **Password** that are used to Log into the CDRS interface on page 43, and a **User email address** that is used to recover the password.

An admin user is created when you Install and register a Cloud DR Server on page 25, using the RecoverPoint for VMs vSphere plug-in.

If you have not responded to the AWS verification email sent after you Define the email address for CDRS password recovery on page 96, a warning icon is displayed next to the **admin** user account **email address**. You can request a new verification email through the AWS console by signing into the console and selecting the **US East (N. Virginia)** region. Open https://console.aws.amazon.com, and select **Email Addresses**. Select the email address that you want to verify, and click **resend**.

Change the CDRS admin user account password

Define the password for the CDRS **admin** user accounts.

Prerequisites

Log into the CDRS interface on page 43.

About this task

The password of the **admin** user is defined when you Install and register a Cloud DR Server on page 25 using the **RecoverPoint for VMs vSphere plugin**.

Steps

1. In the Cloud DR Server user interface, select **Settings** > **Users**.
   The **User Management** page appears.
2. Click the edit (pencil) icon to the right of the **admin** user account.
   The **Edit User Details** dialog box appears.

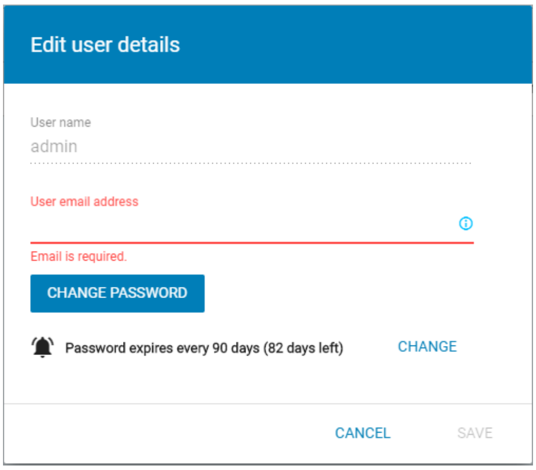

Edit user details

User name
admin

User email address

Email is required.

Password expires every 90 days (82 days left)

3. Click **Change Password**.
4. Enter the new password.

   The password must:

   - Be at least eight characters in length
   - Contain characters of a minimum of three of the following types:
     - English uppercase: A-Z
     - English lowercase: a-z
     - Numeric character: 0–9
     - Special (non-alphanumeric) characters
5. Confirm the new password by entering it again.
6. Click **Save**.

Results

The password is updated, and should be used from now on, when you Log into the CDRS interface on page 43.

Next steps

Use the **RecoverPoint for VMs vSphere plugin** to register the new password with every vRPA cluster that protects a production VM. In the **RecoverPoint for VMs vSphere plugin**, select **Administration** > **Cloud Services**, click the **Edit** icon to the right of the **Cloud DR Server** name, update the value for **CDRS admin user password**, and click **Register**.

Change the CDRS password expiration period

You can change the password expiration period for the CDRS **admin** user.

Prerequisites

Log into the CDRS interface on page 43 as the **admin** user.

Steps

1. In the Cloud DR Server user interface, select **Settings** > **Users**.
   The **User Management** page appears.
2. Click the edit (pencil) icon to the right of the **admin** user account.
   The **Edit User Details** dialog box appears.
3. Click **Change Password**.
4. Select a different expiration period. To set the password to never expire, select **Never**.
5. Click **Save**.

Results

The expiration period of the CDRS **admin** user password is updated. You will be prompted to update the current password when it expires.

Define the email address for CDRS password recovery

Specify the email address to which instructions for resetting the password will be sent, if you should lose your CDRS **admin** account password.

Prerequisites

- Log into the CDRS interface on page 43.
- In the **AWS Management Console**, ensure that the email address that you want to use for password recovery is verified under the **AWS root user** account.

About this task

NOTE: Clicking the Forgot Password? link when you Log into the CDRS interface on page 43 will send an email with the instructions for resetting the password to the User email address that you define.

Steps

1. In the Cloud DR Server user interface, select **Settings** > **Users**.
   The **User Management** page appears.
2. Click the edit (pencil) icon to the right of the **admin** user account.
   The **Edit User Details** dialog box appears.

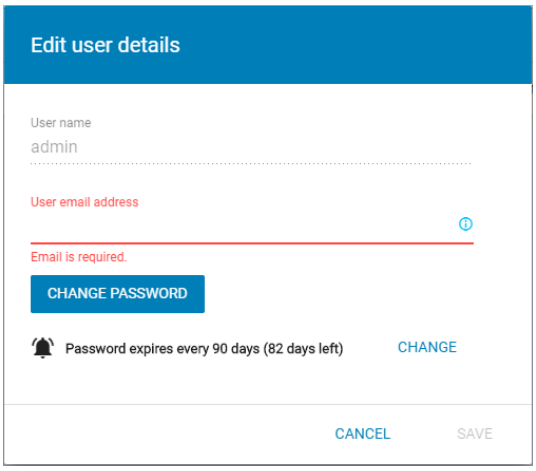

Edit user details

User name
admin

User email address

Email is required.

Password expires every 90 days (82 days left)

3. In the **User email address** field, enter the email address.
4. Click **Save**.

Results

If the new email address exists in the **AWS root user account** and has been verified by AWS, a verification email is sent from AWS to the new email address.

Next steps

Respond to the AWS verification email within 24 hours. After you respond to the AWS verification email, the email address is updated in Cloud DR, and will be used for recovery the next time you click the **Forgot Password?** link, when you Log into the CDRS interface on page 43. The new **User email address** is assigned to the **US East (N. Virginia)** region.

C

Upgrading the cloud solution

This section describes how to upgrade the components of the RecoverPoint for VMs cloud solution, after initial solution deployment and configuration.

Before upgrading the cloud solution, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

ⓘ **NOTE: It is important that you perform the upgrade procedure in the order in which it is presented.**

Topics:

Upload an upgrade package

To upload an upgrade package to the CDRS and CDRA, use the Cloud DR Server **Upgrades** page.

Prerequisites

- The versions of the CDRA and CDRS do not need to be identical, and you are not required to upgrade them at the same time (unless otherwise instructed). When uploading an upgrade package, if the upgrade package version is not supported, you receive a notification.
- Consult the *RecoverPoint for VMs Release Notes* to ensure that the upgrade packages that you upload are for a CDRA/CDRS version that is compatible with the version of RecoverPoint for VMs that you want to upgrade to.

Steps

1. Download the upgrade package (CDRS or CDRA, or both) from online support: https://www.dell.com/support/ (search for "Cloud Disaster Recovery Upgrade Package").
2. From the CDRA **System** menu option, select **Upgrades**.
3. To upload the upgrade package that you downloaded in 1 on page 98, click **Upload Package**.
4. To replace the currently uploaded package with another package, click **Upload Different Package**.

Results

- After uploading an upgrade package for the CDRS, the **Upgrade Cloud DR Server** button is displayed. Upgrade the Cloud DR Server on page 98 provides the steps to upgrade the CDRS.
- After uploading an upgrade package for the CDRA, a message indicates that the CDRA is pending upgrade. Upgrade the Cloud DR Add-on on page 99 provides the steps to upgrade the CDRA.
- If the upgrade package includes both CDRS and CDRA, the package is made available for the CDRA only after the CDRS has been upgraded.

Upgrade the Cloud DR Server

To upgrade a CDRS, use the Cloud DR Server **Upgrades** page. If a recovery operation is in progress, the upgrade process is disabled.

Prerequisites

- Upload an upgrade package on page 98
- Ensure that there is no rapid recovery process running.
- Ensure that you have complied with all of the pre-requisites in the Before you begin on page 14.

- Consult the *RecoverPoint for VMs Release Notes* to ensure that the target CDRS version is compatible with the version of RecoverPoint for VMs that you want to upgrade to.

About this task

NOTE: Do not upgrade the CDRS while the rapid recovery process is running. If you upgrade the CDRS during the rapid recovery process, that process is not monitored after the upgrade (the machine image is lost).

Steps

1. From the CDRS **System** menu option, select **Upgrades**.
2. Click **Upgrade Cloud DR Server**.
3. In the **Cloud DR Server Upgrade** dialog box, click **Upgrade**.

Results

Expect a short downtime during upgrade while the CDRS restarts. You cannot perform recovery operations until the upgrade completes and you restart the browser.

Next steps

Restart the browser, and Log into the CDRS interface on page 43.

Upgrade RecoverPoint for VMs

Use the **RecoverPoint for VMs Deployer** to upgrade the vRPA clusters in your RecoverPoint for VMs cloud solution.

See the *RecoverPoint for Virtual Machines Installation and Deployment Guide* for a detailed description of the RecoverPoint for VMs upgrade process.

Upgrade the Cloud DR Add-on

To upgrade a CDRA, use the Cloud DR Add-on **Upgrades** page.

Prerequisites

- Upload an upgrade package on page 98
- Ensure the CDRA complies with the Virtual machine specifications for Cloud DR with AWS on page 16
- Consult the *RecoverPoint for VMs Release Notes* to ensure that the target CDRA version is compatible with the version of RecoverPoint for VMs that you have upgraded to.

NOTE: CDRA is required only for failback from AWS, or to recover to VMware Cloud on AWS. If you deployed more than one CDRA, remember to upgrade both of them.

Steps

1. From the CDRA **System** menu option, select **Upgrades**.
   The **Upgrades** page displays and provides information about the current version and upgrade status of the Cloud DR Add-on.
2. If an upgrade package is available for the CDRA, click **Upgrade Cloud DR Add-on**.

Results

The CDRA is upgraded to the new version. A short downtime is possible during upgrade while the CDRA restarts. At the end of the upgrade process, the Cloud DR Add-on login page displays.

Next steps

Restart the browser and Log into the CDRA on page 77.

D

Troubleshooting

Use the following information, features and tools to troubleshoot your RecoverPoint for VMs cloud solution.

Topics:

Troubleshooting AWS environments

AWS default limits

The following components in AWS have default limits that may not be appropriate for the Cloud DR environment. For example, if you plan to use more than five VMs and are using elastic IP addresses, you must increase the default limit for the number of elastic IP addresses before performing a disaster recovery.

Table 11. AWS default limits

Component	Default limit
Number of buckets	100
Number of Elastic IP addresses	5
Number of instances per region	20
Number of Internet gateways	5
Number of Instances from the same type in the same region	25

http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html contains information about default service limits and information about how to increase the limits.

AWS encryption

In AWS, policies can be specified for an S3 bucket that requires all objects within the bucket to be encrypted (or non-encrypted) with a specific algorithm or key. Cloud DR does not verify that the policy of the target bucket matches the encryption policy that the user configured for a cloud target. If there is a mismatch between the two, CDRA fails to send the data to the S3 bucket.

In the event of this failure, check the Cloud DR events to determine the issue. Then change the security policy in the cloud target, the target bucket, or the target bucket policy.

AES256 encryption

Cloud DR uses AES256 to encrypt metadata in AWS. As a result, if an S3 bucket policy enforces KMS for all objects within the bucket, the CDRA can upload the user data, but not the metadata.

If this issue exists, edit the bucket policy to allow for AES256 encryption for the metadata folder within the bucket. For example, edit the bucket policy by adding the following:

```
 {
      "Sid": "AllowKMSEverywhere",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::bucket1/*",
      "Condition": {
        "StringEquals": {
          "s3:x-amz-server-side-encryption": "aws:kms"
        }
      }
    },
    {
      "Sid": "AllowAES256InMetadataFolder",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::bucket1/backups/*",
      "Condition": {
        "StringEquals": {
          "s3:x-amz-server-side-encryption": "AES256"
        }
      }
}
```

Error when deploying the Cloud DR Server if AWS Marketplace terms have not been accepted

An error may occur when the Cloud DR is trying to create the EC2 instance for the Cloud DR Server during Cloud DR Server deployment if AWS Marketplace terms have not been accepted. The following error message appears in the log:

```
ERROR [date]
com.emc.cloud_dr.cdr.cdra.cloud_manager.impl.deploy_cdrs.wf.steps.instance.CreateCdrsInstanceTas
klet: Error in Create Cdrs Instance
! com.amazonaws.services.ec2.model.AmazonEC2Exception: In order to use this AWS Marketplace
product you need to accept terms and subscribe. To do so please visit http://aws.amazon.com/
marketplace/pp?sku=aw0evgkw8e5c1q413zgy5pjce (Service: AmazonEC2; Status Code: 401; Error Code:
OptInRequired; Request ID: id)
```

To resolve this issue, accept the AWS Marketplace terms as described in Accept Amazon Web Services Marketplace terms on page 16 and continue with Cloud DR Server deployment.

Incorrect email address when configuring the Cloud DR Server

If you specify an incorrect email address when configuring the Cloud DR Server and are unable to verify the email:

1. Follow instructions for changing the email address at Define the email address for CDRS password recovery on page 96. Then enter and verify the correct email address.
2. Log in to the AWS console and open the Amazon SES console at https://console.aws.amazon.com.
3. Select the US East (N. Virginia) region.
4. Select the incorrect email address and click **Remove**.

Uninstall the cloud solution

To uninstall the RecoverPoint for VMs cloud solution, follow the steps in this procedure.

Prerequisites

NOTE: Failure to perform these steps in the listed order causes undesirable results.

Steps

1. If you installed a CDRA on-premises or on the VMware Cloud on AWS, delete the Cloud DR Add-on appliance from vSphere, as described in VMware documentation.
2. It is important that you clean up cloud-based resources that are no longer needed. From the Amazon Web Services console, perform these tasks in the order presented:

Table 12. Cleaning up cloud-based resources

Task	AWS documentation link
Delete the Cloud Formation stacks from all regions that you used (named **CDRS-DeployStack**, **CDRS-RDSCluster**, **CDRS-RestoreService**).	http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-console-delete-stack.html
Delete the EC2 key pairs that are named **CDRS-KeyPair** and **CDRS-RestoreService**.	http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-key-pairs.html
Delete the IAM role that is named **CDRS-Role**.	http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_manage_delete.html
Delete all S3 buckets that were used as Cloud DR targets.	http://docs.aws.amazon.com/AmazonS3/latest/dev/delete-or-empty-bucket.html ⓘ **NOTE: Perform this step only if the S3 buckets are not being used for purposes other than Cloud DR.**
Unregister AMIs and delete snapshots that Cloud DR Server created for rapid recovery of VMs.	https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/deregister-ami.html Look for AMIs and snapshots where CDRS is displayed in the **Created By** tag name.
Delete the SQS queues that are named **CDRS-RestoreService *<version>*-Events** and **CDRS-RestoreService...*<version>*-Responses**	https://docs.aws.amazon.com/cli/latest/reference/sqs/delete-queue.html

Finding the vRPA cluster management IP

Displays the vRPA cluster management IP of a specific vRPA cluster.

Steps

1. Select **Administration** > **vRPA Clusters** > **vRPA System**
2. Select the vRPA cluster.
3. Note the **vRPA cluster management IP** of the selected vRPA cluster.

Collecting system information

Collecting system information is only relevant in support cases, and should only be performed when instructed to do so by Customer Support. Use the **RecoverPoint for VMs vSphere plugin** to collect system information, and retrieve it from the S3 bucket of the CDRS cloud account, or a specified FTP server.

Prerequisites

Ensure you Enable downloads of Cloud DR logs from AWS on page 18.

About this task

ⓘ **NOTE:**

If there is no connectivity between the vRPA cluster and the CDRS, you cannot collect system information from the CDRS through the RecoverPoint for VMs vSphere plugin. In this case, ask Customer Support to collect the system information for you.

Steps

1. In the vSphere Web Client home page, select **Administration** > **vRPA Clusters** > **Log Collection**.
2. Under **Collection Period**, define a date and time for the start and end of the collection process.
3. Optionally, click **Change to GMT** to change the collection time display to GMT.

   GMT is not adjusted for daylight savings time. Although the system information of the past 30 days is available for collection, only 3 days of system information can be collected at a time.
4. Select the **vRPA Clusters** from which to collect the logs.
5. Optionally, select **Include core files**.

   Core files might be large. Subsequently, including these files in the collection process could substantially increase collection time.
6. By default, **Full system log collection** is selected. If you are instructed to do so by Customer Support, use **Advanced** to select the specific logs that you want to collect.
7. Optionally, select **Copy the output file(s) to an FTP server** and define the FTP server settings.
8. Click **Start**.

Results

Be patient. The collection process can take awhile, depending on the amount of data being collected. After the collection process is complete, the results are displayed.

Next steps

If you selected the **Copy the output file(s) to an FTP server** checkbox, retrieve the output file from the specified FTP server.

Otherwise:

- If you are collecting logs from the vRPA cluster that you are connected to, click the relevant link in the **Output Files on Cloud** column to retrieve the log files.
- If you are collecting logs from a vRPA cluster other than the vRPA cluster you are connected to, you have one of two options:
  - Copy the displayed text into a browser address bar, replacing **`<CDRS-BUCKET-NAME>`** and **`<CDRS-REGION>`** with the names of your AWS bucket and region.
  - Note the displayed log file name(s) and location, and download the log files using your AWS Management Console.

Collecting RecoverPoint for VMs splitter logs

About this task

RecoverPoint for VMs splitter logs are in the ESXi logs. To export the ESXi system logs, use the following procedure.

Steps

1. In the vSphere Web Client, select an ESXi host and click **Actions**.
2. Select **All vCenter Actions** > **Export System Logs...**.
3. In the **Export Logs** screen, specify which system logs are to be exported. If required, select the `Gather performance data` option and specify a `duration` and `interval`.
4. Click **Generate Log Bundle**.
5. Click **Download Log Bundle**.
6. Upload the logs to the FTP site.

   For information on how to upload logs for VMware products, see http://kb.vmware.com/selfservice/search.do?cmd=displayKC&docType=kc&docTypeID=DT_KB_1_1&externalId=1008525

Recovering from a cluster disaster

After a full cluster disaster or a switch disaster, it may take 10 minutes or more for all the components of the vRPA system to restart, reconnect, and restore full operation.

RecoverPoint for VMs licensing

RecoverPoint for VMs supports two types of licensing models; VM-based licensing and socket-based licensing .

> **NOTE: No additional cost is required to protect your production VMs in AWS. If you have an existing RecoverPoint for VMs system that is already licensed, the RecoverPoint for VMs cloud solution uses your existing license.**

VM-based licensing

With VM-based licensing, licenses are based on the number of supported VMs per vCenter server. Only production VMs are counted in the number of supported VMs per vCenter server. Licensing is enforced using the vCenter Server ID.

All vCenter servers must be registered in RecoverPoint for VMs before their licenses can be added. vCenter server registration is performed in the RecoverPoint for VMs Deployer UI. Refer to the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

When you reach the maximum number of VMs that the license supports for each vCenter server, you cannot protect new VMs or enable disabled consistency groups. However, replication of existing VMs and consistency groups continues.

Failover has no effect on the license.

Socket-based licensing

With socket-based licensing, licenses are based on the number of physical CPU sockets in the ESXi servers that host the production VMs. A VM does not 'belong' to a specific socket.

When you reach the maximum number of sockets that the license supports for each vCenter server, you cannot protect new VMs or enable disabled consistency groups. However, replication of existing VMs and consistency groups continues.

As with VM-based licensing, failover does not affect the socket-based license. However, vMotion of production VMs does affect the license and may cause a license violation due to an increase in the number of sockets being used. ESXi servers that host the production VMs are the ones that count in a socket-based license. To avoid license violations, it is a best practice to license all ESXi servers of the ESXi cluster.

Adding a socket-based license to a system with VM-based licenses

When using VM-based licensing, license capacity is measured by the number of VMs. For example, when you view the license capacity in the UI, it may be listed as:

```
Capacity = 30 VMs
```

When using socket-based licensing, license capacity is measured by the number of sockets. For example, the license capacity may be listed as:

```
Capacity = 2 sockets
```

When a socket-based license is installed on a RecoverPoint for VMs system that has VM-based licenses, the system automatically converts VM-based licenses to socket-based licenses at a ratio of 15 VMs per socket. In this case, the license capacity would be listed as:

```
Capacity = 30 VMs (2 sockets)
```

In cases where the ratio does not result in an even conversion, the value is rounded up. For example:

```
Capacity = 31 VMs (3 sockets)
```

Since licenses are applied per vCenter, and not per vRPA cluster, multiple vRPA clusters with VMs or CPU sockets may count towards the same license.

License subscriptions

VM- and socket-based licenses may be installed as subscriptions. Unlike a permanent license, a subscription license has a start date and an end date. The system sends an alert beginning 30 days before license expiration to indicate the number of days remaining. Subscription and permanent licenses may coexist.

You can install a subscription license before its start date. It automatically becomes active on the start date.

Register RecoverPoint by email or phone

If your company is without external connectivity, and you cannot register your RecoverPoint for VMs system online, you can also register by phone.

About this task

- Register the RecoverPoint system after:
  - Installing a RecoverPoint system
  - Connecting RPA clusters in a RecoverPoint system
  - Upgrading a RecoverPoint system
- The registration process is incomplete if valid values are not provided for every field in the post-deployment form.

Steps

1. Gather the required information.
   - Download the post-deployment form:
     a. Access https://support.emc.com
     b. Search for the term *Post-Deployment Form*
     c. Download and fill out the RecoverPoint and RecoverPoint for VMs Post-Deployment Form, for every vRPA cluster
   - Export the RecoverPoint registration information, for every vRPA cluster:
     a. Select **Administration** > **vRPA Clusters**.
     b. Select the vRPA cluster for which you want to export a post-deployment form, and then click **Support**.
     c. In the Registration pane, click the **Export to CSV** button and save the file to the computer.
2. Send the information to the Install Base group:
   - Customers and partners: Email the post-deployment form to the Install Base group at rp.registration@emc.com.
   - Employees:
     - (Preferred) Use the IB Portal at http://emc.force.com/BusinessServices.
     - Call in the information to the Install Base group at 1-866-436-2411 – Monday to Friday (normal Eastern Time Zone working hours).

Glossary

A

account ID
Part of a customer's account settings. The account ID is a unique identifier of a RecoverPoint for VMs customer account.

account settings
The details that comprise a RecoverPoint customer account. The account settings are comprised of:

- One account ID
- One installation ID per installed RecoverPoint for VMs system
- One software serial ID per vRPA cluster
- A RecoverPoint for VMs license key

AMI
Amazon Machine Image (AMI) is a template that contains configuration information which is used to launch an EC2 instance in the AWS environment. In the native cloud solution (AWS cloud), VMware's VMDK format, which is used by VMs, must be converted to AMI format, which is used by AWS cloud. In the VMware Cloud to AWS (VMC) solution, there is no requirement to do a format conversion from VMDK to AMI because a VMware environment exists both on premises and in the cloud.

Asset
A general term that refers to a VM or an application. VMs and applications are considered assets in the Cloud DR solution.

B

bookmark
A label that is applied to a snapshot (PIT) so that the snapshot can be explicitly called (identified) during recovery processes (for example, during image access).

C

call home events
A proactive online service capability that is built into RPAs to enable them to continuously monitor their own health, and the health of the RecoverPoint system, using a pre-defined set of event-filtering rules. If a serious problem arises, the call home event mechanism automatically opens a service request with Customer Support. The service request enables Customer Support to proactively engage the relevant personnel, start working with the relevant parties, or use a configured Secure Remote Services gateway to resolve the issue as soon as possible.

CDRA
Cloud Disaster Recovery Addon (CDRA) manages deployment of on-premises components and CDRS, which runs in the cloud.

CDRS
Cloud Disaster Recovery Server (CDRS) is a virtual server that runs in the customer domain in the cloud. It provides a user interface for disaster recovery testing and failover, and monitors available copies and orchestration activities in the cloud.

NOTE: Multiple on-premises sources (CDRAs and vRPAs) can connect to a single CDRS, but an on-premises source cannot connect to multiple CDRSs.

Classless Inter-Domain Routing (CIDR)
Classless Inter-Domain Routing (CIDR) is a method for IP address allocation and IP routing.

CLI
The RecoverPoint Command Line Interface. Using the RecoverPoint CLI, management and monitoring activities can be run textually, interactively, or through scripts. For information about the command line interface, see the *RecoverPoint Command Line Interface Reference Guide*.

call home events
A proactive online service capability that is built into RPAs to enable them to continuously monitor their own health, and the health of the RecoverPoint system, using a pre-defined set of event-filtering rules. If a serious problem arises, the call home event mechanism automatically opens a service request with Customer Support. The service request enables Customer Support to proactively engage the

relevant personnel, start working with the relevant parties, or use a configured Secure Remote Services gateway to resolve the issue as soon as possible.

cluster control
The process that manages an RPA cluster.

cluster management IP
A virtual, floating IP address assigned to the vRPA that is currently active (runs the cluster control).

NOTE: Unofficially, can be referred to as a floating IP.

compression
The process of encoding data to reduce its size. RecoverPoint uses lossless compression (compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed).

In the RecoverPoint for VMs cloud solution, WAN compression is used to compress consistency group data before transferring it over the WAN.

consistency group
A logical entity that constitutes a container for virtual machines and all their copies, used to replicate virtual machine application data to a consistent point in time.

copy

A logical entity that constitutes all of the data that is replicated and used to protect the production data in the cloud.

E

EBS
Amazon Elastic Block Store (EBS) provides block-level storage volumes for use with EC2 instances.

EC2
Elastic Cloud Compute (EC2) is an Amazon web service that provides resizable compute capacity in the cloud. An EC2 instance is a virtual server in the AWS environment.

ESRS
EMC Secure Remote Support (ESRS) is a server and set of services that allow customer support to remotely access vRPAs to collect system information and provide pre-emptive support.

event
A notification that a change has occurred in the state of a managed device or component. In some cases, the change indicates an error or warning condition for the device or component. Multiple events can occur simultaneously on a single monitored device or service module. A single incident can generate events across multiple system components. Events in RecoverPoint have a level (*Info*, *Warning*, *Error*), scope (*Normal*, *Detailed*, *Advanced*), and a topic (*All*, *Cluster*, *RPA*, *Group*, *Splitter*, *Management*).

F

full sweep
An initialization process that is performed on all of the volumes in a consistency group.

G

group set

A collection of consistency groups. Group sets allow you to manage and perform recovery activities on multiple consistency groups simultaneously. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

H

high-load

A system state that indicates resource depletion during replication. There are two kinds of high-loads in RecoverPoint:

- Permanent high-loads – RecoverPoint stops and waits for a user action in order to come out of high-load (not relevant in the RecoverPoint for VMs cloud solution).
- Temporary high-load – RecoverPoint tries to recover from the high-load and keeps trying until the condition that triggered the high-load changes.

I

image
All of the snapshots that, together, constitute a specific point in time. An image consists of a copy's snapshots in the snap replication datastores, and the data that has already been transferred to the Amazon S3 bucket.

initialization
The process that is used to synchronize the data of the copy volumes with their corresponding production volumes, and ensure consistency. Generally, all synchronization processes are called Initialization.

See also synchronization

initialization snapshot
The first consistent snapshot in a copy journal. Whenever an initialization process completes, this initialization snapshot is created.

Install Base
An database that is used to manage and support equipment installed at customer sites. When new equipment is installed, it is important to update the database with the new information and IDs of the newly installed equipment. The information in the installation base is also used to enable the RecoverPoint system report mechanism.

L

lag
The current RPO of the consistency group. In RecoverPoint, lag starts being measured when a write made by the production host reaches the local RPA, and stops being measured when the write reaches either the target RPA, or the target journal (depending on the transfer_by_non_preferred parameter of the `set_policy` CLI command).

latency
The number of milliseconds or microseconds that it takes for data to get from the local vRPA to the vRPA or journal at the remote vRPA cluster.

link
The communication pipe through which data is transferred between the production and a copy.

See also pipe

load settings
A user-triggered operation that is performed to recreate the system configuration. This operation loads a configuration file that is created from the output of the `save_settings` CLI command, which displays the current system settings in a configuration settings file.

M

maintenance mode
The RecoverPoint for VMs vSphere system enters maintenance mode when undergoing any of the following operations:

- minor version upgrade
- major version upgrade
- replacing an RPA in an existing cluster
- adding new RPAs to existing clusters
- modifying system settings

In maintenance mode, RecoverPoint for VMs can only monitor the system; user-initiated capabilities are disabled.

O

on-premises source
Manages deployment of resources, protects VMs in the cloud, and configures the Cloud DR Server (CDRS). It is possible to have both types of sources (CDRA and vRPA) on premises, each one connecting to the same Cloud DR Server

P

preferred RPA
An RPA that whenever possible, handles replication for the consistency group. If an error occurs in the preferred RPA, in most cases, another RPA at the same RPA cluster handles replication.

See also primary RPA, preferred primary RPA

production
The data that is being replicated and protected.

See also production source, protected copy

R

RDS
Relational Database Service (RDS) is a web service that makes it easier to set up, operate, and scale a relational database in AWS environment.

RecoverPoint for VMs plug-in
The vSphere web client user interface that is used for managing VM replication. The plug-in is installed automatically after the vRPA cluster has been installed.

RecoverPoint for VMs splitter
Proprietary software that is installed on every ESXi host in an ESXi cluster that is involved in RecoverPoint for VMs replication or running virtual RPAs. The splitter splits every write to the VMDK and sends a copy of the write to the vRPA and then to the designated storage volumes. It is installed automatically after you register the ESXi cluster.

Recoverpoint for VMs system
One or more vRPA clusters that have been installed using the RecoverPoint for VMs Deployer.

Rehydration
During protection, Cloud DR initially sends the first full copy of the protected VM to the cloud and afterwards sends only the differences. When you start recovery (for example, test or failover), a temporary Restore Service instance constructs the VMDK file from the raw data chunks that are stored in the Cloud DR target. This process is called rehydration.

replication policy
A user-specified set of parameters driven by business objectives that control system operation during replication.

replication set
A production source and the target(s) at a copy to which it replicates.

replication set volumes
All of the sources that have been added to a replication set.

RPO
Recovery Point Objective. The maximum amount of data, per application, that an organization is willing to lose if there is a disaster. For example, an RPO of 5s means that if there is a disaster, RecoverPoint ensures that no more than the last 5s of data can be lost.

RTO
Recovery Time Objective. The duration of time and a service level within which a business process must be restored after a disaster (or disruption) in order to avoid unacceptable consequences associated with a break in business continuity.

S

S3
Simple Storage Service (S3) is a cloud computing web server that provides scalable, object storage in the AWS environment. Objects are stored in S3 buckets. It is the simplest and cheapest type of storage available from Amazon.

shadow VM
A secondary copy VM that RecoverPoint creates, configures, and manages to allow access to copy VMDK and RDM devices. A copy shadow VM has the .copy.shadow extension at the end of the virtual machine name. User action on copy shadow VMs is not supported.

short initialization
An initialization process that uses marking information to re-synchronize a copy's data with its production sources. Because this initialization process uses delta markers to synchronize the copy with the production, the initialization process is much faster and more efficient. Generally occurs when restarting transfer for a consistency group after a pause in transfer.

See also short init, short resync, short resynchronization, resynchronization, resync

software serial ID
The identification that is used by the install base to support equipment that is installed at customer sites using the system reporting and Secure Remote Services mechanisms. A software serial ID is supplied per RPA cluster in a system installation.

See also SSID

source
The object that RecoverPoint is replicating from. For example, the source RPA or the source copy (for example, production). After failover, the source becomes the target and the target becomes the source.

splitter
Proprietary software that is installed on storage subsystems that splits application writes so that they are sent to their normally designated storage and the RPA simultaneously.

system alerts
A mechanism that allows RPAs to send system events about system components in real-time, to a specified email or the system reports database, via SMTP.

system reports
A mechanism that provides one-way communication between a system installation and the system reports database. This mechanism supports two types of information: system alerts and system reports.

See also SYR

system settings
The output of the `save_settings` CLI command.

See also configuration file, system configuration file, configuration settings

T

target
The object at the copy that the protected data is being replicating to. For example, the target RPA, or the target storage. After failover, the target becomes the source, and the source becomes the target.

throughput
The total amount of writes made by the production hosts and received by the local RPA.

tweak parameters
A configuration parameter that only Customer Support can change. Tweak parameters enable Customer Support to change the hard-coded values of specific internal settings without requiring the RecoverPoint for VMs code to be recompiled or re-loaded onto the vRPAs.

V

volume sweep
An Initialization process that is performed on a specific replication set in a consistency group.

VPC
Amazon Virtual Private Cloud (VPC) is a part of the AWS cloud where you can launch AWS resources in a virtual network that you define.

vRPA
The virtual RecoverPoint Appliance that manages data replication.

vRPA cluster
A group of vRPAs that work together to replicate and protect data.