# An Introduction to VMware Disaster Recovery and Business Continuity

VMware DRaaS

# Table of contents

# An Introduction to VMware Disaster Recovery and Business Continuity

## Introduction

### Audience

This document provides high-level strategic guidance targeted towards Cloud Architects whose focus is on the design of VMware disaster recovery (DR) solutions for on-premises software-defined datacenter (SDDC) or any other VMware Cloud SDDC.

This document does not discuss the configuration and operations of the VMware disaster recovery solutions on any particular cloud provider.

### Overview

Business continuity is a process to ensure that business operations are not affected and in case of a disaster, the downtime to operations is minimized. Disaster recovery is a part of business continuity planning. This document highlights the disaster recovery planning aspect of business continuity with respect to infrastructure components.

Disaster recovery should be one of the primary factors that must be taken into consideration when you are planning an SDDC deployment in either on-premises or in the cloud.  Datacenter availability is a major factor when planning a new datacenter or migrating an existing datacenter to the cloud. Disaster recovery is the process to get the business up and running when a disaster strikes. With a carefully planned solution, we can execute a proactive failover and avoid the disaster altogether and in the event of a disaster, we can perform recovery while minimizing  data loss and downtime.

### Why Disaster Recovery?

What is the need for a disaster recovery solution when organizations already perform backups regularly?

A backup solution restores a single workload or a group of workloads, either to the original location or a new location, with the option to choose multiple point-in-times during restoration. This option of restoring from a backup is time-consuming and may not be optimal for a disaster recovery situation. Traditional backup solutions enable the backup to be saved within the datacenter or on the cloud. However, in case of a data center failure, this solution is not dependable as we do not have a dedicated SDDC from which the data would be recovered. There is no recovery SDDC in the case of a backup solution.

In the VMware disaster recovery solutions, the data from the primary SDDC or the protected datacenter is replicated to the recovery SDDC or a cloud storage continuously (at set intervals). This replicated data can be used to recover workloads quickly in an automated fashion, and hence the time required to recover an application is much faster compared to the traditional backup/restore process which involves restoring VMs manually.

An effective disaster recovery solution also reduces administrative time and complexity with the use of automation, which is critical in cases of a rolling disaster such as a ransomware attack. This solution can also help businesses to meet other objectives such as:

- Infrastructure management will help with inventory and process mapping
- Workload protection
- failover and failback procedures( boot up order)
- Testing and Infrastructure maintenance  aid in risk and impact analysis
- A disaster recovery site readiness contributes to Service-level Agreement (SLA).

## Selecting the Solution

To replicate the data from one site to another, two solutions are currently available:

- Storage-based replication
- Hypervisor-based replication

Traditional recovery methods involved setting up two sites and enabling storage-based replication between these two sites. Replicating data was completely handled by the storage array. Logical Unit Number(LUN) created on the storage array was configured to be replicated over to the secondary site. During disaster recovery, a snapshot of the replicated LUN was created on the recovery site and mounted manually to use the data. This solution lacked automated recovery.

VMware Site Recovery Manager (SRM) filled the gap with the storage array plugin. Not only was the process of creating array snapshots and mounting them automated, but SRM also provided additional granularity on the recovery of the virtual machine. You could now create protection groups to group dependent workloads, create a recovery plan, and set a priority order for the recovery of workloads. However, what lacked was choosing to just replicate an individual VM rather than a complete LUN. Placement of VM on a particular LUN was to be designed based on the need for replication.

Storage-based replication has a dependency on having similar grade storage hardware on the production and recovery site. These requirements forced users to use one particular storage vendor on both sites along with an additional replication license which increased the overall cost of the disaster recovery solution.

VMware introduced Hypervisor/Host-based replication (HBR) to overcome the challenges of storage-based replication. With HBR, you can replicate the data of an individual virtual machine from one site to another. HBR provides the granularity and flexibility of enabling replication on individual virtual machines placed on any supported storage for the hypervisor which addressed the major challenge of the storage agnostic solution. For example, you may choose to have iSCSI or Fibre-attach storage with VMFS on the production site and choose an NFS or VSAN-based solution on the recovery site.

The hypervisor at the production site where the virtual machine is running can perform the initial replication (first full copy) and also track the changes and replicate them over to the recovery site. This is done with the help of a paired replication appliance deployed on both sites. An HBR-enabled machine can also utilize virtualization features such as vSphere HA or vMotion.

One more method of host-based replication is done using the CBT technology on the VM where the changes are replicated to the destination site.

Currently, host-based replication supports two destinations for replicating the data.

1. SDDC Datastore
2. Cloud-based Storage

Note: The difference in these two solution used for replicating the data into the SDDC datastore versus  cloud storage, is cloud storage involves additional rehydration (rebuild the data into the VM format) process. The cloud storage destinations are block storage where the replicated data is saved, and these are rehydrated automatically with the automation available with the replication solutions. However, we must ensure regular failover testing is performed in any of the solutions used to ensure we comply with the  disaster recovery strategy.

## Design Considerations

Consider the following features and properties before you begin designing a recovery solution:

### Architecture

VMware currently offers two distinct disaster recovery solutions:

- **VMware Site Recovery (VSR):** This is a fully-managed DRaaS solution that is delivered with vSphere replication and VMware Site Recovery. With this service, along with enabling the add on, you deploy a vSphere appliance in your on-premises vSphere environment. You can then pair sites and replicate your critical VMs running in the on-premises environment to a Cloud SDDC. The figure below shows an example of architecture where the recovery site is hosted on VMware Cloud.

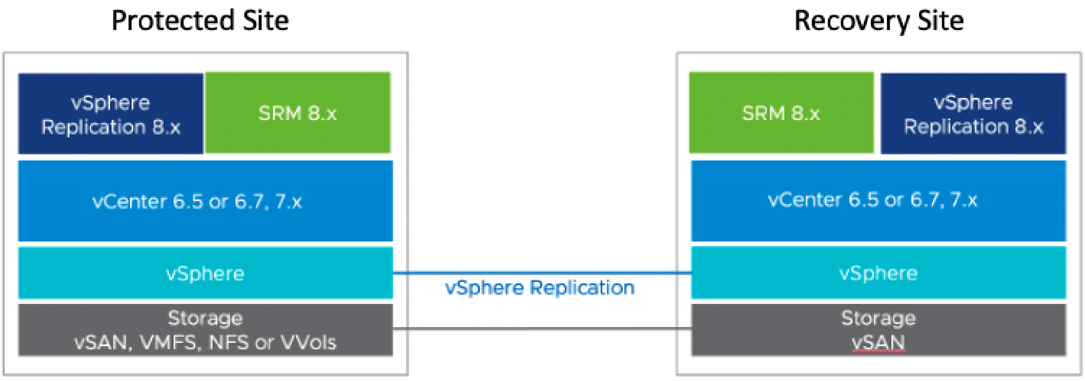

- **VMware Cloud Disaster Recovery (VCDR):** This is VMware's on-demand disaster recovery service that is delivered as an easy-to-use SaaS solution and offers cloud economics to help keep your disaster recovery costs under control. The target SDDC can be created immediately before performing a recovery and not upfront, while also supporting the replications in the steady-state. The DRaaS connector is deployed as a virtual appliance that replicates the data to a Scale-Out Cloud File System (SCFS). This volume is mounted when we perform recovery as a live-mount datastore on the  SDDC. Since the VMs are already in an ESXi supported format recovery is handled at ease. This offering is currently available only with VMC on AWS.

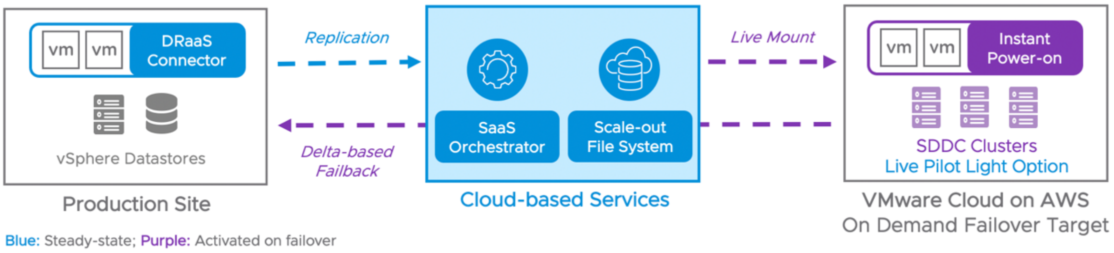

### Use Cases

### On-premises to Cloud (hybrid)

The majority of enterprise workloads currently run in an on-premises datacenter. While the idea to migrate to the cloud is gaining momentum, there can be various reasons to continue using the on-premises datacenter. For workloads which are currently running on-premises, the  disaster recovery SDDC can still exist in the cloud making it a hybrid use case. Currently, only hypervisor-based replication is supported in this model.

### Cloud to Cloud(C2C)

Migrating the workload to the cloud does not eliminate the need for a disaster recovery site. Having a disaster recovery SDDC is still as critical as when running workloads on-premises. There are several high availability features configurable when running production workloads on cloud SDDC enabling resiliency from hardware/zone failure. However, none of these can be assumed to be a replacement for a disaster recovery site.

In both use cases, the primary datacenter is replicating the data to either a  disaster recovery site SDDC’s storage or cloud storage.

### Advantages of a Disaster Recovery Site in the Cloud

- Available on demand
- Dynamic scaling when executing a disaster recovery plan
- Integrated disaster recovery solutions available on the SaaS subscription model
- Eliminate end user need to ensure the disaster recovery site is audited and in compliance with standard requirements (Cloud provider will ensure compliance against standard case)
- Terminate or scale-down recovery site when workloads are failed back to production site thus reducing the cost of ownership (TCO)

### Alternate Uses of Recovery Site/SDDC

With multipoint-replication( usually termed as multi point-in-time which has multiple copies at different intervals) configured on a virtual machine, the data can be recovered to an earlier known,error-free state after a ransomware attack . This reduces the time for recovery and also data loss depending on the RPO configured for their virtual machine.

- Disaster avoidance – With a recovery site setup beforehand, we can implement failover even before disaster strikes. This is useful in cases where we know there could be physical disasters such as flooding, storms, etc.
- Running non-critical workloads on a disaster recovery site – The recovery site is a  datacenter residing in a different geographical location from the primary site.  In the case of such a distributed model, we can run non-critical or temporary workloads on the recovery site, and turn them off in case capacity is needed to failover production workloads, thus reducing the load on the production datacenter.
- Upgrade and patch testing -  With a virtualized disaster recovery solution, the test failover can bring the most recent state of replicated production VM in a test network on a disaster recovery site enabling us to apply the patch/upgrade and validate any effects before being applied to the production workload.
- Ransomware - Ransomware is one of the latest threats to any organization. Without  disaster recovery, there are two ways to deal with ransomware attacks:
    - Pay the attacker with no guarantee of recovering
    - Restore from backup, which may take a long time and result in loss of data.

### Sizing

When designing a recovery SDDC, you must account for the compute, storage, and network requirements that are necessary to keep the critical applications up while the primary site is recovered. Recovery sites can also be used in a distributed model, which means the resources in the recovery site do not need to sit idle all time. Depending on the available resources, we can run the on-premises non-production workload on the recovery site.

**Note -** If you are running the recovery site in distributed mode, you must account for additional resources that will be required for saving the replicated data.

See the table below for the different site considerations for SDDC planning:

| | |
|---|---|
| Account for overhead on each datacenter if used in distributed mode | |
| Network Connectivity Considerations | |
| Perform inventory mapping | |

## Recovery Plan Considerations

The process of recovery from an application perspective requires careful planning. Discuss with your application owners the following:

- Limits on recovery plans, protection groups, etc – How does this affect the overall design and how to optimize it.
- Limits on concurrent recoveries to avoid burst mode - How does this impact disaster recovery events and offer strategies for prioritization of recoveries during a disaster recovery event.
- Restart priority, recovery order, reboot order
- Split-brain breaker(witness)
- Application synchronization (ensure multiple VMs have their replication synchronized to ensure statefulness).
- Application recovery procedures before an application is brought back online.
- Acceptable Performance levels during a disaster recovery  scenario.
- Test plan methodology & frequency.

## Access Management

Creating solution-specific roles with access boundaries is a good security practice. Create roles and permissions to ensure only specific users have the right privileges can execute certain tasks on the shared service during recovery. For example, test-recovery, DNS record updation and planned failover.

## Terminology

- **Recovery Time Objective (RTO):** RTO is the targeted duration of time and a service level in which a business process must be restored as a result of an IT service or data loss issue, such as a natural disaster.
- **Recovery Point Objective (RPO):** RPO defines the maximum acceptable age that the data that can be recovered from the recovery storage in case of a disaster. The lower the RPO, the closer the replica's data is to the original. However, lower RPO requires more bandwidth between the source and target locations, and more storage capacity in the target location depending on the Point-in-time configured on VM.
- **Point-in-Time Instance:** You define multiple recovery points (point-in-time instances or PIT instances) for each virtual machine so that when a virtual machine has data corruption, data integrity, ransomware-encrypted data, or host OS infections, administrators can recover and revert to a recovery point before the compromising issue occurred.