# An Introduction to Disaster Recovery (Site Recovery)

VMware DRaaS

# Table of contents

# An Introduction to Disaster Recovery (Site Recovery)

## Introduction

VMware's Disaster-as-a-Service (DRaaS) for the AWS cloud offering is a part of VMware's Site Recovery portfolio. DRaaS enables customers to protect and recover applications without the requirement for a dedicated secondary site. It is delivered, sold, supported, maintained and managed by VMware as an on-demand service. IT teams manage their cloud-based resources with familiar VMware tools—without the difficulties of learning new skills or utilizing new tools.

### About VMware Site Recovery

VMware Site Recovery Manager (SRM) is a business continuity and disaster recovery solution that helps you plan, test, and run the recovery of virtual machines between sites either designed as protected and recovery or bi-directional sites.

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. VMware SRM is an addon made available as Disaster-Recovery-as-a-Service (DRaaS) with VMC on AWS. SRM works in conjunction with VMware vSphere Replication to automate the process of migrating, recovering, testing, re-protecting, and failing-back virtual machine workloads.

SRM servers coordinate the operations of the VMware vCenter Server™ at two sites. Virtual machine data is configured to be replicated to the other site which can be used later in starting copies of protected virtual machines on the second site during planned maintenance or when a disaster strikes.

The Site Recovery can be configured in two modes:

- Customer's data center and an SDDC deployed on VMware Cloud on AWS
- Between two SDDCs deployed to different AWS availability zones or regions. This option allows VMware Site Recovery to provide a fully VMware managed and maintained Disaster Recovery (DR) solution.

VMware Site Recovery utilizes VMware Site Recovery Manager and vSphere Replication. For the VMware Cloud on AWS instance, this software is automatically installed and configured by VMware when the add-on is enabled.

For the customer site, the customer can install VMware Site Recovery Manager on a Microsoft Windows server or deploy as an appliance (starting With SRM 8.1). vSphere Replication is deployed as an appliance.

VMware Site Recovery utilizes vSphere Replication for transferring data between sites. The following requirements must be met on the customer site for installation:

- VMware vCenter
- One or more vSphere ESXi hosts
- Infrastructure services like DNS, DHCP and Active Directory, must be in place at both the protected and recovery sites.

See the Compatibility matrix for detailed requirements and version compatibility.

## Features and Benefits

- Lower Total Cost of Ownership on recovery site datacenter with a DR site service
- On-Demand DRaaS that works regardless of complexity
- Protect your most sensitive workloads
- Recover fast and minimize downtime
- Eliminate manual process with automated orchestration of site failover and failback with a single-click
- Application-agnostic protection eliminates the need for app-specific point solutions
- 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
- Recovery plans are available on the HTML-based UI
- vSphere Replication integration delivers VM-centric, replication that eliminates dependence on a particular type of storage
- Flexible versioning allows for easier upgrades and ongoing management

## Use Cases

Though the primary use case of a recovery manager is to perform a recovery of a site in case of disaster, it can also perform additional tasks. For all use cases and situations, VMware Site Recovery supports non-disruptive testing of recovery plans in network and storage isolated environments.

This provides the ability to test disaster recovery, disaster avoidance, or planned migrations as frequently as desired to ensure confidence in the configuration and operation of recovery plans.

### Disaster Recovery

Disaster recovery or an unplanned failover is what VMware Site Recovery was specifically designed to accomplish. This is the most critical but least frequently used use case for VMware Site Recovery. Unexpected site failures don't happen often but when they do a fast recovery is critical to the business. VMware Site Recovery can help in this situation by automating and orchestrating the recovery of critical business systems for partial or full site failures ensuring the fastest RTO.

### Disaster Avoidance

Preventive failover is another common use case for VMware Site Recovery. This can be anything from an oncoming storm to the threat of power issues.

VMware Site Recovery allows for the graceful shutdown of virtual machines at the protected site, full replication of data, and ordered start up of virtual machines and applications at the recovery site ensuring app-consistency and zero data loss.

### Upgrade Patch and Testing

The VMware Site Recovery test environment provides a perfect location for conducting an operating system and application upgrade and patch testing. Test environments are complete copies of production environments configured in an isolated network segment which ensures that testing is as realistic as possible while at the same time not impacting production workloads or replication.

### Topologies

VMware Site Recovery can be used in a number of different failover scenarios depending on customer requirements, constraints, and objectives. All of these arrangements are supported and easily configured.

- Active-Active
- Active-Passive
- Bi-directional

In context of DRaaS, there is a different naming used depending on the type of SDDC usage.

- Single site - One on-prem VMC on AWS cloud instance replication
- Multisite - Multiple on-prem or VMC on AWS cloud to One VMC on AWS cloud instance( acts as a dedicated SSDC instance for replication)

## Benefits of Using vSphere Replication with SRM

VMware Site Recovery utilizes vSphere Replication to move virtual machine data between sites. vSphere Replication can utilize any storage supported by vSphere so there is no requirement for storage arrays, similar or otherwise at either site.

- Build Flexible Configurations
- Customize the recovery point objective (RPO) from 5 minutes to 24 hours
- Use multiple point-in-time (MPIT) recovery to revert to previous known states
- Eliminate Storage Lock-In
- Use Microsoft Volume Shadow Copy Service (VSS)
- Use Linux file system quiescing
- Optionally enable data compression to further reduce network bandwidth consumption

### Protection Groups

A protection group consists of the virtual machines that support service or application that together provides a function. For example, an application might consist of a two-server database cluster, three application servers, and four web servers. In most cases, it would not be beneficial to fail over part of this application, only two or three of the virtual machines in the example, so all nine virtual machines would be included in a single protection group.

Creating a protection group for each application or service has the benefit of selective testing. Having a protection group for each application enables non-disruptive, low-risk testing of individual applications allowing application owners to non-disruptively test disaster recovery plans as needed. Note that a virtual machine can only belong to a single protection group. However, a protection group can belong to one or more recovery plans.

### Recovery Plans

Recovery Plans in VMware Site Recovery are like an automated playbook, controlling all the steps in the recovery process. A recovery plan contains one or more protection groups and they 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.

The following configurable options are available in a recovery plan:

### Priority Groups

There are five priority groups in VMware Site Recovery. The virtual machines in priority group one are recovered first, then the virtual machines in priority group two are recovered, and so on. All virtual machines in a priority group are started at the same time and the next priority group is started only after all virtual machines are booted up and responding.

This provides administrators one option for prioritizing the recovery of virtual machines. For example, the most important virtual machines with the lowest RTO are typically placed in the first priority group and less important virtual machines in subsequent priority groups. Another example is by application tier - database servers could be placed in priority group two; application and middleware servers in priority group 3; client and web servers in priority group four.

### Dependencies

When more granularity is needed for startup order dependencies can be used. A dependency requires that before a virtual machine can start, a specific other virtual machine must already be running. For example, a virtual machine named “app01” can be configured to have a dependency on a virtual machine named “DB01” - VMware Site Recovery will wait until “DB01” starts before powering on “app01”. VMware Tools heartbeats are used to validate when a virtual machine has started successfully.

### Shutdown and Startup Actions

Shutdown actions apply to the protected virtual machines at the protected site during the run of a recovery plan. Shutdown actions are not used during the test of a recovery plan. By default, VMware Site Recovery will issue a guest OS shutdown, which requires VMware Tools and there is a time limit of five minutes. The time limit can be modified. If the guest OS shutdown fails and the time limit is reached, the virtual machine is powered off.

Shutting down and powering off the protected virtual machines at the protected site when running a recovery plan is important for a few reasons. First, shutting it down quiesces the guest OS and applications before the final storage synchronization occurs. And second, it avoids the potential conflict of having virtual machines with duplicate network configurations on the same network

A startup action applies to a virtual machine that is recovered by VMware Site Recovery. We have two options to choose from

here. Power and keep powered off. This option would used in accordance on the need of the VM. In some cases we may just want to recover and keep the VM powered off for later use.

### Pre and Post Power On Steps

As part of a recovery plan, VMware Site Recovery can run a command on a recovered virtual machine after powering it on. A common use case is calling a script to perform actions such as making changes to DNS and modifying application settings on a physical server. VMware Site Recovery can also display a visual prompt before or after any step in the recovery plan. This prompt might be used to remind an operator to place a call to an application owner, modify the configuration of a router, or verify the status of a physical machine.

### 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 Site 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.

### Testing and Cleanup

After creating a recovery plan, it is beneficial to test the recovery plan to verify it works as expected. VMware Site 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.

A question often asked is whether replication continues during the test of a recovery plan. The answer is yes. vSphere Replication utilizes virtual machine snapshots at the recovery site 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.

To keep test networks isolated VMware Site Recovery supports two different options. One, automatically created networks on each host. This has the advantage of not requiring any additional networking configuration. However, because this option limits VM connectivity it is best used for testing the function of the recovery plan, not for testing application functionality.

The second network testing option is manually created test network(s) that are configured to duplicate production networks at the recovery site without a connection to the production portion of the network. This is easily possible at the VMware Cloud on AWS site through the tight integration with VMware NSX. See the administration and configuration guide for details about configuring this. This option requires additional configuration upfront and provides the ability to fully test the functionality of both the recovery plan and the application.

At this point, guest operating system administrators and application owners can log into their recovered virtual machines to verify functionality, perform additional testing, and so on. VMware Site 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 clean up workflow is finished, the recovery plan is ready for testing or running.

Note: When a test clean up is performed all accumulated replication data is consolidated with the replica base disks. This can take time and can increase RTO if a disaster happens and the consolidation after test clean up is still running.

### Planned Migration and Disaster Recovery

Testing a recovery plan does not disrupt virtual machines at the protected site. When running a recovery plan, VMware Site Recovery will attempt to shut down virtual machines at the protected site before the recovery process begins at the recovery site. Recovery plans are run when a disaster has occurred and failover is required or when a planned migration is desired.

When you click the Run Recovery Plan on the VMC/SDDC console, you must choose between a planned migration or disaster recovery.

In both cases, VMware Site Recovery will attempt to replicate recent changes from the protected site to the recovery site. It is assumed that for a planned migration, no loss of data, is the priority.

A planned migration will be canceled if errors in the workflow are encountered. For disaster recovery, the priority is recovering workloads as quickly as possible after disaster strikes. A disaster recovery workflow will continue even if errors occur. The default

selection is a planned migration, which includes the following steps:

1. Try to synchronize the virtual machine storage
2. Shut down the protected virtual machines. This effectively quiesces the virtual machines and commits any final changes to disk as the virtual machines complete the shutdown process
3. Synchronize storage again to replicate any changes made during the shutdown of the virtual machines.
4. Replication is performed twice to minimize downtime and data loss.

Note: During a planned migration, after the replication online sync is completed and the source VMs are shut down, power-on method on them is explicitly disabled by SRM. There is no rollback for planned migration. The only workaround to power on the original VMs at the source site is to remove them from vCenter inventory and re-register (and effectively lose all the VM identify - tasks, events, replication configuration, configuration in SRM).

## Re-Protect and Failback

After the failover of the VMs to the DR site is completed and workloads are running as usual, you must ensure that the primary site is up and running and then get the latest copy of these workloads replicated back to production/primary site.

SRM provides a feature called re-protect which is used when the primary site is ready to receive the latest changes of workload VMs from DR site.
Use re-protect to sync the latest data from the recovery site before getting workloads running again on the primary site.

Consider a use case where the threat of rising floodwaters from a major storm is expected at the primary site. Using VMware Site Recovery, you can migrate the virtual machines from the protected site to the recovery site. Fortunately, the floodwater subsides before any damage was done leaving the protected site unharmed. Now that primary site is back, you can now sync the latest changes by reversing the replication direction from recovery site to primary site and after all data is replicated, we can failback workloads to original site.

A recovery plan cannot be immediately failed back from the recovery site to the original protected site. The recovery plan must first undergo a re-protect workflow. This operation involves reversing replication and setting up the recovery plan to run in the opposite direction.

## History Reports

When workflows such as a recovery plan test and clean up are performed in VMware Site Recovery, history reports are automatically generated. These 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 to HTML, XML, CSV, or a Microsoft Excel or Word document.

## Terminology

- Recovery time objective (RTO): The 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): The 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.
- Protected site: Site that contains protected virtual machines. This can be either the customer's data center or VMware Cloud on AWS.
- Recovery site: Site where protected virtual machines are recovered in the event of a failover. This can be either the customer's data center or VMware Cloud on AWS.

Note: It is possible for the same site to serve as a protected site and recovery site when replication is occurring in both directions and VMware Site Recovery is protecting virtual machines at both sites.

## Author and Contributors

Sharath has been working in the IT industry for over 15 years primarily on SDDC and cloud related technologies. Go to https://vmc.techzone.vmware.com/users/sharath-bn for more information.