

# DataStax Enterprise on VMware vSAN™ Deployment Options

---



# CONTENTS

---

VMware and DataStax .....	3
Agility through Storage Policy Based Management (SPBM) .....	4
vSAN FTT=0 with Host Affinity for Applications with Intrinsic High Availability Capabilities	4
vSAN FTT=1	5
DSE Replication .....	5
DSE Rack Awareness .....	5
Node Density.....	6
Option 1: One DSE Node per Host	6
Option 2: Multiple DSE Nodes per Host	6
Multiple vSAN Cluster Deployment.....	7
vSAN FTT=1	7
vSAN FTT=0 without Host Affinity	7
vSAN FTT=0 with Host Affinity	7
Takeaway.....	8
About DataStax .....	9

**VMWARE HCI**

**An Ideal Platform for Cloud-Native Applications**

- ✔ **Predictable and performant infrastructure:** Cloud-native apps tend to scale linearly, and HCI's ability to add both compute and storage resources incrementally more precisely matches resources to needs as the application expands.
- ✔ **Unified management:** HCI provides an ideal platform to manage compute and storage resources for virtual machines as well as containers. VMware HCI uses the same tools to manage cloud-native apps that system admins use to manage their VM environments today.
- ✔ **Lower-cost, high-performance infrastructure:** On the CAPEX side, customers have reported savings up to 75% versus traditional three-tier, all-flash arrays with fiber channel networking. On the OPEX side, hyperconverged infrastructure simplifies operations by prolific use of automation, and studies have shown the management OPEX is up to 58% lower with HCI versus traditional all-flash arrays.

# VMWARE AND DATASTAX

## Hybrid and Multi-Cloud Ready, Enterprise-Grade Availability, Simplified Management

Organizations are leveraging modern data platforms, such as Apache Cassandra™, to power new products and services across industries. For example, Cassandra empowers retailers to provide an omnichannel shopping experience with tailored customer experiences in real time. However, the implementation and ongoing management operations can be a significant barrier to adoption. Organizations are betting their growth on distributed technologies, and they want a scalable full-stack solution that is always available, secure, and simple to deploy and manage.

VMware® and DataStax have partnered to build a hybrid and multi-cloud solution designed to meet the data and availability requirements of the most demanding applications with built-in high availability and failure protection across availability zones, regions, data centers, and clouds in three of the four largest public cloud providers: Amazon, Microsoft, and IBM, as well as a rich ecosystem of hundreds of public cloud providers globally.

Our joint integration simplifies operations and eliminates data and application silos with consistent processes and tooling while offering intrinsic security for both data-at-rest and in-flight software-based encryption that meets strict US federal government standards. DataStax and VMware's hyperconverged infrastructure (HCI), combined with support for on-premises and VMware Cloud, enable organizations to manage both traditional and cloud-native applications with the same interface, delivering vast efficiency gains and consolidation benefits, greater consolidation ratios, a single management toolset, and automated lifecycle management, to name a few.



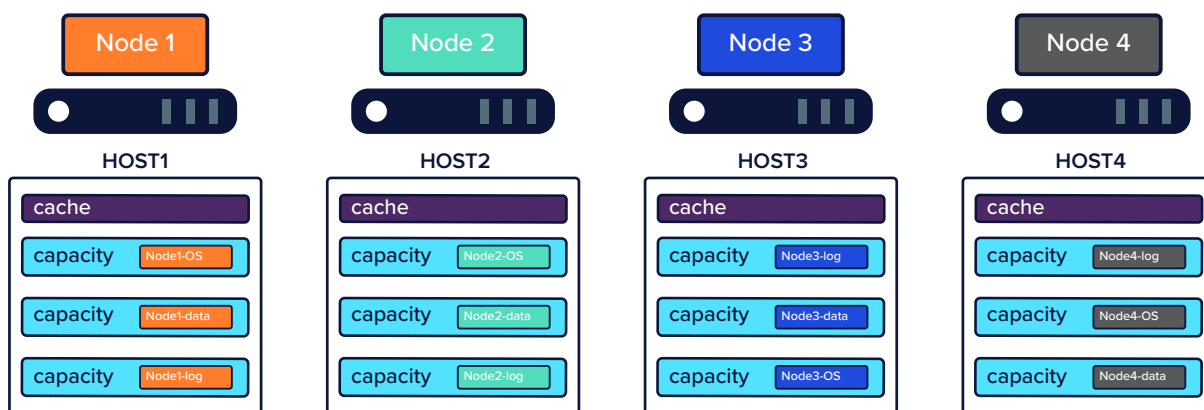
# AGILITY THROUGH STORAGE POLICY BASED MANAGEMENT (SPBM)

VMware vSAN was built around the idea of assigning storage-related settings on a per-VM basis, or even per-virtual disk basis. Administrators can be prescriptive to the specific application needs, based on its role in the data center. They can easily increase the level of failures to tolerate on VM OS disk and log disk, while applying policies for data disks running DataStax Enterprise (DSE), which allows for a level of control that is not possible with other approaches. Administrators can assign key availability and performance capabilities via policy. For instance, the Failures to Tolerate (FTT) capability addresses the key customer and design requirement of availability. With FTT, availability is provided by maintaining replica copies of data to mitigate the risk of a host failure resulting in lost connectivity to data or potential data loss. The FTT policy works in conjunction with VMware vSphere® High Availability (HA) to maintain availability.

	COPIES OF DATA ON VSAN
FTT=0	1
FTT=1	2
FTT=2	3

## VSAN FTT=0 WITH HOST AFFINITY FOR APPLICATIONS WITH INTRINSIC HIGH AVAILABILITY CAPABILITIES

VMs are pinned to a specific host to maintain availability of the VM’s data that is being served up locally, thus reducing additional storage overhead and improving performance due to less writes (one copy) at the vSAN layer.



This configuration provides high performance and storage efficiency. Availability guarantees are provided using replication within the database (Note: mixed storage policies, HA, and Distributed Resource Scheduler (DRS) are not currently compatible).

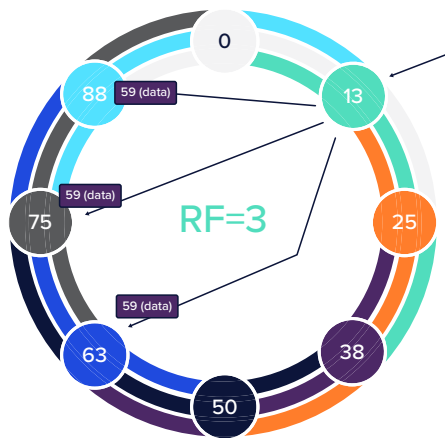
## VSAN FTT=1

The default vSAN storage policy is FTT=1; it allows hypervisor and cluster-level activities such as snapshots, HA, and DRS to be fully aware and compatible with vSAN. Running DSE on vSAN with FTT=1 simplifies deployment while maintaining all operational functionality in the vSphere clusters. vSAN is self-healing and will attempt to re-establish full storage protection compliance with expected limited performance impact to the application. It performs resynchronization actions automatically, all while maintaining a fair balance of resynchronization and guest VM traffic to ensure that DSE can maintain sufficient levels of performance during failure recovery.

With this configuration, both the vSphere stack and application availability are leveraged, applications benefit from vSAN's storage-level protection and vSphere HA's automated recovery from host failures.

## DSE REPLICATION

DSE stores replicas on multiple nodes to ensure reliability and fault tolerance. A replication strategy determines the particular nodes where replicas are placed. The total number of replicas across the cluster is referred to as the replication factor. The commonly used replication factor is 3 and results in three copies of data spread across each data center. All replicas are equally important; there is no primary or master replica.



## DSE RACK AWARENESS

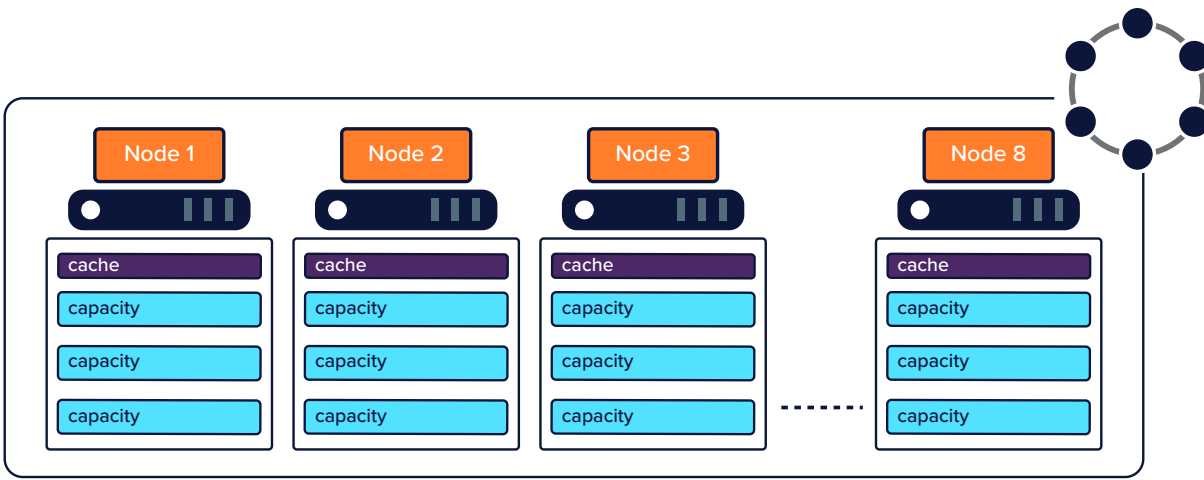
DSE uses racks so that no replica is stored redundantly inside a rack, ensuring replicas are spread throughout different racks in case one rack goes down.

We leverage [rack-aware snitch](#) to ensure that multiple replicas are not stored on DSE nodes of the same ESXi host. With the rack name in the `cassandra-rackdc.properties` file aligned to a physical chassis name, replicas will split across chassis for a given token.

# NODE DENSITY

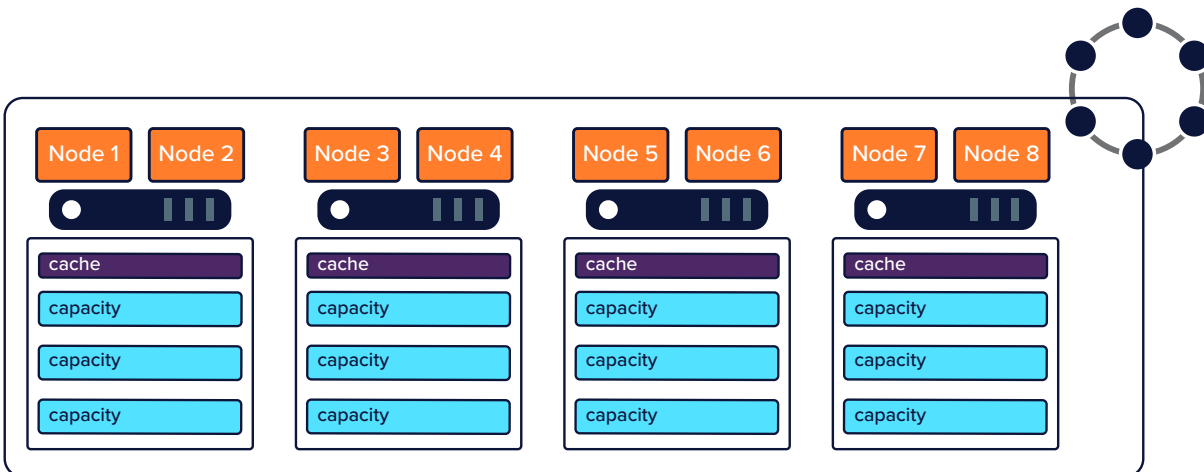
## OPTION 1: ONE DSE NODE PER HOST

In option 1, a DSE cluster is deployed on some ESXi hosts with equal or less than one DSE node per host. Running a DSE cluster in this configuration does not require any changes to the topology or rack configuration, the DSE node fully uses the computing and storage resource of the host, and all deployment and operations are similar to the physical servers. There's no architecture change if customers are migrating their existing physical environment to a virtual environment.



## OPTION 2: MULTIPLE DSE NODES PER HOST

Option 2 shows multiple DSE nodes per host; it offers the server consolidation benefits and provides more flexibility for DSE cluster resource assignments. In this option, DSE rack-aware snitch needs to be configured to ensure DSE replicas are not placed on the same physical host.



For small DSE cluster deployments, we use only one vSAN cluster for management simplicity; for large deployments, multiple vSAN clusters is a feasible option.

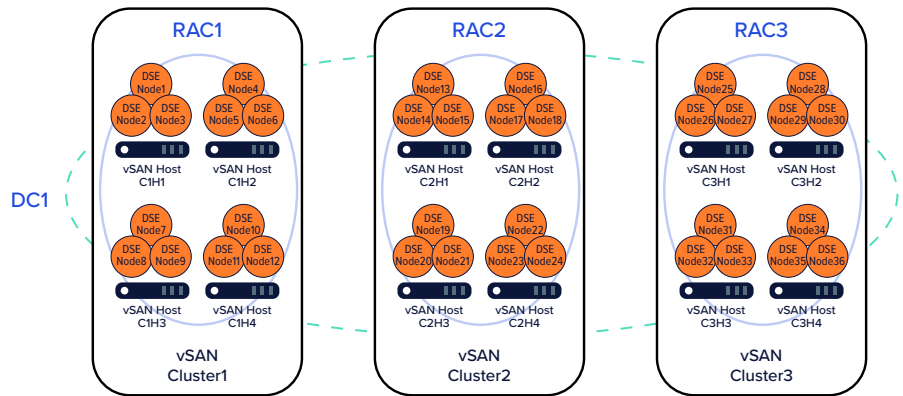
AGILITY PERSONIFIED

- ✓ Adopt and integrate the very latest hardware technologies like 3D XPoint NVMe devices into a cluster.
- ✓ Scale up or out incrementally, as needed by an organization.
- ✓ Maintain full independence of storage from demands of other clusters. Just as with compute and memory, vSAN storage is a cluster resource that remains independent from other clusters.

LEARN MORE ABOUT

- ✓ **Reference architecture:** [DataStax Enterprise on VMware vSAN 6.7](#)
- ✓ **VirtualBlocks:** VMware’s blog site for all topics related to storage and availability
- ✓ **StorageHub:** The one-stop location for all documentation on storage and availability

# MULTIPLE VSAN CLUSTER DEPLOYMENT



As shown in the multiple vSAN cluster deployment diagram, we utilize DSE’s rack awareness, where the DSE object model of a “rack” is actually one or more vSAN clusters.

## VSAN FTT=1

This option is similar to one vSAN cluster deployment, it combines availability from both the vSphere stack and DSE’s built-in resiliency.

## VSAN FTT=0 WITHOUT HOST AFFINITY

In multiple vSAN cluster deployments, data locality is not required for data availability. Application-level availability is guaranteed by ensuring that DSE racks are aligned with vSAN clusters. This deployment scenario requires careful management and will increase the number of virtual machines affected by capacity disk, cache disk, network, or host failures.

## VSAN FTT=0 WITH HOST AFFINITY

This option reduces complexity and increases performance. Host Affinity pins the virtual machine’s data drives to the host it is running on. The Host Affinity feature is currently an RPQ feature and is not fully integrated with vSphere’s features such as DRS and HA.

See the following chart for various configurations in one vSAN cluster and multiple vSAN clusters.

	BENEFITS	TRADEOFFS	AVAILABILITY	CONFIGURATION
FTT=0	Performance (fast). No additional storage compared to Cassandra bare metal	Requires careful planning; disk failures can require more frequent node rebuilds	Available today in all vSAN multi-cluster deployments (minimum 3).	OS, logs, and data configured with FTT=0; storage may not be on the same node as compute
FTT=0 with Host Affinity	Performance (faster). No additional storage compared to Cassandra bare metal	Non-GA solution	Potential future availability in Kubernetes service model	OS, logs, and data configured with FTT=0 with Host Affinity; storage local to VM's
FTT=1	Simple management, HA, V-Motion, DRS available	Minimal tradeoff in performance requires 2x storage	Available today in all vSAN offerings	OS, logs, and data configured with FTT=1
Combined FTT=0 and FTT=1	Probable reduction in manual tasks associated with rebuilds	Requires careful planning and may require node rebuilds on disk failure	Available today in all vSAN offerings	OS configured with FTT=1, logs, and data configured with FTT=0

## TAKEAWAY

DSE contains application-level resilience through replication, and the enterprise-class resiliency features within vSAN provides user flexibility for various application needs. User customizable, storage-level resilience is built directly into vSAN, and is suitable for even the most demanding requirements.

The architecture of VMware vSAN allows it to address agility, consistent performance, scalability, and resiliency requirements that are top needs for data center administrators and application owners alike. Whether the DSE deployments are part of a large and coordinated deployment for a targeted set of applications, or a smaller scale deployment, vSAN provides operational simplicity and flexibility to meet the demands of customers' modern application services.



## ABOUT DATASTAX

---

DataStax delivers the only active everywhere hybrid cloud database built on Apache Cassandra™: DataStax Enterprise and DataStax Distribution of Apache Cassandra, a production-certified, 100% open source compatible distribution of Cassandra with expert support. The foundation for contextual, always-on, real-time, distributed applications at scale, DataStax makes it easy for enterprises to seamlessly build and deploy modern applications in hybrid cloud. DataStax also offers DataStax Managed Services, a fully managed, white-glove service with guaranteed uptime, end-to-end security, and 24x7x365 lights-out management provided by experts at handling enterprise applications at cloud scale. More than 400 of the world's leading brands like Capital One, Cisco, Comcast, Delta Airlines, eBay, Macy's, McDonald's, Safeway, Sony, and Walmart use DataStax to build modern applications that can work across any cloud. For more information, visit [www.DataStax.com](http://www.DataStax.com) and follow us on Twitter [@DataStax](https://twitter.com/DataStax).

© 2019 DataStax, All Rights Reserved. DataStax, Titan, and TitanDB are registered trademarks of DataStax, Inc. and its subsidiaries in the United States and/or other countries.

Apache, Apache Cassandra, and Cassandra are either registered trademarks or trademarks of the Apache Software Foundation or its subsidiaries in Canada, the United States, and/or other countries.