NOVEMBER 2020

VMware Cloud on AWS

Resiliency and Availability Built-in for Peace of Mind

Table of Contents

Introduction

VMware Cloud on AWS provides a seamlessly integrated hybrid cloud platform for running customers' enterprise workloads of today and tomorrow. With VMware Cloud on AWS, customers can start their application and infrastructure modernization journey with minimal disruption to their business. They can rapidly extend and migrate their applications to the cloud in an AWS Region of their choice, live and without the need to change the applications. This cloud service is jointly engineered by VMware and AWS and combines the benefits of VMware's enterprise-class SDDC (Software Defined Data Center) software and the AWS global cloud infrastructure with access to the breadth and depth of 175+ AWS native services.

A key differentiator of VMware Cloud on AWS is that availability and resiliency are designed into the service and the underlying AWS infrastructure. This enables VMware and AWS joint customers to focus on their applications and workloads rather than rethinking availability into the application layer, especially when migration and extension projects have stringent timelines associated with them. VMware and AWS have developed a highly available service that makes VMware Cloud on AWS the easiest and fastest path to the cloud to migrate specific applications, move entire data centers, extend into the cloud, or protect it for business continuity and resiliency. In fact, according to a recent study done by IDC, VMware Cloud on AWS customers that were interviewed experienced an 83% reduction in unplanned downtime.[1]

The core foundation of VMware Cloud on AWS is the VMware Software Defined Data Center (SDDC). The SDDC is based on VMware Cloud Foundation and integrates VMware vSphere, VMware vSAN, VMware NSX, and VMware HCX to create a platform that can be deployed typically in under 2 hours on AWS infrastructure in any one of the 17 Global AWS Regions (including a special controlled instance running in AWS GovCloud (US)). Customers are presented with the VMware Cloud on AWS Console or API that enables them to create an SDDC by selecting the AWS Region and Availability Zones. Once installation completes, customers have a fully functioning and fully managed and supported VMware Cloud on AWS SDDC. Behind the scenes, VMware and AWS manage the software and the infrastructure to provide a seamless, highly available service.

This whitepaper discusses the availability and resiliency capabilities included in a VMware Cloud on AWS SDDC and the different deployment options that are available.

Availability and Resiliency Capabilities

AWS Data Center Architecture

Amazon Web Services has the concept of Regions, which are physical locations around the world containing clustered data centers. AWS calls each group of logical data centers an Availability Zone (AZ). Each AWS Region consists of multiple, isolated, and physically separate AZs within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers many advantages for customers. Each AZ has independent power, cooling, and physical security. They are connected via redundant, high bandwidth, ultra-low-latency, networks. AWS customers focused on high availability can design their applications to run in multiple AZs to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection.

An AZ is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZs in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZs. All traffic between AZs is encrypted, and the network performance is sufficient to accomplish synchronous replication between AZs. AZs make partitioning applications for high availability easy. If an application is partitioned across AZs, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZs are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other.

The AWS control plane (including APIs) and AWS Management Console are distributed across AWS Regions and utilize a multi-AZ architecture within each region to deliver resiliency and ensure continuous availability. This ensures that customers avoid having a critical service dependency on a single data center. AWS can conduct maintenance activities without making any critical service temporarily unavailable to any customer.

Water, power, telecommunications, and internet connectivity are designed with redundancy, so continuous operations can be maintained in an emergency. Electrical power systems are designed to be fully redundant so that in the event of a disruption, uninterruptible power supply units can be engaged for certain functions, while generators can provide backup power for the entire facility. People and systems monitor and control the temperature and humidity to prevent overheating, further reducing possible service outages. AWS teams run diagnostics on machines, networks, and backup equipment to ensure they're in working order now and in an emergency. Routine maintenance checks on data center equipment and utilities are part of regular operations.

Partition Placement Groups

VMware Cloud on AWS leverages AWS Partition Placement Groups (PPG). With this feature, EC2 instances are spread across logical partitions that do not share underlying hardware, including the physical rack, to minimize the impact of host and rack failures.

VMware Cloud on AWS automatically allocates Partition Placement Groups within an SDDC cluster. Instances in a cluster are placed on a best-effort basis in separate logical partitions that do not share underlying hardware. Placement happens automatically for every new SDDC, cluster or host add operation. Customers simply benefit from the increased availability of Partition Placement Groups; there is no configuration or effort required.

Figure 1 shows an example placement for a 3-host cluster.

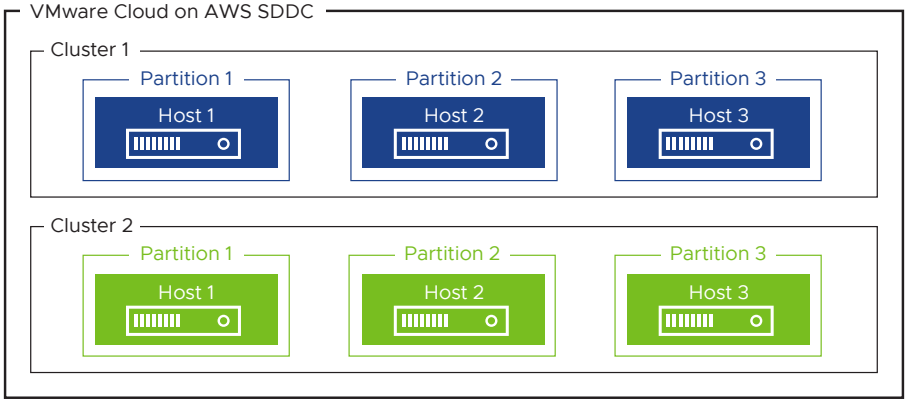

Figure 1. Example placement for a 3-host cluster

When hosts in a cluster are fully placed in separate partitions, a failure, such as a Top-of-Rack switch going down, impacts a single host instead of multiple hosts in a cluster. This aligns with VMware vSAN SPBM policies. With the SPBM Failures-To-Tolerate (FTT) setting, a cluster can recover from one or two host failures. Note that a single Top-of-Rack switch failure can impact multiple clusters in an SDDC, but with Partition Placement Groups only a single host per cluster is affected and the single host failure does not impact workloads. For example, if Partition 2 in the figure above experiences a Top-of-Rack switch failure, there will be two host failures in the SDDC, but only one host failure in each cluster and each of those failures doesn't impact workloads. As a result, AWS Partition Placement Groups and VMware vSAN SPBM combine to provide a solution that is very resilient to rack failures.

Partition Placement Groups is an AWS feature that provides customers with added protection against single points of failure. With the collaboration between VMware and AWS, the VMware Cloud on AWS service brings this capability automatically and transparently to customers.

See AWS documentation for more information on *Partition Placement Groups*.

Auto Remediation

VMware Cloud on AWS SDDCs are deployed using bare metal instances in the global AWS infrastructure. VMware and AWS work together to keep SDDCs healthy and available:

- VMware regularly delivers software bug fixes, new features, and software upgrades.
- AWS delivers a robust global cloud infrastructure that is constantly improving.
- AWS performs host planned maintenance for hardware repair/replacement across the fleet.
- VMware and AWS work closely to detect and repair any failures that do occur and to improve software, hardware, and processes to reduce the chances of future failure.

To minimize the impact of host failures, VMware Cloud on AWS implements an Auto Remediation service that monitors every host across all customers and springs into action whenever there is a problem. The Auto Remediation service works in conjunction with predictive features such as vSphere HA and DRS along with the resilient storage capabilities of vSAN to automate recovery from failed hosts, minimizing downtime and impact to customer SDDCs. Unlike in an on-premises deployment, there is no need to maintain hot spares in the cloud. VMware and AWS jointly manage the pool of instances in every Region and AZ to ensure Auto Remediation has sufficient spare capacity to automatically replace a faulty host with a healthy new one. This happens without the need for customer intervention.

Learn more about *Auto Remediation with VMware Cloud on AWS*.

Managed Storage Policies

VMware Operations manage the vSAN cluster, but that is not to say that customers have no control over storage, thanks to Storage Policy-Based Management (SPBM). SPBM is the declarative control plane that manages all data services inside vSAN. It empowers VMware Cloud on AWS customers to mix and match different classes of data logically via policy.

By default, any new VMs, VMDKs, containers, etc. are associated with the workload policy integrated into the service. VMware Operations configures and manages this policy based on cluster scale and type. Customers are free to create a custom policy for any reason. These policies are stored within the vCenter Server and can be applied to any new or existing VM/VMDK, even a Kubernetes container. Later, if the requirements change, customers can either update the policy - updating any associated objects all at once - or create a new policy and then change the association. Storage Policy Based Management empowers rational data management regardless of scale.

Getting started with this powerful system is a simple process. Every vSAN policy contains an Availability rule - and may optionally include advanced settings should they be warranted. Focusing on the Availability declaration, this is a two-step process. First, the site disaster tolerance controls which AZs should have a copy of the data. Then, the failures to tolerate configures how resilient to failure each copy is. This combination empowers customers to mix and match business-critical workloads alongside less essential workloads, with the ease of policy change controlling if a VM and its data should be synchronously replicated between AZs or pinned to a particular AZ..

Find out more about *SPBM for VMware Cloud on AWS* and *vSAN Policies*.

Elastic DRS

In addition to a robust and highly available infrastructure, the VMware Cloud on AWS service also helps ensure that applications have the capacity and resources they need to perform optimally. VMware and AWS work closely together to ensure sufficient capacity in every AWS Region and AZ across the globe. As a result, customers do not need to overprovision their SDDCs by purchasing more instances than they need. Instead, customers can elastically add hosts in minutes when needed.

VMware has also developed the Elastic DRS (EDRS) feature to automate the scaling of SDDCs when storage capacity, memory, or CPU thresholds are reached. Customers have a choice of four policies:

1. Default storage scale-out
2. Optimize for best performance
3. Optimize for lowest cost
4. Optimize for rapid scale-out

Each of the Elastic DRS policies has a different purpose, depending on customer needs. SDDC clusters always scale out if the 75% storage capacity threshold is reached. The automatic scaling leads to consistent vSAN performance and is necessary to ensure availability and support of the SDDC. Automatic scale-out and scale-in based on memory and CPU thresholds provides optimal resources for applications.

Learn more about *Elastic DRS* and *Rapid Scale-Out* for VMware Cloud on AWS

vSphere Availability

vSphere High Availability (HA) and the Distributed Resource Scheduler (DRS) are enabled on all vSphere clusters within the SDDC by default and fully managed by VMware. These capabilities combined result in a more resilient and well-balanced cluster.

vSphere HA ensures the availability of virtual machines by restarting them on healthy hosts within the cluster if a host failure occurs. This works in conjunction with the VMware Cloud on AWS Auto Remediation service outlined above. vSphere HA is configured with Admission Control enabled reserving a percentage of cluster resources for failover capacity as well as host failure and isolation enabled. Virtual machine and application monitoring are also enabled.

vSphere DRS is used to manage resources such as CPU and memory to automatically resolve resource overcommitment and guarantee or limit resources to ensure performance and efficient usage. Virtual machines will be migrated to balance the cluster for optimal usage and performance. Resource pools are created by default to isolate management and customer workloads, and resources to management VMs are guaranteed via reservations. Customers can create additional resource pools to guarantee or limit CPU and memory usage across their workloads to meet their needs. vSphere DRS is configured with a migration threshold set to level 3 (default) to avoid excessive vMotion operations.

Network Availability

A key attribute of any solution's availability is the network, and VMware Cloud on AWS is no exception. There are two primary components utilized to deliver a high performance, scalable and resilient network which customers require for their workloads. The first is the underlying AWS network infrastructure that provides connectivity from the underlying hosts to the top of the rack network switches and beyond. AWS has a highly engineered networking solution that has been proven over the years and every SDDC ultimately relies on this underlying network to communicate.

VMware and AWS have partnered closely to engineer a robust connectivity model between each SDDC and the AWS network in a way that is easy to consume, operationally sustainable, and innovative in the way access to native AWS services is delivered. The perimeter of the SDDC is the Edge router which is powered by VMware NSX. The Edge router is a virtual machine that enjoys the protection mechanisms of VMware vSphere including resource reservations and vSphere HA. The Edge router itself is deployed in an Active/Standby pair with state synchronization between the two nodes. In the event of a failure, the Standby Edge router assumes all of the routing and security functions without the need for TCP sessions to be re-established. NSX Edge software is also updated with only a brief disruption as part of the *SDDC upgrade process*.

A unique VMware and AWS specific integration is the use of an AWS Elastic Network Interface (ENI) to provide high speed connectivity to an AWS VPC that is paired to every SDDC called a Connected VPC. Customers can implement native AWS services like S3, EC2, and more in the Connected VPC and allow them to have high speed access to VMs in the SDDC over the ENI through the VMware NSX Edge router. Every host in the SDDC is connected to the ENI so that this critical path is available if the host with the active NSX Edge router fails and the Standby takes over, and connectivity across the ENI is maintained.

As of VMware Cloud on AWS Release 1.12, a new option for connectivity was introduced, VMware Transit Connect. Transit Connect is a VMware managed implementation of a highly available AWS Transit Gateway (TGW). Transit Connect simplifies complex network tasks through the use of SDDC Groups to allow customers a fast and easy way to interconnect SDDCs, native AWS VPCs and back to their on-premises data centers. VMware Transit Connect "plugs into" an SDDC through the VMware NSX Edge router and as such is able to take advantage of the same high availability attributes that the Active/Standby design employs. The VMware Managed Transit Gateway (VTGW) component of Transit Connect is in and of itself highly available in the same way that the native AWS TGWs are and uses the underlying AWS infrastructure to provide resilient network connectivity between SDDC Group members.

Learn more about *VMware Transit Connect*.

Deployment Options

Single AZ SDDC

The basic deployment option for an SDDC is to deploy all bare metal hosts in a single AWS AZ. An SDDC has one or more clusters and each cluster can have from 2 to 16 hosts in a production environment. To comply with the VMware Cloud on AWS SLA terms, clusters up to 5 hosts must have SPBM configured with FTT=1 while clusters of 6 or more hosts require FTT=2. This ensures that the cluster can tolerate host failures without data loss.

An SDDC deployed in a single AZ is susceptible to a failure that impacts the entire AZ, but VMware Cloud on AWS availability and resiliency capabilities minimize the impact of the more localized failures.

- **Auto Remediation**: Auto Remediation actively monitors the SDDC and remediates or replaces hosts when necessary.
- **Partition Placement Groups**: Rack level failures, such as Top-of-Rack switch failures, do not impact more than one host per cluster.
- **Managed Storage Policies**: Replication or erasure coding protect against data loss in the event of a host failure. By default, the managed policy automatically changes the managed objects to be protected.
- **Elastic DRS**: Clusters scale out and scale in as needed to maintain an optimal environment for applications in the SDDC.

While these combined capabilities drastically improve the robustness of an SDDC deployed in a single AZ, the most resilient deployment model stretches the SDDC clusters over two AZs, increasing availability significantly.

Multi-AZ SDDC

Multi-AZ SDDCs, or stretched clusters, is the preferred deployment for highly available VMware Cloud on AWS SDDCs. In traditional cloud deployments, if a service requires additional resiliency, traditional applications need to be refactored to obtain high availability. With Stretched Clusters, using native vSAN, customers deploy the bare metal AWS instances of an SDDC across two AZs, creating a cluster that can survive the loss of an entire AZ. Both AZs can run virtual machines. The data belonging to the virtual machines is synchronously mirrored between AZs to maintain accessibility if an AZ is offline.

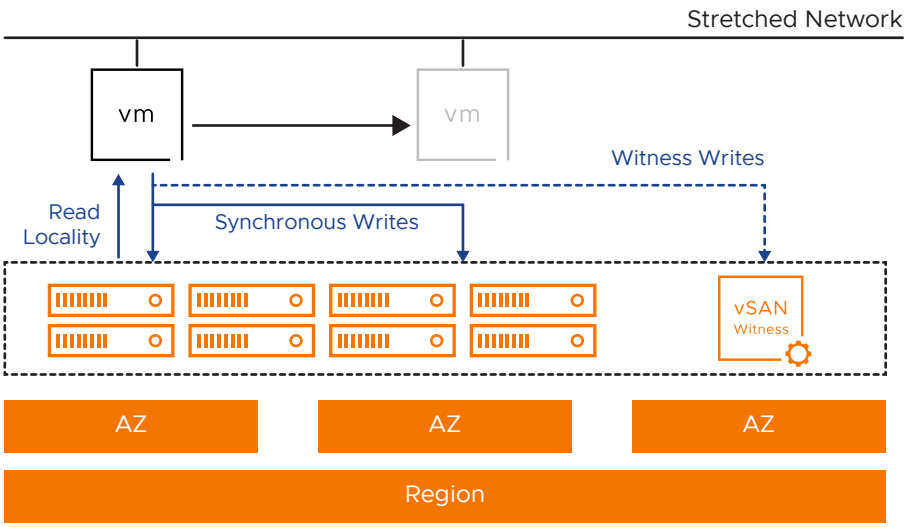

Figure 2. 8 host stretched cluster

Figure 2 shows an 8 host Stretched Cluster with 4 hosts in one AZ and 4 hosts in the other. There is also a fully managed witness node in a third AZ that is used as a tiebreaker to guard against split brain scenarios.

The underlying vSAN Stretched Clusters technology synchronously replicates I/O writes between the AZs. If two or more hosts fail in an AZ or an AZ fails entirely, the service automatically recovers management and customer workloads to continue operating in the surviving AZ. With VMware Cloud on AWS, existing on-premises applications can be migrated as-is to the cloud and they can immediately benefit from the added protection of Stretched Clusters.

The underlying vSAN Stretched Clusters technology synchronously replicates I/O writes between the AZs. If two or more hosts fail in an AZ or an AZ fails entirely, the service automatically recovers management and customer workloads to continue operating in the surviving AZ. With VMware Cloud on AWS, existing on-premises applications can be migrated as-is to the cloud and they can immediately benefit from the added protection of Stretched Clusters.

The Stretched Clusters feature is implemented and managed by the VMware Cloud on AWS service. There are some configuration options:

1. Customers must specify that they would like an SDDC to be configured with Stretched Clusters enabled. As shown in Figure 3, this is a simple checkbox when the user deploys an SDDC.

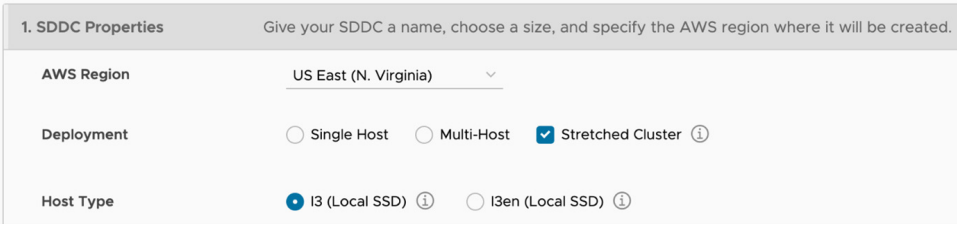

Figure 3. SDDC deployment

2. Two AZs and the related network subnets are selected. (Figure 4)

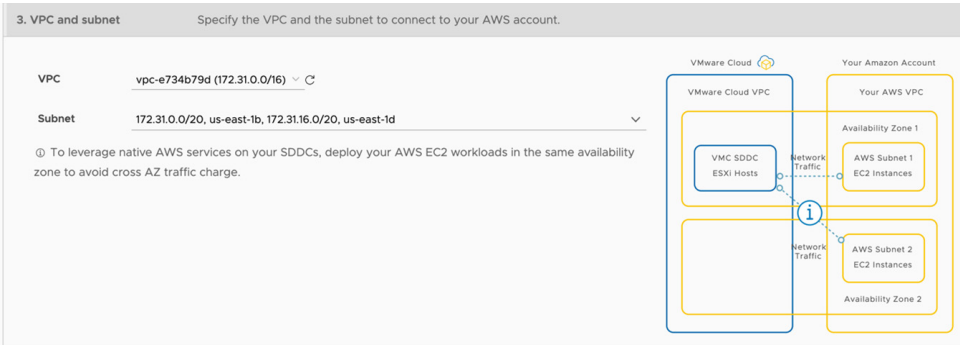

Figure 4. SDDC deployment

3. SPBM policy can be customized, although users are often better off relying on automatic policy management. (Figure 5)

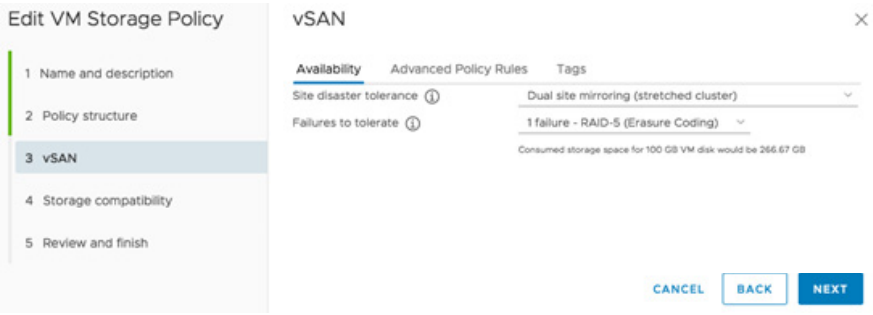

Figure 5. SDDC deployment

In figure 5, the Site disaster tolerance policy specifies Dual site mirroring. This enables synchronous replication and must be enabled on workloads that should be protected by Stretched Clusters. Failures to tolerate (FTT) is set to 1 with RAID-1 or with RAID-5 Erasure Coding. There is no need to configure FTT=2 in a Stretched Cluster because Dual site mirroring already provides protection against double failures. The lower AZ level FTT setting reduces overhead and makes more capacity available to customer workloads. Assuming at least four hosts are deployed in each AZ, erasure coding can be used instead of mirroring to reduce capacity consumption while still maintaining a failures to tolerate setting of one within each AZ.

The availability benefits of Stretched Clusters cannot be overstated. The VMware Cloud on AWS service offers a 99.9% SLA for SDDCs deployed in a single AZ. A Stretched Cluster offers a significantly higher 99.99% *SLA* because it protects against entire AZ failures. Deploying applications across multiple AZs is a best practice that VMware and AWS recommend for all customers. While AZ failure is extremely rare, it can happen. When such an event occurs, production workloads must be protected with Stretched Clusters to avoid service interruption or data loss. This resiliency allows businesses to focus on application requirements, capabilities, and performance instead of infrastructure availability. Stretched Clusters management and availability are described in more detail *here*.

Understand which Data Protection Solutions are supported via the *VMware Compatibility Guide*.

	Availability	Storage Policy	Mirroring	Elastic DRS	Partition Placement Groups
SINGLE AZ SDDC	99.90%	<=5 Hosts FTT=1 >= 6 Hosts FTT=2	Mirroring within AZ	Enabled	Enabled
MULTI AZ SDDC (STRETCHED CLUSTER)	99.99%	FTT=1 with RAID= 1 or RAID=5	Dual Site Mirroring	Enabled	Enabled for each AZ

Table 1. Deployment options summary

Business Continuity

Data Protection

Protecting workloads is critical in safeguarding data from corruption, data loss, and compromise due to several factors such as human error, ransomware, etc. Data protection for VMware Cloud on AWS is driven by partner solutions, with most of the common solutions being supported. VMware offers a 'VMware Ready for VMware Cloud'- a partner certification program for certified solutions and provides software vendors with the highest level of product endorsement. Often these common solutions are already being used on-premises by customers to back up their existing workloads; leveraging the same tools to backup workloads in the cloud provides a consistent operations model.

One key benefit of VMware Cloud on AWS is that data protection solutions can leverage the elastic network interface (ENI) to provide high speed, low latency, access to AWS native storage. This allows traffic to traverse the VPC interconnect in the same AZ with no ingress or egress charges when using an EC2 based backup appliance and S3, EFS, or Glacier as a backup storage target.

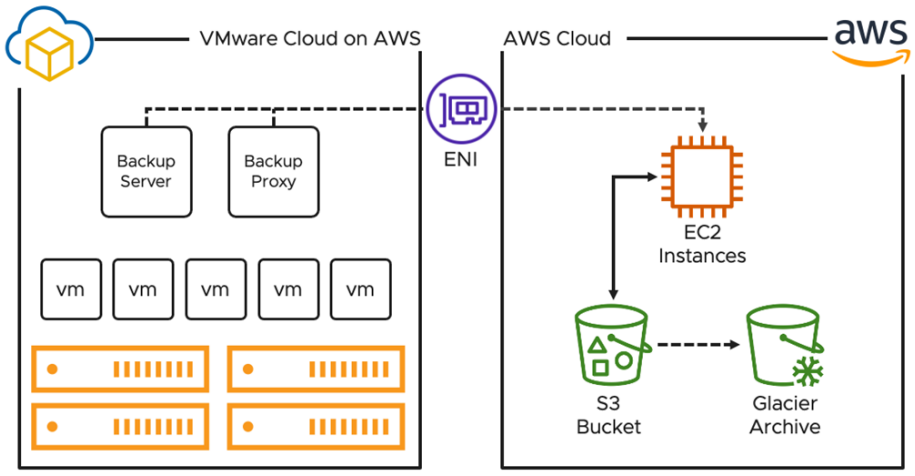

Figure 6. Placeholder caption

Disaster Recovery

VMware offers two in-house DRaaS solutions supported on VMware Cloud on AWS.

1. VMware Cloud Disaster Recovery

VMware Cloud Disaster Recovery offers easy to use, on demand disaster protection and recovery, delivered as a SaaS service, with cloud economics. It combines cost-efficient cloud storage with simple SaaS-based management for IT resiliency at scale, and customers benefit from a 'pay when you need' failover capacity model for DR resources.

Leveraging the elasticity of cloud computing, VMware Cloud Disaster Recovery spins up VMware Cloud on AWS infrastructure only during a DR testing or failover event. It utilizes a highly efficient cloud storage layer for storing backups, lowering DR costs. It delivers fast recovery with zero copy and no rehydration of data from cloud storage to VMware Cloud on AWS hosts where the recovered VMs can be immediately powered-on. For longer term use, and achieving production-level performance, workloads can be migrated (using storage vMotion) to vSAN storage in the SDDC. Using optional pilot light clusters makes the recovery time even faster. VMs are maintained in their native VMware vSphere format, eliminating the need for brittle and time-consuming VM disk format conversions. Instant power-on of VMs is very powerful for rapid identification of the best recovery point when recovering from a ransomware attack.

VMware Cloud Disaster Recovery can protect a very broad set of IT services in a cost-efficient manner, with fast recovery capabilities (On-demand DRaaS). Learn more about VMware Cloud Disaster Recovery *here*.

2. VMware Site Recovery

VMware Site Recovery offers Disaster Recovery as-a-Service (DRaaS) for VMware Cloud on AWS. VMware Site Recovery can protect mission critical IT services that require very low RPO and RTO (Hot DRaaS). Automated workflows deploy the Site Recovery components inside the SDDC while VMware maintains and supports them. This further extends availability for customer workloads by protecting VMware Cloud on AWS SDDCs operating in one AWS Region with a failover target in another AWS Region.

VMware Site Recovery enables customers to create recovery plans that fully automate and orchestrate failovers, allowing IT teams to offload manual tasks during the recovery process. At the heart of the DR solution is Site Recovery Manager (SRM) and vSphere Replication (VR). This proven DR tool helps customers reduce risks during critical times. VMware Site Recovery has extensive built-in testing capabilities and enables customers to perform frequent non-disruptive DR tests that automatically generate detailed reports, reducing the exposure to disasters. Learn more about VMware Site Recovery *here*.

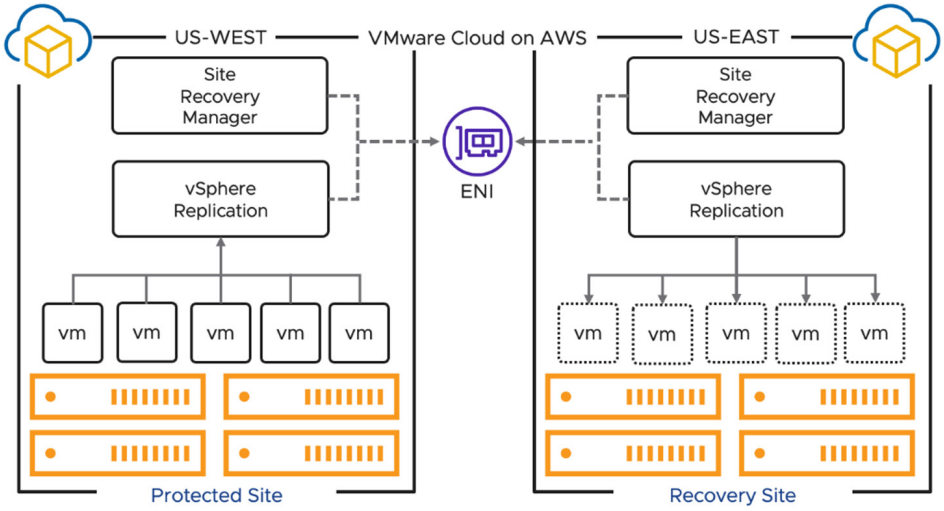

Figure 6. Placeholder caption

Customers can combine the benefits of vSAN stretched clusters and DRaaS options to further enhance resiliency. Stretched clusters across AZs protect against a wide variety of host and AZ failures. DRaaS adds rapid recovery options for scenarios such as ransomware attacks. This combination of services and automation provides maximum availability and rapid disaster recovery—all without the need to manage the underlying infrastructure.

Migration

VMware HCX

VMware HCX, an all-in-one solution for workload mobility is included with VMware Cloud on AWS. It abstracts and removes the boundaries of the underlying infrastructure, focusing on the workloads. Another boundary VMware HCX removes is that it supports several versions of vSphere going back to vSphere 5.0 to the most current vSphere 7.0 Update 1. VMware HCX has built-in WAN optimization, deduplication, and compression to increase efficiency while decreasing the time it takes to perform migrations. It can leverage a dedicated network connection such as AWS Direct Connect or your internet connection (100 Mbps). The established network tunnel is secured using suite B encryption. HCX has a single click option to extend on-premises networks (L2 stretch) to VMware Cloud on AWS, and once the workloads have been migrated, there is also an option to migrate the extended network. VMware Cloud on AWS customers have several migration types to choose from based on their required service level agreements, ranging from Replication Assisted vMotion (zero downtime), Bulk Migration (minimal downtime), and Cold Migration (downtime). Also included is the ability to create mobility groups based on workload or application dependency mapping for wave migrations.

Conclusion

VMware Cloud on AWS delivers a seamlessly integrated hybrid cloud solution that extends on-premises vSphere environments to a VMware Software-Defined Data Center (SDDC) running on Amazon Elastic Compute Cloud (Amazon EC2) elastic, bare-metal infrastructure that is fully integrated as part of AWS. It provides a highly resilient and available platform with capabilities that enable customers to focus on their applications rather than the infrastructure. In summary, key highlights Include:

- Robust AWS global infrastructure
- Constant monitoring, failure detection and auto-remediation of the infrastructure
- Partition Placement Groups to minimize the impact of rack level failures
- vSAN Stretched Clusters to protect against AZ level failures
- Per-VM storage policies to manage availability levels and capacity consumption Elastic DRS to automatically scale clusters and optimize application performance
- vSphere HA to automate workload recovery from host failures Reliable networking due to AWS network infrastructure and VMware NSX
- Robust ecosystem of data protection and disaster recovery solutions
- VMware HCX facilitates migration of workloads to the cloud

Combined, these capabilities make VMware Cloud on AWS the easiest way for customers to migrate to the cloud, with built-in availability and resiliency. For more information, please contact your VMware or AWS representative or visit our website at *https://cloud.vmware.com/vmc-aws/*.

1. IDC White Paper, sponsored by VMWare, The Business Value of Running Applications on VMware Cloud on AWS in VMware Hybrid Cloud Environments, doc #US46919520, October 2020

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com