VMware Cloud Disaster Recovery - Implementation Practices Guide

VMware DRaaS

Table of contents

VMware Cloud Disaster Recovery - Implementation Practices Guide

Overview

Introduction

This guide will cover many of the implementation practices identified that will help improve the setup and operation of VMware Cloud Disaster Recovery.

VMware Cloud Disaster Recovery provides a framework for protecting VMware workloads - either in on-premises vCenter environments or in VMC on AWS SDDCs There are protected to a cloud-based repository (SCFS) of Recovery Points that can be used to fail over VM operations to a Recovery SDDC running in VMware Cloud on AWS.

This document will evolve over time as new practices and areas of consideration are added. It is suggested to use the dynamic html link to the guide whenever possible to reference the content.

Purpose of This Guide

This guide takes you through several areas of using VMware Cloud Disaster Recovery in a variety of situations. The goal is to provide insights and criteria that can be applied to your specific deployment scenarios.

For more information, see the Product Documentation.

Audience

This guide is intended for IT administrators and product evaluators who are familiar with VMware Cloud Disaster Recovery. Familiarity with VMware vSphere and VMware vCenter Server as well familiarity with networking and storage in a virtual environment, Active Directory, identity management, and directory services is assumed.

Prerequisites

This guide assumes a basic working knowledge of VMware Cloud on AWS, VMware Cloud Disaster Recovery, vSphere and vCenter environments.

Protecting Workloads

This section will focus on the elements of VMware Cloud Disaster Recovery that pertain to protecting the VM workloads in preparation for disaster recovery failover needs.

Protecting workloads with VMware Cloud DR involves setting up the Protected Site(s) and defining Protection Group policies.

These are documented in more detail in the product documentation here:

- Protected Sites
- Protection Groups

Implementation considerations for these two operational setup tasks are discussed below.

Protected Sites

Understanding how the data flows from the Protected Site DRaaS Connector(s) to the Scale-out Cloud File System (SCFS) will help in designing and implementing a good site topology.

Each VMware Cloud DR instance can support multiple Protected Sites and each Protected Site can support multiple DRaaS Connectors and multiple vCenter Server registrations. The current configurations limits on these are found in the online ConfigMax tool. Each DRaaS Connector has a recommended limit of VMs from the VCenter inventories that it will need to handle.

These limits need to be considered to provide optimal resiliency and performance considerations.

Some other logistic considerations are outlined below.

Practice – network connectivity for replication traffic

Note that VMware Cloud DR does not carry or configure user traffic to failed over VMs. Please see the VMC on AWS documentation for user traffic options.

The replication connection can be configured over a public interface (either internet or DX) or over a private DX method:

- public internet including AWS Public VIF Direct Connect (DX) – with this option VMware Cloud DR securely replicates VMs via SSL over the public internet to the SCFS service. In the use of an AWS public VIF traffic will enter your private side of the connection and exit the public side of the AWS network.
- private interface via Direct Connect (DX) – alternatively customers can use an AWS Private VIF connection if they wish to keep all replication off public infrastructure either over the internet or within the AWS network.

Recommendation: use the network connection method that provides your organization the appropriate bandwidth and security attachment method. The available bandwidth may be greater than the actual used bandwidth based on a number of other architectural choices (e.g., concurrency) implemented in the Protected Site setups.

Practice – connector instances

While a single connector appliance per Protected site is supported, it is always highly recommended that you deploy at least two connectors. Multiple connectors within a site will improve load balancing and increase the availability of the protection tasks should something happen to one of the connectors.

There is a limit to the amount of additional bandwidth achieved with additional connectors from a given site. Current practices show that ~4-5 connectors from a single Protected Site will likely achieve the maximum network flow bandwidth for that site.

A single VMware Cloud DR instance has an upper limit on the total number of Protected sites or DRaaS Connectors supported.

Recommendation: check the current ConfigMax levels when designing the overall topology of sites and number of connectors planned so that you stay within the bounds of each deployed VMware Cloud DR instance.

Practice – Protected Site topology considerations

A Protected Site is a combination of one or more DRaaS Connectors and one or more registered vCenter servers. The VMs within the inventory of the vCenter server(s) will get protected through the associated DRaaS Connector operations. The DRaaS Connectors gets deployed to an ESXi host/cluster managed by the vCenter server. In some situations, a singe vCenter may be managing multiple physical locations.

The DRaaS Connectors within a Protected Site are stateless with respect to the VMs that they are tasked to protect. For sites where a single vCenter may have multiple physical sites under its control, a connector could end up accessing a VM in a different cluster than where the connector is actively running.

Recommendation: for more complicated Protected site configurations, work with the VMware Cloud DR support team to pin your connectors to the clusters from which they are replicating the VMs.

Practice – PCI Compliant SDDC

- Protected Site SDDC – due to a requirement to access the NSX-T API's during the configuration of the protected site, you will need to work with the VMware Cloud GSS Support Team to have it temporarily turned off. Once the SDDC is protected you are then able to turn PCI compliance back on.
- Recovery Site SDDC – deploying a new SDDC - due to a requirement to access the NSX-T API''s during the deployment of VMware Cloud DR, it is recommended that VMware Cloud DR be deployed and attached to the Recovery SDDC before turning on PCI compliance.
- Recovery Site SDDC – attaching an existing SDDC – due to a requirement to access the NSX-T API''s during the attachment, you will need to work with the VMware Cloud GSS Support Team to have it temporarily turned off. Once the SDDC is attached you are then able to turn PCI compliance back on.

Recommendation: for more complicated PCI compliance configurations, work with the VMware Cloud GSS support team to configure the NSX-T configuration.

Protection Groups

First let’s cover some Protection Group (PG) basics. Each PG policy defines:

- an inventory of VMs based on selection criteria (name patterns, folder location, vCenter tags)
- a choice of snapshot method
    - either Standard Frequency for 4-hour RPO or High-Frequency for 30-minute RPO – note that Standard Frequency snapshots also support guest level quiesce capabilities as outlined in the documentation
- a set of schedules for taking snapshot recovery points and replicating data to the SCFS
- a retention period for each schedule for how long to keep the recovery points in the SCFS inventory

Practice – inventory considerations and DR plans

VMs can belong to more than one PG and the VMs will be processed based on each active PG policy. When VMs are included in multiple PGs, there may be unnecessary snapshot schedules and overlap that are not readily apparent. A PG definition can support up to 2000 total snapshots and thus support a robust schedule / retention plan to avoid creating multiple PGs just for scheduling.

Also, each VM in the PG inventory must be processed by a DR plan Recovery step. These steps include:

- Recover all VMs in the PG as a single Recovery step
- Recover individual VMs from a specific PG
- Recover all remaining VMs, files and groups as a “catch-all” for the DR plan

The granularity applied when building PGs can have a direct effect on the flexibility and processing of DR plans.

Recommendation: do not include VMs in multiple PGs unless there is a specific reason to do so.

If running a DR plan – either for test or failover – and there is already a VM in the Recovery Site that matches a VM in the running plan, there will be a DR plan error. The error could be ignored with manual administrator response.

Recommendation: if protecting more than one site to the same Recovery site, make sure that VM naming conventions are unique across and between sites to avoid inventory conflict should more than one site need to use the Recovery site at the same time.

Practice – inventory mappings to the Recovery SDDC

PGs are the basic building blocks for the inventory specified for failover (and failback) for DR plans. All VMs in a PG are processed by the Orchestrator when that PG is included in a DR plan. In the construction of the DR plan, you will need to establish mappings (folders, resources, networks) for each VM in the PG inventory. The source side of the mapping will be discovered by the Orchestrator, the destination side targets will need to be created in the Recovery SDDC before the mapping can be established in the DR plan.

Note that currently, DR plan mappings are a 1-1 relationship. You cannot map multiple source (Protected Site) items onto the same (Recovery Site) destination items.

Recommendation: create the Recovery SDDC mapping targets (networks, folders, and resources) before

constructing the DR plan, otherwise you may need to exit the plan construction simply to add another desired folder, resource group or network target for the desired mappings.

Practice – what to protect – what not to protect

VMware Cloud Disaster Recovery protects VMs that are in the vCenter inventory that has been registered for a Protected Site and whose VMDK files reside on a supported datastore and VM configuration. For the most part, this covers a broad range of application VMs, but there are some considerations for what to exclude from a VMware Cloud DR snapshot protection method.

It is always good to check the latest documentation and release notes, but the following list identifies many situations that are not covered by VMware Cloud DR snapshots:

- RDMs, Guest Attached, independent disks, and non-Datastore disks
- Service VMs included Tanzu containers
- Windows 11 VMs
- Any VM configured with vTPM
- Shared disks and clustered VMs

Recommendation: Review the inventory that is required for disaster recovery and identify alternate methods (e.g., application level, HCX, etc.) that can be used to enable DR for VMs that are required but not supported with VMware Cloud DR methods.

Practice – Protection Group scope, size, and granularity

Outside of scalability concerns for the total number of VMs to protect, the first choice in designing PG policies is the number and purpose of Protection Groups. In a small environment, it might make sense to put everything into a single or very small number of protection groups.

The size of the environment to protect is driven by a few

- Scope – what is the role of the VMs in the inventory selection
- Size – how many VMs (and vmdk) are being protected
- Granularity – what is the recovery set size – this can affect certain aspects of protection and failover concurrency

Recommendation: Keep the number of protection groups to the minimum needed to satisfy business requirements. Unnecessary creation of more protection groups can add complexity. In other words, keep things as simple as possible as this leads to more reliable disaster recovery.

In summary we have:

Single Protection Group

Pro	Con

Multiple Protection Groups

Pro	Con

Concerns:

- Replication failures impact only the VM’s in the effected PG. Unaffected PG’s will continue to replicate.

There are situations where a single VM or just a few VMs per PG makes a lot of sense, or is a requirement. For example, typical VMs that fit in this category are:

- MsSQL, MS Exchange or MS AD VMs that may all require VSS and the rest of the fleet doesn’t need VSS or High

Frequency Snapshots are in use.

- All VM’s for a single application or possibly multiple app’s serviced by a single other VM (e.g., DB server).

Practice – inventory selection method

When a PG is created, it will have a snapshot schedule or schedules associated with the PG. Every time the schedule runs, it evaluates the inventory specification providing a dynamic inventory capability. Consider how to maximize this auto-protect coverage of PGs and minimize chances that a newly created VM does not get included in the desired PGs.

As stated earlier, there are 3 methods for including a VM into a PG policy. They are: Naming Pattern, Folder Location, and vCenter Tags. You can apply one or more of these in the Protection Group policy setup to provide the proper inventory selection for each scenario.

Naming Pattern

Pro	Con

Folder Location

Pro	Con

vCenter Tags

Pro	Con

NOTE: there is a concern that VM’s which get included in multiple PG’s could lead to excessive snapshot processing and inventory mapping issues when using the overlapped PGs for DR plans.

Recommendation: review Protection Group inventories and schedules and adjust as needed. Snapshot recovery point inventories can also be reviewed from the Virtual Machine view in the Orchestrator.

Practice – snapshot method

VMware Cloud DR provides 2 types of snapshot methods that can be configured for a Protection Group.

Standard Frequency Snapshots – this method allows for a minimum supported frequency of 4 hours, uses VMware vSphere Storage APIs – Data Protection, and provides the capability to quiesce the guest VM through VMware Tools.

Standard Frequency Snapshots

Pro	Con

High Frequency Snapshots

Pro	Con

Quiesce Note: – https://docs.vmware.com/en/VMware-Cloud-Disaster-Recovery/services/vmware-cloud-disaster-recovery/GUID-DBF1DFD5-F956-4ED9-AF06-95664D3AA89D.html

Recommendation: apply the snapshot method that best suits the VM protection and recovery needs for snapshot frequency and guest OS quiesce integration.

Practice – snapshot schedules

When defining the Protection Group policies, the schedule structure is the next most important consideration. The schedules define the RPO (frequency) and the retention (depth of recovery points) you will have for each PG.

What is the ideal number of schedules per Protection Group? The answer varies based on specific use cases and SLA needs, but the goal should be to use the least number of schedules to accomplish the business goals and to try to avoid multiple PGs just to address schedule needs for the same VMs.

It is also good to note here that schedules within a PG that occur at the same time will get combined into a common recovery point. For example: if the hourly snapshot occurs at the same time as the daily snapshot (e.g., midnight) then only one recovery point will be established that addresses both point in time references and will be retained for the longer of the coinciding schedules.

Backup schedules default to midnight, 00 and :30 minutes but you can stagger PG schedules to avoid triggering large numbers of simultaneous backup requests.

What we see in practice are typically 3 levels of schedules as follows:

- Operational recovery – shorter frequency and shorter retentions – for example 30 minutes to dailies
- Ransomware recovery – slightly longer frequencies and retentions – these will typically span weeks or months based on malware dwell time analysis objectives
- Compliance recovery – longer term – both frequency and retention – these tend to be monthly and kept for years

A good practice is to use multiple schedules to give retention depth, simplicity, and efficiency. Here is an example of how what those schedule/retention policies might look like and for what use case?

- Operational DR
    - 30min-12hour @ 1 week – Critical Systems
    - Daily @ 1-2 week
    - Weekly @ 1-2 month
- Ransomware Recovery
    - 30min-12hour @ 1 Week – Critical Systems
    - Daily @ 1-2 Month
    - Weekly @ 6-9 Months
- Long Term Archive / Compliance
    - Daily @ 1 Month
    - Weekly @ 6 Months
    - Monthly @ 1+ Year

With the ability to retain up to 2,000 recovery points per Protection Group, you can see that just a few schedules provide a robust solution and leave plenty of room for expansion if needed. For example. The following 4 schedules combine all 3 needs and consume less than 500 recovery points (even fewer based on some schedule sequence overlap):

- 30 minutes for 1 week = 336
- Daily for 2 months = 60
- Weekly for 6 months = 26

- Monthly for 3 years = 36

Recommendation: avoid constructing Protection Groups simply to enable another schedule while at the same time define the minimum number of schedules to enable the desired protections SLAs.

Recommendation: evaluate business SLAs against snapshot schedule frequency and set schedule to levels that provide the desired RPO – avoid excessive snapshots and replication traffic if not required – in other words don’t set schedule for every 30 minutes if every 2 hours is sufficient.

Recommendation: you can set recovery point retention for longer than periods if unsure – it is easier to delete unwanted recovery points to reclaim space than to go back in time to enable deeper retention.

Recommendation: for environments with several or more Protection Groups, consider staggering the schedule times to allow better interleaving of snapshot and replication traffic over time.

Summary

This document will evolve over time as new practices and areas of consideration are added. It is suggested to use the dynamic html link to the guide whenever possible to reference the current content.

Additional Resources

For more information about VMware Cloud Disaster Recovery, you can explore the following resources:

- VMware Professional Services
- VMware Knowledge Base
- VMware Product Interoperability Matrices

Changelog

The following updates were made to this guide:

Date	Description of Changes
2022/09/2	

Author and Contributors

- Mike McLaughlin – https://vmc.techzone.vmware.com/users/mike-mclaughlin
- Max Daneri – https://vmc.techzone.vmware.com/users/max-daneri
- Don Rush