VMware Cloud Disaster Recovery Technical Overview

VMware DRaaS

vmware by Broadcom © VMware LLC.

Table of contents

vmware by Broadcom © VMware LLC.

vmware by Broadcom © VMware LLC.

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 from both disasters and ransomware attacks. It efficiently replicates VMs to a Scale-out Cloud File System (SCFS) 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 provide 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 needing 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.

Recovery 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 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 even many years old, are instantly powered on from the Scale-out Cloud File system that is already "Live Mounted" 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, minimizing egress charges, before 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.

© VMware LLC.

- 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.
- Isolated Recovery Environment (IRE): A dedicated environment for the recovery, analysis, validation and remediation of VMs separate from the production environment.
- Next-Gen Anti Virus (NGAV): Modern cyber threat tools (software) capable of performing behavior analysis of running workloads as well as traditional vulnerability and malware signature detection.

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

Features and Benefits of VMware Ransomware Recovery (add on solution)

- Quickly identify next-gen ransomware strains with embedded behavior analysis
- Safely iterate recovery points in a managed, on-demand Isolated Recovery Environment
- Leverage an integrated, step-by-step guided ransomware recovery workflow
- Easily assign VM network isolation policies to prevent lateral movement of ransomware
- Deploy a single product to address the entire ransomware recovery operation
- Boost collaboration between Infrastructure and Security teams

© VMware LLC.

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.

© VMware LLC.

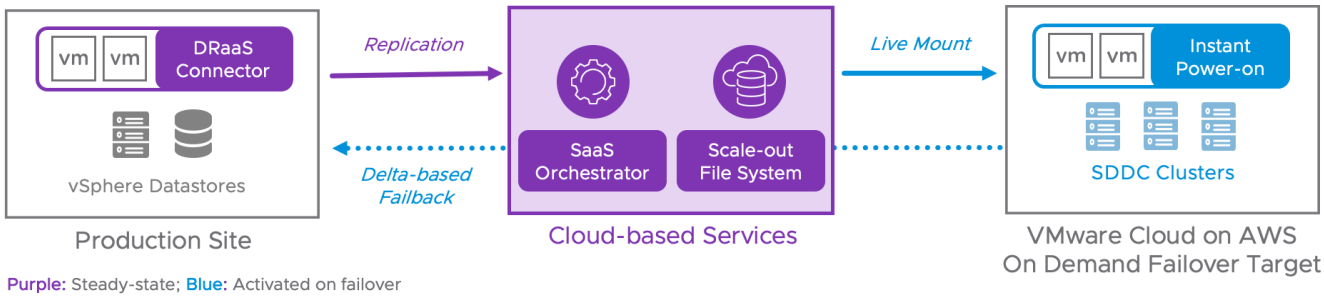

Purple: Steady-state; Blue: Activated on failover

Figure 1. VMware Cloud DR architecture.

Ransomware Services

SDDC / Isolated Recovery Environment (IRE)

The VMC Software-Defined Data Center (SDDC) environment is the core of the recovery site solution. This is the case whether it is being used as a full-site disaster recovery site for VMware Cloud DR, or as the special purpose Isolated Recovery Environment (IRE) for VMware Ransomware Recovery scenarios. The SDDC can be deployed 'Just-in-time' (on-demand) or can be deployed as a minimal-sized always-on configuration (also known as Pilot Light). In either case, the configuration can take advantage of cloud elasticity and scale to the desired size (compute resources or storage capacity) to serve the appropriate recovery needs.

NOTE: For ransomware recovery situations where the VMs are being iterated through in the guided workflow, the final size of the required SDDC being used as the IRE may be smaller than the SDDC size than would be required to support a full site disaster failover.

NSX Advanced Firewall

VMware Cloud on AWS includes the NSX – Networking and Security – features needed for basic network management of the SDDC within the VMC and broader internet environment. These networking capabilities can easily be leveraged for a more secure and connected configuration of the SDDC for production use, disaster recovery scenarios, as well as ransomware recovery needs. The integration of particular features such as NSX Advanced Firewall support will provide the network isolation functionality discussed later in this guide.

NOTE: For the Advanced Firewall capabilities used to provide some of the integration, there may be additional per-use charges incurred during ransomware plan operations – either for testing or recovery.

Carbon Black Cloud

Another service available through the VMware Cloud portal is the Carbon Black Cloud Endpoint and Workload Protection platform. This advanced cybersecurity solution provides improved capabilities against modern threats with the addition of behavioral analysis of running workloads. This is critical in the detection and remediation of ransomware attacks and will be used in the overall Ransomware Recovery solution described in this guide.

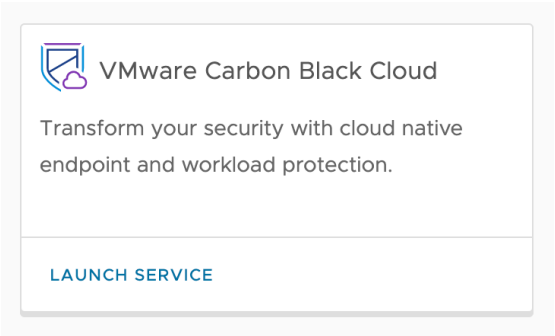

Next Gen Anti-Virus (NGAV)

The latest malware attacks are getting more sophisticated and using techniques not easily detected by simple system scans that only look for known vulnerabilities or existing virus signatures. The new attack vectors exploit fileless attacks that are only observable when a system is running and exhibits an undesirable behavior. NGAV solutions need to include behavioral analysis as part of the validation methods applied. We'll look more closely at the integration of the Carbon Black Cloud capabilities later in this guide.

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.

© VMware LLC.

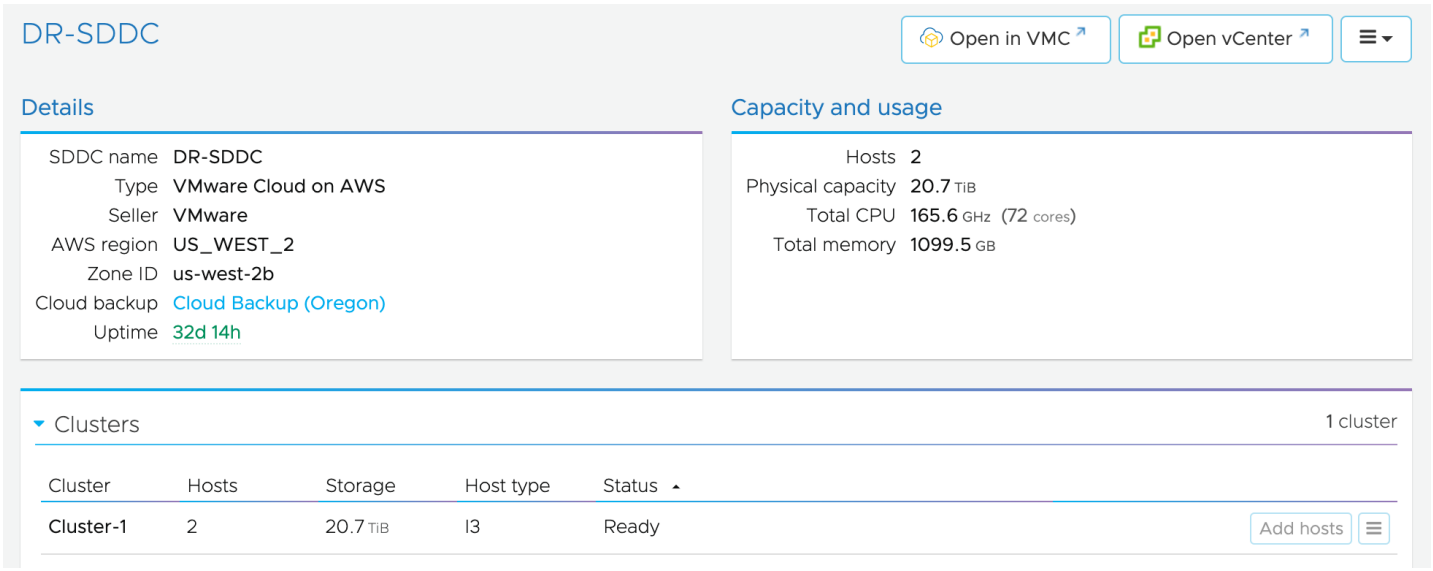

Figure 2. 2-node pilot light SDDC.

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.

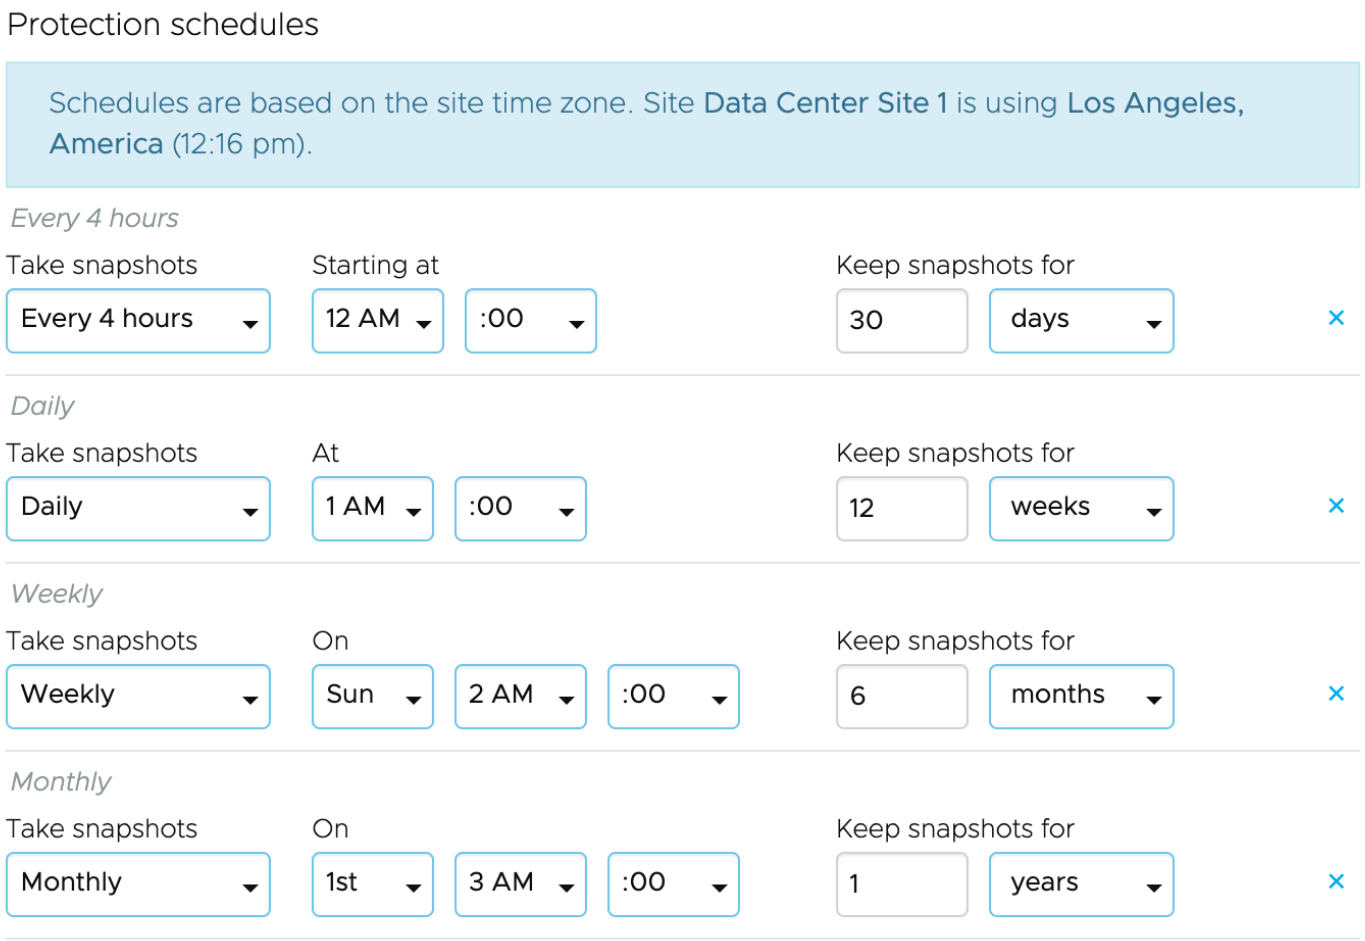

Figure 3. Protection schedule and retention policy.

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

Site Disaster 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.

Site Disaster 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.

© VMware LLC.

Ransomware Recovery Setup

Before covering the details of the Ransomware Recovery solution, let's review some of the key differences between a ransomware recovery and a more traditional site disaster recovery. Many of the differences are outlined in the documentation and other sources, but a simple synopsis is worth including here:

1. The production site is still available, but the content (applications and VMs) has been compromised – the goal is to get the production site operational
2. Finding and validating the best recovery point(s) typically involves an iterative and piecemeal workflow as opposed to a more linear runbook style of recovery, and bringing order to a more complicated process is essential
3. Performing recovery tasks quickly and in a controlled environment is critical to minimize the overall recovery time and reduce the risk of reinfection of the applications being handled

We will cover these topics and a few other related considerations in the remaining parts of this guide.

The figure below outlines one of the fundamental challenges of ransomware recovery. The point of detection and impact to the systems – shown in the vertical red lines on the right side of the timeline – may be separated by days, weeks, or longer from the initial attack point when the ransomware was introduced – shown as the skull on the left side of the timeline.

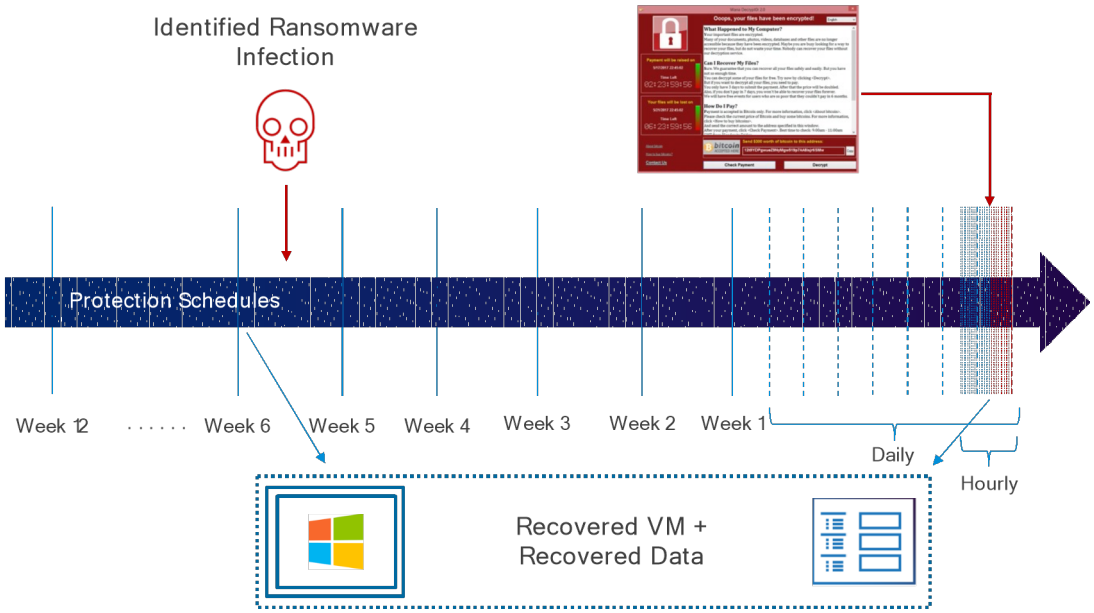

551x302

It may be necessary to iterate over the recovery points identified in the protection schedules to find the best usable VM and data set(s) for the recovery and validation tasks.

Preparation

Before the ransomware recovery processes can be accomplished, organizations need to do pre-work to ensure everything will run as it should when the recovery time comes. This pre-work includes scoping and building the necessary Protection Groups and Recovery plans and ensuring they are tested and ready for use.

Testing Considerations

Testing Recovery plans is critical to improving readiness. It allows organizations to identify and address potential deficiencies to prepare for the worst, regardless of an unplanned outage or ransomware attack. During testing tasks, the Recovery SDDC, or IRE, needs to be in place and operational. Organizations may require multiple classes of networks in the IRE to bring infected VMs back into inventory safely. All of these multiple classes of networks need to be validated, cleaned, and integrated into recovery mode operations. The need for network isolation can be met with NSX Advanced Firewall methods incorporated into the IRE framework with the Ransomware Recovery setup.

Organizations can consider using different mappings for test/failover, especially while planning development stages. Folder mapping can be leveraged to provide easier identification (e.g., quarantine folder) of inventory items during testing or partial failovers.

Another goal during testing is determining what additional service or utility VMs, specific to what ransomware may impact, are needed in the IRE when provisioned. These can be failed over as initial actions during early response activities. They can be

included in the specific DR plans where they may be used, either manually or automatically, with script VMs.

NOTE: The status of Recovery Points used during testing phases can be captured in the Badges construct provided with the Ransomware Recovery capabilities. Badges are covered in a later section.

© VMware LLC.

Ransomware Recovery Workflows

While recovery from ransomware has many similarities with traditional site disaster recovery, there are also important differences that necessitate developing specialized recovery workflows. This section covers the main stages relevant to ransomware recovery scenarios:

1. Response – declaring a ransomware disaster and initiating the recovery workflow tasks
2. Recovery (phase 1) – recovering and validating individual VM recovery points into the IRE and preparing them for eventual recovery back into the Production site
3. Recovery (phase 2) – failing back the remediated production workloads from the IRE back into a clean production site

Response

Once ransomware has been detected in the production (protected) environment, there are many steps that need to be taken to safeguard the remaining infrastructure and respond to the attack. In addition to the tasks dictated by the security team and the ransomware defense systems, there are actions organizations can take before they start the VM recovery tasks, which will ultimately improve the overall success of the recovery.

Isolated Recovery Environment

To be prepared, organizations must work off the premise that their backup data has been infected. Most likely, they will not know the precise moment of an infection, so they often assume they should restore their backup data from a point in time before it was encrypted in the production environment. Unfortunately, that backup copy may be impacted and could re-introduce ransomware into that environment, causing more harm than good.

An alternative to simply trying to restore backups into the production environment is to restore backup data to an Isolated Recovery Environment (IRE). This provides a safer environment, so the ransomware can be fully remediated before migrating virtual machines back into a production environment. An IRE offers a staging area for restored virtual machines isolated from other networks. This means the remediation process can occur without external ransomware triggers and without the risk of (re)infecting different workloads.

If an IRE is not already prepared and running, the organization should deploy and configure the Recovery SDDC that will be used for the IRE.

The application networks defined during the preparation stage should now be constructed, and their connectivity isolation should be verified. For a ransomware recovery, deploying a new, “clean” SDDC for use as the IRE to prevent reinfection is best. The VMware Cloud DR management interface provides the mechanisms to perform this deployment. Provisioning and configuring the SDDC will take approximately two hours before it can be used as the IRE.

If the SDDC is already in operation (e.g., in Pilot Light mode) and connected to the protected site that has been infected, then it may need further analysis and processing by the security and networking teams (e.g., separation, validation, cleaning) before it can be used as the IRE.

Once deployed, the IRE may be used to run VMs that have been fully recovered until the entire recovery process (phase 1) is complete.

With VMware Cloud DR, the recovery SDDC can also be scaled up after it is created to accommodate fluctuations in incoming recovery workloads. Organizations need enough compute capacity to service the VMs, and enough storage capacity to hold the VMs. Running the workloads directly from the SCFS datastore during ransomware recovery can avoid the storage vMotion process for the SDDC vSAN datastore.

In general, running testing from the SCFS to expedite plan run times is helpful by eliminating the storage vMotion background step. When running an actual ransomware recovery failover, the workloads will be left on the SCFS. This approach also reduces some of the SDDC scalability requirements that are driven by storage consumption needs.

The creation and deletion of an SDDC for recovery efforts can be performed in the VMware Cloud DR UI. Minimizing the number of tools and interfaces needed helps simplify the entire disaster recovery workflow and reduces risk or downtime associated with the cloud provisioning processes.

Once the remediation process is complete, workloads can be migrated back into the original production site without fear of reintroducing the ransomware. When finished with the IRE, it can be deleted. The automated processes of creating and deleting an SDDC help minimize ransomware recovery times and the cost of maintaining an IRE. Organizations create and pay for an IRE only when they need it.

Another option is to augment the IRE with a greenfield VMware Cloud on AWS environment as a DR site to recover workloads and

run applications once remediated. This option removes the pressure on the IT organization to return an impacted data center to service quickly. It buys the organization some time, so they can conduct the appropriate forensics on the existing infrastructure without prematurely eliminating points of inspection in a rush to recover. Sometimes the only way to recover VMs from an on-site backup would be to delete the infected machines from the datastore to make storage space available for a restore. Once deleted, all the forensic data on those VMs is lost. A greenfield, clean operating environment eliminates this necessity, allowing organizations to make workloads and applications available in the cloud to keep business running while they investigate and remediate the attack.

Review and Organize the DR plans

As part of the initial response to ransomware, it is good practice to check the status of the recovery components and overall recovery plan compliance. In cases where the recovery SDDC was just deployed, recovery plans must be activated and checked. Organizations will also need to resolve any issues with the recovery plans that may prevent a successful ransomware recovery procedure. For instance, if the protected site has been segmented or disconnected from VMware Cloud DR and the DR site, organizations can expect to receive a DR plan compliance alert notifying them the protected site is unreachable. This will not affect the failover process since the DR plans work primarily from the VMware Cloud DR (UI and SCFS) and VMware Cloud on AWS (SDDC) elements.

The DR plans control the VM power on state at recovery time. Under normal disaster recovery operations, having the VMs automatically powered on may be desirable The VMs will be powered on for ransomware recovery as part of the recovery. The VMs will be powered on for ransomware recovery when they are recovered to the IRE.

Review snapshot status

As described earlier, the Protection Groups run automatically on schedules pre-defined by the organization’s policies. At this point in the ransomware response, organizations should review the current snapshot status and make any appropriate adjustments. For example, performing a manual snapshot for protection groups whose next cycle is hours away might be helpful. The snapshots are independent and immutable in the SCFS, so there is no harm in capturing a more current recovery point that may benefit data extraction or forensics later.

If the protected site where the ransomware has been detected is disconnected from the VMware Cloud DR cloud components, organizations can stop the snapshot and replication schedules. This may also help control alerts from the environment while other ransomware tasks are being performed. The objective is to confirm where things stand with respect to the available recovery points and any automated actions that VMware Cloud DR is serving.

NOTE: During a ransomware recovery operation, the Protection Group (PG) expiration task is suspended for any PG involved in the recovery plan. This will generate warnings at the Protection Group level that can safely be ignored while working on the recovery. It also prevents any snapshots from inadvertently being expired that may have proven helpful in the other VM or data recovery tasks.

Recovery Point Selection

Part of ransomware recovery involves locating a valid restore point that an organization can use as the best recovery option. This recovery point must be either before the infection happened or before the data has been encrypted/compromised so that is can still be used.

As a result, it’s not as simple as restoring the most recent backup data, because it is possible the ransomware has been in the environment or encrypted the data without being detected for several hours, days, or possibly even weeks. This means organizations may need to use an older restore point to be sure the data isn’t compromised. To accomplish this, the recovery solution must have robust data retention capabilities that can support hundreds of restore points, ranging from just a few hours to possibly several months in age.

Finding a valid recovery point might require multiple recovery operations before the organization locates data that is not compromised (encrypted) and still viable. Even if the organization locates a copy of the data that is not encrypted, it is possible that the recovery point still contains the ransomware. As a result, organizations may need to quickly restore and validate alternate recovery points, from a handful to a dozen or more, to be certain they have a clean one they can use to recover.

There are some additional capabilities added to the ransomware recovery workflow UI around recovery point selection. The list of available recovery points is shown graphically in the timeline also showing any previously applied badges marking the status of certain points as shown in the figure below.

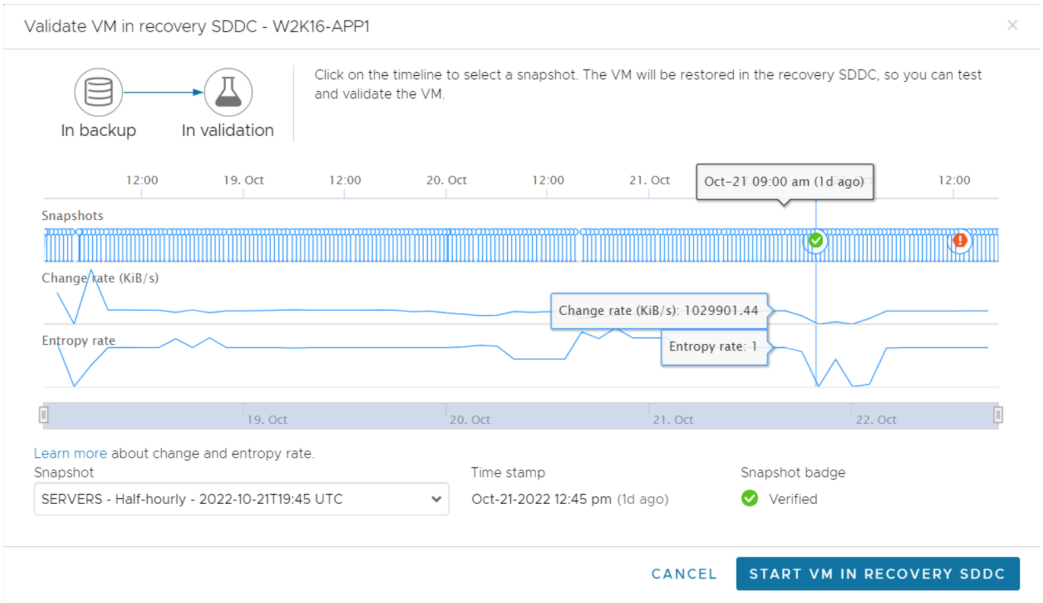

In addition to this badge status tracking, the timeline also shows data change rate and snapshot entropy data for each snapshot taken along the protection timeline. By reviewing this information, recovery admins can make better choices about where to begin the validation process.

VMware Cloud DR provides an instant virtual machine power-on capability that can simplify and accelerate this restore and validation process. Virtual machines can be started directly from the Scale-out Cloud File System (SCFS) with no need to assemble or migrate data into a usable image before powering on a virtual machine. If the selected recovery point is not ideal, an alternate recovery point from the SCFS inventory can be easily selected and brought into service for validation. This capability reduces the amount of time spent when iterating over different recovery points, which results in overall faster recovery times.

Naturally, the restoration and use of data from an older recovery point means the organization will lose any changes that occurred to that data after that particular backup. The main objective is to find a valid recovery point that minimizes this data loss. This means finding the most recent copy of a virtual machine that can be remediated and used as the baseline for the recovery. The file and folder restore capabilities can be used to extract subsets of the virtual machines data set into this recovery point as needed to reduce overall data loss.

As with any restore process for data protection and disaster recovery, recovery point validation procedures should be well documented and practiced routinely to minimize recovery times when an actual ransomware attack occurs.

Recovery Point Iteration

One of the final activities of the response phase is to identify the recovery points that will be used for the actual recovery operations. The security team can use their tools to help identify recovery points available in VMware Cloud that are potential candidates, investigating if malware has infected the systems or if there are other concerns that could render the points unusable. The details of their investigations can be recorded in the recovery detail annotations as part of their record. Organizations can quickly iterate over alternate recovery point candidates, since they are mounted directly from the SCFS and can be quickly examined in the IRE.

Guided Workflow

VMware Ransomware Recovery provides a simple, well defined workflow for each VM to be processed. This workflow encompasses four basic steps as shown in the figure below:

When a Recovery Plan is activated for Ransomware Recovery, the VM inventory, as defined by the members of the Protection Groups used in the plan, is determined and placed into the VM list starting in the backups stage, ready for processing through the workflow.

The first step of the guided workflow has already been discussed above – selecting the initial recovery point for validation. During the remainder of the guided workflow, there are several other activities that are involved before a VM can be properly staged for

recovery.

Security and Vulnerability Analysis

As previously mentioned, organizations need to proceed with recovery on the assumption that backup data is also infected with ransomware. As a result, they cannot simply recover an older copy of an infected virtual machine back into production, for fear that it will reinfect the environment. Instead, they will need to ensure the data and applications are clean before moving the virtual machines back into production.

The ideal solution and environment for conducting this analysis and remediation activity will include the following items:

- An IRE(s) to stage recovered workloads into a controlled and safe environment for remediation
- A virtual machine recovery solution that has robust retention capabilities to give the organization multiple recovery points that can be quickly deployed in the IRE(s)
- Tools to detect and remove a wide variety of ransomware variants from workloads both statically and during operational (running) conditions
- Documented procedures to help teams quickly:
    - Locate valid recovery points for remediation
    - Extract data from a recovered workload
    - Iterate on the restore process

NOTE: To ensure organizations benefit from the latest list of known ransomware, CBC will use updated signature lists at the time of recovery to detect any new variants that might not have been known when the backup was originally performed. CBC also includes malware behavioral detection to further improve its ability to uncover and address new strains of file-less ransomware.

Behavioral Analysis

Once the selected recovery point is brought into the IRE and instrumented with the security sensors, it will begin running, constrained by the default network isolation level of “Quarantined + Analysis” settings of the integrated NSX Advanced Firewall settings. In addition to the vulnerability and malware signature scans, the VM will be monitored with behavioral analysis. The vulnerability and malware scans are fixed tasks based on the current content and configuration of the VM. The behavioral analysis is a bit more open ended and should run for a period of time and under varying conditions. These variations could be in network isolation changes and interoperability with other VMs within the IRE.

Guest File Restore

When a limited amount of data (files) from an alternate recovery point is desired, the files can be extracted from SCFS recovery points with guest file restore utilities. These utilities allow a recovery administrator to view the VM contents of specific recovery point snapshots and extract specific data sets from those backup images into a local download location or a cloud based (S3) location, without having to bring the VM into the SDDC/IRE inventory. The extracted files are presented as a zip archive that can be downloaded, unpacked, validated, and merged into the working recovery copy of the VM.

The exact mechanics of guest file restore and subsequent data validation (cleaning) tasks varies from case to case. Sometimes it can be useful to have an intermediary IRE VM (e.g., one for Windows and one for Linux guests) equipped with all the required ransomware utilities, stationed in the IRE for initial data extraction and preparation. From here, the data can be safely merged into the baseline VM with significantly less risk.

Badges

During the validation workflow steps, either testing or actual recovery, it is possible to assign a predefined badge to the specific recovery point under consideration. The badge is a user applied attribute for the specific recovery point. The predefined values for the snapshot badge are shown in the figure below.

© VMware LLC.

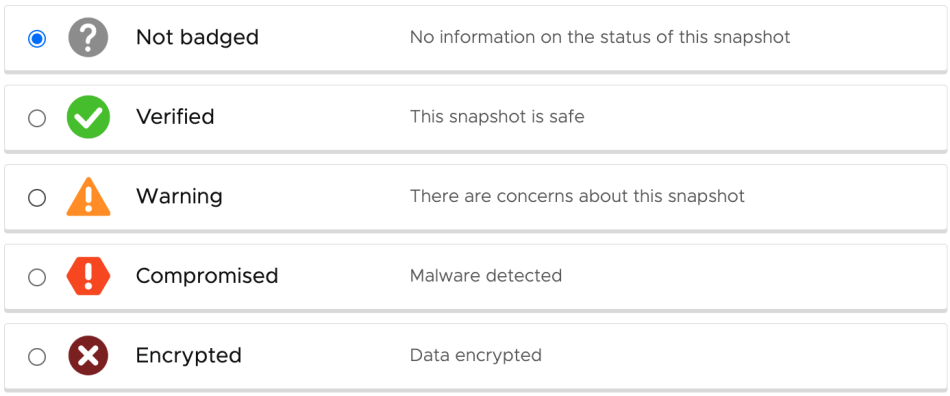

As was noted in the recovery plan testing section, applying these badges to snapshots evaluated during testing or practice runs will end up in the snapshot timeline (covered in the Recovery Point Selection section) and provide more information and guidance to the rest of the team when dealing with actual recovery tasks.

User Annotations

During the iterations of the workflow for an active recovery plan, the recovery administrator can add notes to the workflow for each iteration conducted. These notes improve communications, status and results across the team and further facilitates a comprehensive divide-and-conquer approach to scaling the recovery tasks.

Network Isolation / Firewalls

When a VM is initially brought into the IRE and started up, and the Carbon Black Cloud (CBC) sensors are enabled, it begins in a "Quarantined + Analysis". In this state, the VM can only access or be accessed by the NGAV tools – all other north/south or east/west network traffic is limited through NSX Advanced Firewall rules. With the integration provided in VMware Ransomware Recovery, the network isolation control is easily managed with "push-button" selection options.

This short video on Network Isolation goes into more details about the network isolation capabilities available and the leverage of the NSX Advanced Firewall capabilities.

Final Recovery

Stage the Recovery Point

Once the team – both infrastructure and security – has completed the validation of the VM and performed the necessary patching, remediation, data updates, and malware removal, then that instance of the VM can be staged for final recovery back to the protected site. A new staging snapshot can be created in the SCFS that contains any updates to the VM made during validation and also has the security sensors removed to avoid any conflict when recovered.

Prepare the Protected Site

One objective of any recovery from cloud-based restores is to minimize the cloud egress impact — both in cost and transfer time. To support this, VMware Cloud DR and VMware Ransomware Recovery will attempt to restore the original protected site VM back to the same recovery point that was used for the recovery task. To make this functionality possible, organizations are advised to not delete the VMs from the original protected site. Instead, they are asked to simply power them off once the recovery begins, or sooner, if warranted for site security.

If for some reason, the original site is unavailable for the final recovery step, an alternate recovery site can be supported. In this case, the configuration of the alternate site will need to be mapped the same as the original site.

NOTE: If an alternate site is used for the final recovery, there will be additional time (and bandwidth) involved as the baseline VM data will not be present and need to be transferred from the SCFS to the alternate site.

Run the Recovery workflow process

The final step of the guided recovery workflow for each VM is to recover that staged VM from the SCFS back to the original site. The orchestration of the failback activities can be monitored from the UI.

Once all the recovery plans used to process the VMs to the VMware Cloud on AWS IRE have been completed, everything should be back to its original protected site (production) location. Organizations can then use the IRE for forensics, as needed, or simply begin cleaning up the remaining items, such as the temporary VMs and other services, and decommissioning or scaling down the

VMware Cloud on AWS SDDC environment.

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.

© VMware LLC.

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.

VMware Cloud Disaster Recovery can be extended to include additional services, features and capabilities that are purpose-built to address the complexities of recovering VMs that have been impacted by modern ransomware attacks.

© VMware LLC.

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

© VMware LLC.

Copyright © 2005-2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.