TECHNICAL WHITE PAPER: **January 2026**

# vSAN ESA with VMware Cloud on AWS

## Technical Deep Dive

vmware by Broadcom

## Table of contents

## Introduction

With version 1.26 of VMware Cloud on AWS SDDC, vSAN Express Storage Architecture (ESA) has been released. vSAN ESA is now the default for all new VMware Cloud on AWS SDDCs and all new secondary clusters that are deployed in a single availability zone and updated to version 1.26.

Let's look at the capabilities, supported use cases, deployment scenarios, and benefits of vSAN ESA on VMware Cloud on AWS

## vSAN Express Storage Architecture Overview

vSAN ESA was introduced with vSphere 8 and enhanced in the vSAN 8 U1, U2, and U3 updates. It is designed to utilize the full capabilities of modern hardware and delivers high performance and scalability.

vSAN ESA leverages single-tier storage architecture utilizing high-performance NVMe devices opposite to two-tier vSAN Original Storage architecture (OSA) where cache and capacity tiers are separate. Using a single-tier architecture increases the overall capacity available and narrows the impact of drive failures.

With other improvements including a new log-structured file system, improved compression technique, redesigned encryption process and native snapshots vSAN ESA provides lower TCO, faster, more reliable storage architecture with improved capacity and performance without compromise. You can find a comprehensive review of new Express Storage Architecture features and benefits on the vSAN Express Storage Architecture (ESA) page.

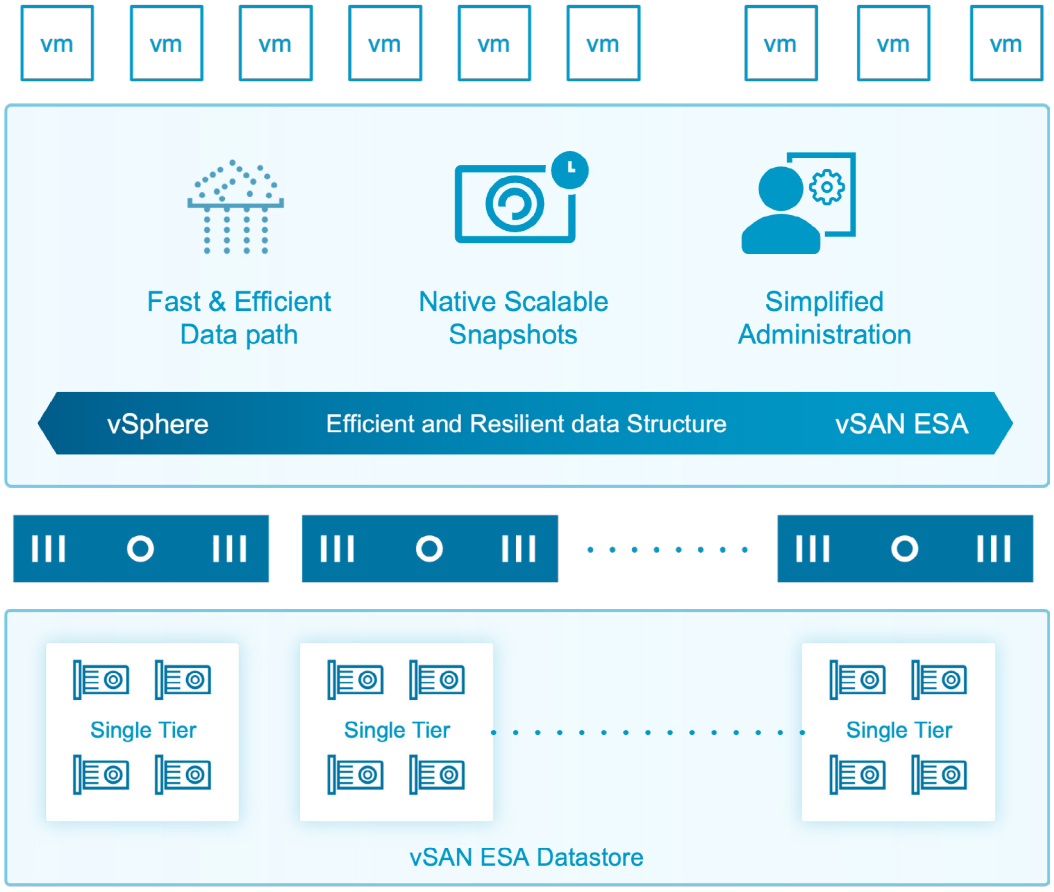

## vSAN ESA Implementation Details on VMware Cloud on AWS

### Prerequisites and Limitations

Clusters with new vSAN architecture are supported with i4i.metal host type only. vSAN ESA is available in VMware Cloud on AWS 1.26. With this release, we support the following use cases:

- New SDDC deployment (both the primary and secondary clusters),
- New secondary clusters once SDDC is upgraded to 1.26,
- Available for only SDDCs in a single availability zone. Stretched Cluster is not currently supported.

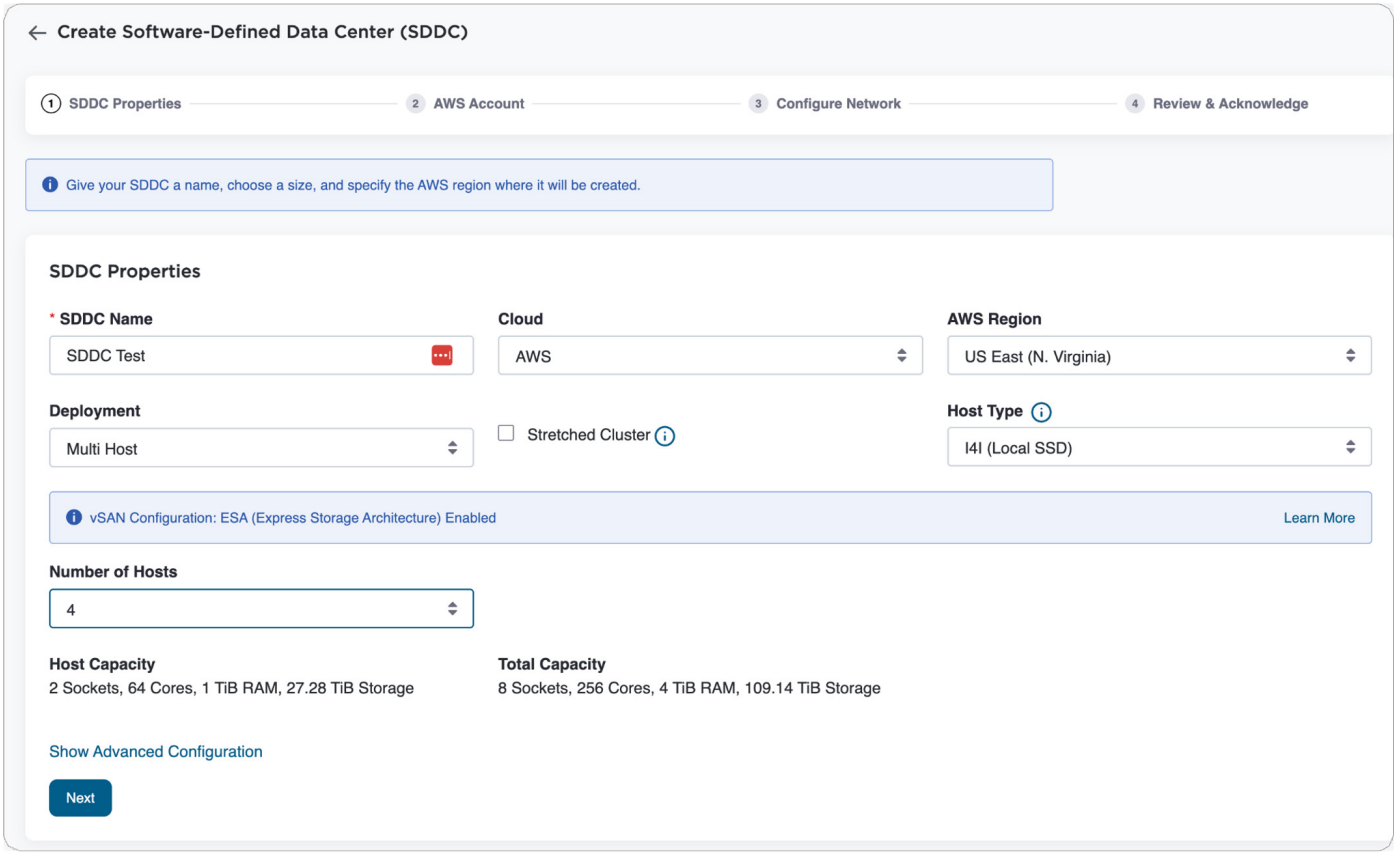

## vSAN ESA Benefits with VMware Cloud on AWS

With vSAN Express Storage Architecture on i4i hosts VMware Cloud on AWS customers are getting instant capacity benefits and storage performance increase compared to vSAN Original Storage Architecture.

### Capacity

vSAN ESA removes the need for dedicated cache devices making the capacity of all NVMe drives usable by the vSAN datastore.

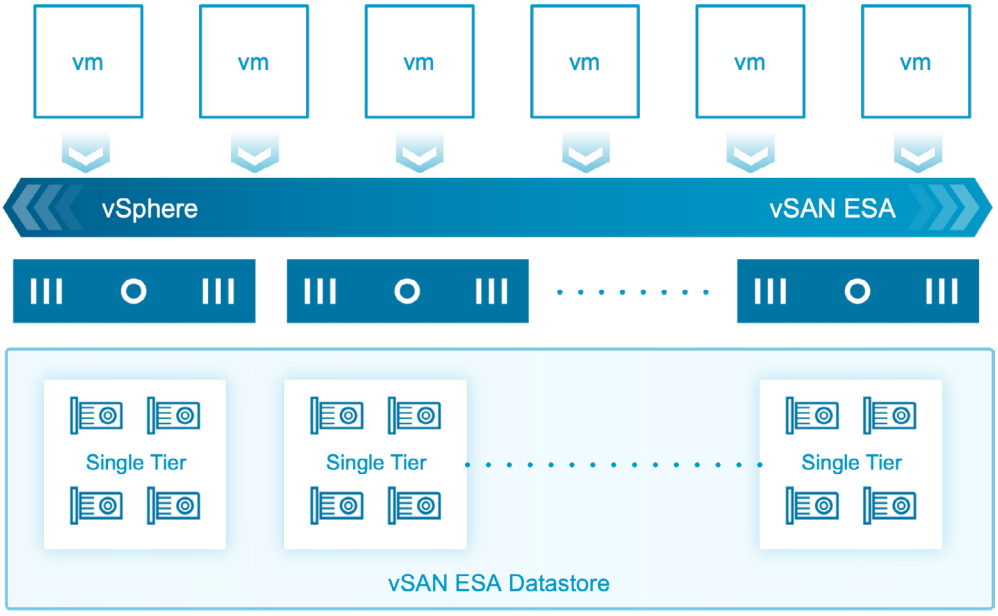


Let's illustrate the capacity gains by example: when deploying a vSAN ESA cluster with i4i nodes customers would get 27.28 TiB (8 disks*3.41 TiB ) raw capacity, compared to 20.46 TiB (6 disks* 3.41 TiB) using traditional vSAN OSA with 2 disk groups effectively claiming 2 disks for cache on each host i4i giving impressive 33% increase for raw capacity available.

vSAN RAID-5 policy is available for clusters as small as 3 nodes and RAID 6 is enabled starting from 6 nodes during deployment of vSAN ESA cluster in SDDC bringing additional capacity gains specifically for small cluster deployments.

### Compression

With the new Express Storage Architecture compression process is now executed on the ESXi host where the virtual machine is running minimizing CPU and storage writes amplification required and decreasing network traffic.

Putting compression "up the stack" allowed it to move its configuration to VM or even virtual disk level differing from vSAN OSA where it was a datastore level setting.

The improved compression algorithm can reduce the 4KB block size to as little as an 8:1 ratio which is 4 times better compared to vSAN OSA which could theoretically achieve a 2:1 compression ratio.

Compression is enabled by default and recommended for most of the workloads but can be disabled using the "No Compression" policy (i.e. use cases where applications are performing their own compression) saving CPU cycles.

Changing the compression policy on existing objects will affect only subsequent writes and will not change the compression state of existing data retroactively. This removes performance impact compared to vSAN OSA where changing compression configuration forced re-writing all data to the vSAN datastore.

### Performance

vSAN ESA is designed to deliver higher performance for both read and write workloads. This is due to a number of factors, including the use of a log-structured file system and reduced write amplification.

Log-structured file system (LFS) implementation in vSAN ESA allows very fast ingesting of incoming small-writes and writing them to durable log which is tied to specific vSAN object and mirrored as a component across hosts depending on the FTT policy for the object. This is called the "performance leg" and allows to provide write acknowledgment to the guest virtual machine with minimal latency.

When the object's durable log becomes full LFS writes data to the "capacity leg" which is another branch of vSAN object data. Write is performed in full stripe and aligns with the assigned policy for the object.

Such an approach reduces CPU and I/O amplification required and eliminates traditional RAID-5/6 trade-off for capacity over performance. You can find a detailed overview of the vSAN log-structured file system in the article: RAID-5/6 with the Performance of RAID-1 using the vSAN Express Storage Architecture.

Native snapshots capability delivered with vSAN ESA Architecture allows to minimize performance degradation regardless of the number of snapshots taken. The new approach is using a highly efficient lookup table instead of a traditional chain of base and delta disks. vSAN ESA LFS performs incoming writes with the appropriate metadata pointers providing the intelligence behind which data belongs to which snapshot. This speeds up snapshot operations which are one of the most intensive yet frequent activities and makes it much more efficient compared to traditional vSAN OSA architecture. You can find more details in the article: Scalable, High-Performance Native Snapshots in the vSAN Express Storage Architecture.

Data at rest encryption is enabled by default on each VMware Cloud on AWS SDDC cluster and can't be turned off. For vSAN ESA this data service similar to compression is performed on an ESXi host level before writing data to disks thus making less impact on the performance when compared to the OSA. More details on encryption improvements can be found in the article: Cluster Level Encryption with the vSAN Express Storage Architecture

### Less Impact Of a Disk Failure

Using a single-tier architecture increases the overall capacity available and narrows the impact of drive failures. Eliminating the vSAN OSA disk group concept where a failing cache drive led to the failure of the whole group narrows the failure domain in vSAN ESA to a discrete disk reducing the amount of data affected in case of the error.

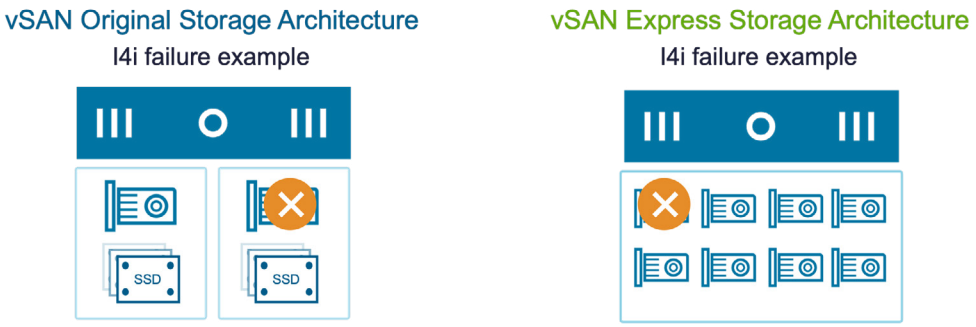


## Summary

vSAN Express Storage Architecture is one of the key features introduced with the new 1.26 release of VMware Cloud on AWS.

Supported on primary and secondary clusters running on i4i hosts, customers could gain immediate benefits from modern vSAN architecture getting more capacity and performance and lowering TCO.

With better compression, native snapshots and encryption which requires less overhead, vSAN ESA is a no-brainer choice for customers running workloads on VMware Cloud on AWS.

Item No: vSAN ESA with VMware Cloud on AWS 1/26