Well-Architected Design: VMware Cloud Disaster Recovery Planning and Preparation

Table of contents

Well-Architected Design: VMware Cloud Disaster Recovery Planning and Preparation

Introduction

VMware Cloud Disaster Recovery safeguards your virtual machines, whether they are hosted on-premises or in a VMware Cloud on AWS Software-Defined Data Center (SDDC), by replicating them to the scale-out cloud file system and restoring them to a VMware Cloud on AWS SDDC.

Scope of Document

There are several aspects to consider while implementing a VMware Cloud Disaster Recovery software-as-a-service (SaaS) solution that encompasses all infrastructure components. This design addresses VMware Cloud Disaster Recovery planning and preparations considerations.

Summary and Considerations

Use Cases	
Prerequisites	
General Considerations/Recommendations	
Performance Considerations	
Network Considerations/Recommendations	VMware Cloud Disaster Recovery network requirements.
Cost Implications	If the protected site is a VMware Cloud on AWS SDDC, the egress replication data will be subject to network costs depending on AWS pricing. Please see the VMware Cloud Disaster Recovery pricing page for further information on the costs involved.
Document Reference	
Last Updated	April 2023

Recovery Objectives Considerations

A disaster recovery solution is designed to assist businesses recover their IT infrastructure, systems, and data following a disruptive event. The recovery objective of a disaster recovery solution refers to the specific recovery goals that an organization aims to achieve after a disaster. This includes the recovery time objective (RTO) and the recovery point objective (RPO).

- Recovery Time Objective: The RTO is the targeted duration of time and a service level in which a business process must be restored because of an IT service or data loss issue, such as a natural disaster.
- Recovery Point Objective: RPO defines the maximum acceptable age that the recovered data can have. The lower the RPO, the closer the replica's data is to the original. However, setting a 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 can define multiple recovery points (point-in-time instances or PIT instances) for each virtual machine. If a virtual machine experiences data corruption, data integrity, or host OS infections, administrators can recover and revert to a recovery point before the compromising issue occurs.

Consider an example where a company has a critical application that must be operational within 2 hours of a disaster, and the maximum tolerable data loss is 30 minutes.

In this case, the RTO for the application is 2 hours, which means that the company must have a disaster recovery plan in place that can recover the application within 2 hours of a disaster. The company will need to ensure that all the required infrastructure, applications, and data are available and can be restored within 2 hours.

The RPO for the application is 30 minutes, which means that the company must have a disaster recovery plan in place that can restore the data up to the last 30 minutes before the disaster occurred. The company must ensure that the data is frequently backed up and can be restored up to the last 30 minutes before the disaster occurs.

It's critical to understand that RTO and RPO are not simply numbers, but they represent the actual time and data that a company can afford to lose during a disaster. Therefore, it's crucial to perform regular tests and ensure that the disaster recovery plan can meet the RTO and RPO requirements.

Infrastructure Planning

Greenfield and brownfield planning are two approaches to consider when implementing the VMware Cloud Disaster Recovery service. Greenfield planning involves deploying new VMware Cloud on AWS SDDC infrastructure that is designed explicitly for disaster recovery. On the other hand, brownfield planning involves integrating the disaster recovery service with the existing VMware Cloud on AWS SDDC infrastructure.

Regardless of the approach, proper planning is critical to the success of implementing VMware Cloud Disaster Recovery. This includes understanding the organization's recovery objectives and identifying the appropriate replication technology. Additionally, the failover plan must be designed to minimize downtime and ensure data availability. By considering both greenfield and brownfield planning approaches and correctly planning the implementation, organizations can ensure a successful deployment of the VMware Cloud Disaster Recovery service.

The scope of this design only covers planning for a VMware Cloud on AWS SDDC as this is the only supported recovery target for VMware Cloud Disaster Recovery.

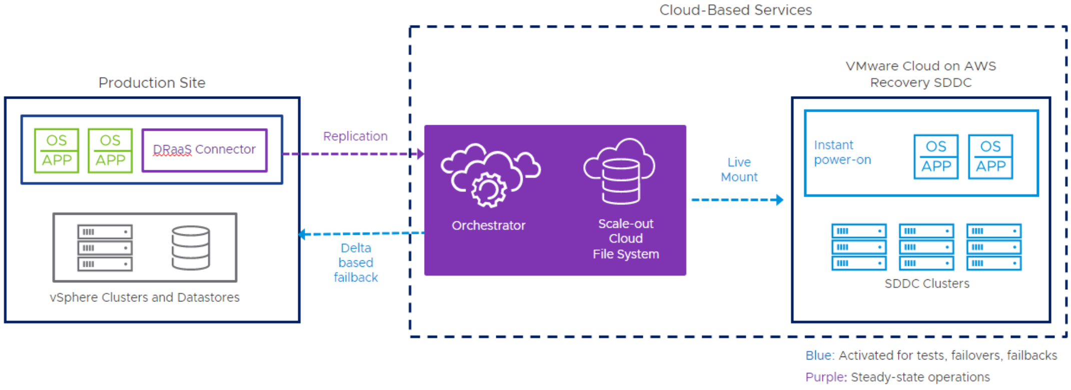
Cloud-Based Services
Production Site
OS
APP
OS
APP
DRaaS Connector
vSphere Clusters and Datastores
Replication
Delta based failback
Orchestrator
Scale-out Cloud File System
Live Mount
VMware Cloud on AWS Recovery SDDC
Instant power-on
OS
APP
OS
APP
SDDC Clusters
Blue: Activated for tests, failovers, failbacks
Purple: Steady-state operations

Figure 1 – Disaster Recovery Overview

Deployment Models

VMware Cloud Disaster Recovery offers two target SDDC deployment models for disaster recovery: On-Demand and Pilot Light.

On-Demand

The on-demand disaster recovery deployment model is presented as an easy-to-use SaaS offering with cloud economics to help keep disaster recovery costs as low as possible. The target SDDC is built immediately before executing a recovery plan instead of having a pre-built stand-by target, while still allowing steady-state replications. The recovery SDDC is not created during the replication process in this architecture. The recovery SDDC is only built when a test or scheduled failover is planned.

One thing to keep in mind when selecting the on-demand mode of consumption is the time required during a recovery process. The recovery time will include the time required for the VMware Cloud on AWS SDDC deployment and configuration as well as the VM recovery process itself. The decision to use this strategy must be thoroughly evaluated and assessed throughout the planning stage. Capacity requirements and availability should also be considered in the event of a broader outage.

Pilot Light

In a pilot light deployment model, businesses can protect their critical applications and data in the cloud, without the need to maintain a full disaster recovery site. This approach also allows businesses to easily test their disaster recovery plan and ensure their critical applications and data can be recovered quickly in the event of a disaster. In a pilot light deployment, you provision a minimal footprint of failover capacity, which can be scaled up during the recovery process. To ensure faster recovery times, VMware recommends a minimum of a two-node SDDC for pilot light configurations.

This SDDC may be scaled based on resource requirements during a test or actual failover. The required number of nodes in a pilot light deployment is determined by the workload and application SLA running on them. To determine the number of minimum hosts in your recovery SDDC, evaluate the application needs and their SLA.

Design Considerations for VMware Cloud Disaster Recovery On-Demand and Pilot Light

Component	On-Demand	Pilot Light
Recovery Time Objective (RTO)	RTO includes recovery SDDC deployment as well as recovery plan execution. RTO may take several hours.	RTO is the total amount of time required for recovery plan execution. For Pilot Light, this does not include the SDDC deployment. RTO is shorter than an hour.
Recovery Point Objective (RPO)	RPO configuration for the protected workload is the same for On-demand and Pilot light.	RPO configuration for the protected workload is the same for On-demand and Pilot light.
Infrastructure	All infrastructure components are provisioned and deployed when a disaster is declared.	Minimal infrastructure components are pre-deployed and maintained, with additional infrastructure components provisioned as needed.
Cost	Lower cost due to no active footprint.	Higher cost as a result of a small always-on failover capacity.
Scalability	Limited scalability due to no pre-deployed infrastructure components. However, high scalability is possible once SDDC is provisioned.	High scalability, with the ability to quickly provision infrastructure resources as needed.
Use Cases	Best suited for smaller organizations with less critical workloads and limited budgets.	Best suited for larger organizations with more critical workloads and larger budgets, and/or more frequent testing requirements. Suited for organizations who wish to use the recovery SDDC as an extended datacenter by running lower priority workloads and enabling connectivity from primary to secondary SDDC via a network extension or VPN or Direct Connect.

Topology Design

Inter/Intra Region Design

When planning a VMware Cloud Disaster Recovery deployment, Availability Zone (AZ) and Region design, as well as consumption type are critical considerations. If the protected site is also a VMware Cloud on AWS SDDC, plan the SDDC deployment to meet your AZ or region-level availability requirements. Since VMware Cloud Disaster Recovery components must communicate with the recovery SDDC, ensure that it is in the same region as the cloud file system.

Topologies

VMware Cloud Disaster Recovery supports various topologies to meet the requirements of different organizations. These topologies include:

- On-Prem to Cloud (O2C) Topology: This topology involves deploying a disaster recovery environment in a VMware Cloud on AWS SDDC in a specific region, while the production environment is a vSphere datacenter on-premises. This is a commonly used topology by customers with an existing environment who seek to have a cloud-based datacenter for disaster recovery purposes.

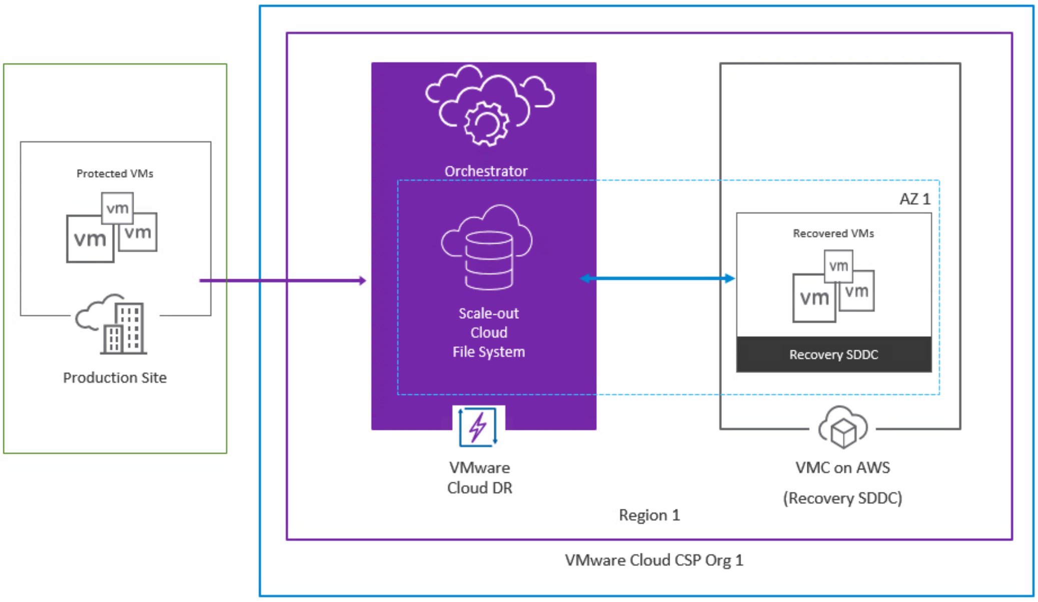

Figure 2 - On-Prem to Cloud (O2C) Topology

- Cloud to Cloud (C2C) Topology: This topology is used for a Greenfield deployment, which involves creating a new production SDDC on VMware Cloud on AWS, or for a Brownfield deployment that has been migrated into the VMware Cloud on AWS SDDC. In this topology, both the production and recovery SDDCs are deployed in different Availability Zones (AZ) or regions to ensure high availability.

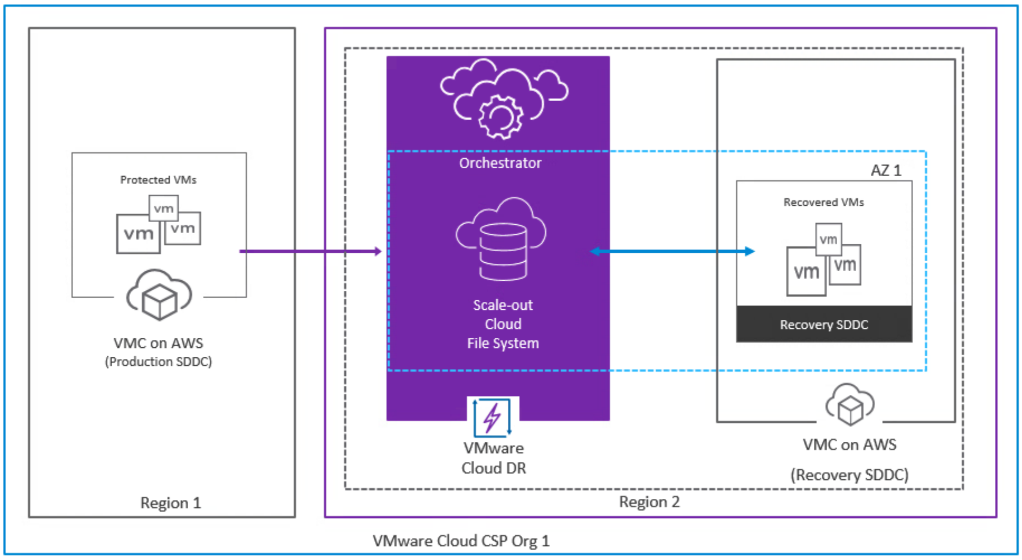

Figure 3 - Cloud to Cloud (C2C) Topology

- Multi-Region Topology: In this topology, the production environment is spread across multiple VMware Cloud on AWS regions, and the disaster recovery environment is also distributed across multiple regions. This provides high availability and can withstand regional disasters, but it can be more complex to manage.

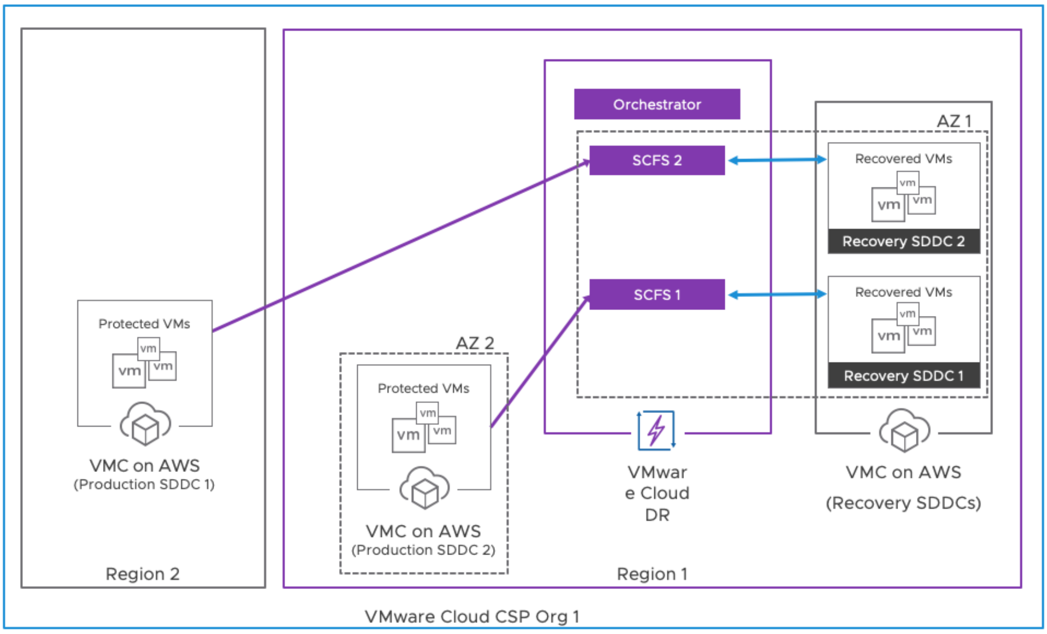

Figure 4 - Multi-Region Topology

- Hybrid Topology: This topology involves having protected workloads in both on-premises infrastructure and cloud-based infrastructure for disaster recovery. This is a good option for organizations that want to leverage their existing VMware Cloud on AWS SDDC infrastructure for disaster recovery.

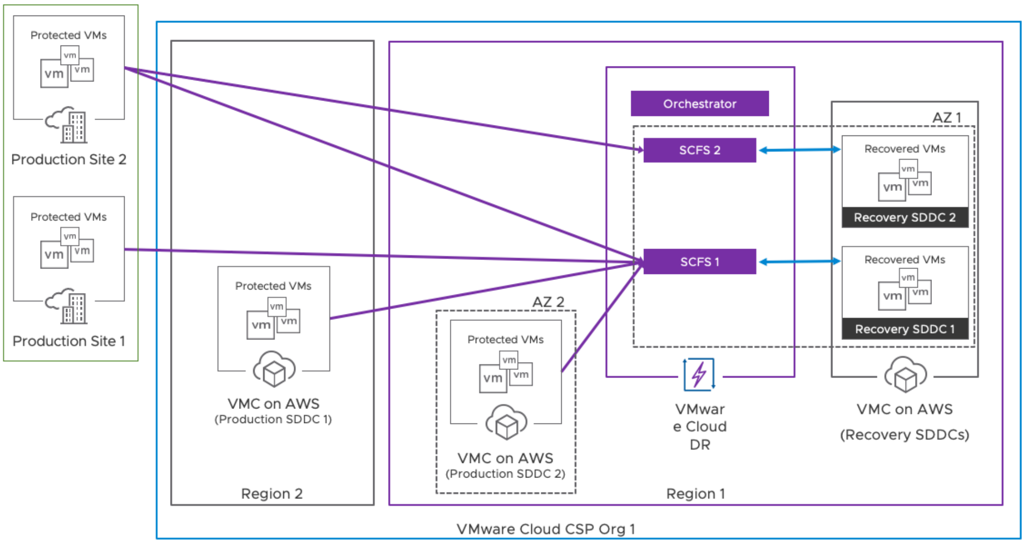

Figure 5 - Hybrid Topology

VMware Cloud Disaster Recovery supports various topologies through its cloud-based disaster recovery solution. It allows organizations to choose the topology that best matches their needs, while also providing replication, failover, and failback capabilities to ensure data availability and minimum downtime. It can also support folder/file level restore. By supporting different topologies, VMware Cloud Disaster Recovery can accommodate organizations of all sizes and types, providing a disaster recovery solution that meets their specific requirements.

SDDC Planning

Using a VMware Cloud on AWS SDDC recovery site, you must create and size the SDDC in the same way that you would any other production SDDC to support your most critical/lowest RTO workloads.

As VMware Cloud on AWS is the only supported recovery cloud/environment, it's important to correctly size the host and storage resources to ensure optimal performance and cost efficiency.

Host Sizing

- Before sizing the disaster recovery site, you must evaluate the current infrastructure and list all resources that must be protected and recovered.
- If you have multiple on-premises datacenters, build the system to use either one dedicated SDDC as a recovery site or multiple recovery SDDC or dual mode (each site as a production and recovery site)
- Deployment Type
  - Single with a single host (test use case only)
  - On-demand (also known as "just-in-time")
  - Pilot Light with cloud bursting

Design Considerations for SDDC Host Commitments:

Type	Design Consideration	Implication
Single Host SDDC	A single host VMware cloud on AWS SDDC is deployed for the purpose of developing and testing disaster recovery plans, after which the recovery SDDC can be destroyed to save on recurrent expenses.	A single host SDDC provides no data security, no production-level SLAs, and is automatically decommissioned after 60 days. A single-host deployment should only be used for testing and is not supported for production.
On-demand	VMware cloud on AWS SDDC is created only when needed	SDDCs are created throughout the recovery process, and the time it takes to build them should be factored into recovery planning. Elastic DRS enables the addition of extra hosts as needed.
Pilot Light	VMware Cloud on AWS supports a minimum of 2-node SDDC for pilot light which enables always on recovery site (For more details, see VMware Cloud Sizer User Guide)	You can scale up the number of hosts while restoring higher-priority workloads. 15 minutes per additional host is required for full failover capacity*. The time it takes to add new hosts might vary and may be shorter than the values given above, which are a conservative average based on testing. These times are subject to change as VMware continues to enhance VMware Cloud on AWS.

Storage Sizing

SDDC sizing is required while running workloads during test and planned failover events, storage sizing in VMware Cloud Disaster Recovery can be considered in two ways.

- Cloud file system - The cloud file system is a distributed file system that allows businesses to replicate their on-premises data to the cloud in real-time. It uses a log-based replication approach to ensure that all changes made to the data are captured and replicated to the cloud. This allows businesses to recover their data quickly and easily in the event of a disaster. During sizing of the cloud file system, determine the total amount of replicated data and the retention duration. (For more details, see the VCDR size) .
- vSAN Storage - VMware vSAN is a software-defined storage solution that aggregates local storage devices into a single pool of storage resources, providing businesses with a highly scalable and cost-effective storage solution. Sizing the vSAN in accordance with the workload recovered during a test or scheduled failover is optional for higher performance needs. Create storage policies with consideration for the various workloads and map the necessary resources on the recovery plan.

Example Size Planning

For example, assume a customer needs to protect a workload that requires 16 virtual CPUs, 64 GB of memory, and 1 TB of storage.

Workload Requirements	Value
Number of vCPUs	16
Memory (GB)	64
Storage (TB)	1

Host Sizing: To determine the number of hosts required, we can use the following calculation: (Number of Virtual CPUs / Number of CPUs per Host) + (Memory Required / Memory per Host) = Number of Hosts In this example, assuming the hosts have 36 virtual CPUs and 512 GB of memory, the calculation would be: (16 / 36) + (64 / 512) = 0.44 + 0.13 = 0.57 hosts, which would round up to 1 host.

Host Sizing	Value
Number of CPUs per host	36
Memory per host (GB)	512
Number of Hosts needed	1

Storage Sizing: To determine the amount of storage required, we can use the following calculation: (Storage Required x RAID) / Number of Hosts = Total Storage Required In this example, assuming a RAID 5 configuration, the calculation would be: (1 TB x 1.33) / 1 = 1.33 TB, which would round up to 2 TB to provide some headroom.

Storage Sizing	Value
RAID Configuration	RAID 5
Total storage required	2TB

Overall, based on this example, a customer would need at least one host with 16 virtual CPUs and 64 GB of memory, and 2 TB of storage to meet the workload requirements. However, this is just an example, and actual host and storage sizing will depend on various factors, such as the number of workloads, their resource requirements, and performance needs.

Additional consideration while accounting for overhead when using each SDDC in distributed mode:

- Include the snapshot and swap space.
- Add DRaaS connector overhead depending on the number of connector VMs deployed for replication.

Solution Interoperability

In protected locations, several VMware or VMware-certified solutions, such as HCX, Site Recovery Manager, and others may be present. Make sure that any deployed solutions are VMware Cloud Disaster Recovery compatible. A simple example can be a backup application. When backing up virtual machines in a vSphere environment, many backup solutions use the snapshot mechanism. Backup and VMware Cloud Disaster Recovery replication configurations are both configurable on virtual machines. Configure these schedules to avoid overlap.

Network connectivity

Network connectivity plays a crucial role in ensuring the success of disaster recovery . It refers to the ability of a system to communicate with other systems or networks to transfer data and access resources. Internet and Direct Connect (DX) are the two ways supported by VMware Cloud Disaster Recovery for replicating data from a protected site to the cloud file system.

The DRaaS connectors communicate with the VMware Cloud Disaster Recovery SaaS orchestrator over an encrypted tunnel across the internet or Direct Connect for data transfer and metadata operations utilizing port 443. In both On-Demand and Pilot Light consumption models, the DRaaS connector from the protected site sends the data to the SaaS orchestrator to be saved in the cloud file system.

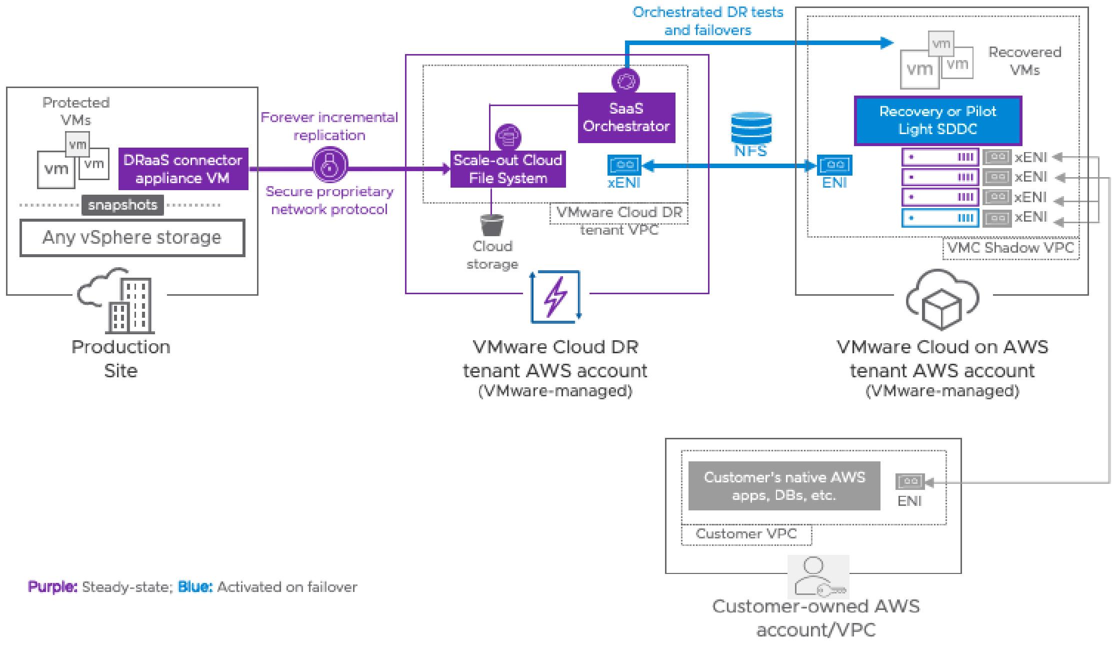
Orchestrated DR tests and failovers
Protected VMs
vm
vm
vm
DRaaS connector appliance VM
snapshots
Any vSphere storage
Production Site
Forever incremental replication
Secure proprietary network protocol
Scale-out Cloud File System
SaaS Orchestrator
xENI
VMware Cloud DR tenant VPC
Cloud storage
VMware Cloud DR tenant AWS account (VMware-managed)
NFS
Recovered VMs
Recovery or Pilot Light SDDC
ENI
xENI
xENI
xENI
xENI
VMC Shadow VPC
VMware Cloud on AWS tenant AWS account (VMware-managed)
Customer's native AWS apps, DBs, etc.
ENI
Customer VPC
Customer-owned AWS account/VPC
Purple: Steady-state; Blue: Activated on failover

Figure 6 - Communications Flow

Internet

The VMware Cloud Disaster Recovery DRaaS connector establishes a secure tunnel connection to the cloud file system. The Internet connection is the VMware Cloud Disaster Recovery's default setup parameter. There are ports that the DRaaS Connector must use to send replication traffic to the cloud file system, and these must be configured on the firewall. See VMware ports portal for more information.

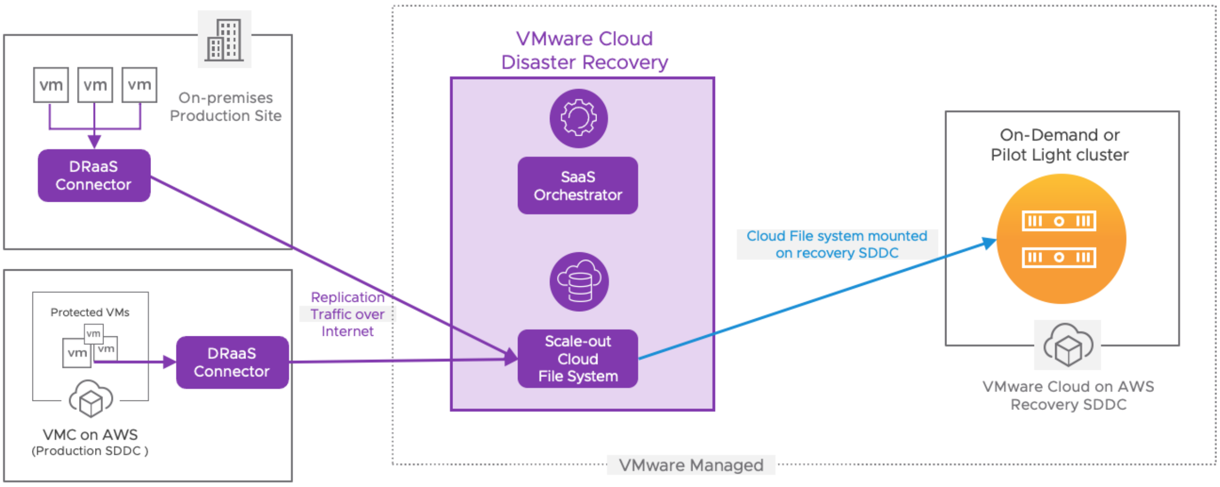
vm
vm
vm
On-premises Production Site
DRaaS Connector
Protected VMs
vm
vm
vm
DRaaS Connector
VMC on AWS (Production SDDC)
Replication Traffic over Internet
VMware Cloud Disaster Recovery
SaaS Orchestrator
Scale-out Cloud File System
Cloud File system mounted on recovery SDDC
On-Demand or Pilot Light cluster
VMware Cloud on AWS Recovery SDDC
VMware Managed

Figure 7 - Cloud File System Connectivity

Direct Connect

VMware Cloud Disaster Recovery supports utilizing Amazon Web Services (AWS) Direct Connect (DX) public or private virtual interfaces (VIFs) for on-premises protected site networks. See AWS Direct Connect to set up a DX connection.

You can deploy an AWS Direct Connect to create or attach a public VIF to facilitate communication between the DRaaS Connector and VMware Cloud Disaster Recovery components.

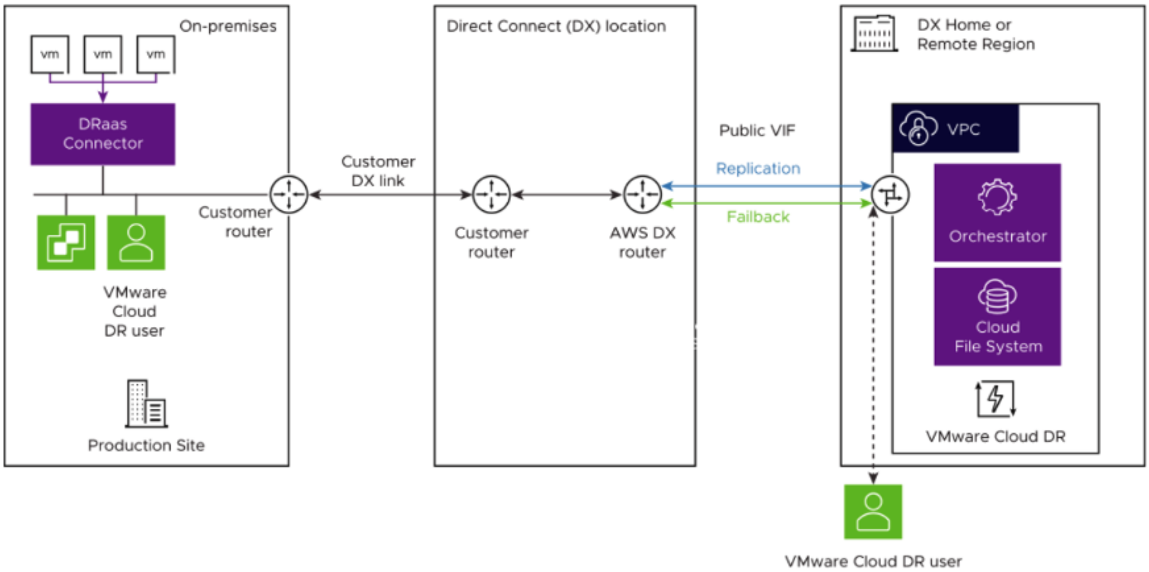

Figure 8 - AWS Direct Connect via public VIF

Once your Direct Connect connection is established, you create a private virtual interface to connect private IP addresses to the VMware Cloud Disaster Recovery Virtual Private Cloud (VPC).

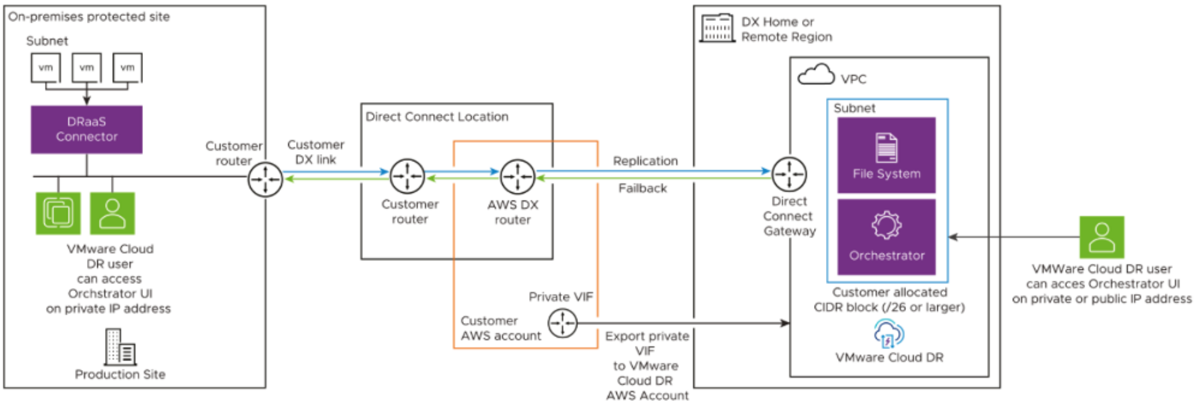

Figure 9 - AWS Direct Connect via public VIF

VMware Cloud Disaster Recovery Network Design Considerations

Consideration	Internet	Direct Connect
Bandwidth	Limited by internet connection speed	Direct connect provides dedicated, high-speed bandwidth
Latency	Latency may be higher due to internet traffic	Lower latency due to dedicated, private connection
Security	May require additional security measures, such as VPN or SSL/TLS encryption	Direct connect provides a secure, private connection
Reliability	Internet connections may be less reliable and more prone to outages	Direct connect provides a more reliable and stable connection
Cost	Generally lower cost, as it does not require dedicated infrastructure or connections	Higher cost due to the need for dedicated infrastructure and connections
Use Cases	Best suited for smaller organizations or less critical workloads that do not require high-speed or dedicated connections	Best suited for larger organizations with critical workloads that require high-speed, reliable, and secure connections

The specific considerations may vary depending on the organization's specific needs and requirements. It's important to carefully consider each factor and weigh the pros and cons of each option.

Additional Network Connectivity Considerations includes:

- Replication Objectives like RTO and RPO configuration are dependent on the network.
- Network Compression
- ISP selection, network bandwidth, and redundancy
- IP migration if any (public IP)
- Name record update

Recovery Plan Considerations

These are the primary components of a disaster recovery strategy.

- Limits on VMs under protection groups, recovery plans, etc. - Discuss the effects on the entire design as well as the best ways to optimize with application owners.
- Concurrent recoveries are regulated to prevent burst mode.
- Discuss the impact on disaster recovery events and offer strategies for distributing VMs across recovery plans in a disaster recovery event and the priority sequence of workload startup with application owners.
- Consider configuring options like restart priority, recovery sequence, and split-brain breaker (witness) in your recovery plan.
- Perform Inventory Mapping.
  - VM resource pool and folder inventory mapping
  - Datastore mapping
  - Network mapping
  - Swap datastore configuration

Configure Role and Permission for Recovery Management

- Set up permissions and responsibilities for recovery management.
- Establish a user account with the least privilege access for the DRasS connector to connect to vCenter Server.
- When a dedicated user is assigned to carry out recovery-related actions on the shared service, take roles and permission into consideration. Update a DNS record as an example during recovery.

- Control the handling of shared services including DNS, DHCP, and domain authentication.

The following table can be used as a planning checklist both during and after deployment to achieve an optimized deployment. Refer to the examples used in the guide to plan the requirements and validate the same post deployment to ensure it matches the requirements.

Planning Consideration	Planning Phase	Post deployment
Recovery SDDC Deployment		
Defined Recovery Objectives		
Review the Supported Workloads		
Assess Network Requirements **• DRaaS connector to ESX on protected site- required minimum 20ms**.		
Size the SDDC **Validate the configuration's maximum.**		
Load balancing techniques.		
Firewall rules **• DRaaS connector - VMware auto-support server**		