

# VMware Cloud on AWS vSAN Stretched Cluster Configurations

Multi-AZ deployments  
to ensure VM availability

TECHNICAL WHITEPAPER

SPEAK WITH A  
CLOUD SPECIALIST

 +1 855 532 4734  
 [getcloud@factioninc.com](mailto:getcloud@factioninc.com)



The Stretched Cluster feature on VMware Cloud on AWS is an optional deployment model that provides true high availability (HA) across Availability Zones (AZ), with a 0-RPO, fully self-healing touchless recovery using elastic (DRS). As a managed provider, Faction can deploy this architecture on behalf of the customer, and it requires no re-architecture of existing workloads. This configuration deploys clusters in two Availability Zones inside of the selected AWS Region.

**The minimum configurations for the Stretched Cluster is as follows:**

- » A minimum of 6 Hosts (3 Hosts per AZ)
- » 16 Hosts maximum, must be added in 2 Host increments
- » Only one stretched cluster per SDDC deployed
- » You cannot attach a stretched cluster to an existing SDDC or vice versa
- » 3 Availability Zone minimum
  - 2 AZ's for VMC hosts
  - 1 AZ for vSAN witness node appliance (managed by VMware transparently)
- » HA deployment, not a DR deployment

This is a great solution for customers who have need highly available infrastructure that can survive the total failure of an Availability Zone. The stretched cluster allows for synchronous writes between the vSAN storage offering "0" RPO for the environment. A witness host appliance is automatically deployed in a 3rd AZ and transparently managed in case communication between the vSAN instances is interrupted.



# Deployment Guidelines

The configuration of the service is done during the deployment of the SDDC. Based on how the deployment is configured, the vCenter instance will present certain attributes that define how the clusters are deployed. Some of the configuration screens are as follows:

## VMware Cloud on AWS Console

The feature is selected during the configuration of the SDDC (a “checkbox” when you select the AWS region). Once the SDDC is deployed, you will notice the SDDC is designated as Multi-AZ and lists two availability zone.

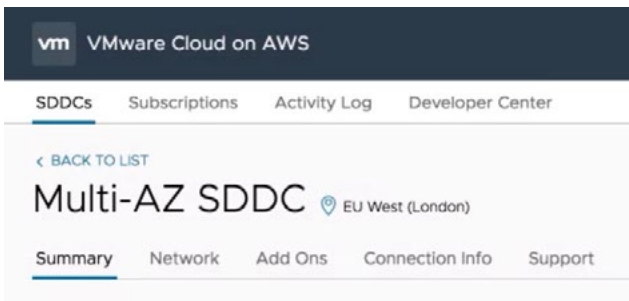


Figure 1

## Deployed AZ / Domain

When the hosts are selected on the console and you specify “Multi-AZ,” the specific AZ domains will be represented as “Fault Domains” on the vCenter cluster. There will be two fault domains listed.

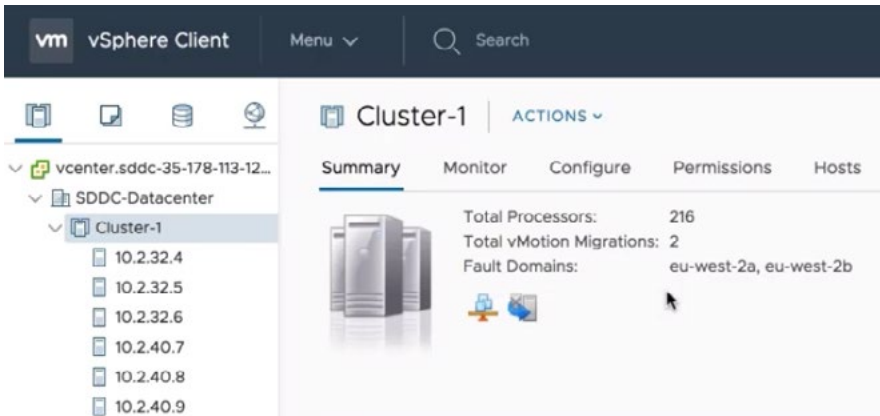


Figure 2

## ESXi Host Domain

The ESXi hosts are balanced between the two availability zones in a stretched cluster. However, the host will reside in one of the AZs. By selecting the “Summary” tab for a host in vCenter, the Fault Domain will be listed on the page.



## VM Designation to an AZ

When a VM is created in vCenter, the New Virtual Machine wizard allows the customer to select which SDDC in which AZ they want designate. VMware will honor that request when possible and set each VM to a specific AZ. The Fault Domain can be viewed under the summary tab for that virtual machine.

## Recovery Considerations

In the event of an Availability Zone failure, vSphere's HA function immediately begins powering up workloads in the other Availability Zone; this operates just like a hardware failure in an on-prem vSphere cluster. Coupled with VMware Cloud on AWS Elastic DRS, which will automatically add hosts to an existing cluster that needs more resources, this allows for a 100% fully self-healing cloud architecture, even in the event of an Availability Zone failure.

Customers may additionally leverage application-level clustering for certain workloads, such as databases, by creating clusters which span Availability Zones; e.g., SQL Server Always-On Availability Groups, with at least one node in each Availability Zone.

## vSAN Stretched Cluster Pricing on VMC

VMware Cloud on AWS charges you for the additional hosts deployed in the adjoining AZ. In addition, there are charges for the witness host and transit charges between the AZs:

1. The stretched cluster add costs \$0.82/hr on top of the normal VMC host prices, this is for the vSAN witness host
2. There is a cross-Availability Zone network charge (\$.01/GB) for network traffic between zones; this applies to both workloads exchanging traffic between zones, as well as every write to disk, since vSAN replicates each write synchronously to the other Availability Zone.



## Conclusion

The vSAN Stretched Cluster deployment on VMC is a standard deployment model that provides a highly available topology for VMware Cloud on AWS environments that can seamlessly survive the failure of an entire Availability Zone. It operates in the same way on AWS and is managed on a single vCenter instance making the management efficient for the customer. Native to vSphere, customers enjoy a true HA environment for their Virtual Machines in an automated environment without any human intervention.

