



Feature Brief: Cluster Types and Sizes

VMware General

Table of contents

Feature Brief: Cluster Types and Sizes	3
Introduction	3
Software-defined Data Center	4
Types of Clusters	5
Single Host SDDC Starter Configuration	5
Multi-Host SDDC Clusters	5
Stretched Clusters	7
Failure Tolerance	8
Adding or Removing Hosts	9
Summary	10

Feature Brief: Cluster Types and Sizes

Introduction

VMware Cloud on AWS supports several different types of clusters that accommodate use cases ranging from evaluation and experimentation to critical business applications. The ability to quickly add or remove resources in response to shifting demand makes this cloud infrastructure service powerful, flexible, and cost-efficient.

This Feature Brief explains the types of clusters and the permitted transitions so that those responsible for managing and sizing cloud resources can make architectural decisions to benefit from this dynamic capability.

Software-defined Data Center

When you create a new software-defined data center (SDDC) you initially decide if it will consist of non-stretched clusters that are deployed in a single AWS Availability Zone (AZ), or if increased resiliency is warranted to protect against an AWS availability zone failure and stretch clusters across two AZs. This document focuses on non-stretched SDDC clusters, to learn about the differences, please see the [Stretched Clusters Feature Brief](#).

An SDDC can contain up to 10 clusters by default, but you can request an increase to 20 clusters per SDDC by contacting VMware support. Additional configuration maximums can be found on the [SDDC Configuration Limits site](#).

The first cluster that is created runs the SDDC management components, such as the vCenter Server appliance, NSX Manager and NSX Edge, in addition to customer workloads. If you create subsequent clusters in an SDDC, these management components will also manage the additional clusters - no additional management components are deployed into the new clusters.

If you are anticipating a deployment with over 30 hosts or 3000 VMs, we recommend the advanced configuration option “SDDC Appliance Size” of “Large” (see Figure 1). For a large deployment there are also additional recommendations in the [Designlet: VMware Cloud on AWS Management Cluster Planning](#).

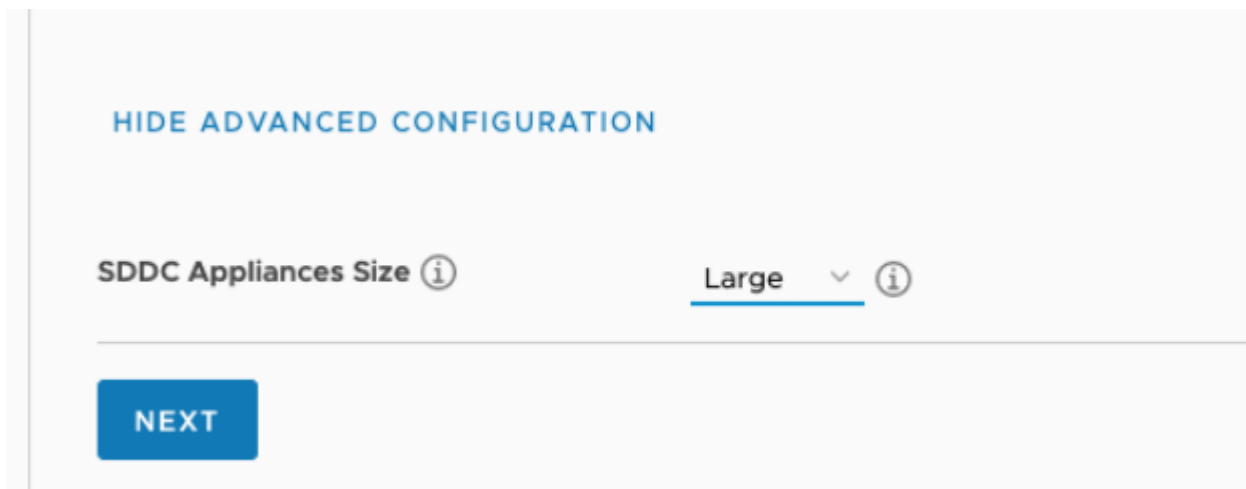


Figure 1 Advanced configuration option enabling Large SDDC Appliance Size during deployment

In case you created the SDDC with a medium appliance configuration and find that you need additional management cluster resources, you can change the SDDC Appliance Size to large. See [Upsize SDDC Management Appliances](#).

VMware Cloud on AWS also supports a feature named Custom CPU Core Count. The Custom CPU Core Count capability allows you to select a reduced number of CPU cores to run per host with respect to the default number of cores for the host type. The extra cores are not visible to the host hypervisor software. This allows to potentially reduce the costs of running applications licensed per-core depending on the vendor (e.g. not applicable for MS SPLA). This function is available from the second cluster onwards, as all host CPUs in the initial SDDC cluster (named “Cluster-1”, which hosts the SDDC management appliances) are enabled and required.

Types of Clusters

A SDDC supports different types of clusters, namely:

- Single Host (although not technically a “cluster”)
- Multi-Host (2 - 16 hosts) in a single AZ or stretched over two AZs

For the first deployment you choose the number of hosts and optionally for Multi-Host to create a stretched cluster..

An SDDC can contain multiple multi-host clusters (but only one Single Host cluster) and each host in a cluster must be of the same host type. The [SDDC Host Types Feature Brief](#) explains the different types of physical hosts available for use with VMware Cloud on AWS. All clusters within a SDDC must be either stretched or non-stretched (no mix and match within a single SDDC; requires additional SDDC). The size of the cluster itself can be changed after deployment - the specifics are covered in the following sections.

Single Host SDDC Starter Configuration

The single host option provides a fully functional SDDC with VMware vSphere, vSAN, and NSX for up to 60 days. Single host SDDCs are therefore ideal for proof of concepts or pilots. A single host can be either an i3 or an i4i host type.

Customers can optionally convert a single host SDDC into a multiple host SDDC at any time during the 60-day operational period. After conversion, the SDDC will be suitable for production workloads because data will be replicated across multiple hosts.

Region	EU (Frankfurt)	Clusters	1
Type	VMC on AWS SDDC	Hosts	1
Availability Zones	eu-central-1a	Cores	64

CPU	Memory	Storage
185,6 GHz	1 TiB	19,98 TiB

Figure 2 A single host SDDC shown in the VMware Cloud on AWS console

A lot of the capabilities of a production-sized cluster are offered so that customers can evaluate and experiment with VMware Cloud on AWS for the lowest possible cost.

Features that do not require more than one host are included in the Single Host SDDC offering, including hybrid operations between on-premises and VMware Cloud on AWS. However, any operations or capabilities that require more than one host would not work. For example, High Availability (HA) and stretched clusters across two AWS AZs. Due to the nature of a single host, there is no data redundancy (FTT=0), meaning that if your host fails, data would be lost. Single Host SDDCs are not patched or upgraded during their 60-day lifespan. At the end of 60 days without conversion to multiple hosts, these clusters are automatically terminated, and all virtual machines and data are securely removed.

For these reasons, Single Host SDDCs are not covered by a service level agreement and are not to be used for production applications. It is recommended to validate the [SDDC Single Host limitations](#) before deployment.

Multi-Host SDDC Clusters

Multi-Host SDDC clusters provide fully redundant data replication and are suitable for production usage since they are covered by

the VMware Cloud service level agreement (SLA) and receive lifecycle management by VMware.

Each Multi-Host cluster can consist of 2 to 16 hosts depending on resource requirements.

Starting with SDDC version 1.24 you can choose between [vSAN Express Storage Architecture \(vSAN ESA\)](#) or vSAN Original Storage Architecture (vSAN OSA) for non-stretched clusters of i4i hosts. vSAN ESA leverages the characteristics of newer hardware to deliver improved capabilities and performance.

Stretched Clusters

If your disaster recovery plan needs to take into account the "failure of an AZ" scenario without data loss (RPO=0) and low recovery time objective (RTO) you have the option to use Stretched Clusters within a VMware Cloud on AWS SDDC. Compared to non-stretched clusters, the [SLA for the SDDC infrastructure](#) increases to 99.99% with six hosts (3:3, three hosts per AZ) and more.

For this purpose, a vSphere cluster is distributed across two availability zones (with a hidden VMware-managed "witness" host in a 3rd AZ). Our vSAN technology provides a single datastore for customer workloads in the cluster that's replicated across both AZs. If one AZ fails, virtual machines are brought up automatically in the other one.

A stretched cluster is already possible with two hosts (1:1, one host per AZ) and can be scaled out to 16 hosts (8 hosts per AZ).

Some points to consider for Stretched Clusters:

- Stretched clusters are not supported with [vSAN ESA](#)
- Stretched clusters are not supported with external storage
- A stretched cluster that has been scaled out to four or six hosts cannot be scaled in. Stretched clusters with more than six hosts can be scaled both out and in.

More details can be found in the [Stretched Clusters Feature Brief](#) and the [product documentation](#).

Failure Tolerance

Depending on the number of hosts in a cluster, workloads may be protected against varying degrees of failure in one or more elements of the underlying infrastructure. This resiliency is specified via the vSAN storage policy, which uses the nomenclature “failures to tolerate,” or FTT, to indicate the extent of protection.

FTT defines the number of host and device failures that a virtual machine can tolerate. You can choose to have no data redundancy, or select a RAID configuration optimized for either performance (Mirroring) or capacity (Erasure Coding).

More details regarding available vSAN policies can be found in the [product documentation](#).

To be [eligible to receive any SLA Credits](#) for an SLA Event, you must meet the following requirements:

- For non-stretched clusters, you must have a minimum configuration for all VM storage policy. Numbers of Failures to Tolerate (FTT) = 1 when the cluster has 2 to 5 hosts, and a minimum configuration of FTT = 2 when the cluster has 6 to 16 hosts. This is not dependent on RAID levels.

The following applies to stretched clusters:

- For stretched clusters with four hosts or less, spanning across more than one availability zone, you must have a minimum configuration for all VM storage policy Site Disaster Tolerance (PFTT) = Dual Site Mirroring.
- For stretched clusters with six hosts or more, spanning across more than one availability zone, you must have a minimum configuration for all VM storage policy Site Disaster Tolerance (PFTT) = Dual Site Mirroring and Secondary level of failures to tolerate (SFTT) = 1. This is not dependent on RAID levels.

Adding or Removing Hosts

Adding or removing hosts is a common operation that can be performed by cloud administrators either interactively via the web console, or through automation using the REST API or PowerCLI. The following demo shows how to manually add a host:

Alternatively, Elastic DRS (EDRS) is a policy-based service that can add or remove hosts in a production cluster as the demand for resources fluctuates. EDRS can be tuned to prioritize performance or cost optimization. EDRS requires a minimum of two hosts, so it does not apply to single-host SDDCs. In two-host SDDCs and stretched clusters with fewer than six hosts, only the Elastic DRS Baseline policy is available. The EDRS Baseline policy adds hosts after storage utilization reaches 80% or vSAN component utilization reaches 85%. To learn more, see the [Elastic DRS Feature Brief](#) and the [product documentation](#).

There are certain restrictions pertaining to host addition and removal operations:

- Single hosts can be scaled up to 2- or 3-host-clusters before the 60-day lifespan ends
- Non-stretched clusters with two or more hosts can be expanded and contracted from 2 to 16 nodes as needed
- A stretched cluster that has been scaled out to four or six hosts cannot be scaled in. Stretched clusters with more than six hosts can be scaled both out and in.

Summary

VMware Cloud on AWS provides a range of cluster types in order to provide the proper resources and service levels to meet customer use cases. It's easy to get started with a single host SDDC and then expand it to multiple hosts that offer a service level suitable for production workloads. By allowing addition and removal of hosts according to resource requirements, customers can avoid over-provisioning and over-spending on cloud infrastructure services.

