VMware Cloud Disaster Recovery Technical Overview

VMware BC/DR

Table of contents

VMware Cloud Disaster Recovery Technical Overview

Introduction

VMware Cloud Disaster Recovery is a VMware delivered disaster recovery as a service (DRaaS) offering that protects on-premises vSphere and VMware Cloud on AWS workloads to VMware Cloud on AWS. It efficiently replicates VMs to a scale-out cloud file system that can store hundreds of recovery points with recovery point objectives (RPOs) as low as 30 minutes. This enables recovery for a wide variety of disasters including ransomware. Virtual machines are recovered to a software-defined data center (SDDC) running in VMware Cloud on AWS. VMware Cloud Disaster Recovery also offers fail-back capabilities to bring your workloads back to their original location after the disaster is remediated. Built for IT infrastructure professionals responsible for IT services and their availability, it provides highly reliable, low TCO, easy to use disaster recovery with fast recovery capabilities.

VMware Cloud Disaster Recovery is an add-on feature to VMware Cloud on AWS. VMware Cloud on AWS integrates VMware's flagship compute, storage, and network virtualization products—VMware vSphere, VMware vSAN, and VMware NSX—along with VMware vCenter Server management. It optimizes them to run on elastic, bare-metal AWS infrastructure. VMware Cloud on AWS and VMware vSphere provides the same architecture and operational experience on-premises and in the cloud. VMware Cloud Disaster Recovery extends VMware Cloud on AWS to provide managed disaster recovery, disaster avoidance, and non-disruptive testing capabilities to VMware customers without the need for a secondary site or complex configuration.

VMware Cloud Disaster Recovery utilizes a SaaS Orchestrator to coordinate the VMware SDDC as virtual machines at the protected site are brought into inventory at the recovery site during failover. Using the data replicated from the protected site, these virtual machines assume responsibility for providing the same services.

VMware Cloud Disaster Recovery can protect virtual machines between a customer's data center, and an SDDC deployed on VMware Cloud on AWS. VMware Cloud Disaster Recovery can also protect virtual machines between two SDDCs deployed to different AWS availability zones or regions. The second option allows VMware Cloud Disaster Recovery to provide a fully VMware-managed and maintained Disaster Recovery solution.

Migration of protected inventory and services from one site to the other is controlled by a recovery plan that specifies the order in which virtual machines are started up, the resource pools to which they are allocated, and the networks they can access. VMware Cloud Disaster Recovery enables the testing of recovery plans, using a temporary copy of the replicated data and isolated networks in a way that does not disrupt ongoing operations at either site. Customers can use multiple recovery plans to migrate individual applications or entire sites, providing finer control over what virtual machines are failed over and failed back. This also enables flexible testing schedules.

VMware Cloud Disaster Recovery with VMware Cloud on AWS provides on-demand DR for all VMware workloads and provides recovery alternatives for ransomware attacks and other disasters. In a steady state, customers only pay for replicas stored in the VMware Cloud Disaster Recovery Scale-out Cloud File System. VMware Cloud Disaster Recovery keeps your data safe and secure, and continuous compliance checks of DR plans allow you to execute failover and fail-back confidently.

In the event of a disaster, VMware Cloud Disaster Recovery can be used to provision VMware resources in the form of an SDDC in VMware Cloud on AWS. The stored replicas, which could be minutes old or up to 7 years old, are instantly powered on by mounting the Scale-out Cloud File system directly to the SDDC, resulting in low RTO. Fail-back is fully automated as well. Once the disaster is over, with the click of a button, changed data is compressed and encrypted, which minimizes egress charges and is automatically sent back to the source data center.

Terminology

- Recovery time objective (RTO): Targeted amount of time a business process should be restored after a disaster or disruption in order to avoid unacceptable consequences associated with a break in business continuity.
- Recovery point objective (RPO): Maximum age of files recovered from backup storage for normal operations to resume if a system goes offline as a result of a hardware, program, or communications failure.
- Failover: Method of recovering applications and services to a recovery site when the primary site experiences a failure or disaster.
- Failback: Restoring applications and services from a recovery site back to the primary site after a failover and recovery has occurred.
- Protected site: Site that contains protected virtual machines.
- Recovery site: Site where protected virtual machines are recovered in the event of a failover.
- Protection group: A collection of protected virtual machines that are backed-up and failed over as a group.

- Disaster Recovery (DR) plan: Documented process to recover applications and services in the event of a disaster. A recovery plan in VMware Cloud Disaster Recovery includes one or more protection groups.
- Inventory mappings: Protection groups, VMs, files, vCenter(s), all vCenter folders, compute resources, virtual networks, and IP addresses (individual or ranges) can all be mapped within a DR plan.

Features and Benefits

Features and Benefits of VMware Cloud Disaster Recovery

- Provides familiar features and functionality with enhanced workflows to reduce time to protection and risk
- An easy-to-use disaster recovery/secondary site that is supported and maintained by VMware. This lowers capital costs and makes it easier to protect more virtual machines faster
- Policy-based application-agnostic protection eliminates the need for app-specific point solutions
- Automated orchestration of site failover and failback with a single click reduces recovery times
- Frequent, non-disruptive testing of recovery plans ensures highly predictable recovery objectives
- Enhanced, easy to use, consolidated protection workflow simplifies replicating and protecting virtual machines
- Centralized management of recovery plans replaces manual runbooks
- VM-centric, replication that eliminates dependence on a particular type of storage
- Flexible versioning allows for easier upgrades and ongoing management
- Consistent operating environment on-premises and in the cloud
- Lower DR costs from an on-demand data center in the public cloud
- RPOs as low as 30 minutes to minimize data loss
- Instant restart of any VMware workload from cost-effective cloud-optimized storage, after a ransomware attack or other disaster
- Pure SaaS service with no hardware deployment on-premises and no hardware or software maintenance

Architecture Components

All VMware Cloud DR components (SaaS Orchestrator and Scale-out Cloud File System, are deployed and managed by VMware in an AWS account dedicated to each tenant. Authentication and access controls are unified via the VMware Cloud Services platform. VMware Cloud on AWS SDDCs can be managed directly via the cloud console. The DRaaS Connector is deployed by the DR administrator into the protected site and then managed and monitored by the cloud based services.

VMware Cloud on AWS

VMware Cloud on AWS provides an on-demand VMware Software-Defined Data Center (SDDC), which is used by DRaaS as a cloud DR target. SaaS orchestrator, a cloud DR orchestrator and component of VMware Cloud Disaster Recovery, can provision an SDDC with different trade-offs in runbook RTO and prices (see VMware Cloud SDDC Deployment Modes below). A provisioned SDDC incurs hourly charges. Upon DR test completion, the SDDC can be decommissioned in the VMware Cloud Disaster Recovery UI. VMware Cloud Disaster Recovery performs automated network configurations for both AWS and VMware Cloud to make backups available for spin-up in SDDC. The SDDC is managed in the familiar vCenter interface.

SaaS Orchestrator

SaaS orchestrator is a DR orchestration service that runs in AWS and executes DR plans from new or old replicas. SaaS orchestrator provisions and monitors SDDCs in VMware Cloud on AWS. The SaaS Orchestrator automatically checks your plan for health and compliance every 30 minutes, so you can be confident your DR plan is going to work when you need it.

DRaaS Connector

DRaaS Connector is a downloadable, lightweight virtual appliance that enables customers to protect any VMware workload in just minutes with no new software or infrastructure to deploy. DRaaS Connect enables VMware Cloud Disaster Recovery to orchestrate failover from a VMware Cloud SDDC in one AWS AZ to another AZ or from any on-premises vSphere infrastructure, including SAN, NAS, vSAN, vVol, or local storage to VMware Cloud on AWS.

Designed as a distributed architecture, the DRaaS Connector virtual appliance linearly scales throughput, and parallel processes egress and ingest streams as more virtual appliances are added to the cluster. DRaaS Connect uses VMware APIs for Data Protection (VADP) to create snapshots of the virtual machine disk file and Changed Block Tracking (CBT) to query only for changed blocks, eliminating the need to install in-guest agents across the virtualized infrastructure. VMware APIs for Data Protection (VADP) is VMware's data protection framework that enables centralized and efficient backup of vSphere virtual machines. VADP leverages the snapshot capabilities of VMware vSphere to enable backup without requiring downtime for virtual machines. As a result, backups can be performed non-disruptively without requiring extended backup windows and downtime to applications and users associated with backup windows.

DRaaS Connector delivers crash snapshots of VM disk files. After snapshots are taken, DRaaS Connect compresses and encrypts the new data before sending it to the Scale-out Cloud File System.

After the disaster, when the original recovery site is back online, VMware Cloud Disaster Recovery and the DRaaS Connector orchestrate failback as well. Failback occurs in a similar fashion using VADP within the VMware Cloud SDDC, with unique backup data being further optimized with data compression, which reduces AWS egress bandwidth charges.

DRaaS Connectors can be deployed to multiple on-premises or VMware Cloud on AWS sites to create a fan-in topology to a single Scale-out Cloud File System.

Scale-out Cloud File System (SCFS)

Backups are encrypted and stored in the native vSphere VM format in a highly efficient cloud storage layer called the Scale-out Cloud File System (SCFS) instead of primary storage in a VMware Cloud on AWS SDDC. This harnesses the benefits of cloud storage economics. The Scale-out Cloud File System is optimized, encrypted and cataloged. A unique capability that enables quick RTOs is the ability to "Live Mount" the Scale-out Cloud File System. Live Mount means the ability for hosts in VMC on AWS to boot VMs directly from snapshots stored securely in the Scale-out Cloud File System which is backed by Cloud Native storage, and it acts as an NFS Datastore for the failover SDDC. It supports both short-term as well as long-term retention of immutable snapshots.

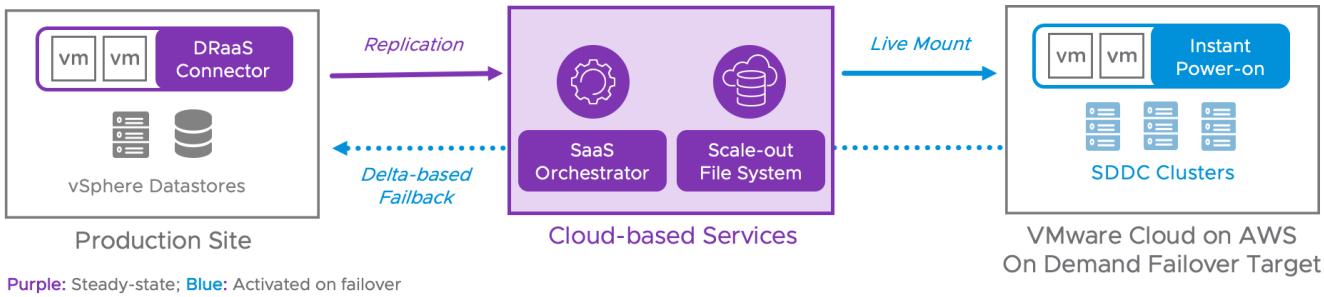

Purple: Steady-state; **Blue:** Activated on failover

Figure 1. VMware Cloud DR architecture.

Topologies/Deployment Options

Just-in-Time Deployment

Just-in-time deployment of a cloud DR site presents an attractive alternative to continuously maintaining a warm standby cloud DR site. With just-in-time deployment, the recurring costs of a cloud DR site are eliminated in their entirety until a failover occurs, and cloud resources are provisioned.

The on-demand nature of public clouds allows DRaaS to drastically reduce the operating costs of DR by deploying the bulk of the DR infrastructure programmatically following a DR event. During steady-state operation, DRaaS maintains a minimal, low-cost AWS cloud footprint to accommodate cloud backups with no ongoing charges for the cloud DR site. The backups are sent to the cloud backup site, and after some processing, land in a cost-effective compressed form. In the just-in-time deployment mode, a cloud DR site is created only following a disaster. VMware Cloud SDDC, a Cloud DR site with a significantly larger server footprint and associated costs, is only deployed immediately before executing a DR plan.

To make that possible, VMware Cloud Disaster Recovery leverages the space and cost efficiencies of the Scale-out Cloud File System. The protected site replicates VMs or protection groups in their forever-incremental format to the Scale-out Cloud File System, which stores them in a compressed and native format. During regular operation, costs are minimized.

Following a DR event, VMware Cloud Disaster Recovery deploys a new VMware Cloud on AWS SDDC and orchestrates the failover to this SDDC as part of a DR plan execution. This process uses a fast, high-bandwidth network link from VMware Cloud on AWS SDDC to the Scale-out Cloud File System to access the replicated VMs. The Scale-out Cloud File System and its connection to the SDDC make the recovered VMs immediately runnable with no rehydration. This capability ensures low RTO restarts, whether they are recent or years old. The recurring charges for the Cloud DR site start accumulating only after the SDDC deployment.

Ahead-of-time vs. just-in-time provisioning of SDDC is a trade-off between costs and RTO. With ahead-of-time SDDC provisioning, SDDC creation latency is eliminated. Just-in-time SDDC provisioning dramatically lowers the costs but increases the RTO by deploying SDDC only in the event a failover.

Ahead-of-Time Deployment

In cases where a DR site has the secondary function of executing non-DR workloads during regular operation, an SDDC can be provisioned before failover.

If the sole purpose of the Cloud DR site is to take over workload execution in the event of a disaster and it remains otherwise unutilized, further significant cost savings are possible with the just-in-time deployment.

Pilot Light with Cloud Bursting Deployment

In Pilot Light mode, DRaaS enables a smaller subset of SDDC hosts to be deployed ahead of time for recovering critical applications with lower RTO requirements.

This deployment model allows organizations to reduce the total cost of cloud infrastructure by keeping a scaled-down version of a fully functional environment always running in warm-standby while ensuring that core applications are readily available when a disaster event is triggered.

With Pilot Light mode, DRaaS presents an option for administrators to add extra SDDC hosts through Cloud Bursting and failover the remaining applications. Expanding the SDDC by adding hosts happens in minutes, providing a lower RTO for all applications than the just-in-time deployment RTO at a fraction of the cost of the ahead-of-time deployment. A full SDDC deployment is a more time-consuming operation with a higher RTO impact than SDDC expansion. Pilot Light mode is an efficient solution with a range of options to balance costs and RTO.

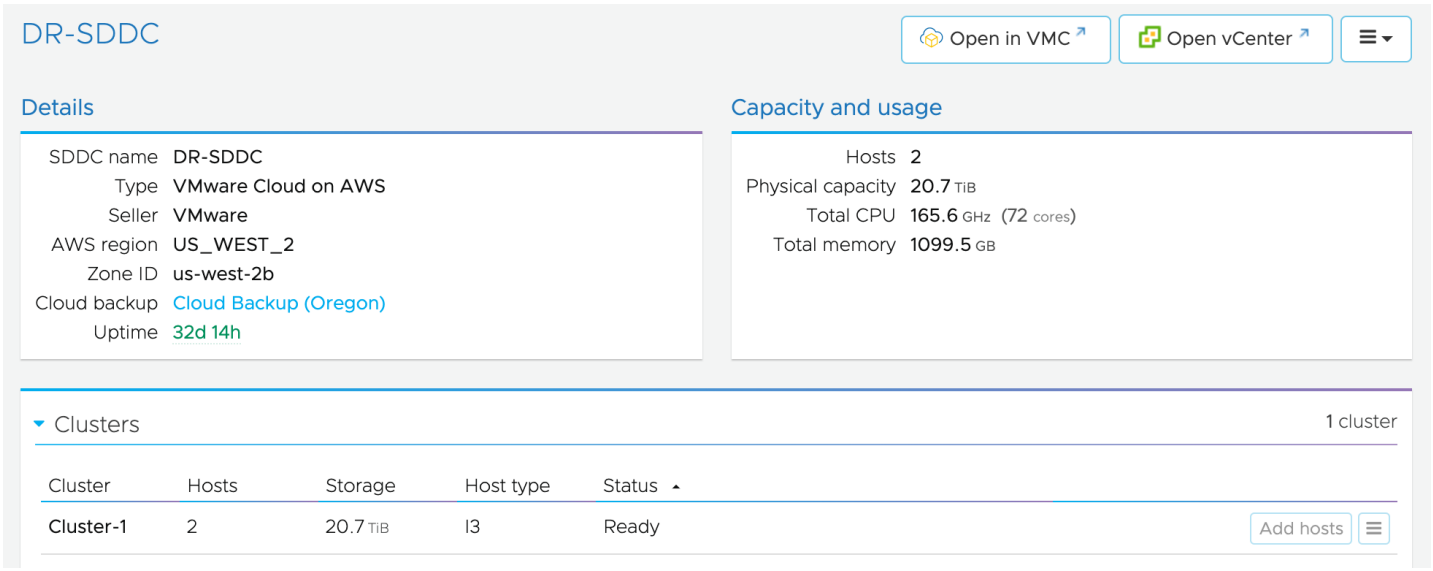

Figure 2. 2-node pilot light SDDC.

Deployment and Configuration

The process of deploying and configuring VMware Cloud Disaster Recovery is simple and logical. This document will cover these steps at a high level. For detailed installation and configuration instructions please see the VMware Cloud Disaster Recovery Installation and Administration Guides.

VMware Cloud on AWS

To start configuring VMware Cloud Disaster Recovery requires an AWS account and the appropriate VMware Cloud on AWS license and credits. The AWS account will then be linked to VMware Cloud on AWS. Depending on the deployment model a VMware Cloud on AWS SDDC may also be required. Activating the service requires a request to VMware through your account team. The specific details for these steps are detailed in the VMware Cloud Disaster Recovery Product Documentation.

Deploy the DRaaS Connector

The SaaS Orchestrator will act as a guide through the process of deploying and configuring the DRaaS Connector(s) as required and getting them connected with the SaaS Orchestrator. The SaaS Orchestrator will also provide guidance for setting up and configuring the Protected Site.

Create Protection Groups

Protection groups are a way of grouping virtual machines that will be recovered together. Often, a protection group will consist of virtual machines that support a service or application. Organizing protection groups by service or application allows testing and failover to be more granular and flexible. VMware Cloud Disaster Recovery supports protecting virtual machines located on any datastore supported by vSphere, local storage, VMFS, NFS, vSAN, or vVols.

A protection group contains virtual machines whose data will be replicated by the DRaaS Connector to the Scale-out Cloud File System following the same protection policy. The protection policy defines the frequency when snapshots are taken and how long the recovery point is retained in the cloud-based Scale-out Cloud File System.

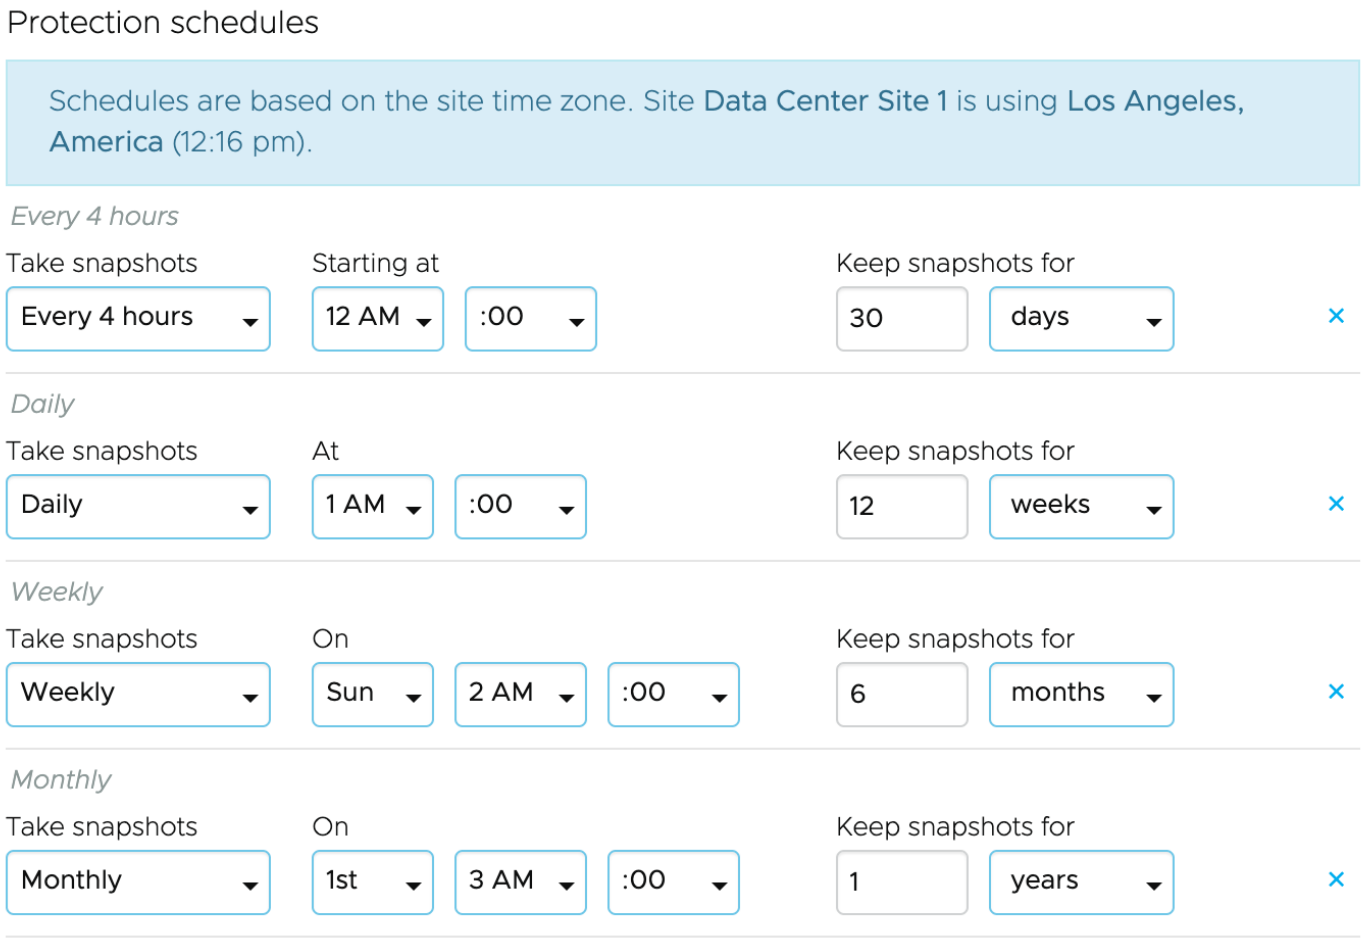

Figure 3. Protection schedule and retention policy.

After the Protected Site is configured and connected to the SaaS Orchestrator protection groups can be created and virtual machines (VM) can start being replicated to the Scale-out Cloud File System. The first snapshot and replication may take a little longer to complete as this is the first copy of data transmitted offsite. Subsequent snapshots will be incremental and much smaller/quicker.

Each VM can belong to more than one Protection Group. And protection groups can belong to more than one recovery plan. Each protection group is specific to a protected site and therefore a vCenter within that site.

VMs can be added to a protection group through a number of different flexible options. VMs can be added to protection groups by naming pattern, by folder or by tag. Schedules can be set for the frequency and timing of replications. Replication frequency can be set to a minimum of 30 minutes (RPO=30 minutes). Multiple schedules can be defined to cover near and longer-term recovery requirements. Each schedule has an associated retention period allowing for flexibility of retention for differing requirements (eg. ransomware, disaster recovery, etc).

Once the Protection Group is set up, the policy runs automatically based on the schedule. Each point-in-time snapshot recovery point is stored in the Scale-out Cloud File System, offsite, on cloud efficient storage. Each recovery point is an immutable snapshot independent of the others in the collection – the delta changed blocks transmitted from the DRaaS Connector are transformed into a synthetic full representation of the VMs in the Protection Group. This underlying full VM representation is what makes the Live Mount capability for quick recovery possible.

Inventory Mappings

There are multiple types of inventory mappings in VMware Cloud Disaster Recovery: Resource mappings, folder mappings, and network mappings. Since virtual machines are being moved from one vCenter to another, these mappings provide default settings for recovered virtual machines. For example, a mapping can be configured between a network port group named “Production-100” at the protected site and a network port group named “Production-200” at the recovery site. As a result of this mapping, virtual machines connected to “Production-100” at the protected site will, by default, automatically be connected to “Production-200” at

the recovery site. Networks to be used only during testing can also be configured in these mappings.

Creating the inventory mappings requires an SDDC. This SDDC can be temporary or part of a pilot-light configuration. Once the SDDC is available, it should be configured with the desired networks, resource groups, and folders to hold the workloads defined by the Protection Groups

Recovery Plans

Recovery Plans in VMware Cloud Disaster Recovery are like an automated run book, controlling all the steps in the recovery process. The recovery plan is the level at which actions like failover, planned migration, testing, and failback are conducted. A recovery plan contains one or more protection groups and a protection group can be included in more than one recovery plan. This provides for the flexibility to test or recover an application by itself and also test or recover a group of applications or the entire site. Active plans have continuous DR health checks performed every 30 minutes and their current status is displayed in the list.

Sequencing

The orchestration events will go in the order of the recovery steps and the recovery steps can be relative to taking action on the entire protection group that is part of the scope of the plan or individual virtual machines that are in those protection groups that are in the scope of that plan, or they may be special actions that are built into the plan orchestration.

Here are some examples of those special actions:

- Wait for user input. Some examples of this might be to remind an operator to place a call to an application owner or modify a network configuration required by the failover.
- Delay a period of time so that something can happen in the environment
- Run a script on the script virtual machine. Some common use cases are calling a script to perform actions such as making changes to DNS or modifying application settings on a virtual machine.

Each step also has pre and post-actions available so that as a virtual machine is being brought into inventory and powered on the recovery plan can take steps before and after that recovery action.

To determine sequencing, VMware Cloud Disaster Recovery doesn't just power on a virtual machine and then move on to the next one. VMware Tools heartbeats can be used to validate when a virtual machine has started successfully and only after that is the next virtual machine in sequence powered on.

Startup Actions

A startup action applies to a virtual machine that is recovered by VMware Cloud Disaster Recovery. Powering on a virtual machine after it is recovered is the default setting. In some cases, it might be desirable to recover a virtual machine, but leave it powered off. Startup actions are applied when a recovery plan is tested or run.

IP Customization

The most commonly modified virtual machine recovery property is IP customization. The majority of organizations have different IP address ranges at the protected and recovery sites. When a virtual machine is failed over, VMware Cloud Disaster Recovery can automatically change the network configuration (IP address, default gateway, etc.) of the virtual network interface card(s) in the virtual machine. This functionality is available in both failover and failback operations.

Workflows

Testing and Cleanup

After creating a recovery plan, it is beneficial to test the recovery plan to verify it works as expected. VMware Cloud Disaster Recovery features a non-disruptive testing mechanism to facilitate testing at any time. It is common for an organization to test a recovery plan multiple times after creation to resolve any issues encountered the first time the recovery plan was tested.

The SDDC used for testing can be deployed just-in-time if provisioning time is acceptable - about 2-3 hours – or always on in pilot-light mode with a minimal footprint and then scale as needed to accommodate failover workload.

A question often asked is whether protection and replication continues during the test of a recovery plan. The answer is yes. VMware Cloud Disaster Recovery Manager utilizes snapshots as part of the recovery plan test process. This approach allows powering on and modifying virtual machines recovered as part of the test while replication continues to avoid RPO violations.

By default, the plan will take the latest good replication point that is stored on the Scale-out Cloud File System. Previous recovery points can be selected as an alternative based on the retention periods specified in the protection group policy details. Each protection group in the plan could have a different recovery point available.

During functional testing runs, there is the option to leave workloads running on the Live Mount datastore, potentially saving time during the test cycle. For performance-related testing, the workloads can be migrated fully into the SDDC, as they would during an actual failover event.

When the VMs are powered on, guest operating system administrators and application owners can log into their recovered virtual machines to verify functionality, perform additional testing, and so on. VMware Cloud Disaster Recovery easily supports recovery plan testing periods of varying lengths - from a few minutes to several days. However, longer tests tend to consume more storage capacity at the recovery site. This is due to the nature of snapshot growth as data is written to the snapshot.

When testing is complete, a recovery plan must be “cleaned up”. This operation powers off virtual machines and removes snapshots associated with the test. Once the cleanup workflow is finished, the recovery plan is ready for more testing or running.

Planned Migration and Disaster Recovery

Running a recovery plan differs from testing a recovery plan. Testing a recovery plan does not disrupt virtual machines at the protected site. There are no dependencies between the protected site and the recovery site when it comes to recovery.

The first step for recovery is to ensure that an SDDC is deployed or to get one deployed if required. This SDDC could be a “just in time” SDDC, it could be a pilot-light SDDC or it could be a fully provisioned cloud site. Whatever makes the most sense based on requirements. For “just in time” and on-demand recovery there isn’t an always-on SDDC running in the cloud. In these situations, it will take approximately two hours to provision and prepare the new SDDC.

For a faster time to recovery, a pilot-light configuration could be used where a few hosts, a minimum of two, are already deployed and running in the cloud to begin supporting the failover and then scale to the site size required to fully support the workloads

A final option is to have the cloud-based SDDC fully deployed and ready for full workload failover. Customers have the flexibility to determine the right match between recovery SLAs and cloud compute cost economics

After the SDDC is in place the recovery point(s) to failover to can be chosen. The recovery point could be the last good replication point or something hours, days, or even weeks older if that is required by circumstances (eg. ransomware, data corruption).

Once the recovery plan has finished the failover and recovered the virtual machines, there is the choice to commit the plan and continue running at the recovery site or to roll back. The rollback process is similar to cleaning up after a test. The recovered VMs are powered off and the SDDC is returned to the state it was in prior to executing the plan.

Failback

After the disaster has been resolved, returning back to normal operations is just as easy as failing over in the event of a disaster.

Simply select the desired plan, duplicate it and then reverse its direction. Once the new plan is created, run through the health checks to make sure that everything's ready to failback. Changes may need to be made to the plan or the environments depending on what happened while operating in the cloud or resolving the on-prem datacenter. The health check process will provide guidance on what needs to be addressed. Then the failback plan can be executed.

The failback process uses change block tracking to minimize the amount of data that needs to replicate back to the on-prem site through the Scale-out Cloud File System back to the DRaaS Connector.

At the end of the failback, all virtual machines are restored to the same point in time that the cloud instance was last running. At that point, the related cloud compute resources are no longer needed, and the VMware Cloud on AWS SDDC could be reduced in size or even eliminated, depending on requirements.

It is important to note here that a failback operation is a planned activity and there will be some downtime of the applications. This will occur during the snapshot and replication stages of this process and that will depend on how much has changed during the DR operation period as well as on network bandwidth.

Reporting and Health Checks

When workflows such as a recovery plan test and cleanup are performed in VMware Cloud Disaster Recovery, history reports are automatically generated. These run reports document items such as the workflow name, execution times, successful operations, failures, and error messages. History reports are useful for a number of reasons including internal auditing, proof of disaster recovery protection for regulatory requirements, and troubleshooting. Reports can be exported as PDF files.

VMware Cloud Disaster Recovery conducts automatic, regular, continuous DR health check operations. All active/ready plans are checked every 30 minutes for consistency pertaining to

- The primary site configuration and health
- The failover site configuration and health
- The DR Plans orchestration steps
- General VMware Cloud DR component health

Based on administrator-driven alert settings, any DR Plan issues will automatically generate email alerts to the proper resources.

Summary

VMware Cloud Disaster Recovery with VMware Cloud on AWS is a comprehensive cloud-based disaster recovery service that protects VMware vSphere environments on-premises and in the cloud. It leverages the execution and operational efficiencies of a single integrated data stack to automate and orchestrate all aspects of DR. The solution is much more streamlined and significantly less resource-intensive than legacy DR solutions, resulting in lower RPO and RTO for cloud and on-premises environments.

VMware Cloud Disaster Recovery delivers a single disaster recovery service for enterprises without contract management overhead. The solution has a single provider and billing, and enterprises can use VMware Cloud Disaster Recovery for automated and user-defined DR plans with failover and failback from the public cloud.

VMware Cloud Disaster Recovery uses a Scale-out Cloud File System to optimally utilize the benefits of cloud storage, while the integrated data and orchestration stack enables consistency checking of the entire environment, which dramatically reduces errors when a disaster occurs.

With VMware Cloud Disaster Recovery, customers receive a complete solution that delivers comprehensive support, simplified purchasing, and billing, which eliminates the cost and friction of multiple point solutions. Customers receive everything needed for improved, on-demand disaster recovery for all VMware workloads with optimized RPO and RTO at a lower cost in a single solution.

Next Steps

Automate and Orchestrate Your DR Plans with VMware Cloud Disaster Recovery

Make VMware Cloud Disaster Recovery a part of your vSphere deployments and improve your virtual machine availability and reduce your risk. Take the VMware Cloud Disaster Recovery Hands-on Lab today and see how simple it is to get the benefits of automated and orchestrated protection of your critical virtual machines as an an integrated part of your IT platform.

Additional Resources

For more information about VMware Cloud Disaster Recovery, please visit the product pages. Below are links to documentation and other resources:

- Product Documentation
- FAQ
- Hands-on Lab

Providing Feedback

VMware appreciates your feedback on the material included in this guide and in particular, would be grateful for any guidance on the following topics:

How useful was the information in this guide? What other specific topics would you like to see covered?

Please send your feedback to docfeedback@vmware.com, with “VMware Cloud Disaster Recovery Overview” in the subject line. Thank you for your help in making this guide a valuable resource.

About the Author

Cato Grace is a Senior Technical Marketing Architect at VMware. He works on business continuity and disaster recovery solutions in the Storage and Availability group. Cato started as a VMware customer in 2005 and has also worked as a VMware partner. He has worked in Technical Marketing at VMware since 2013.

- Cato blogs here:
- Follow Cato on Twitter: @vCatoGrace