

Simplify your Infrastructure with Hyperconverged Infrastructure (HCI)

Technical Overview

Presenter: Patrick Fagen
Presales Engineer



Agenda

- Enterprise Storage Overview
- 2020 Storage Concepts
- VxRail Overview
- VMware's Storage Solution
- vSAN Deep Dive
- VxRail Failure Scenarios
- Learn more from Sterling

Sterling Client to Cloud™

Sterling, a leader in secure technology solutions, is fully capable and prepared to offer our customers complete customized solutions from *Client to Cloud*.

Core Competencies

- Datacenter, Cloud, Security, Managed Services, Client Services

Contract Vehicles:

- CIO-CS, SEWP-V, ADMC-3, GSA, NASPO, MHEC

Proven Execution:

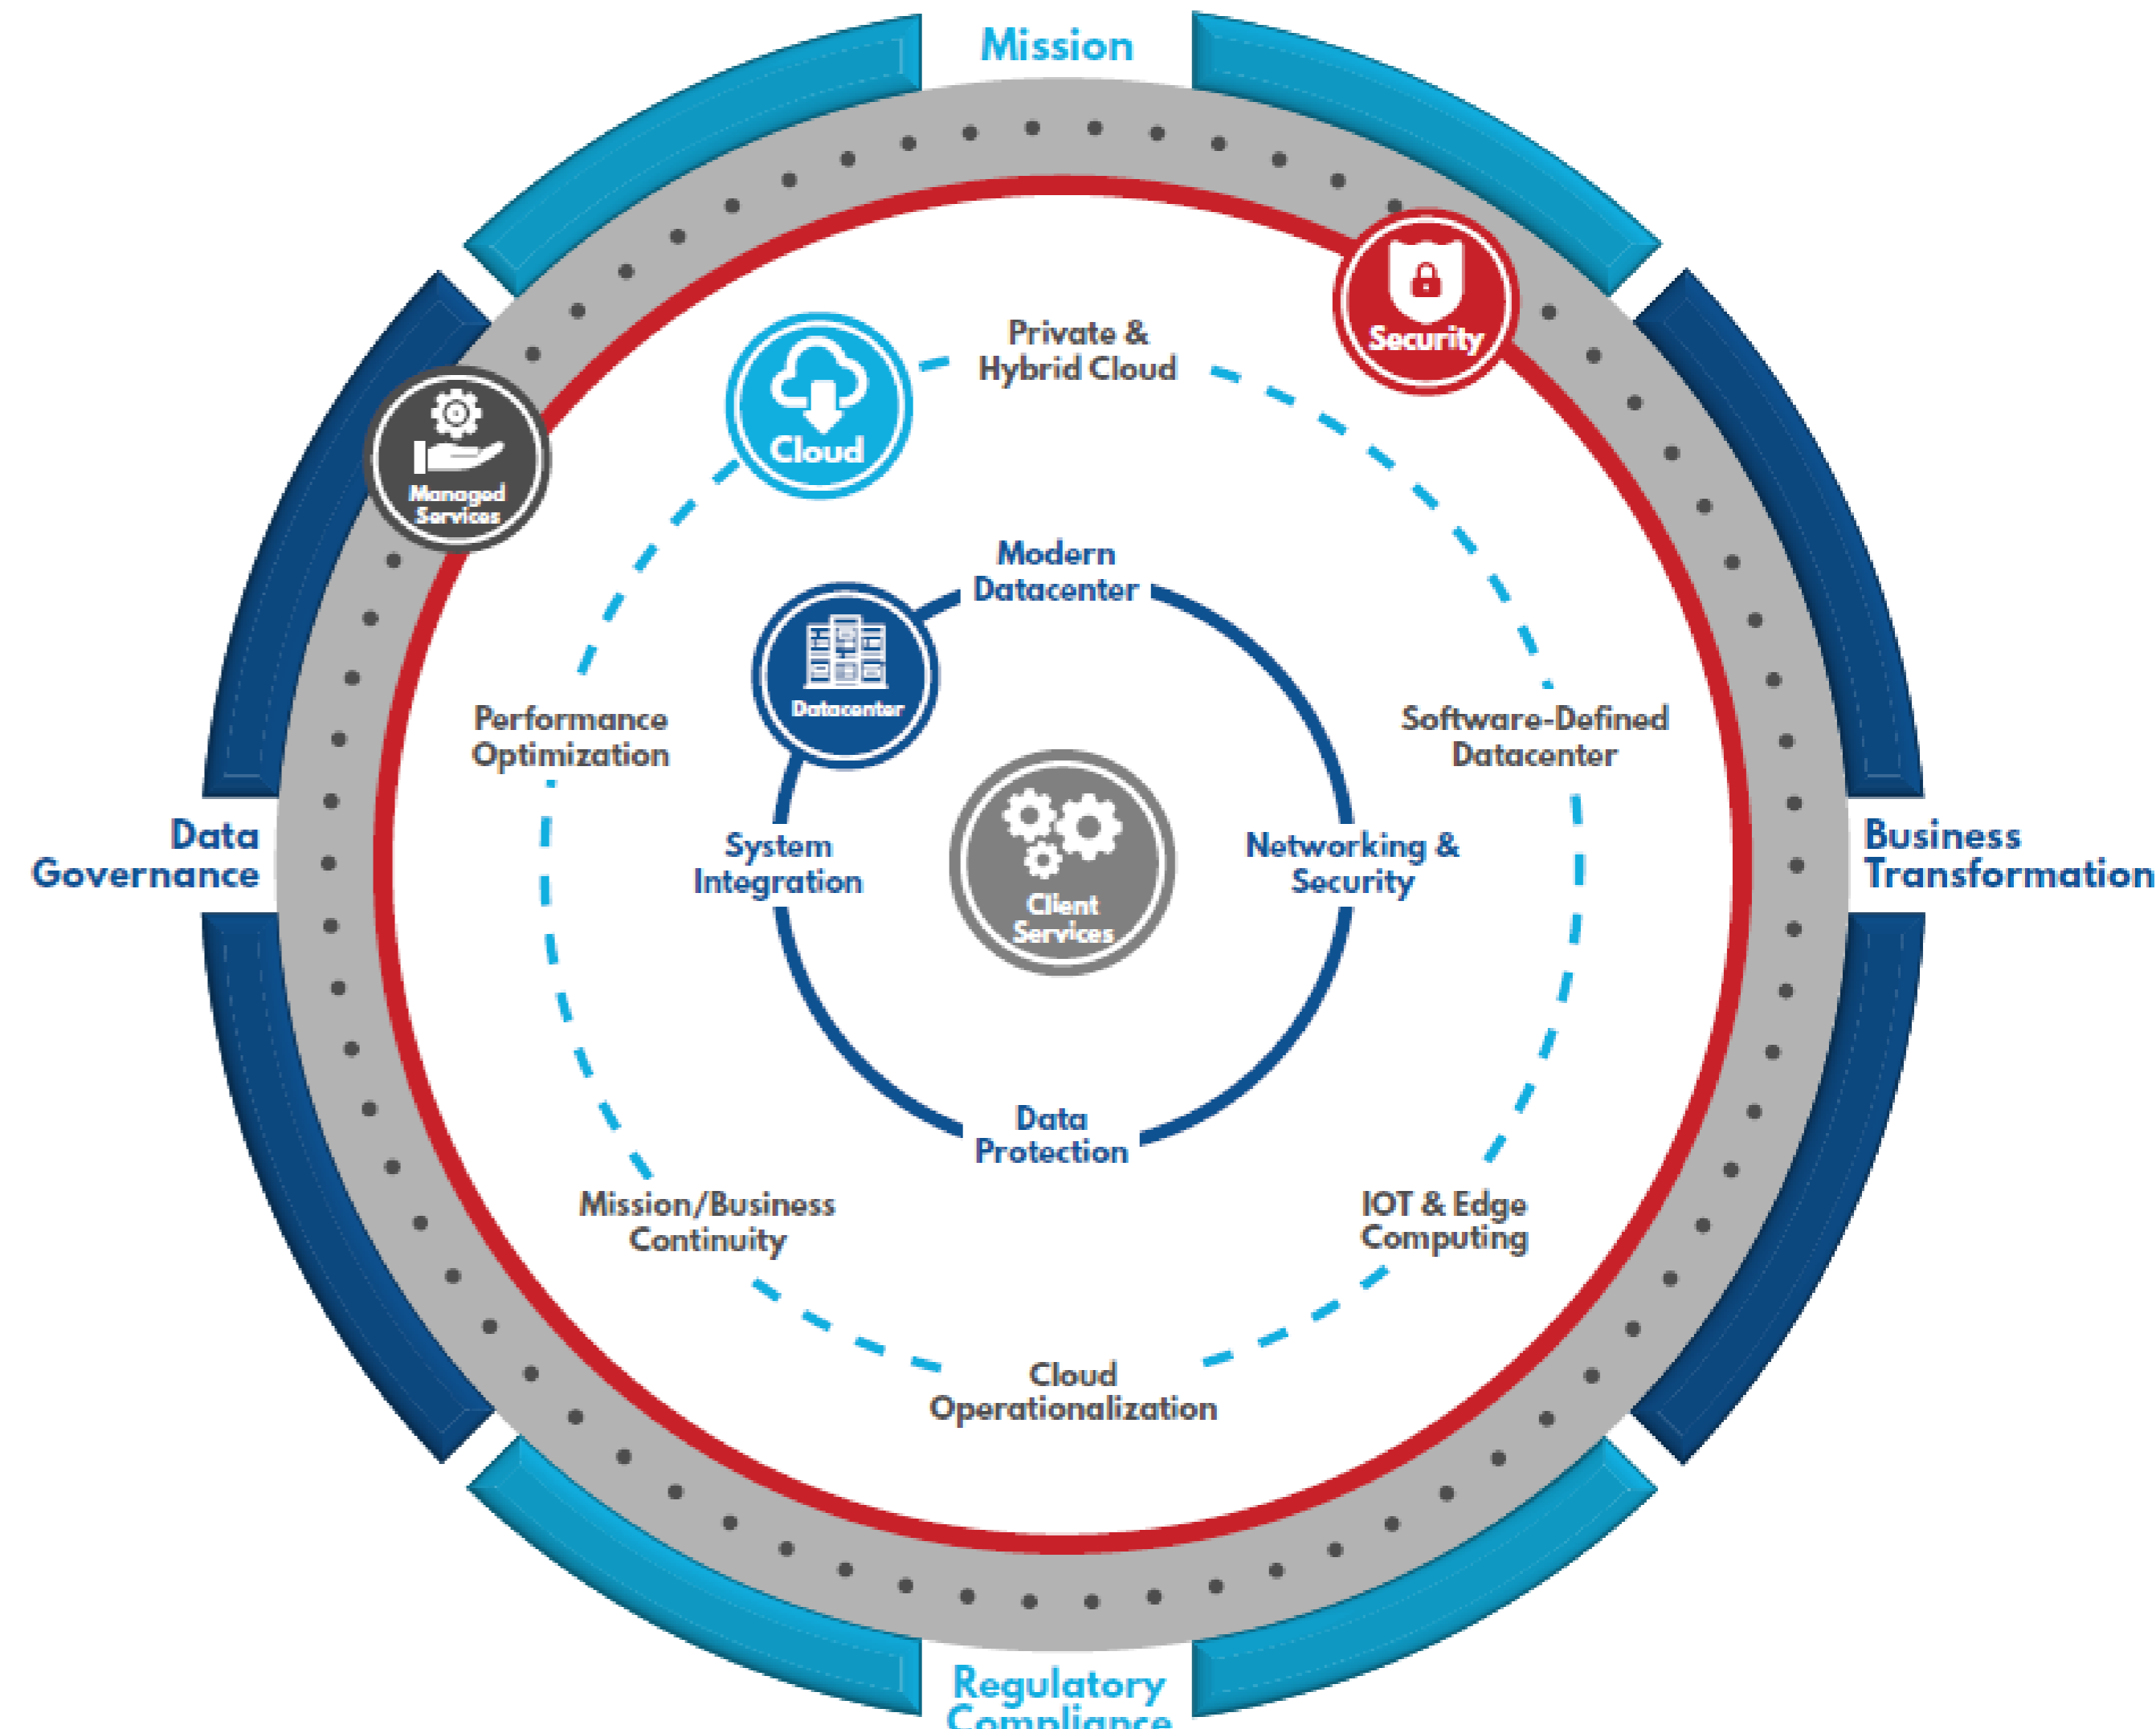
- Two decades of proven performance in the Federal Government, State and Local Government, Education, and Commercial Sectors

Certifications and Demo:

- Multiple VMware, AWS, Azure, Oracle, and Google Cloud competencies and certifications.
- On-Premises multicloud datacenter, fully operational for Demo's and customer engagements utilizing VMware vCF 4.0 to include a full substantiation and integration of the VMware Tanzu portfolio.

Robust Secure Supply Chain:

- O-TPPS Certified, CMMC Level 3 Certified, Complete SCRM plan




It's time to rethink traditional hardware infrastructures!




Traditional Data Centers are Unable to Keep Up

But public clouds have their own challenges



Traditional Architectures



Controlled



Secured



Compliant



Performant

Public Clouds



Innovation




On-Demand



Time to Market




Ease of Scale



Complex Life cycle Management



Limited automation impacts developer agility



Under-utilized Investments



Shadow IT limits Visibility & Control



Enterprise Compliance



Cost Premium

Traditional Apps

Ops Driven
Cost Focus


Cloud Apps

Apps Driven
Agility Focus



Hardware-Based Infrastructure Can Be Challenging

NOTE: This approach limits the ability to deploy a future-proof Private Cloud!




Complex

Rely on manual, specialized processes to deploy and operate




Slow

Long lead time to provision IT and application resources



Inefficient

Unplanned downtime, wasted capacity



Non-Compliant

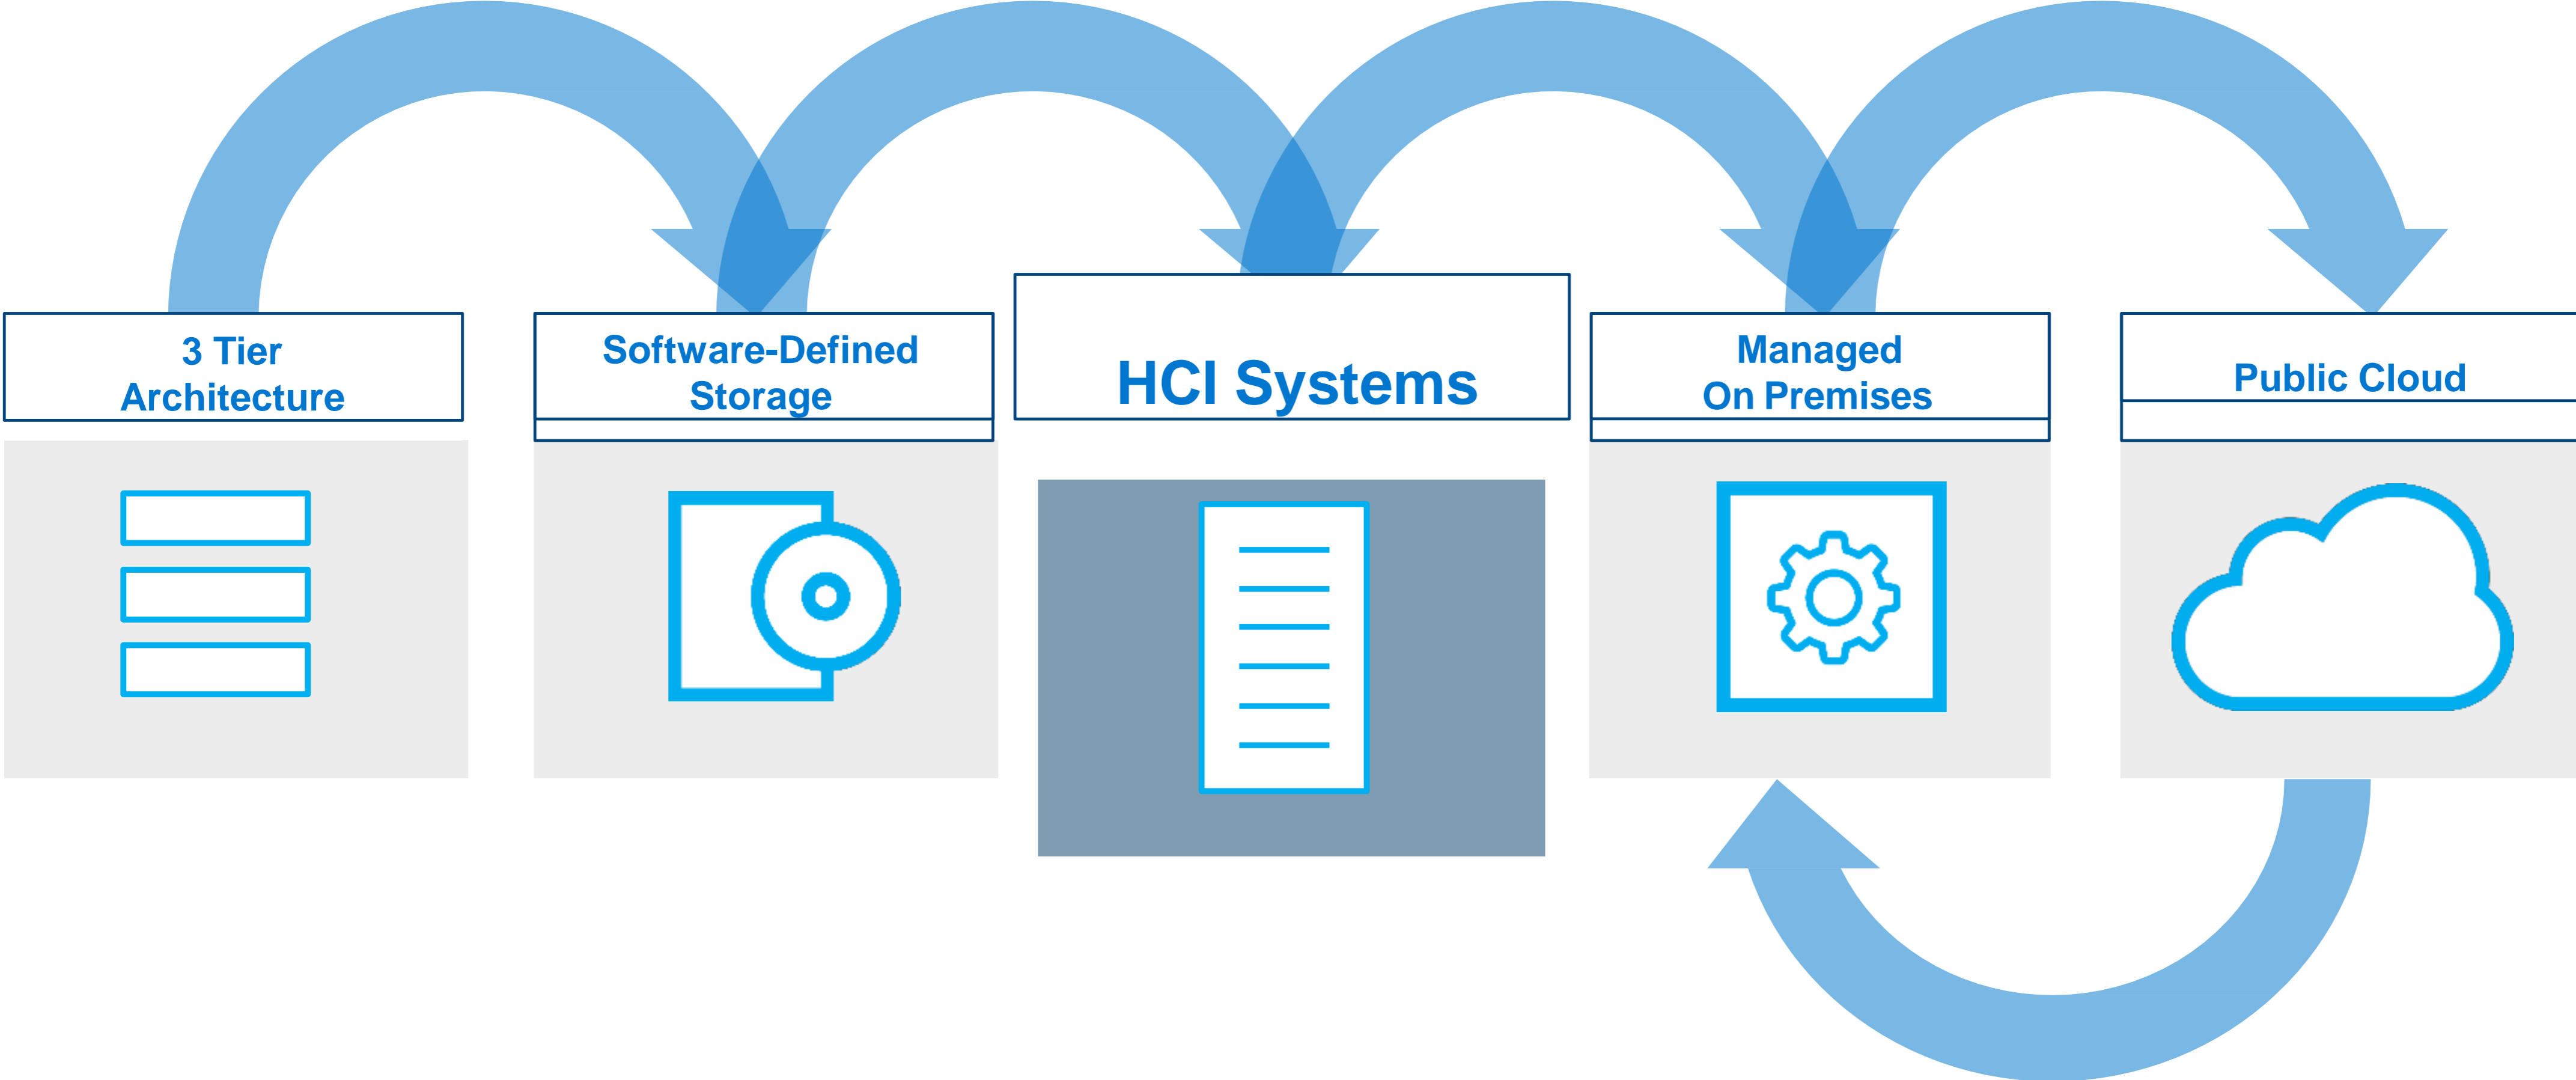
Security, visibility and lack of governance



Siloed

Hard to maintain and manage across clouds

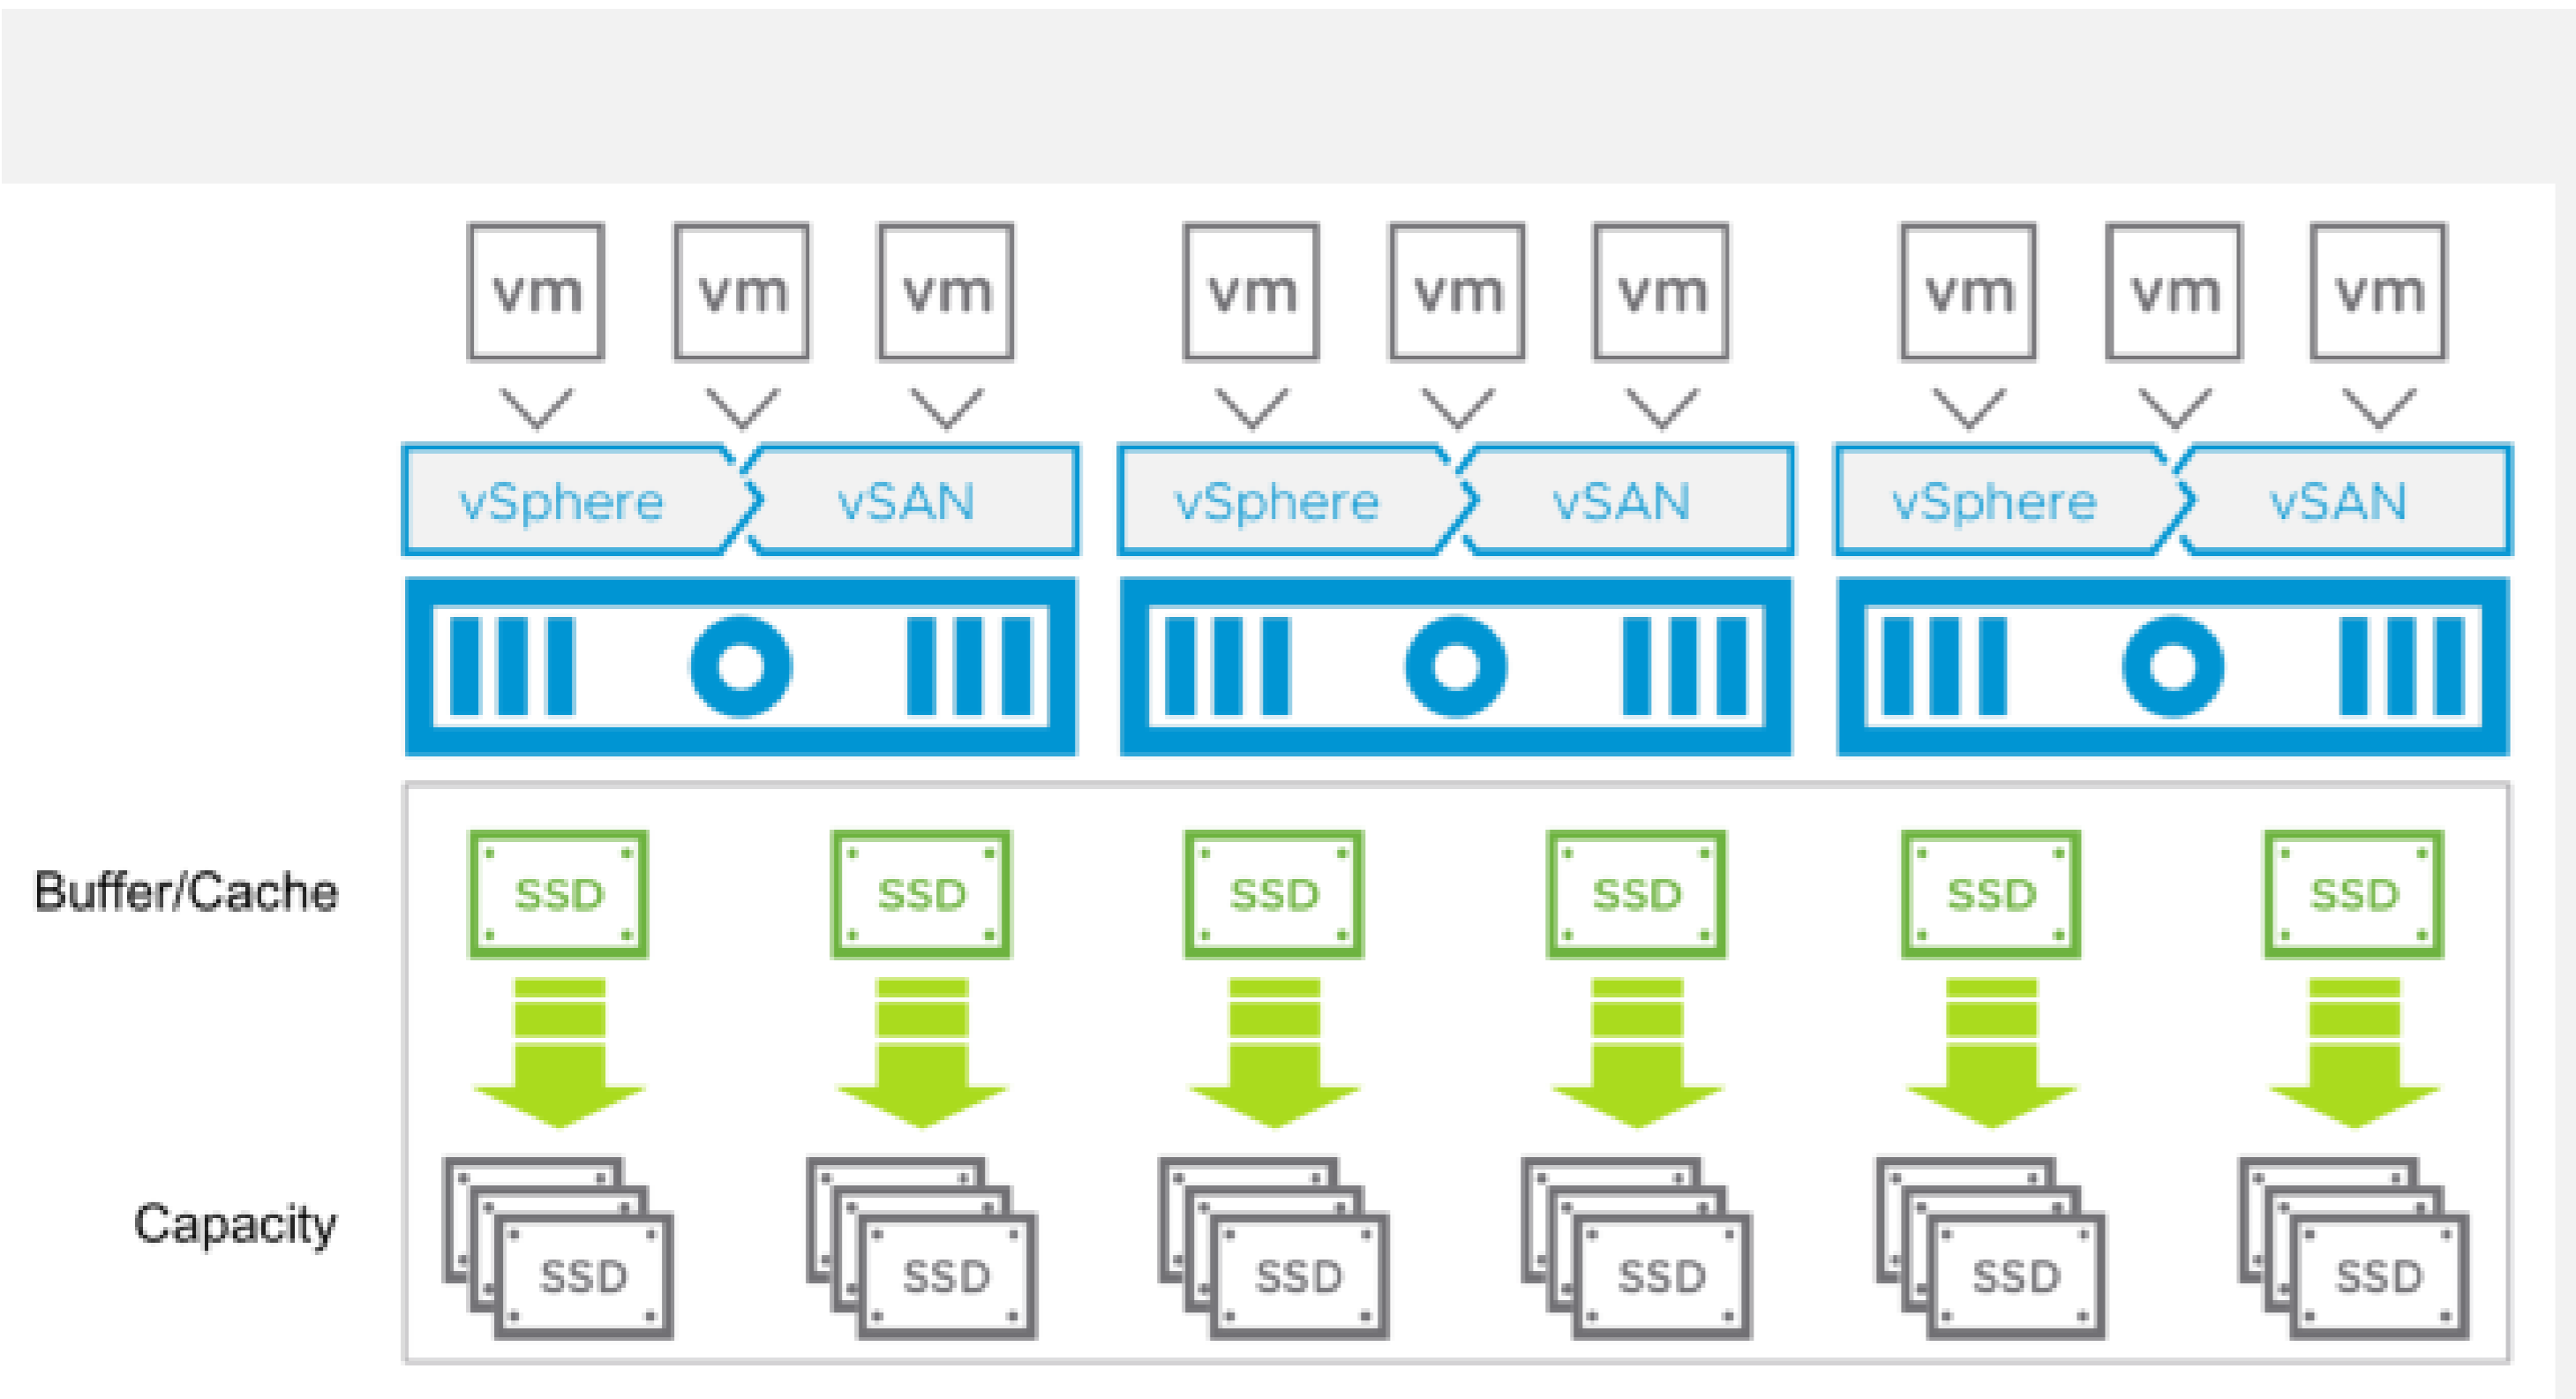
Rapidly Changing Infrastructure



What is Hyperconverged?

Storage and compute for VMs are delivered from the same x86 server platform running the hypervisor.

Industry-standard components are used with direct-attached storage devices, from spinning disks to SSDs to latest NVMe technology.



Key building blocks for HCI

Software-defined storage

- Abstracts storage functions from hardware
- Virtualizes direct-attach storage into shared pool
- Automates provisioning, load balancing

x86 servers


- High-performance processors, large memory
- Flash media delivers consistent, predictable performance

Virtualization

- Abstracts compute and network functions
- Enables physical resources to be shared
- Greatly improves utilization, mobility and security


High-speed Ethernet

- Connects nodes together to create cluster
- Enables HCI to deliver high IOPS and reduced latencies



Customer Examples

Problems with status-quo



2020 Enterprise Storage Vision

Software is eating the world....
Including Storage Arrays



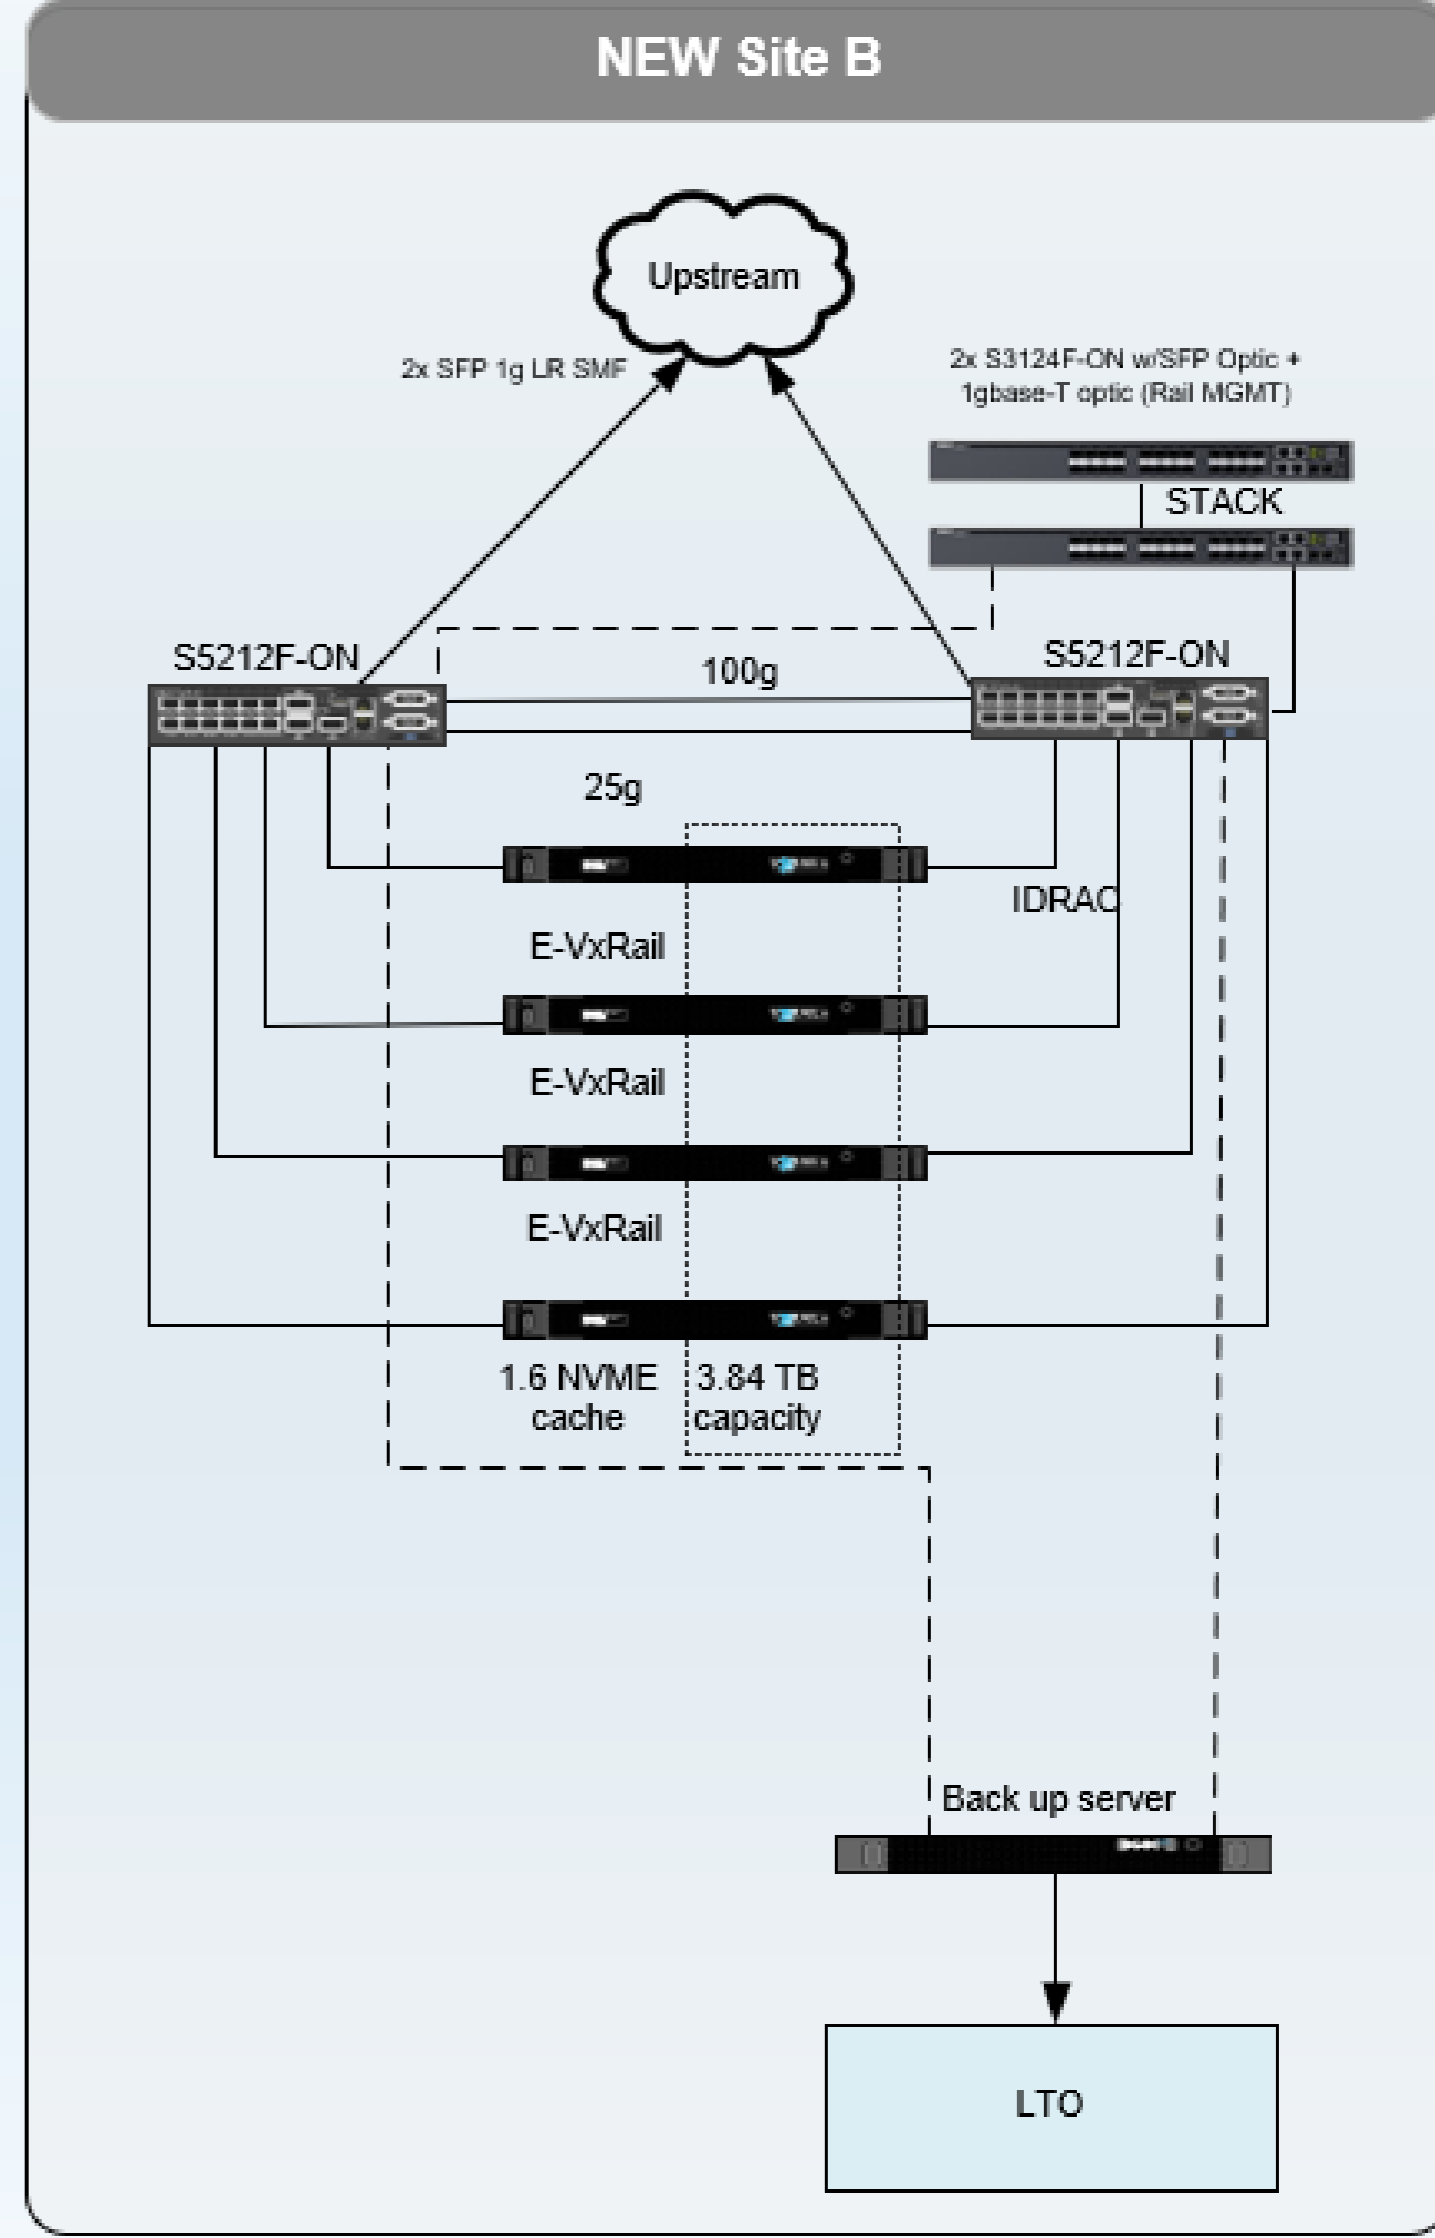
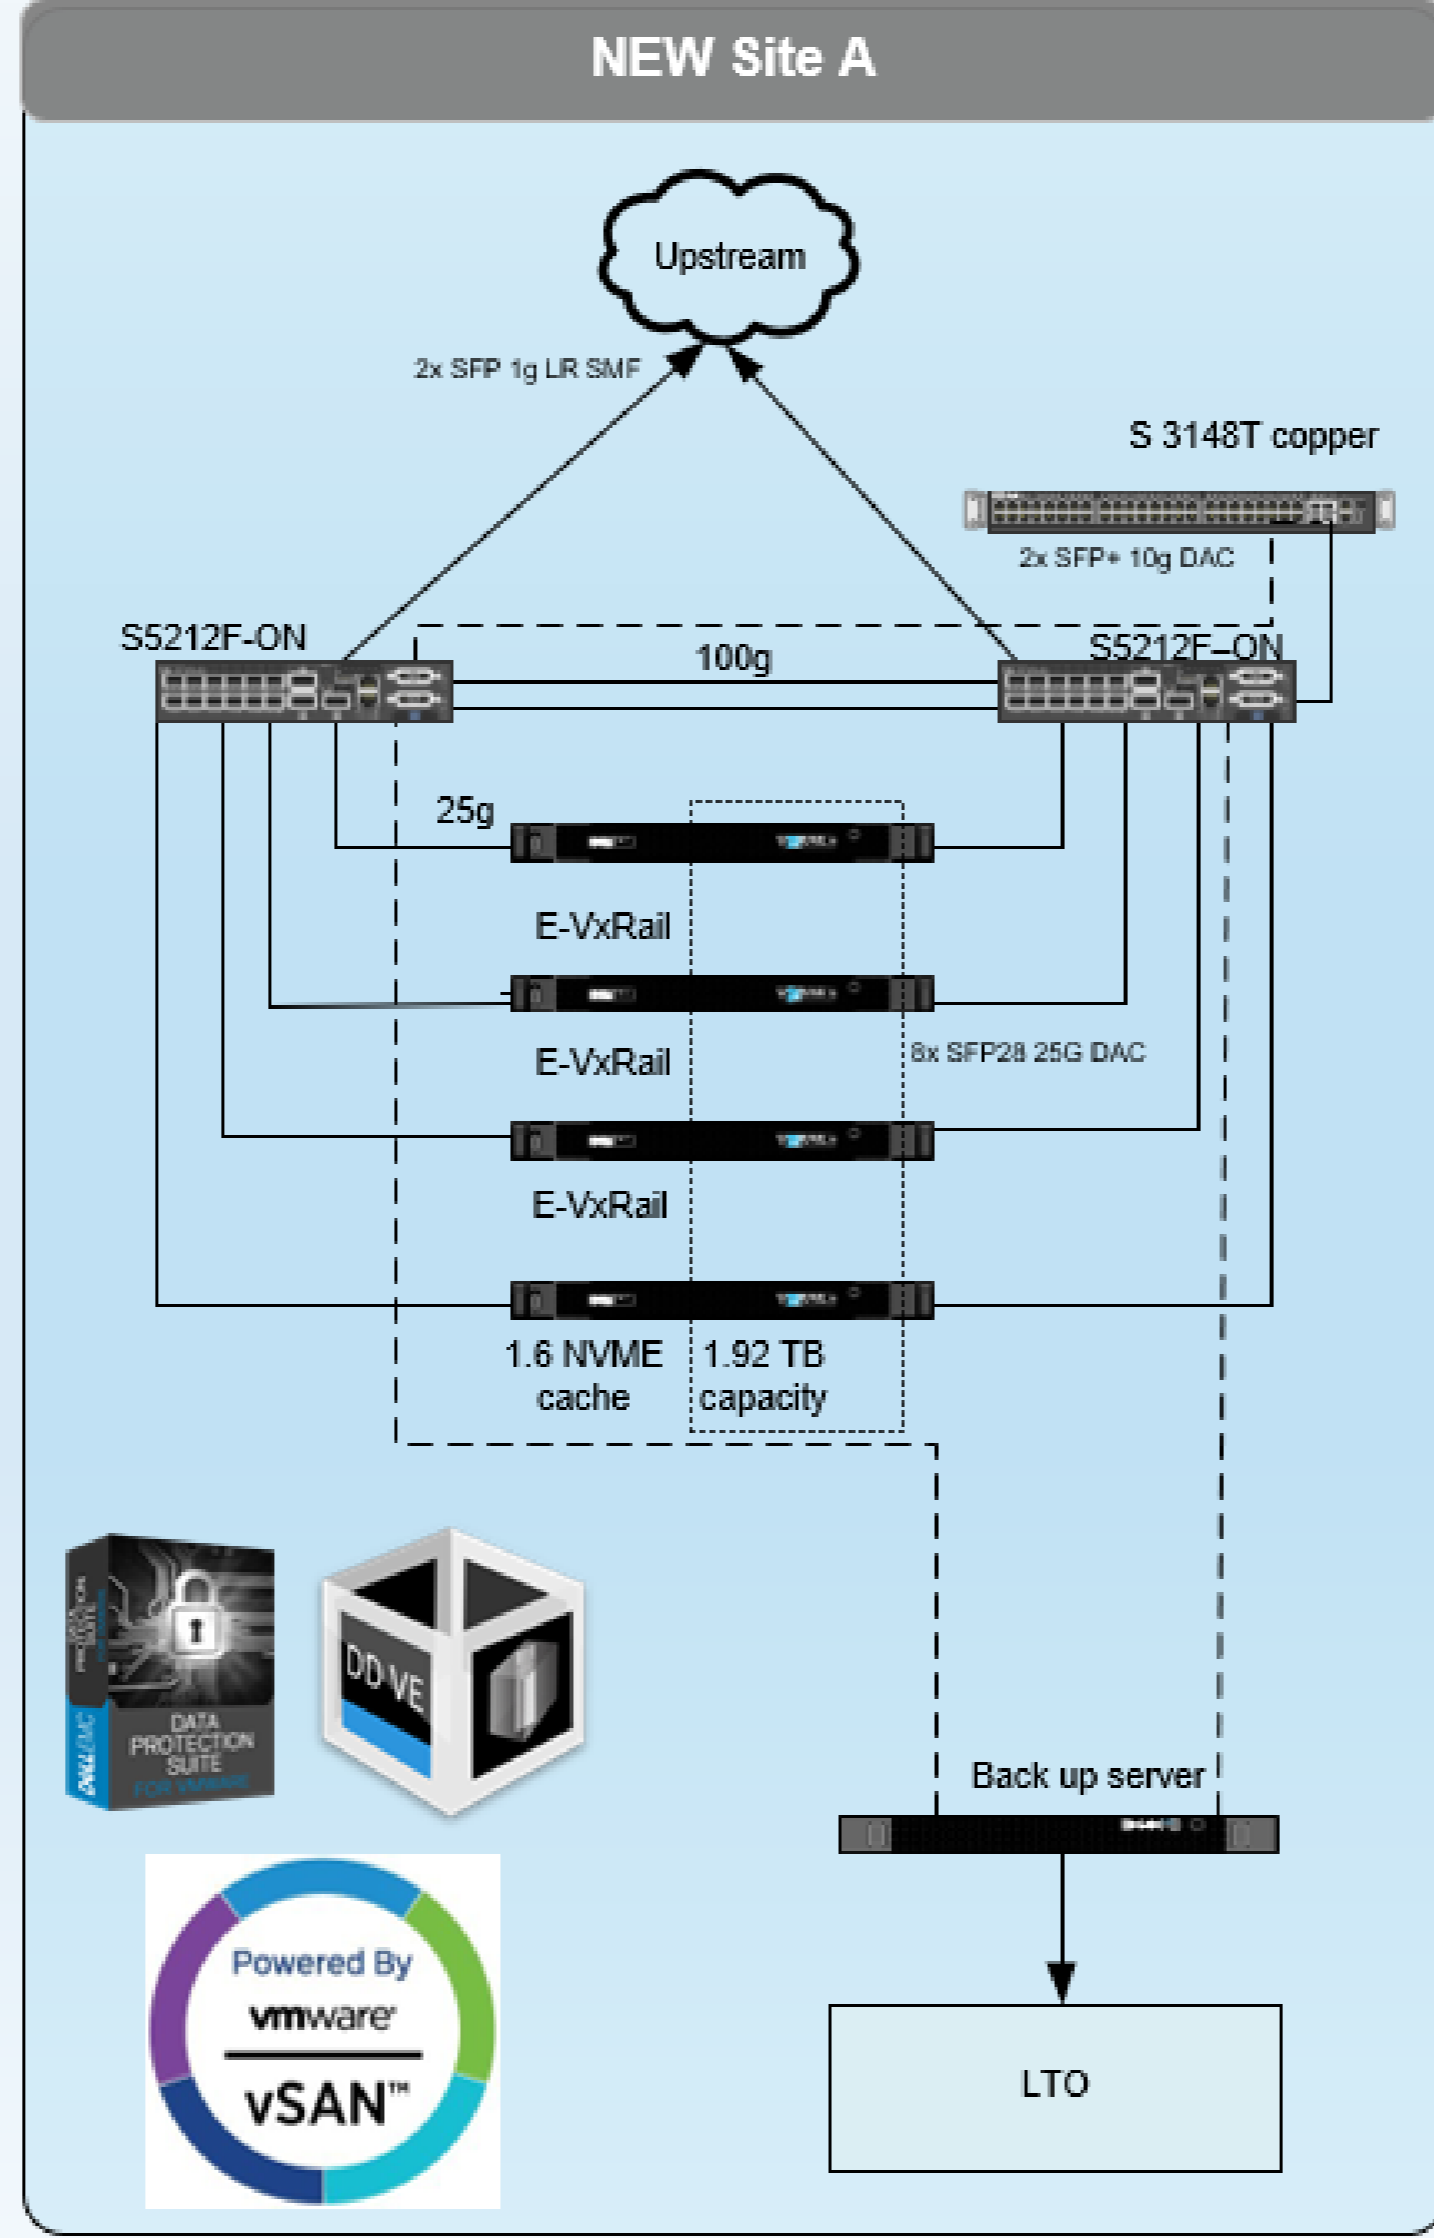
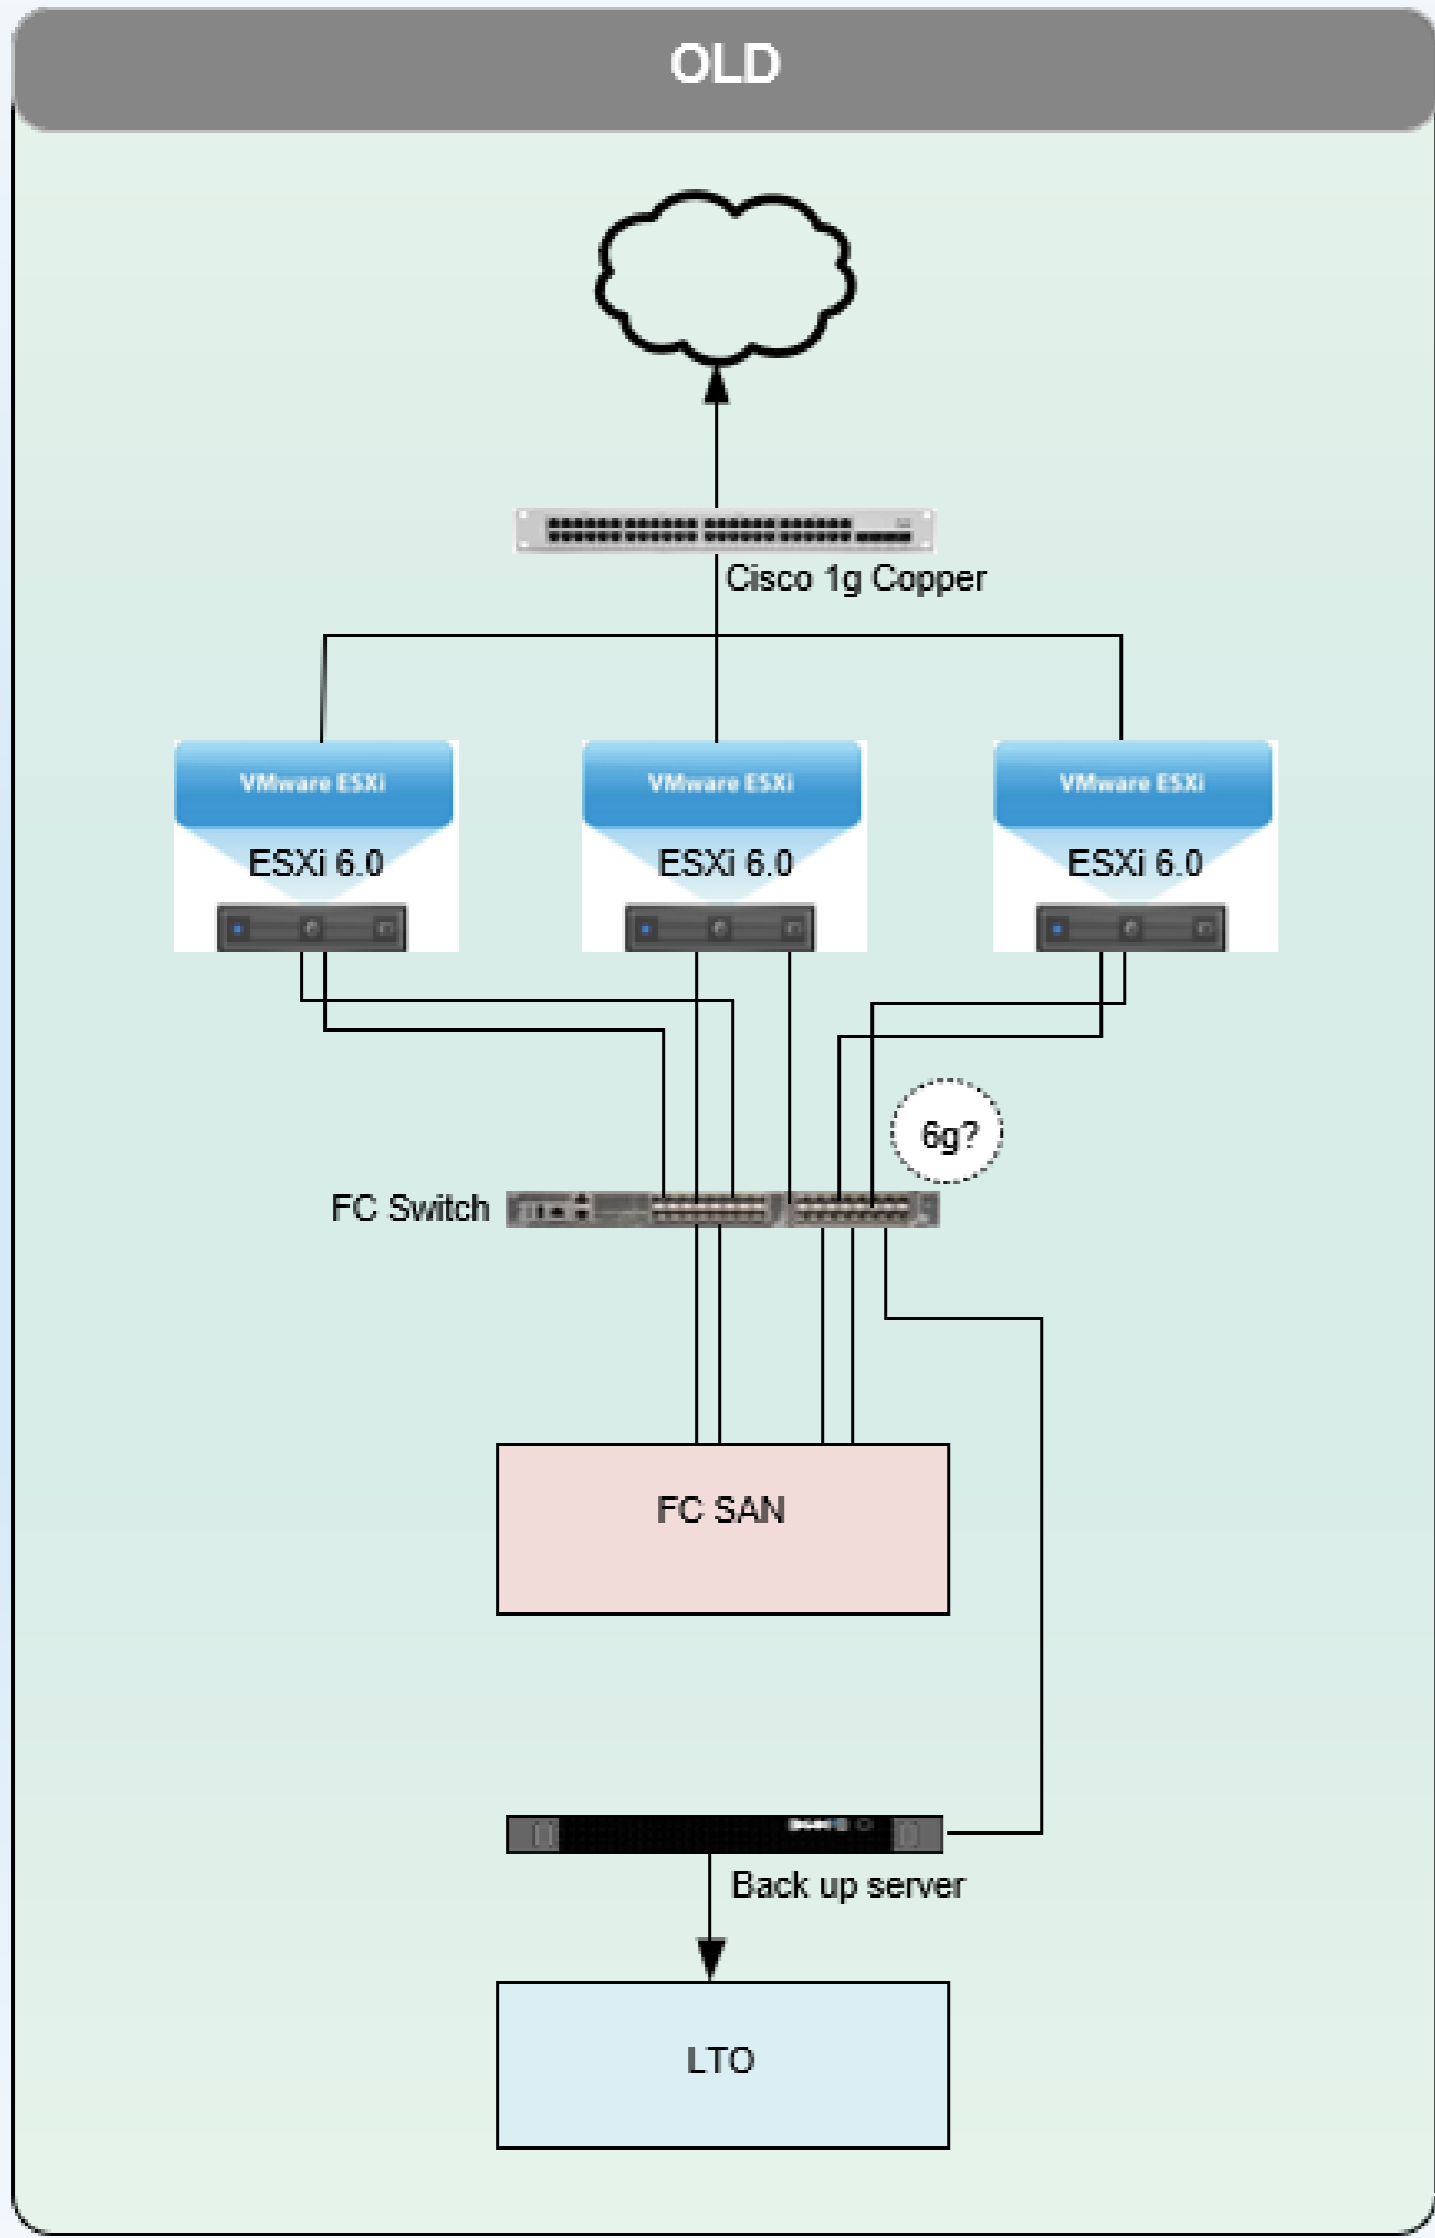
Agility
Public cloud speed, efficiency, and economics within the data center

Scalability
Start small and easily scale up or scale out while maintaining performance levels

Simplicity
Simplify operations with software-driven automation and lifecycle management


Improve efficiency

Lower OPEX costs

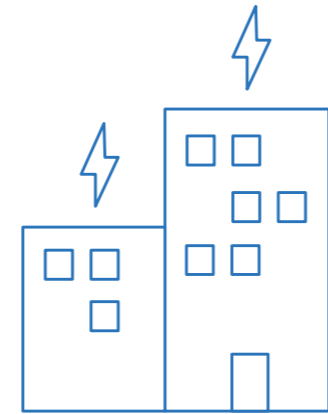


Lifecycle management: Typical challenges

New!



Overburdened IT staff



Disruption to business



Budget efficiencies

- Limited resources to research upgrade options & optimal components
- Testing diverts time & resources from strategic activities
- Patching and updating is a continuous process


- Unsuccessful migrations or upgrades could result in outages or need to go back to prior state
- Need to schedule downtime for upgrades or migrate workloads
- Incremental Change management requests can delay updates

- Disproportionate amount of resources spent on managing infrastructure
- Administrative time to research and upgrade drains OpEx



Introducing: Dell EMC VxRail

The fastest & simplest path to IT outcomes in VMware environments



Turnkey experience

Fully loaded, secure VxRail HCI with Full Stack integrity to create the most seamless and automated VMware experience



Lifecycle management

Pre-tested, validated, engineered for automated end-to-end lifecycle management and non-disruptive upgrades... so you're always in a good known state



Highly differentiated

Only jointly engineered HCI system with VMware, the first network configuration automation, and backed by single support, Future Proof and Flex on Demand programs

What is Inside VxRail?



VMware SDDC

- Choice of vSAN
- VMware Cloud Foundation
 - vCenter server
 - vRealize Suite Ready
 - vSphere Ready*

VxRail HCI System Software

- VxRail Manager
 - VxRail ACE (Analytical Consulting Engine)
 - RESTful APIs
- Automation & orchestration services
 - Backend services and downstream connectors

Data Protection Options

- RecoverPoint for VMs
- VMware vSphere replication

*Compatible with a broad range of customer-supplied vSphere licenses




What is a node?

VxRail Hyperconverged, Self-contained Infrastructure

What is in a node?




Processor




RAM




Redundant power
& cooling



Network Connectivity
(1Gbps/10Gbps/25Gbps)



GPU (V Series)








All-flash or hybrid
disk packs



VxRail on latest generation Dell EMC servers

Updated!

G Series	E Series	P Series	V Series	S Series	
					
Compute dense	Low profile	Performance optimized	Memory Dense	Storage dense	
G560/F	E560/F/N	P570/F	P580N	S570	
<p>Supporting latest generation Dell EMC PowerEdge servers based on Intel® Xeon® Gen 1 and 2 Processor Scalable Family Higher core counts, faster clock frequency, more memory channels, faster memory, higher endurance and redundant boot devices</p>					
2000W or 2400W 10GbE Optane & NVMe cache Mixed-use SAS cache	1100W or 1600W 10GbE or 25GbE Optane & NVMe cache Mixed-use SAS cache FC HBA GPU- T4	1100W or 1600W 10GbE or 25GbE Optane & NVMe cache Mixed-use SAS cache FC HBA	1600W 10GbE or 25GbE NVMe cache NVMe capacity FC HBA	2000W 10GbE or 25GbE Mixed-use SAS cache FC HBA Dual Socket Only GPUs- P40, RTX6000, V100, T4, M10	1100W 10GbE or 25GbE Hybrid Only Mixed-use SAS cache FC HBA

VxRail Configuration Flexibility for Your Workload

G, E, P, S, V Series based on the latest Dell EMC PowerEdge servers

Updated!

Processor

Choice of 37 Intel® Scalable® Gen 2 processors
Choice of 6 Intel® Scalable® Gen 1 processors

From 4 to 56 cores per system


RAM

24 DIMM slots

16GB RDIMM
32GB RDIMM
64GB LRDIMM, RDIMM
128GB LRDIMM

Power supply

1100W	100-240V AC
1600W, 2000W, 2400W	200-240V AC
1100W	48V DC



Options vary by series

Storage

Cache Drives: Optane 375GB, NVMe 1600GB
SAS 400GB, 800GB, 1600GB, 7.68TB

Capacity SSDs (SAS & SATA) : 1.92TB, 3.84TB
HDDs: 1.2TB, 1.8TB, 2.4TB, 2.0TB 4.0TB 8.0TB

Capacity NVMe: 1TB, 4TB

Base networking

SFP28, SFP+, RJ45

2x 25GbE

4x 10GbE

2x 10GbE

4x 1GbE (4x 10GbE auto-negotiate)

Optional add-on NICs, FC HBA

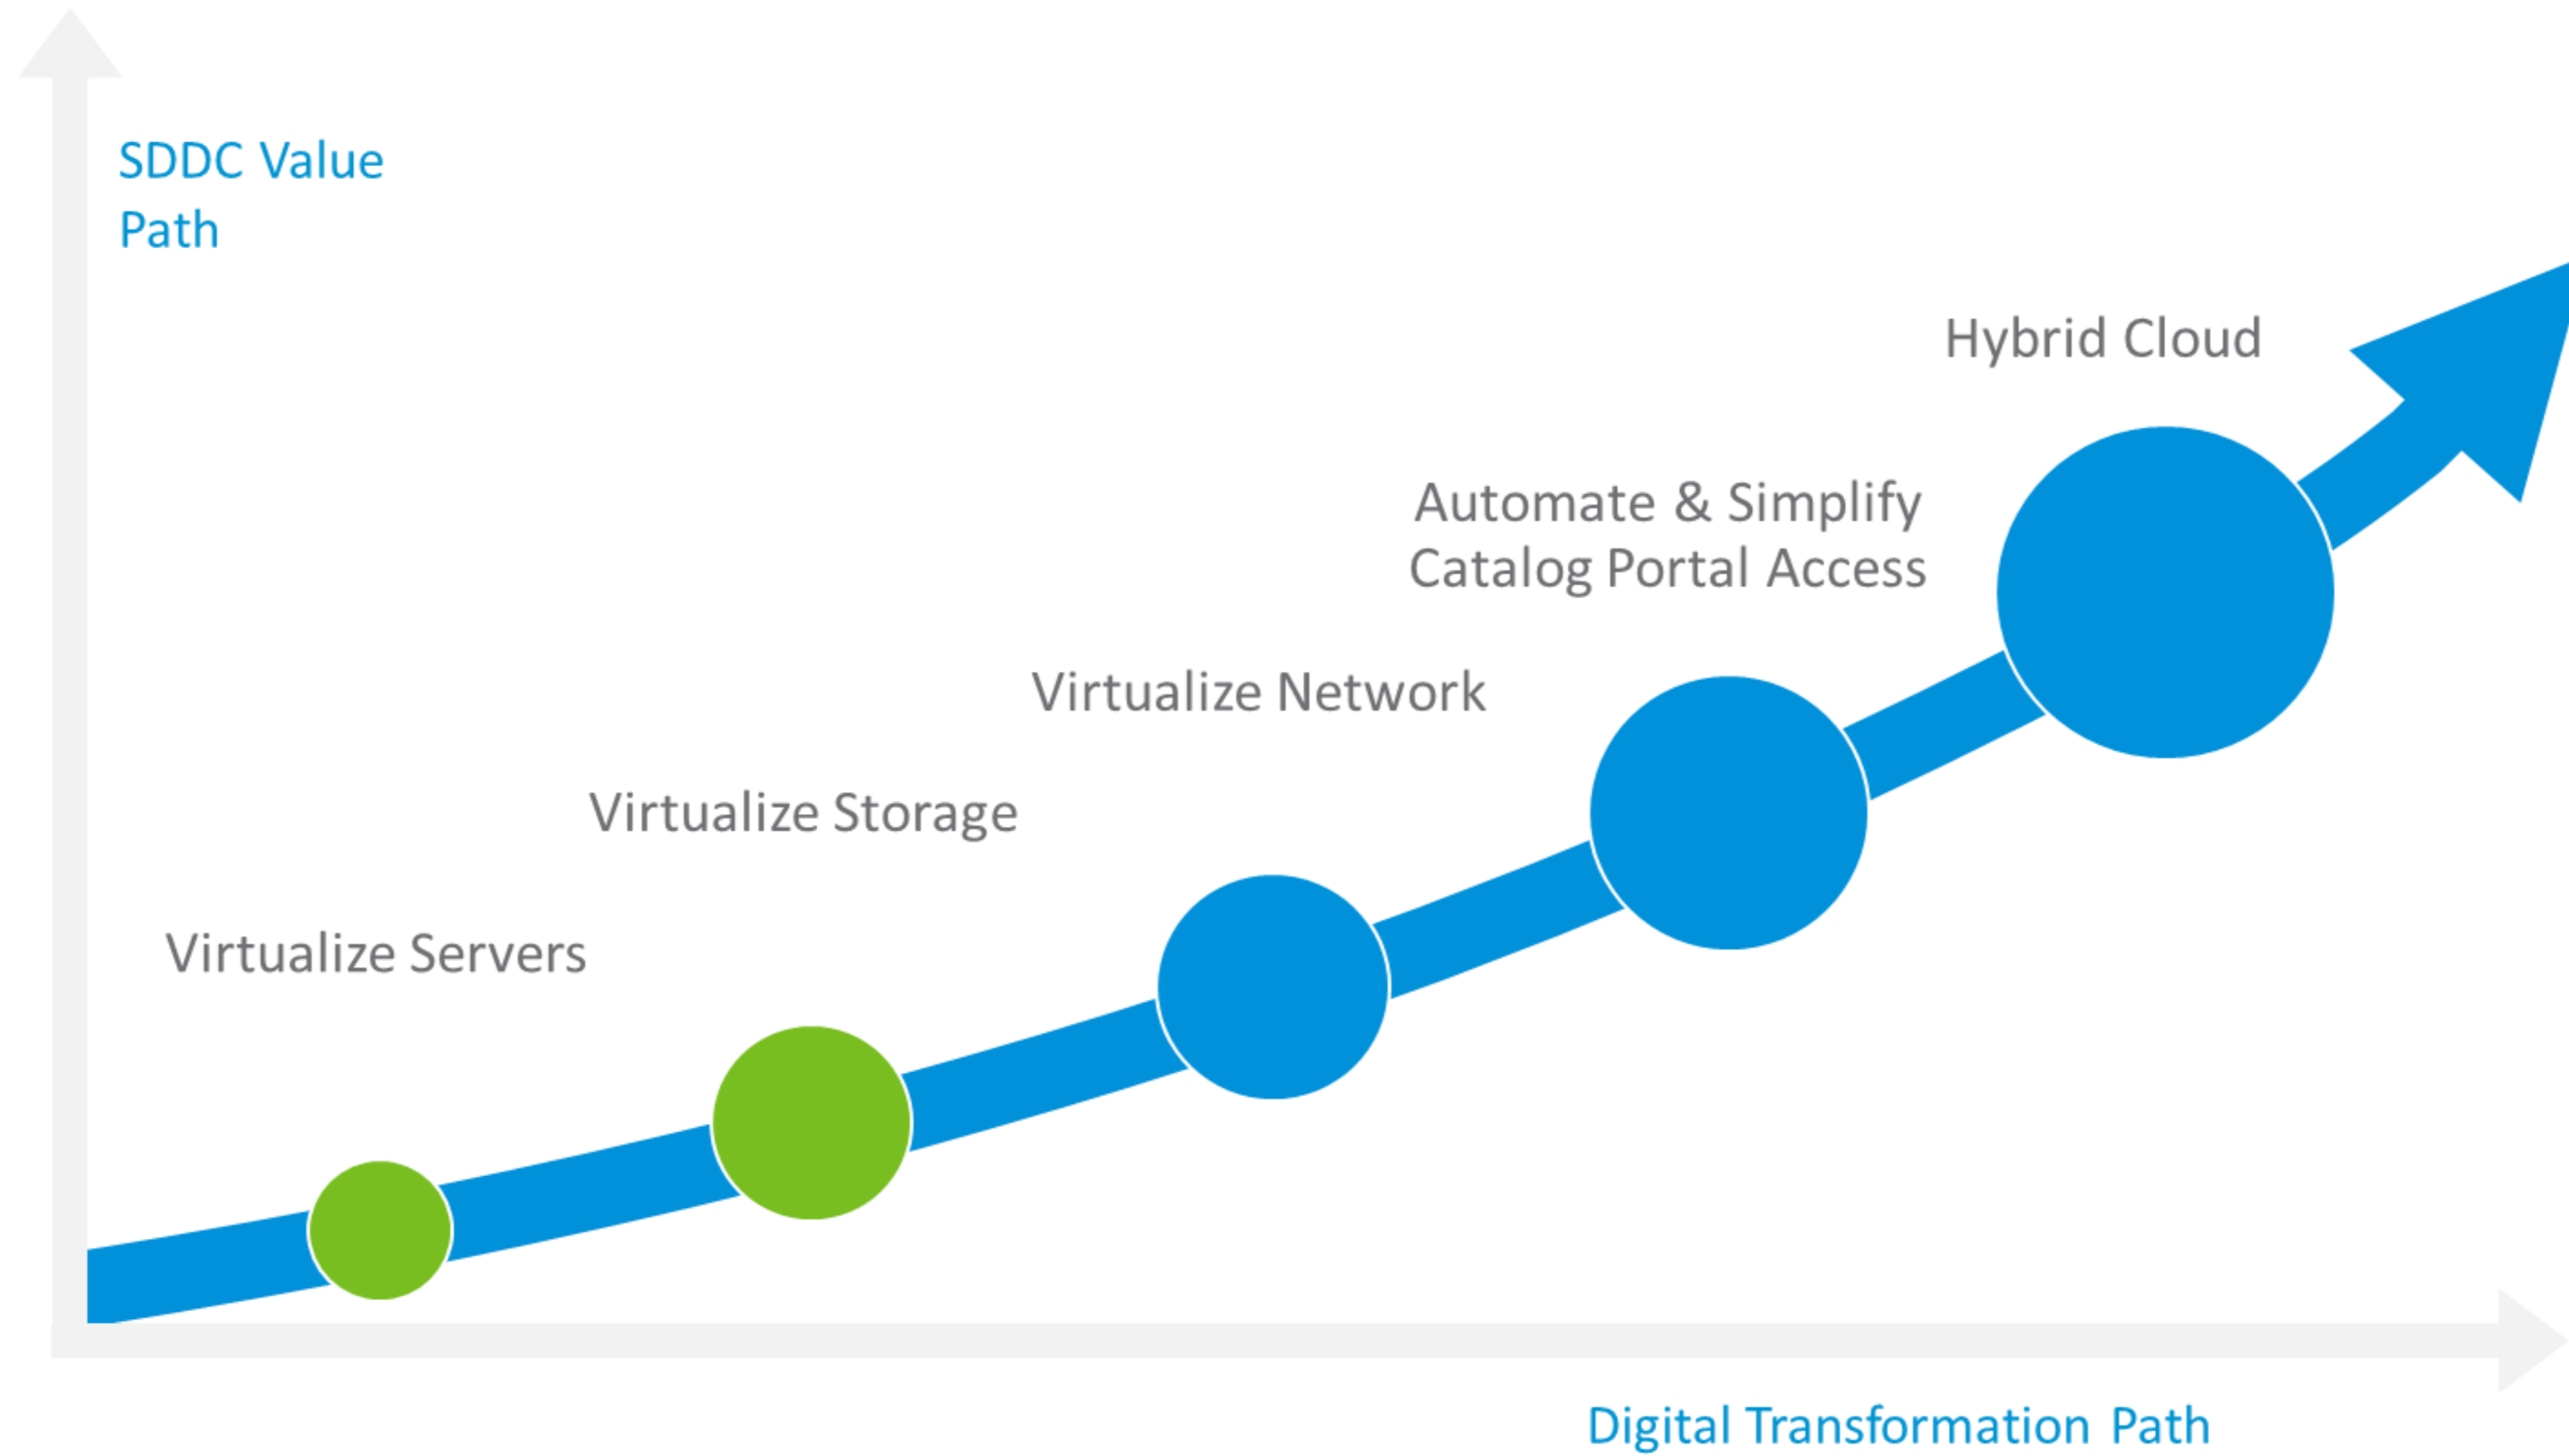
GPUs

NVIDIA P40, V100, T4, M10

Note: GPU SW & drivers
sold separately

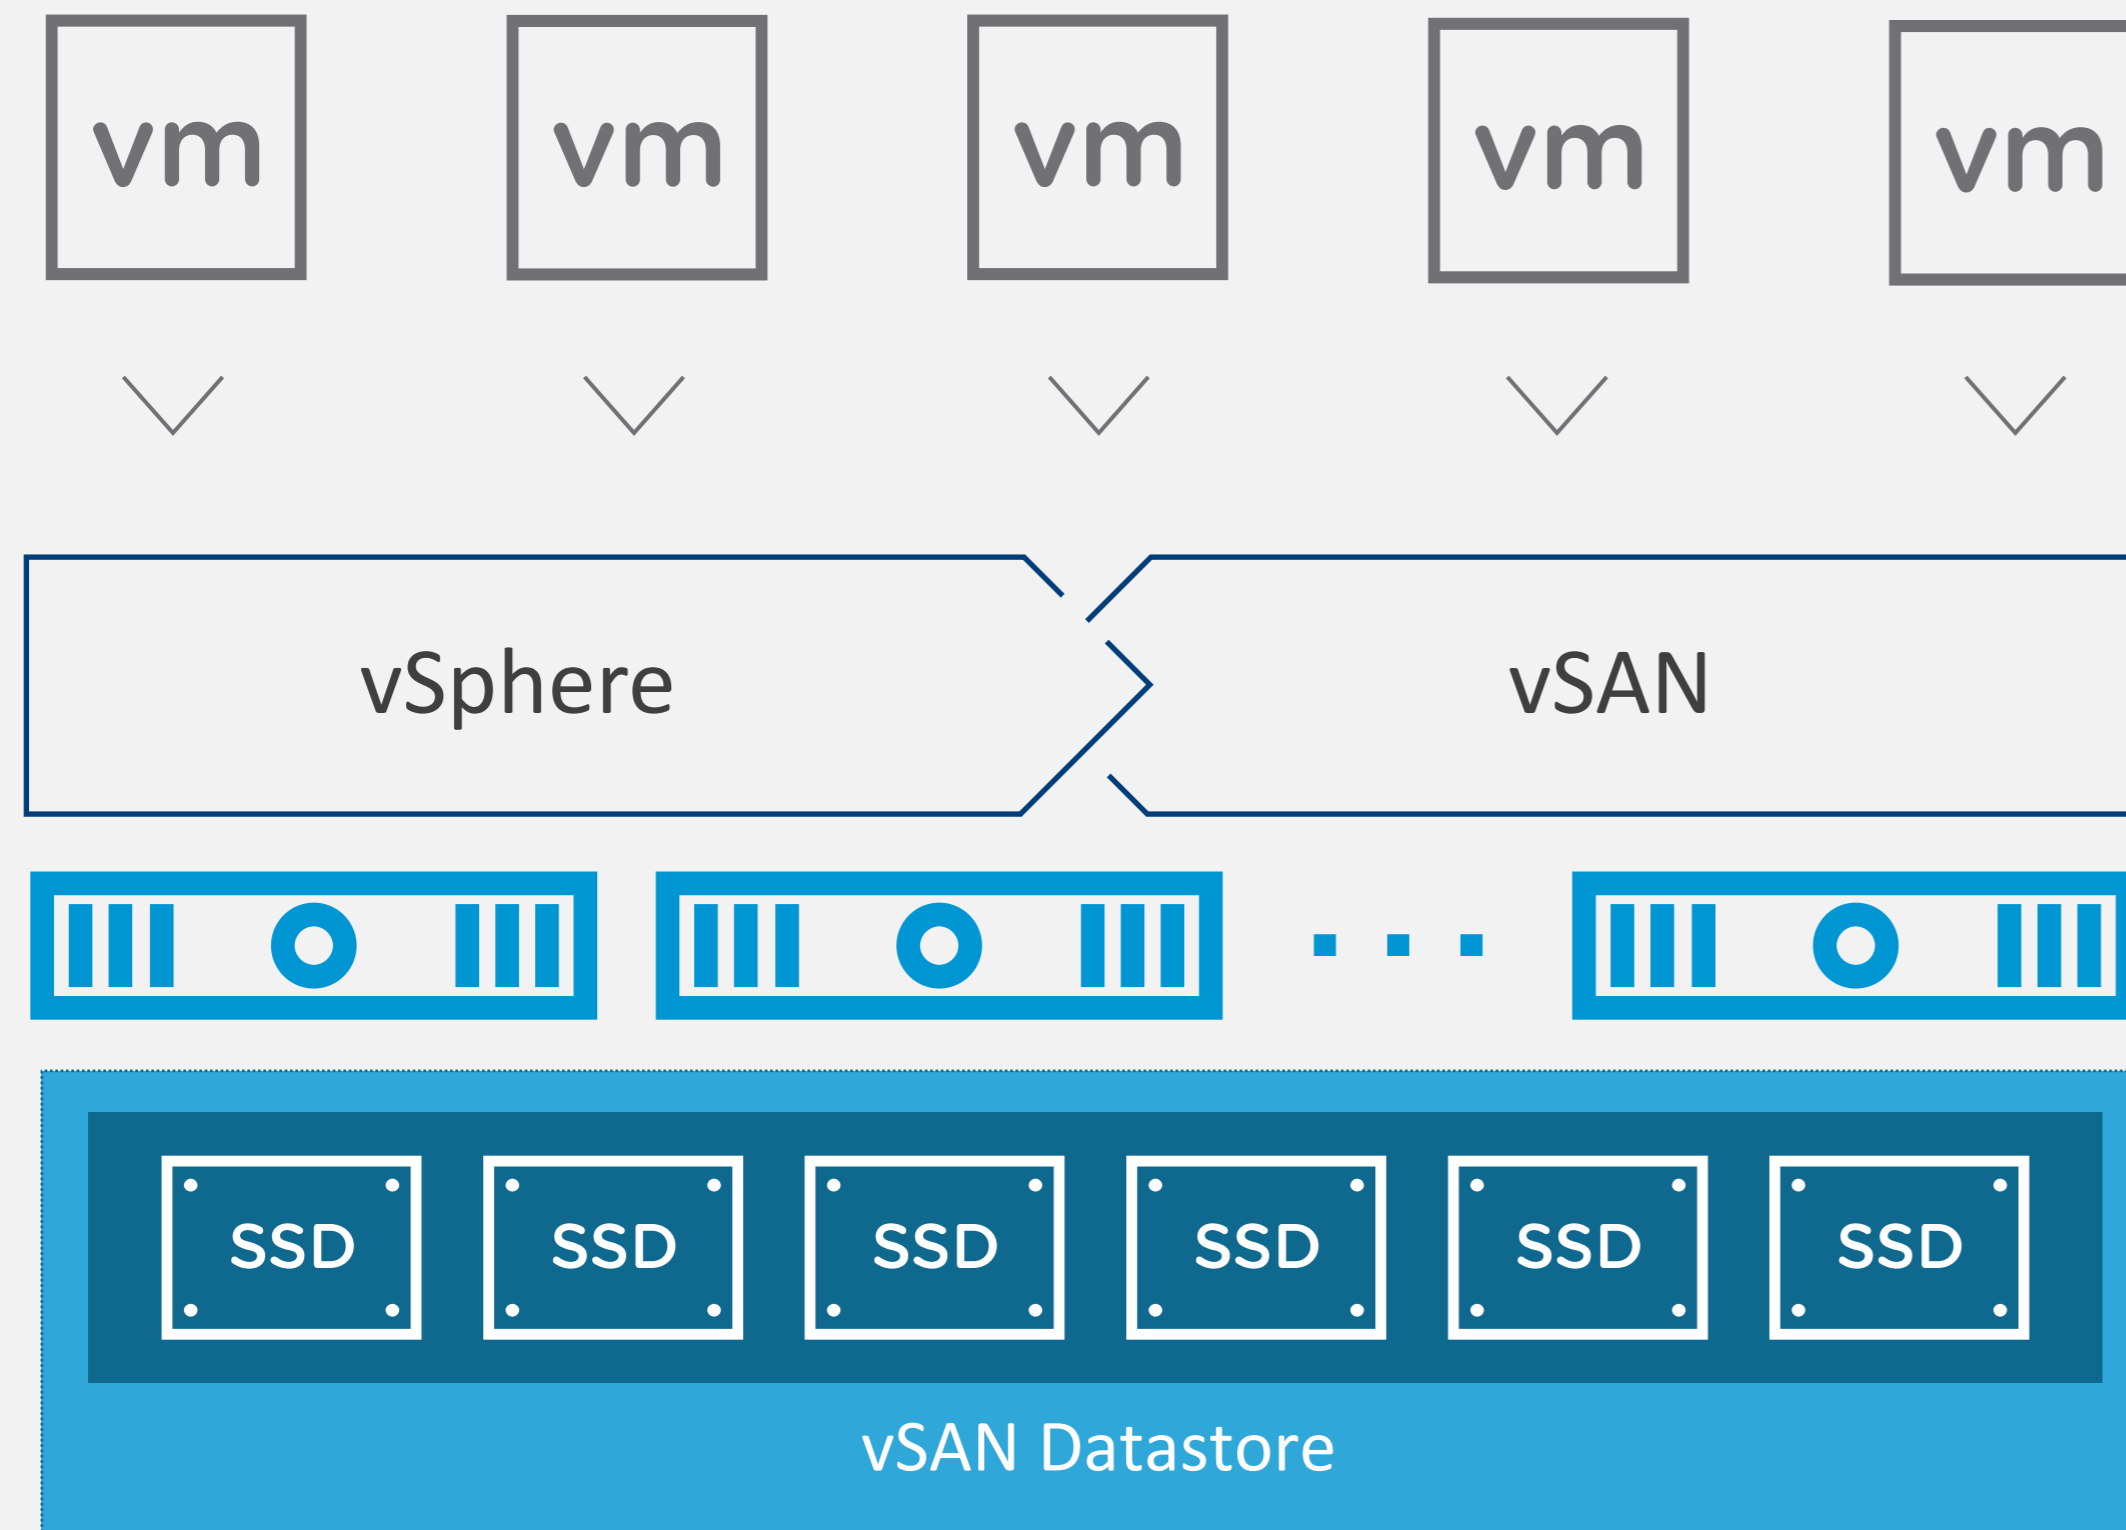


Software Defined Everything



Enterprise Storage in a Native vSphere Architecture

VMware vSAN



Runs on any standard x86 server

Integrated into hypervisor

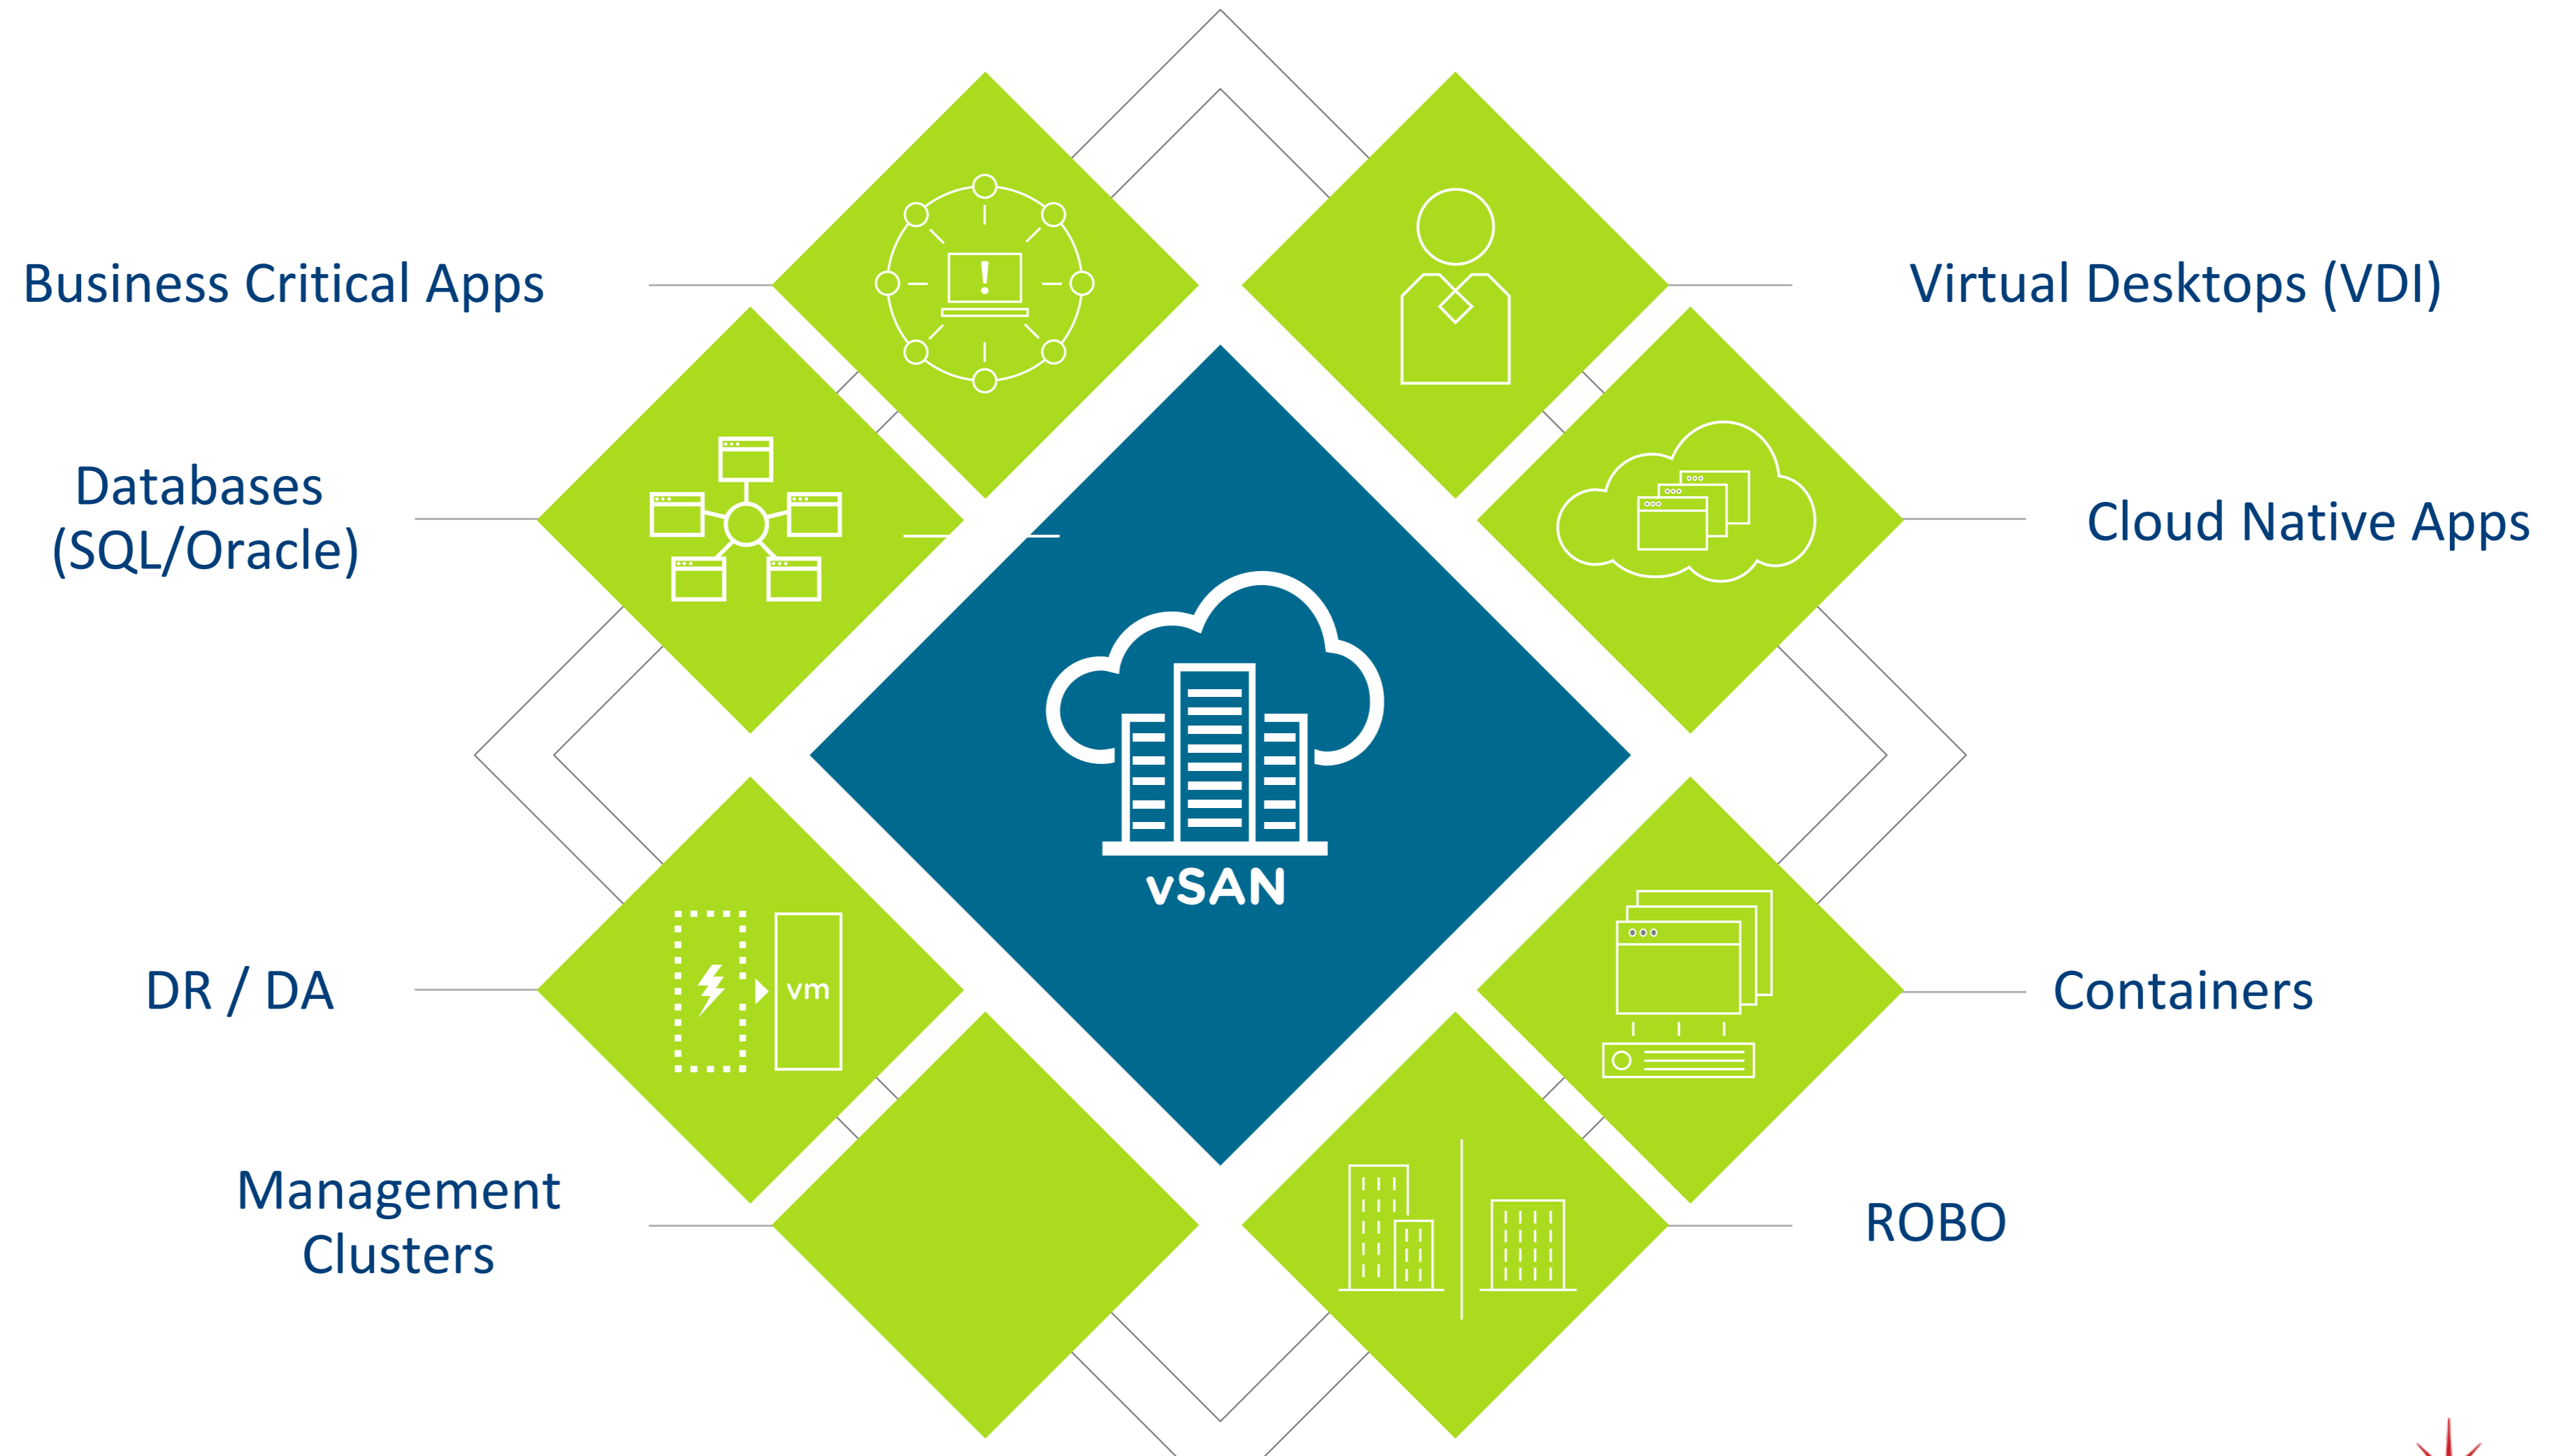
Pools HDD/SDD into single **cluster-wide shared datastore**

Easily scalable

Managed through VM storage policies


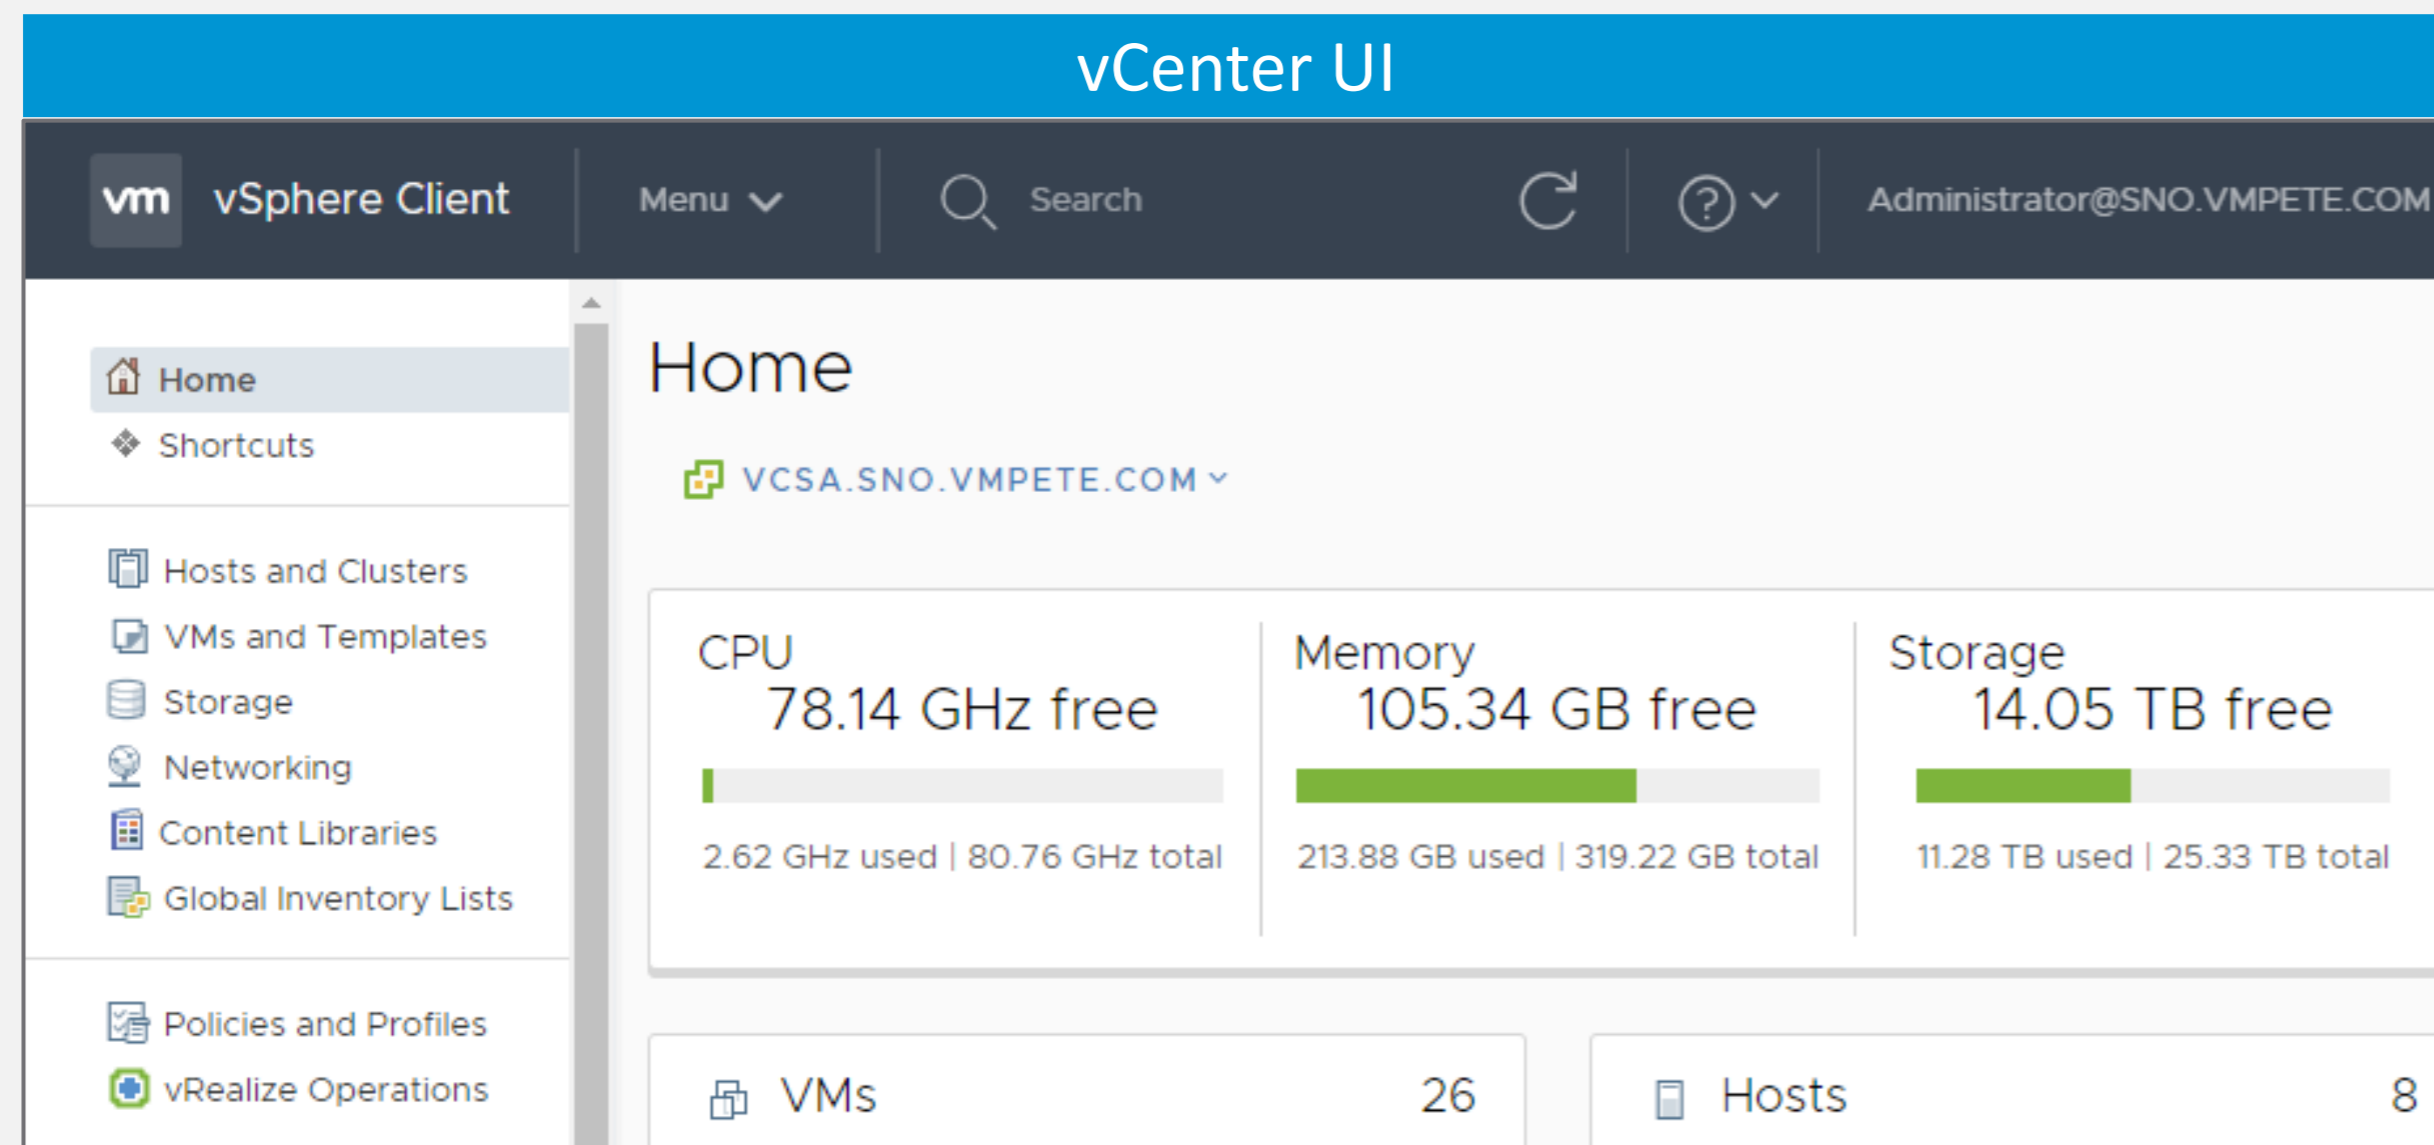
Supporting a Broad Variety of Use Cases

A platform capable of running traditional and next generation applications




One Management Plane – A Common Framework


Compute and storage management in one spot, using software they already know



On Premises Cluster



Remote Site Cluster



VMware Cloud on AWS Cluster

Customers enjoy **familiarity**

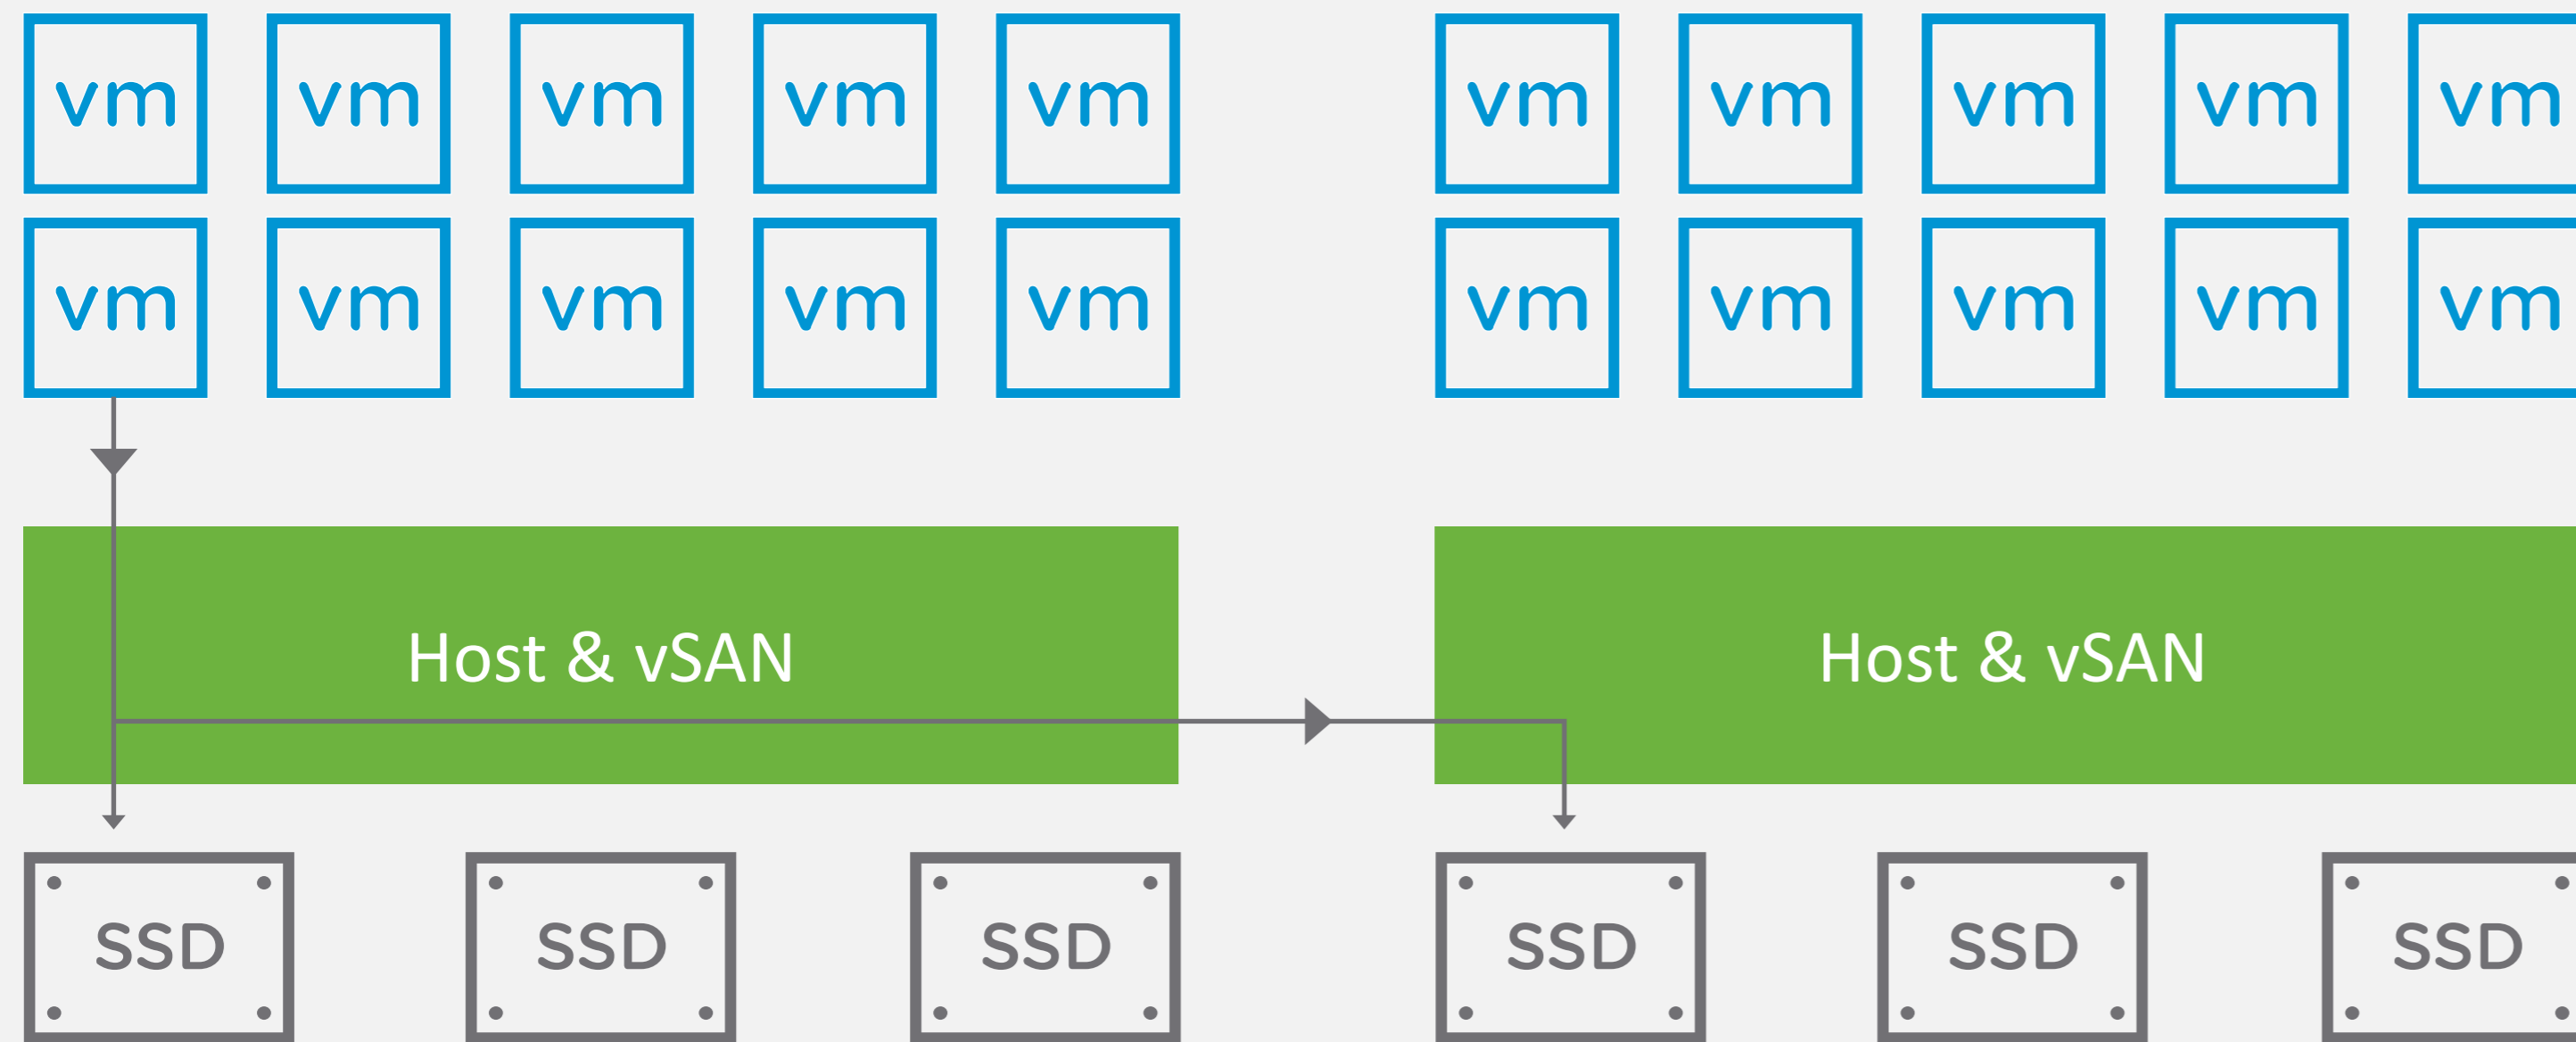
vSAN/vSphere allows them to:

- Maintain familiarity
- Evolve without risk
- Reduce TCO
- Scale to tomorrow

Introduces new functionality into **software they already use**

The vSAN Difference

Hypervisor integration for supreme levels of efficiency, and integration



Committing a Write I/O in vSAN



vSAN **embedded** into vSphere

Simplified, most efficient I/O path

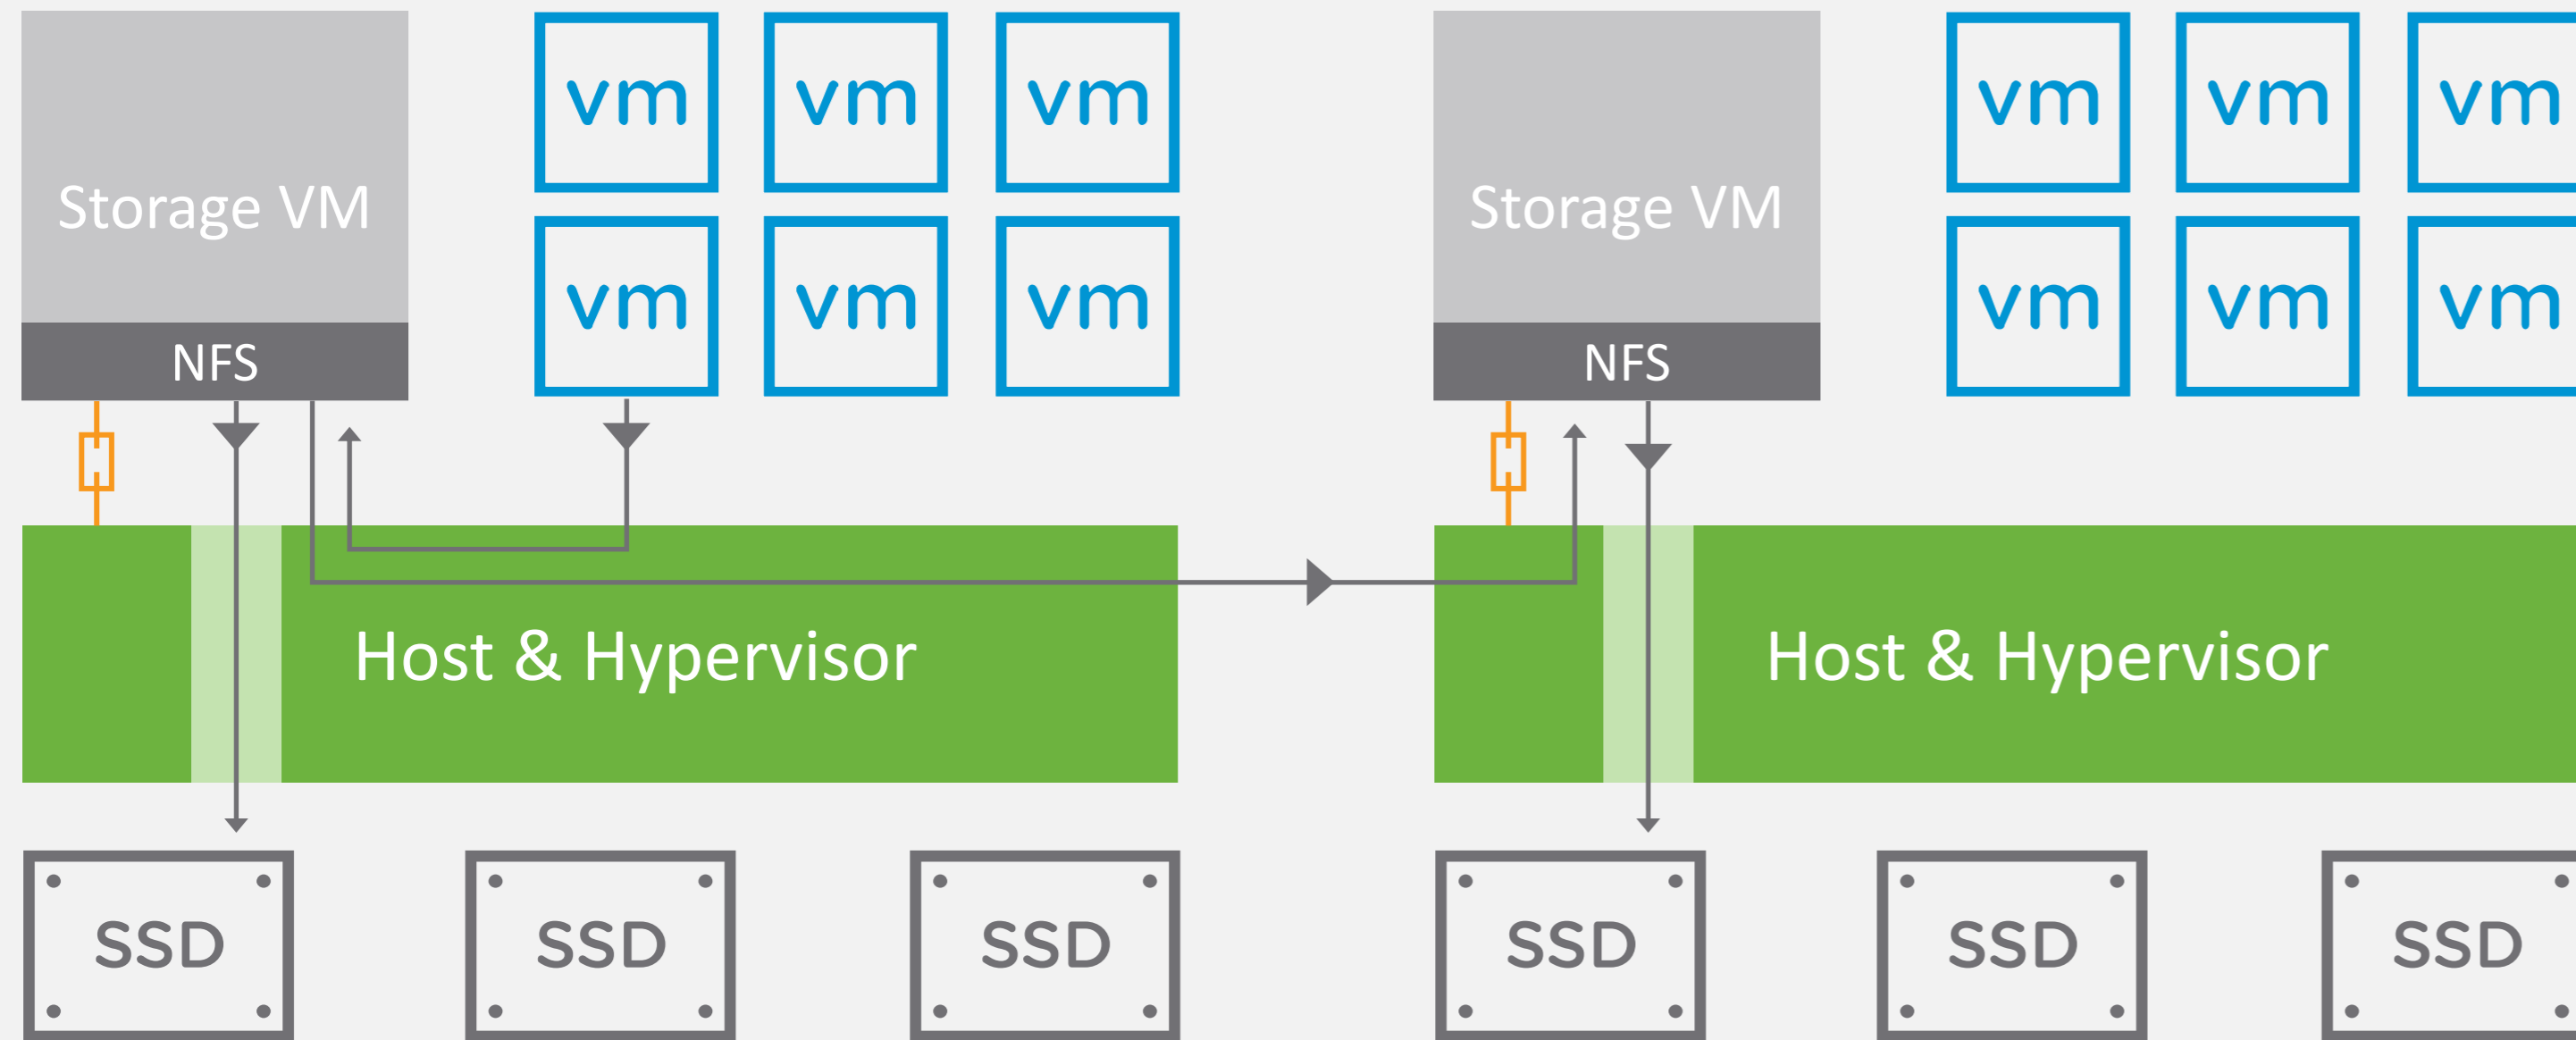
Minimal Host CPU and I/O overhead

More VMs per host, with more consistent performance

Awareness of hypervisor activities

The vSAN Difference

Other HCI solutions using non-integrated VMs for storage processing



Committing a Write I/O in other HCI solutions

 Compute Utilization

Resource-intensive storage controller VMs on every host

More hops, context switching, queues, and locks

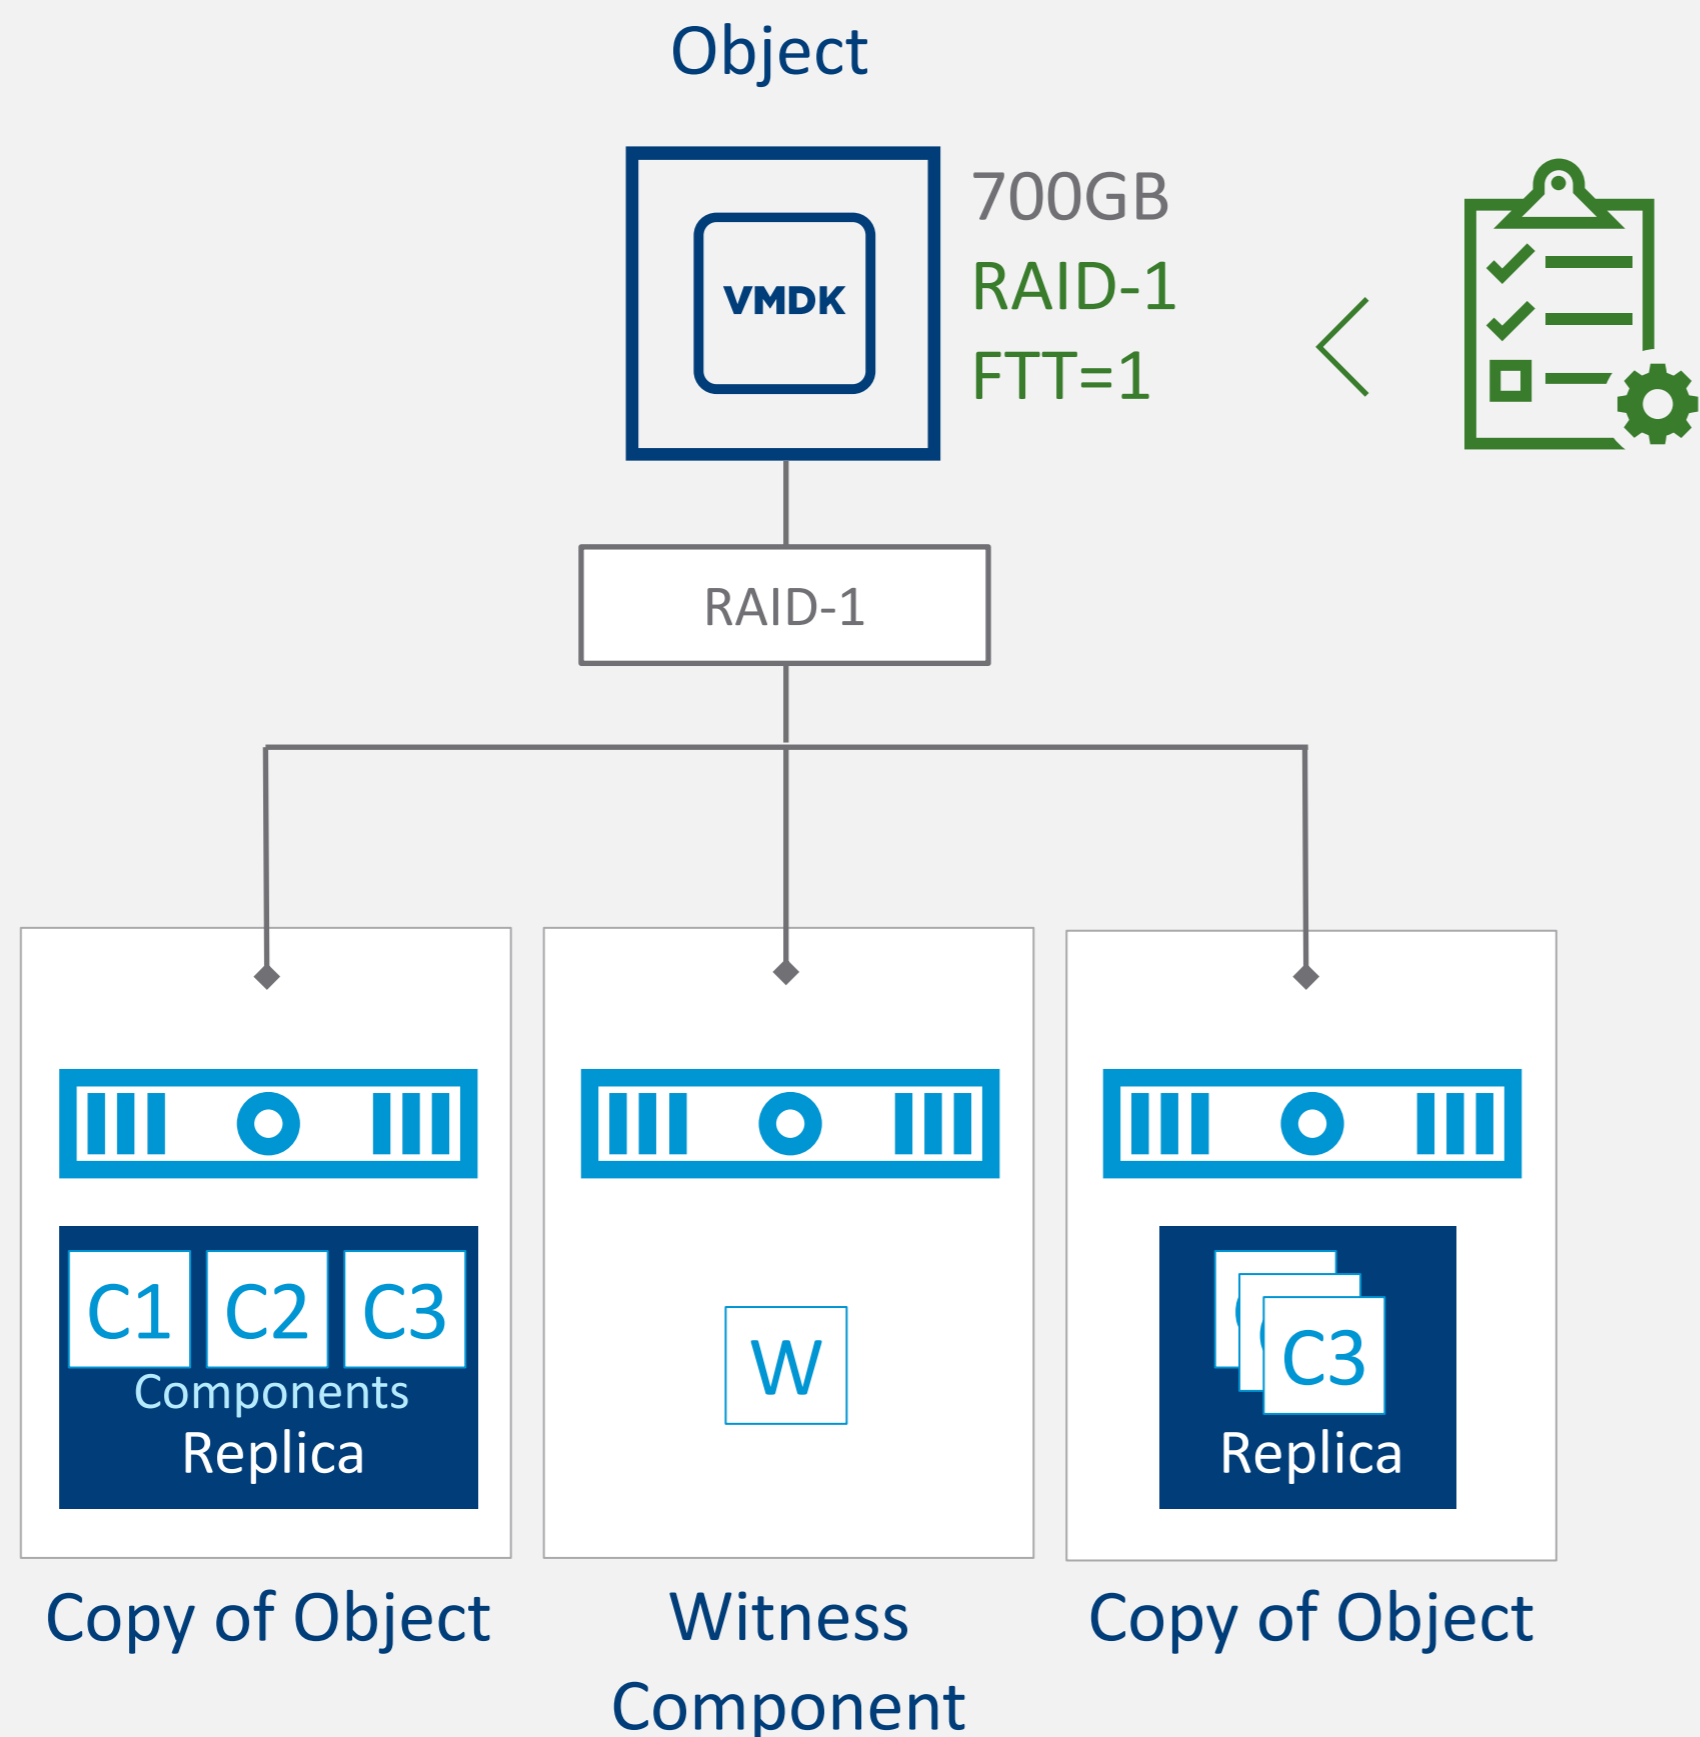
Host CPU and I/O **amplification**

Fewer VMs per host, with less consistent performance

Unaware of hypervisor activities

Modern Object Based Storage for vSphere

vSAN objects and components



The vSAN datastore is an **object** store

Set availability and performance requirements per object

Each object made up of one or more **components**

Data (components) is distributed across cluster based on VM storage policy

Modern Object Based Storage for vSphere

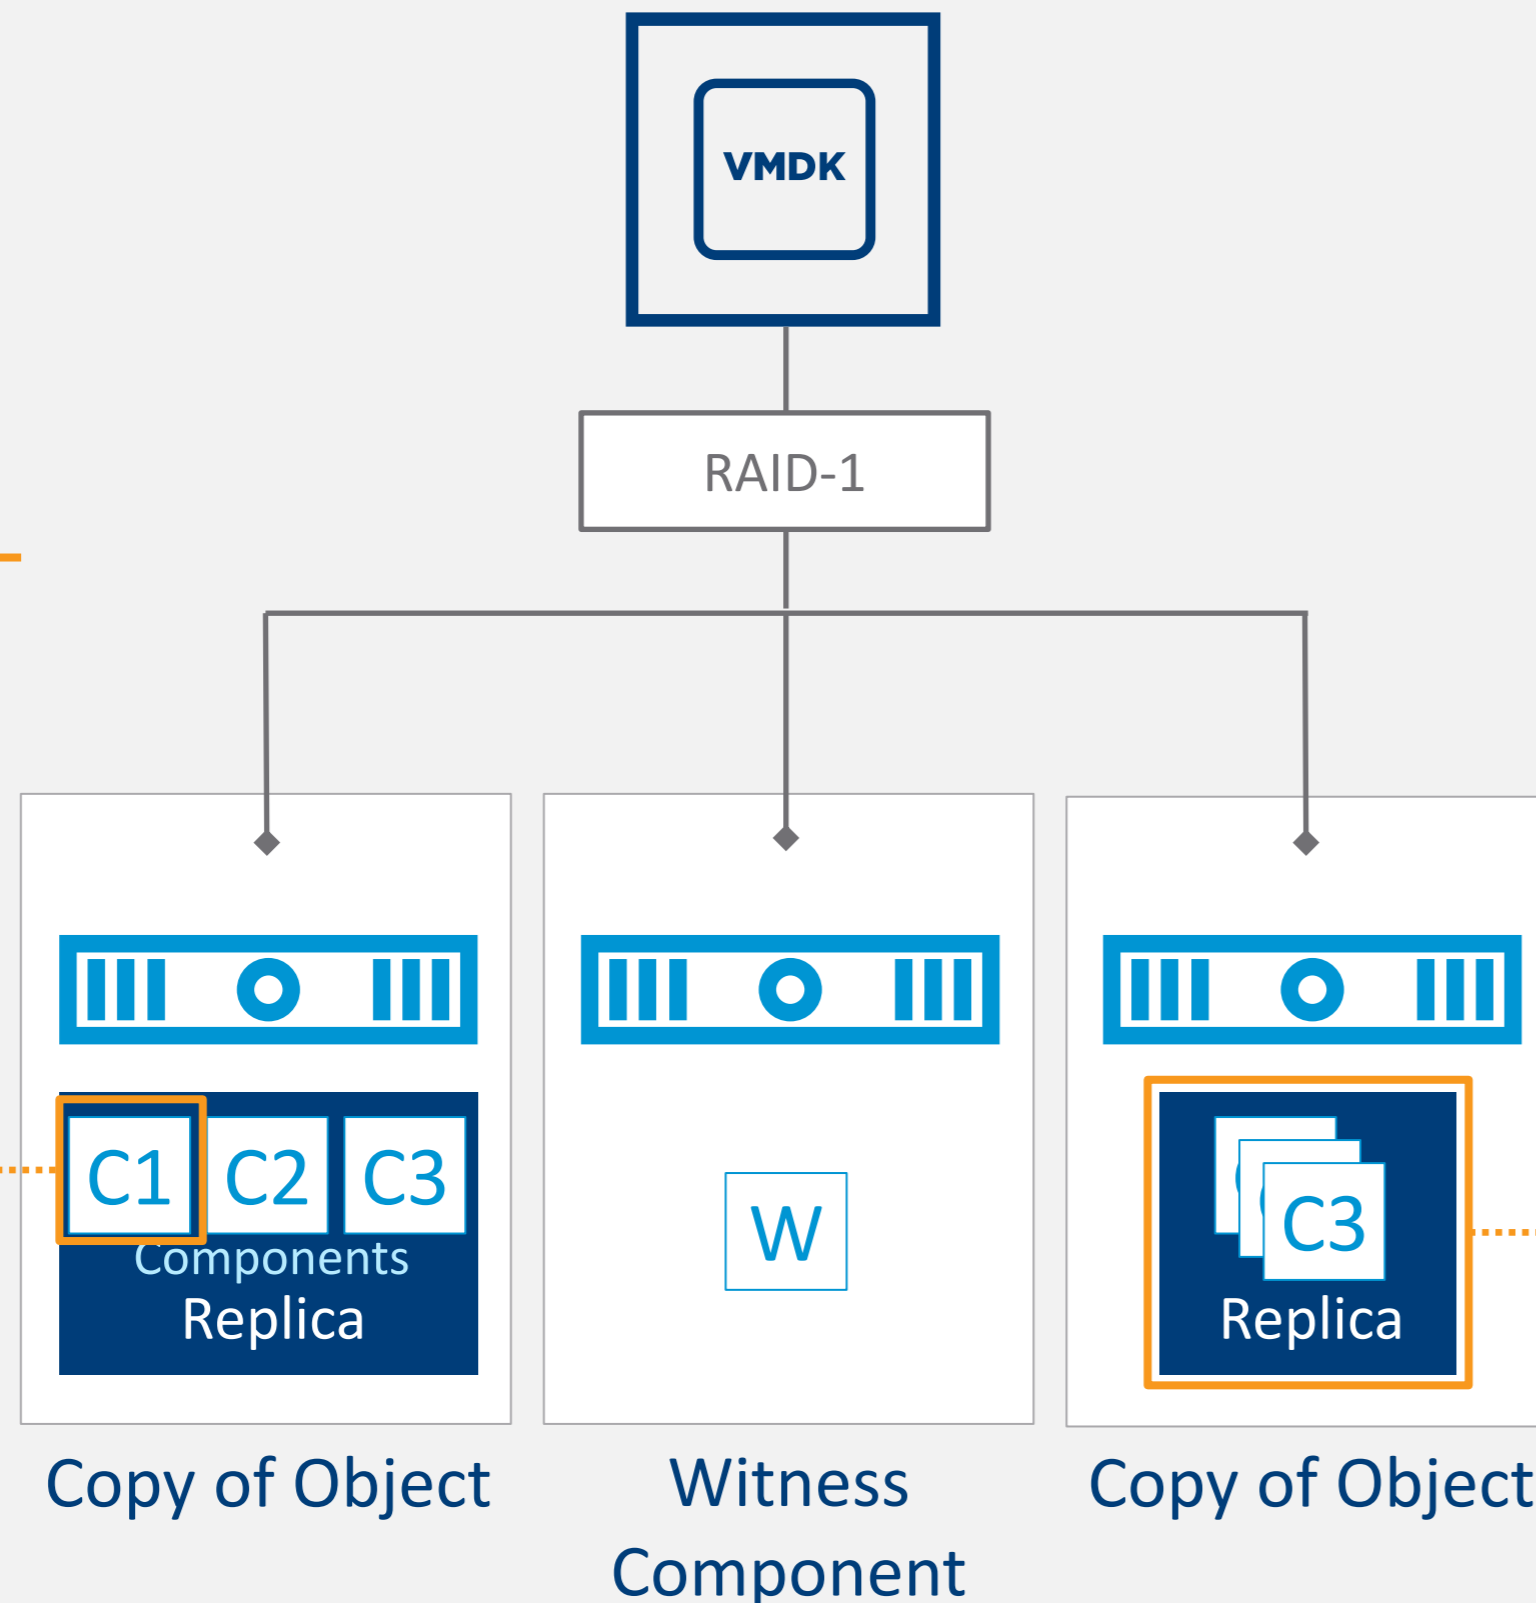
vSAN objects and components

Components


Max size: 255GB

May be split due to policy settings or environmental conditions

Object



Object Types

-  VMDK
-  VM Home namespace
-  VM swap
-  Snapshot delta
-  Snapshot

RAID tree consist of a **leaf** that makes up given object

Components **dispersed** across hosts in a cluster

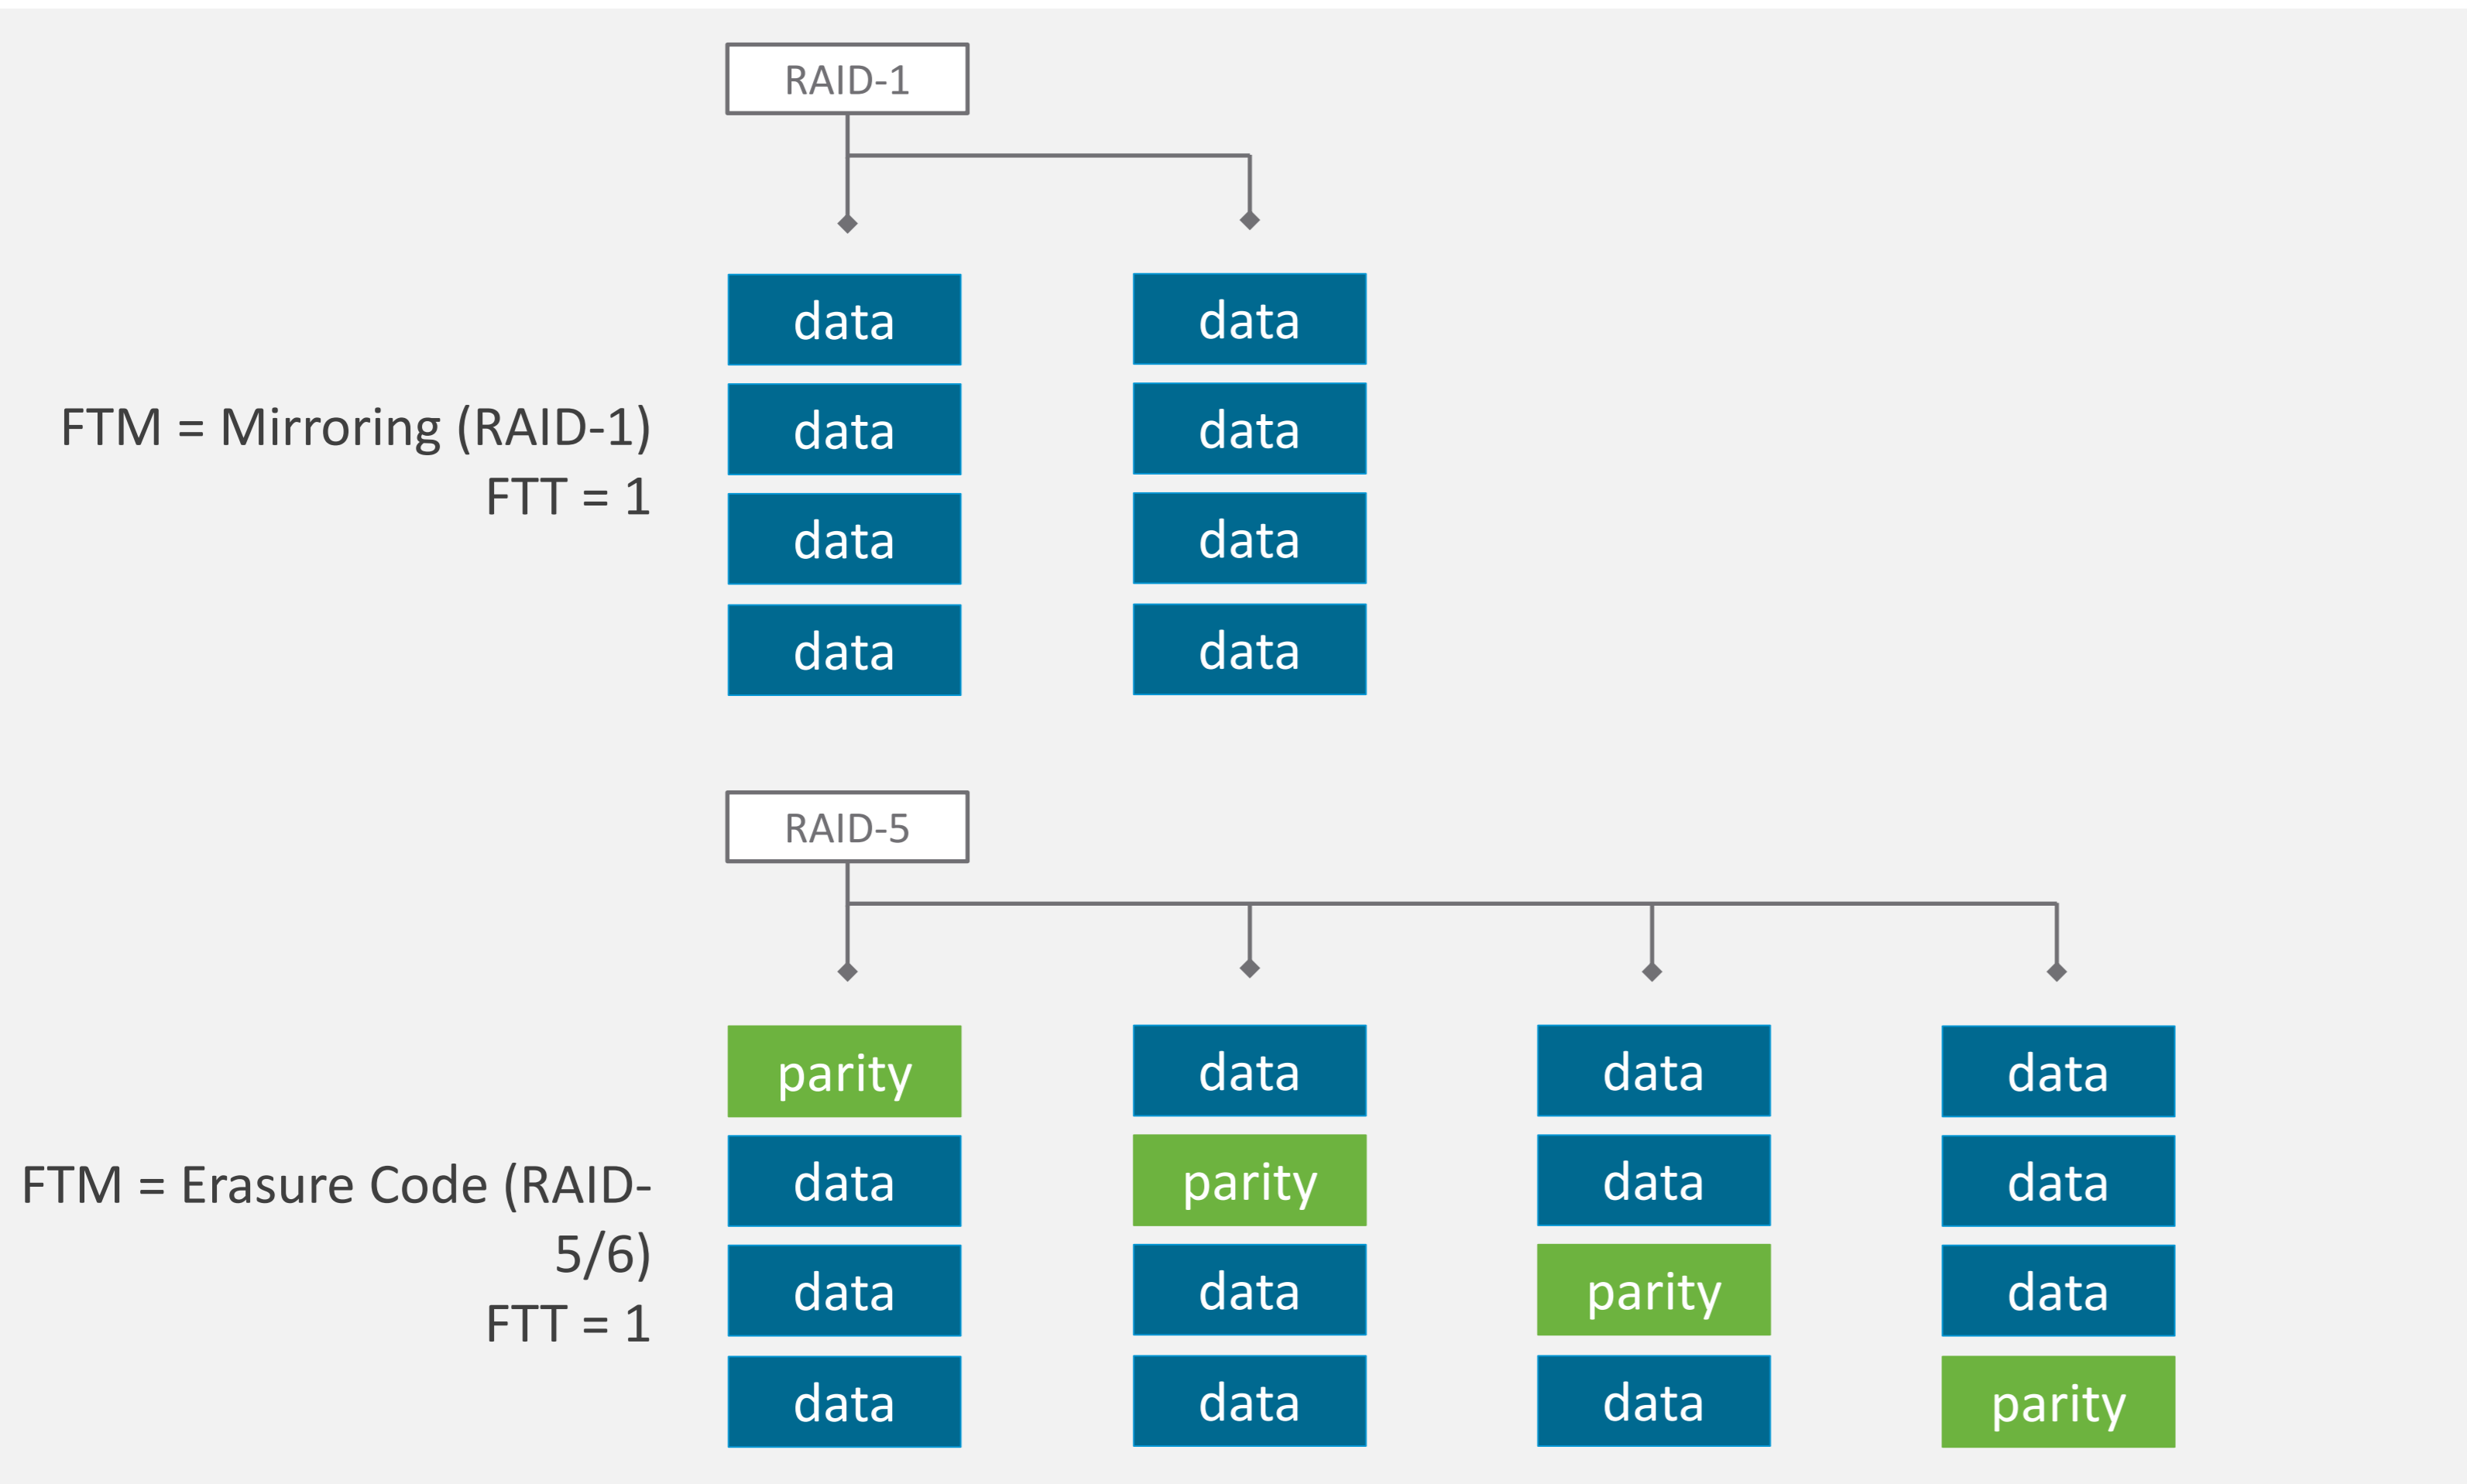
vSAN determines **placement** of object components

Adheres to **assigned policy** of object

“Witness” components used to determine **quorum**

Option with Data Protection Levels and Schemes

Basic concepts around “failure tolerance method” and level of “failures to tolerate”



Failure Tolerance Method (**FTM**) defines **data layout approach**

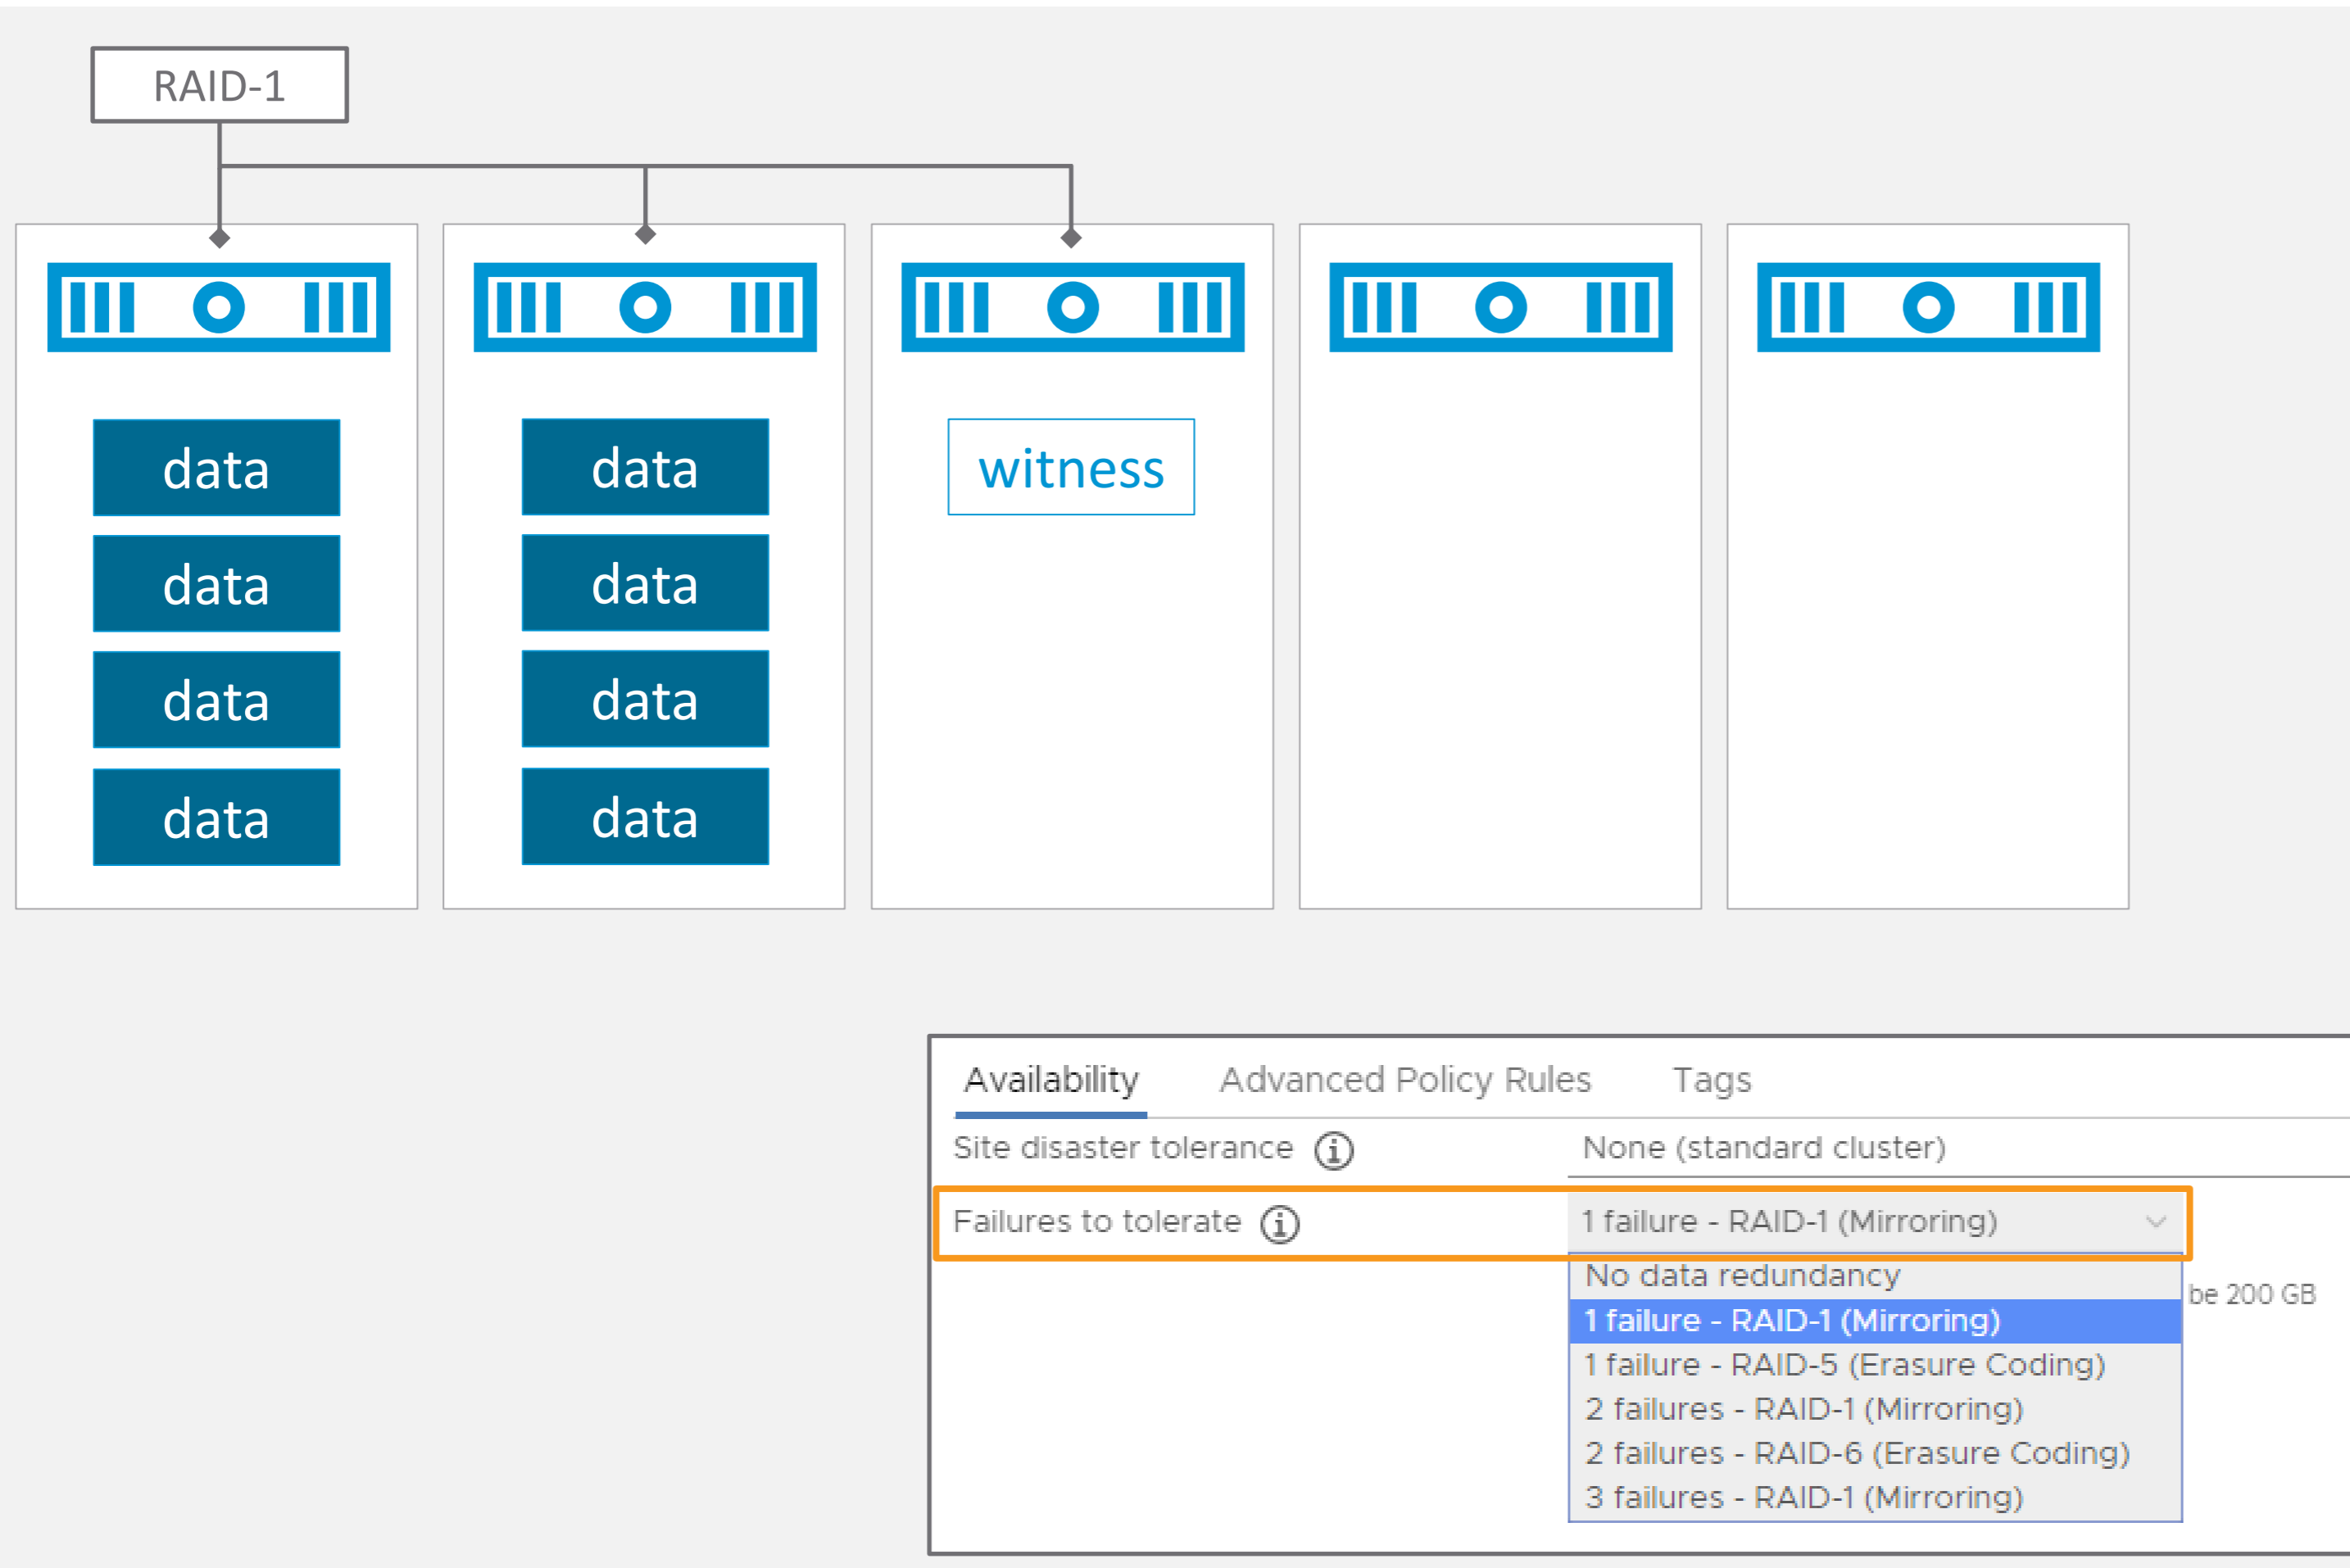
- Mirroring (RAID-1)
- Erasure Coding (RAID-5/6)

Level of Failure to Tolerate (**FTT**) defines **level of resilience**

- FTT=1 – accessible with one failure
- FTT=2 – accessible with two failures
- FTT=3 – accessible with three failures

Object Based Storage – A Better Way for Data Protection

Setting failures to tolerate (FTT) to 1 with RAID-1 mirroring



Alternative FTM to RAID-5/6 erasure coding

Data **mirrored** to another host

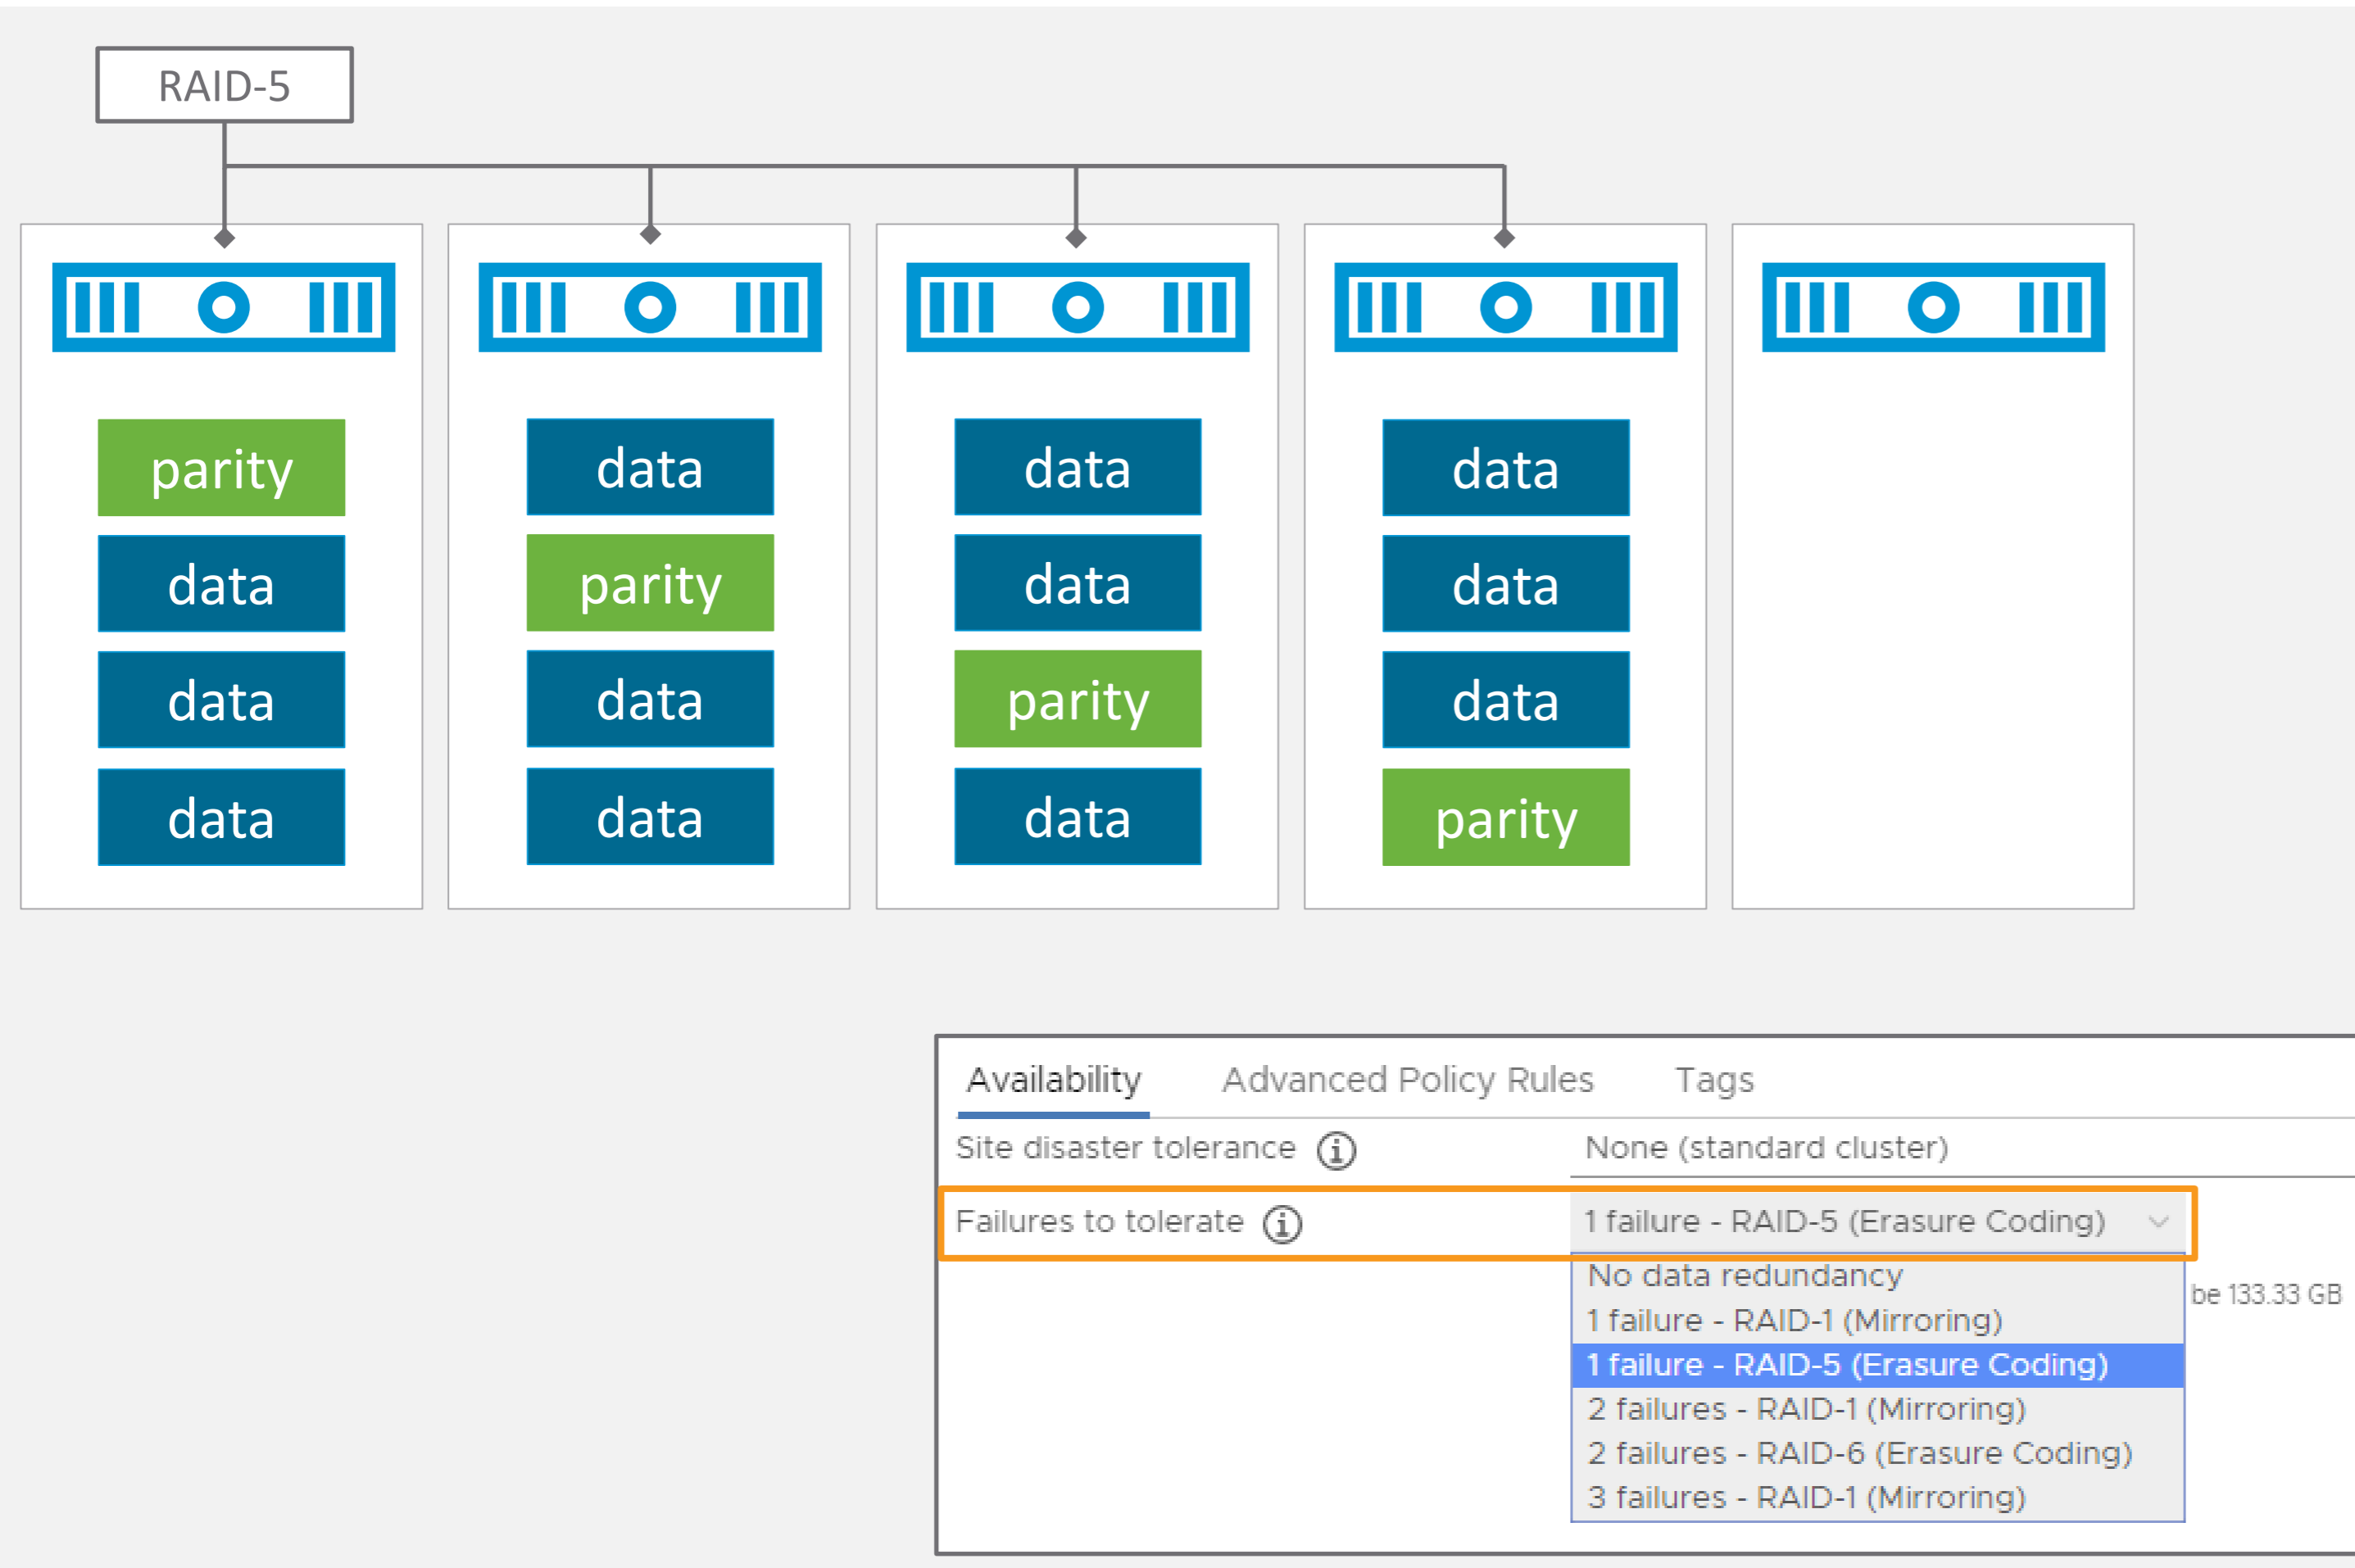
Witness needed to determine quorum

Requires **fewer hosts** but not as space efficient as RAID-5/6

Additional hosts needed to support greater than FTT 1 or maintenance operations

Object Based Storage – A Better Way for Data Protection

Setting failures to tolerate (FTT) to 1 with RAID-5 erasure coding



Alternative to RAID-1 Mirroring

Data with parity **striped** across hosts

For erasure coding, **FTT 1 implies RAID-5**, FTT 2 implies RAID-6

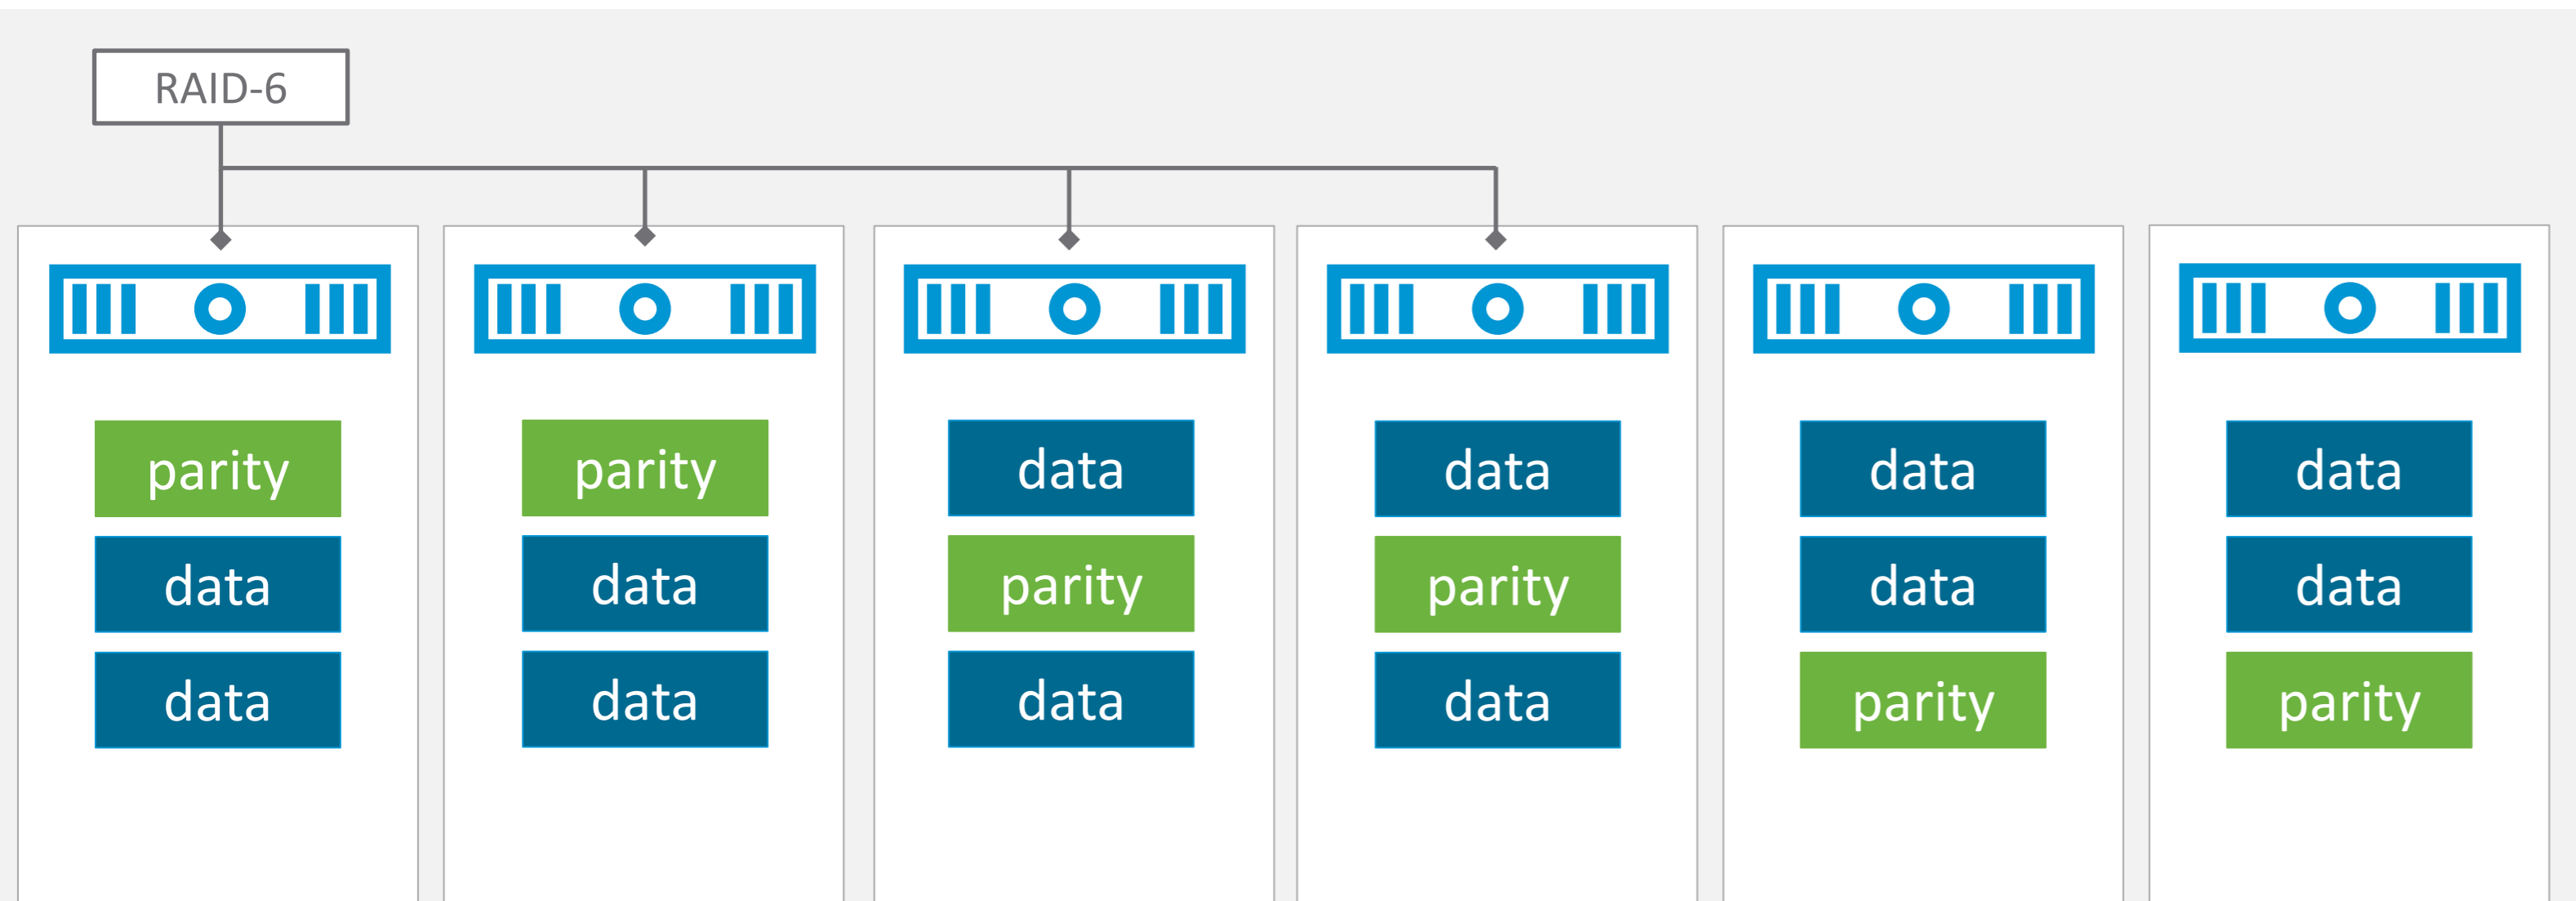
Guaranteed **space reduction**

- 30% savings with RAID-5
- 50% savings with RAID-6

Additional hosts needed to support greater than FTT 1 or maintenance operations

Object Based Storage – A Better Way for Data Protection

Setting failures to tolerate (FTT) to 2 with RAID-6 erasure coding



Availability	Advanced Policy Rules	Tags
Site disaster tolerance ⓘ		None (standard cluster)
Failures to tolerate ⓘ		2 failures - RAID-6 (Erasure Coding) ▾
		No data redundancy
		1 failure - RAID-1 (Mirroring)
		1 failure - RAID-5 (Erasure Coding)
		2 failures - RAID-1 (Mirroring)
		2 failures - RAID-6 (Erasure Coding)
		3 failures - RAID-1 (Mirroring)

2 hosts can fail without data loss
(Implied FTT=2)

6 hosts **minimum**

7+ hosts desired to recover
resiliency level upon failure

50% reduction in **overhead**
compared to mirroring.

- 20GB disk consumes 60GB with RAID-1, FTT=2 (3x)
- 20GB disk consumes 30GB with RAID-6, FTT=2 (1.5x)

vSAN Cluster Host Count Matters

Levels of resilience depend on quantity of hosts within a vSAN cluster



RAID-1
FTT=1



RAID-1
FTT=2



RAID-1
FTT=3



RAID-5
(FTT=1)



RAID-6
(FTT=2)



Data services and functionality **dependent** on cluster **host count**

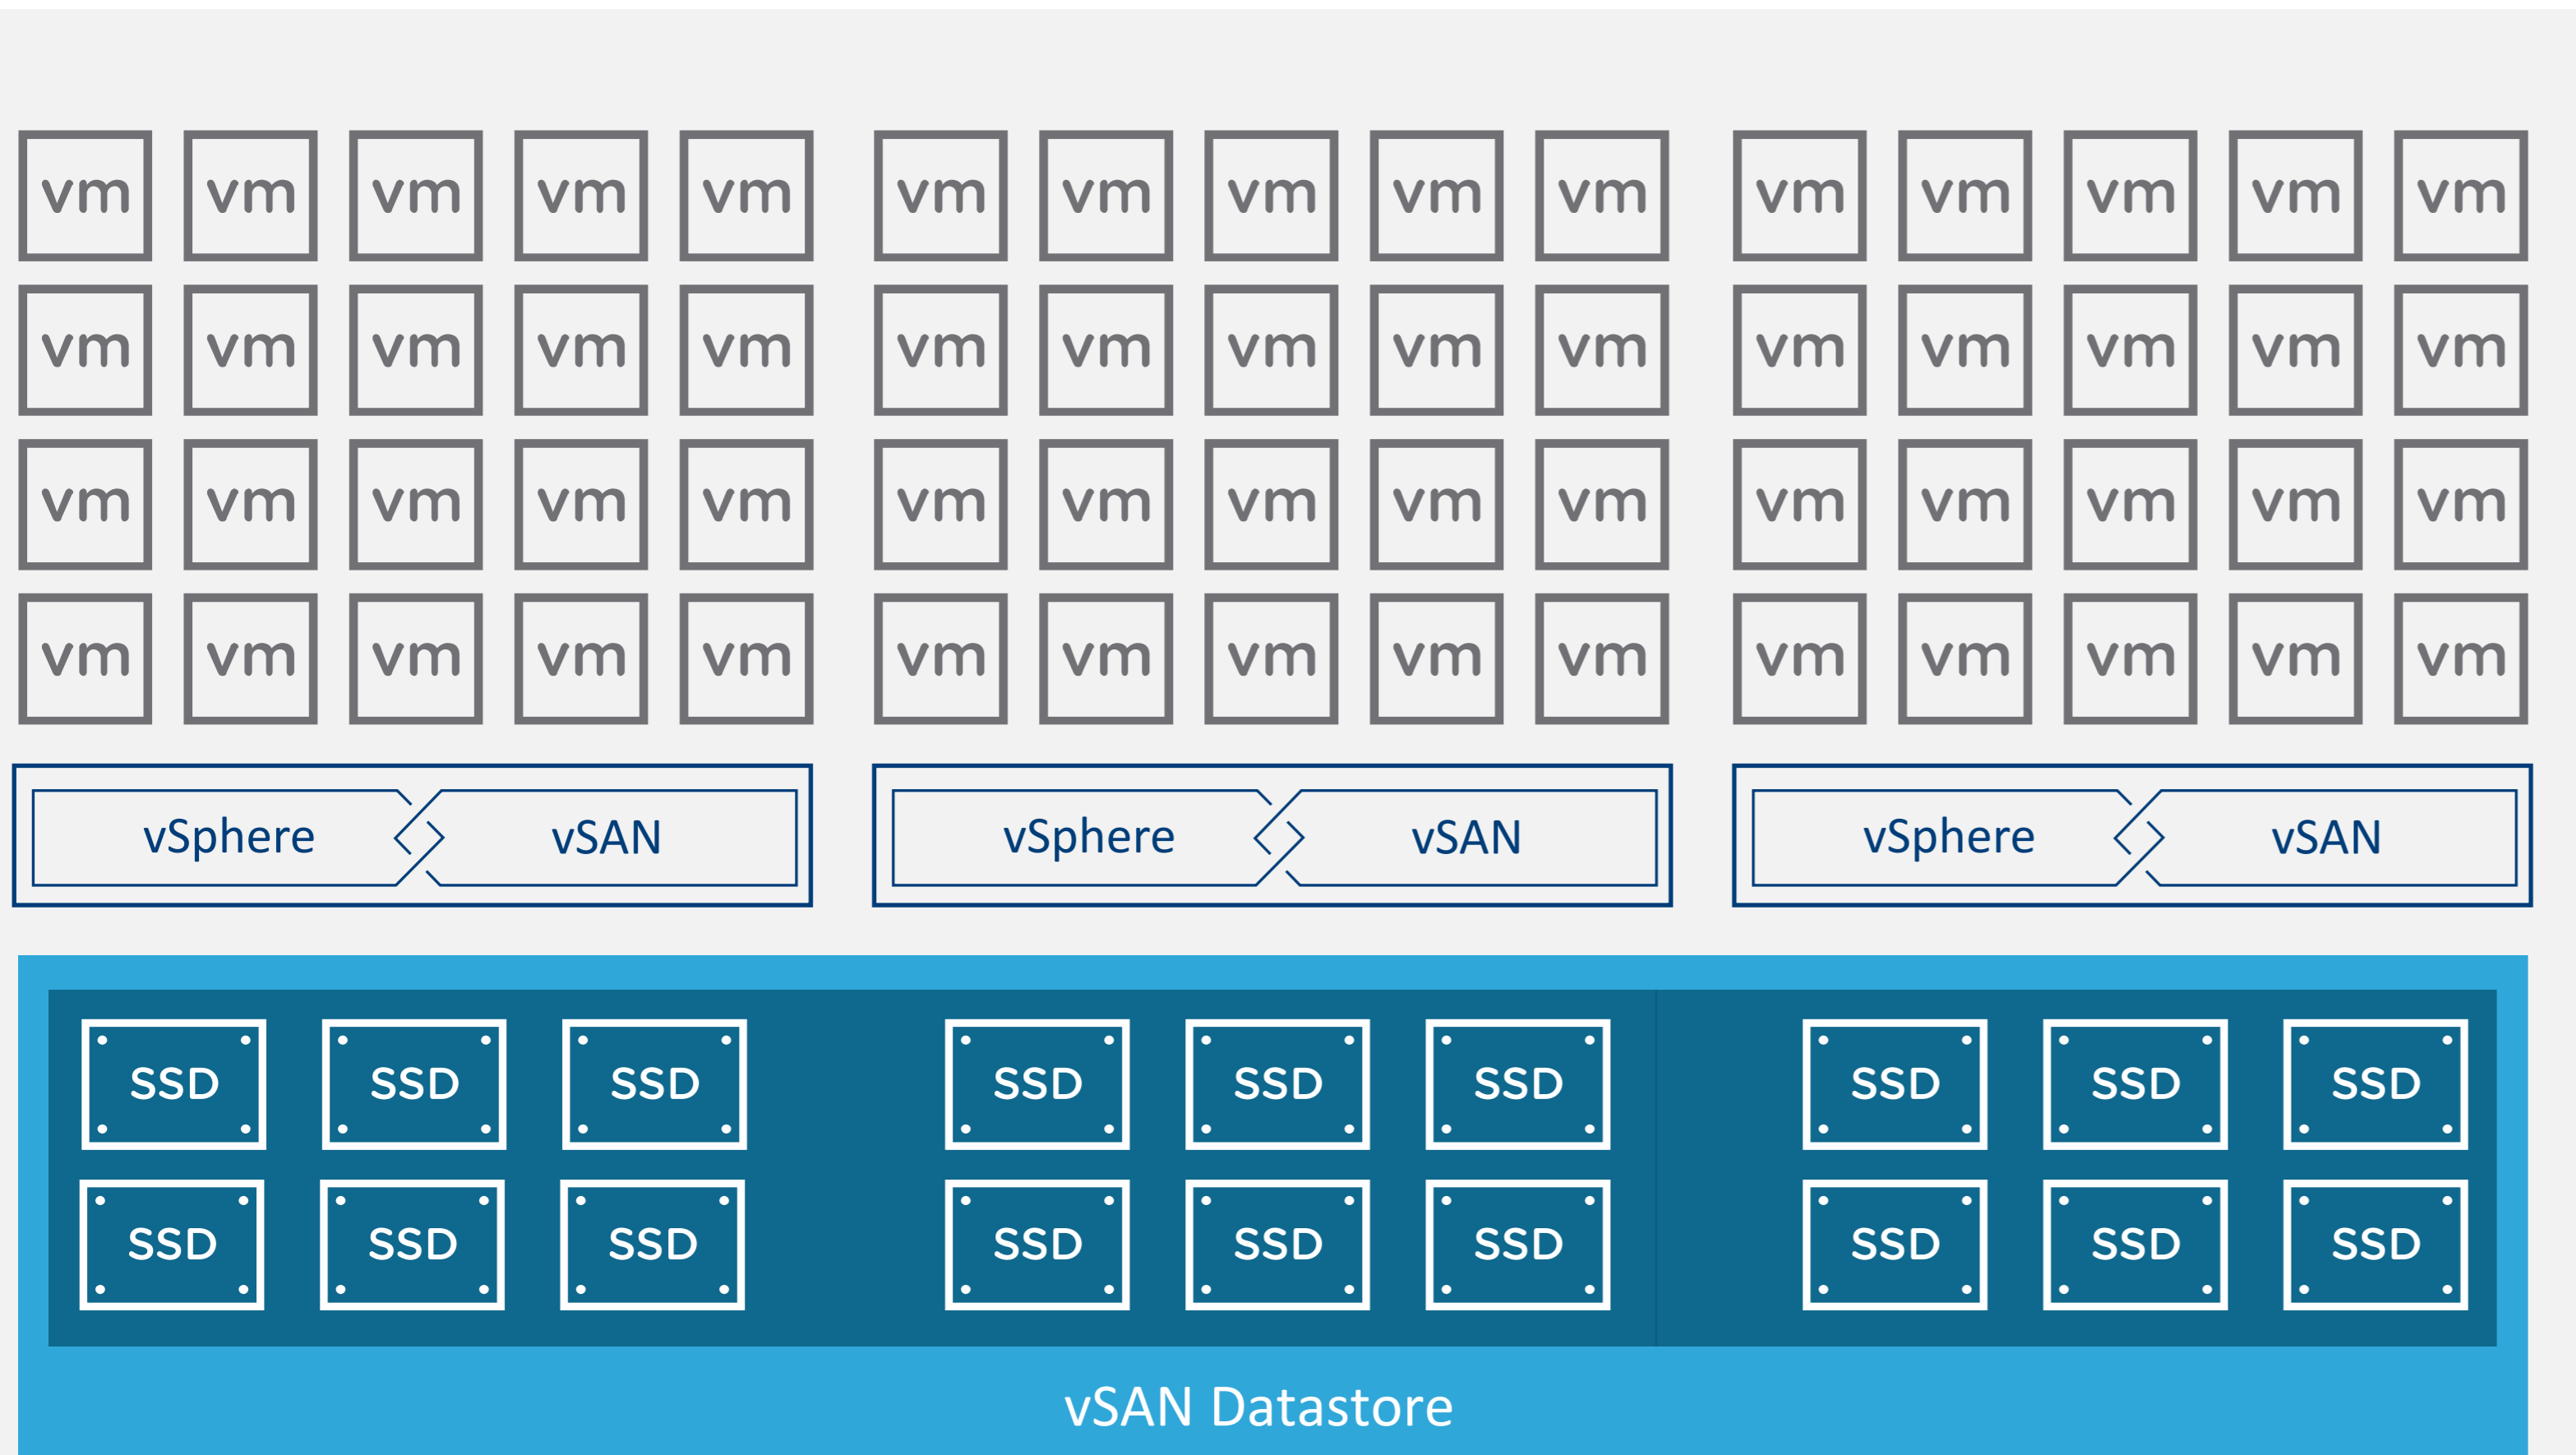
Ensure cluster is sized for N+1 or N+2 in **compute and storage capacity**

Associated penalties in performance and capacity with each option

Remind the customer this is a per **VM** or per **VMDK** setting, thanks to **SPBM**

The vSAN Difference

Scale UP and OUT for maximum agility



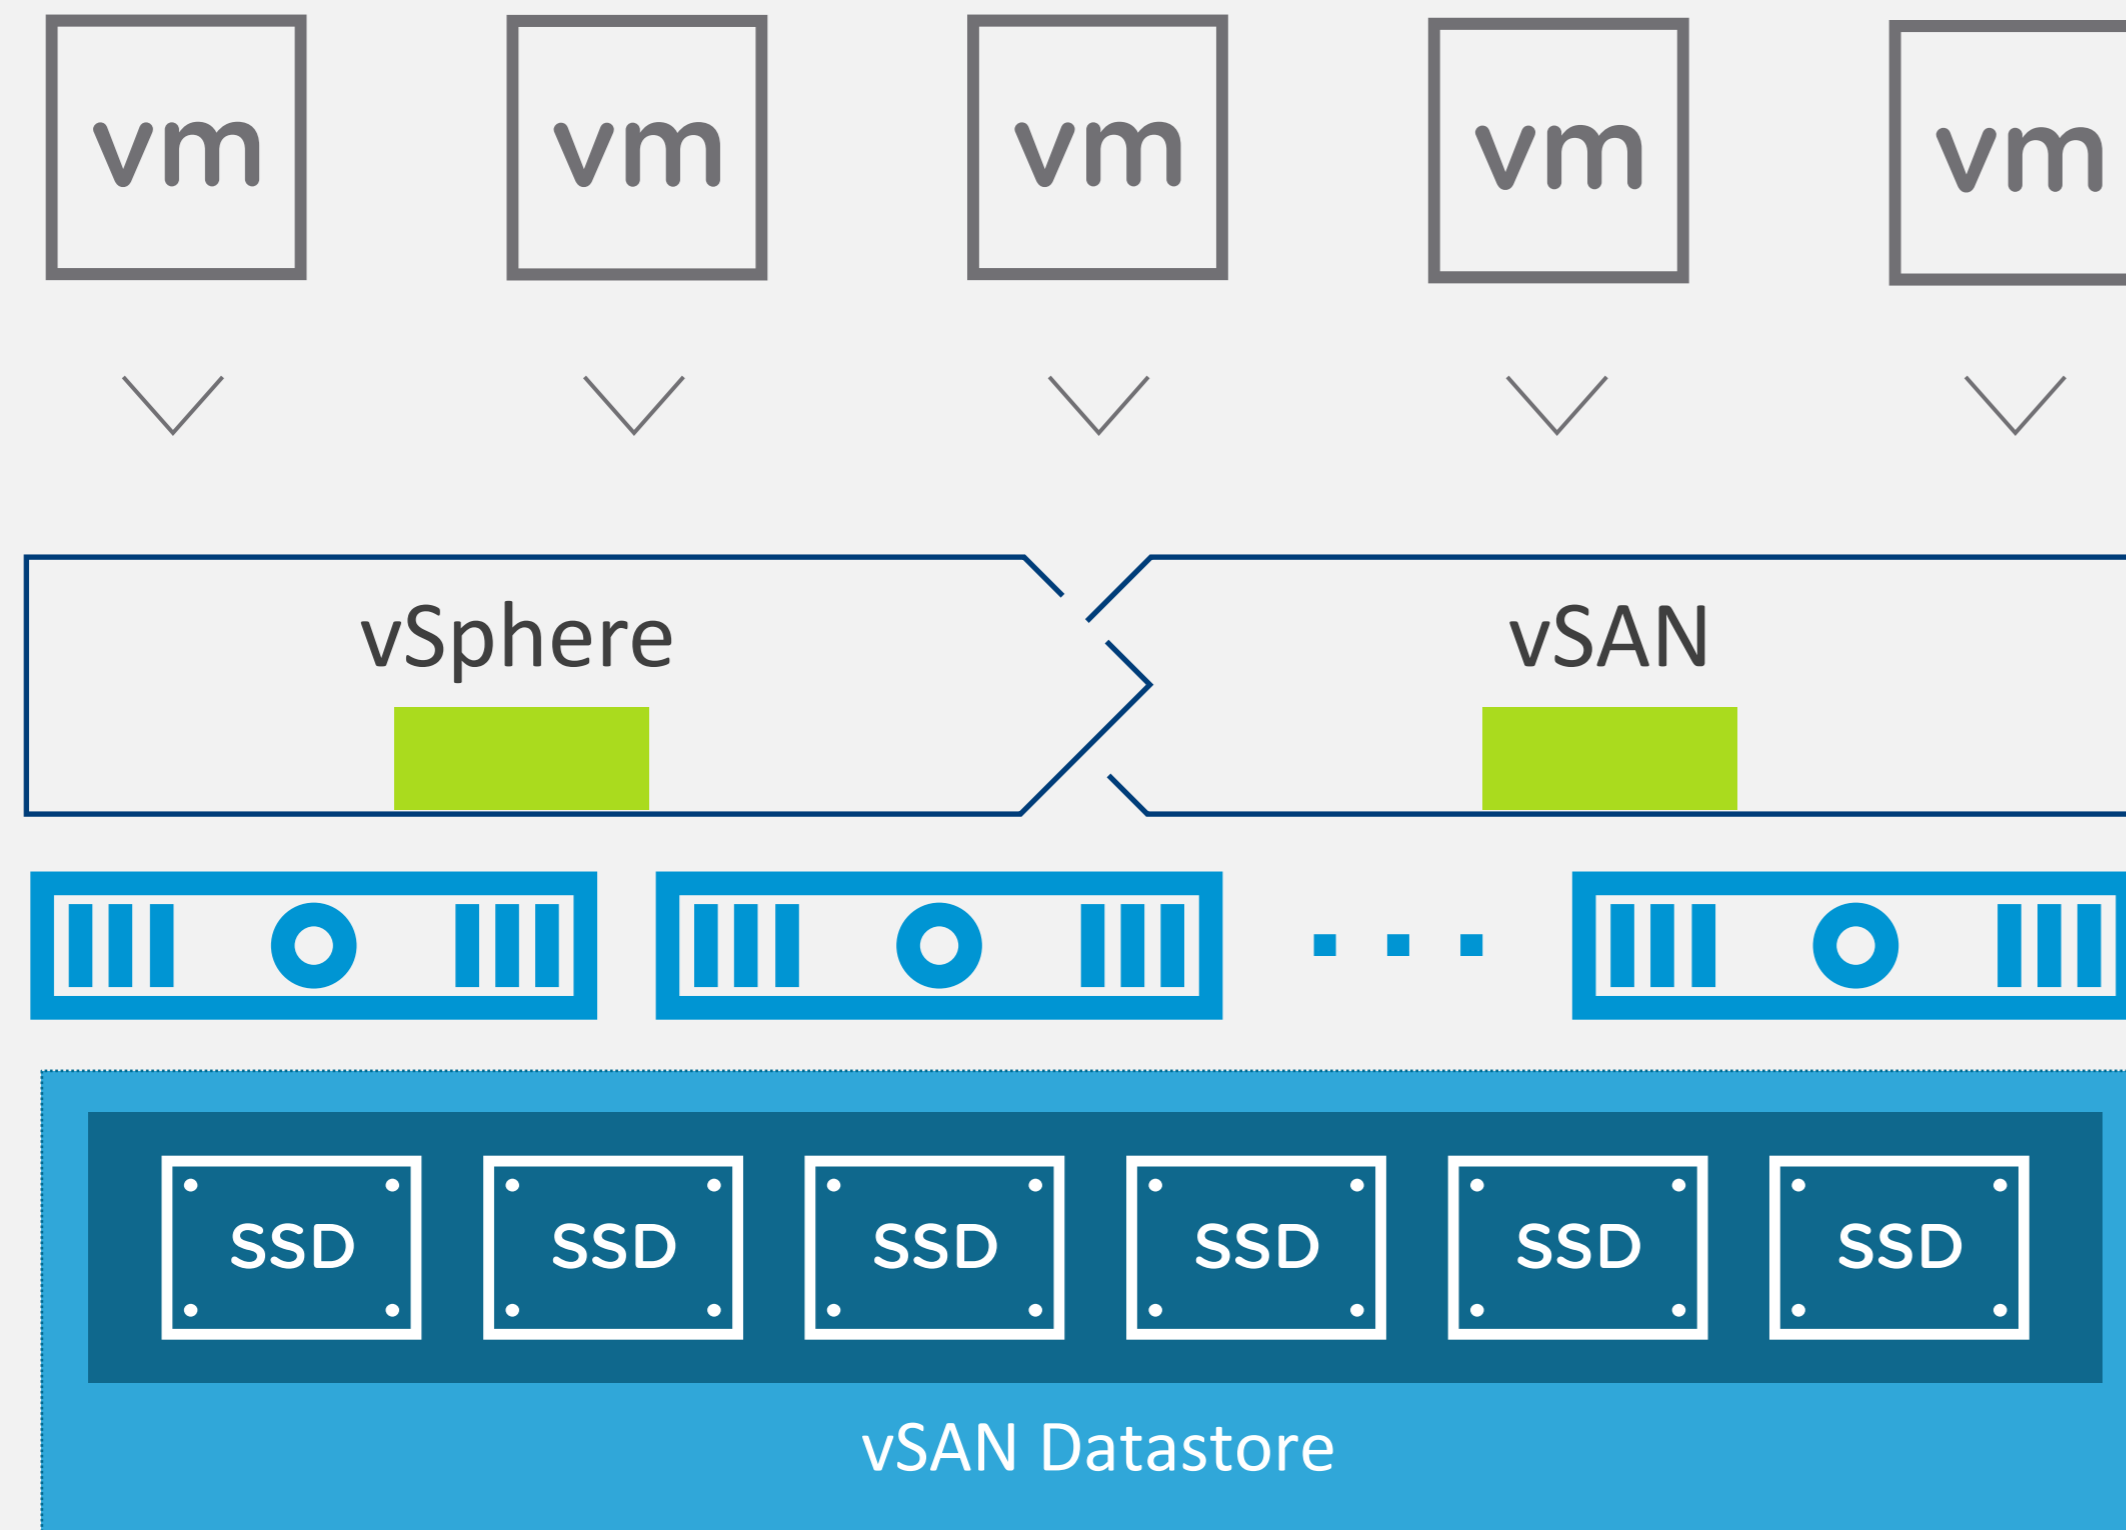
Add capacity the way you want

Scale **UP** by adding drives

Scale **OUT** by adding hosts

Meet Security Requirements with a vSAN-powered Environment

vSphere and vSAN FIPS 140-2 validation



Improved Security in vSphere with FIPS 140-2 validation

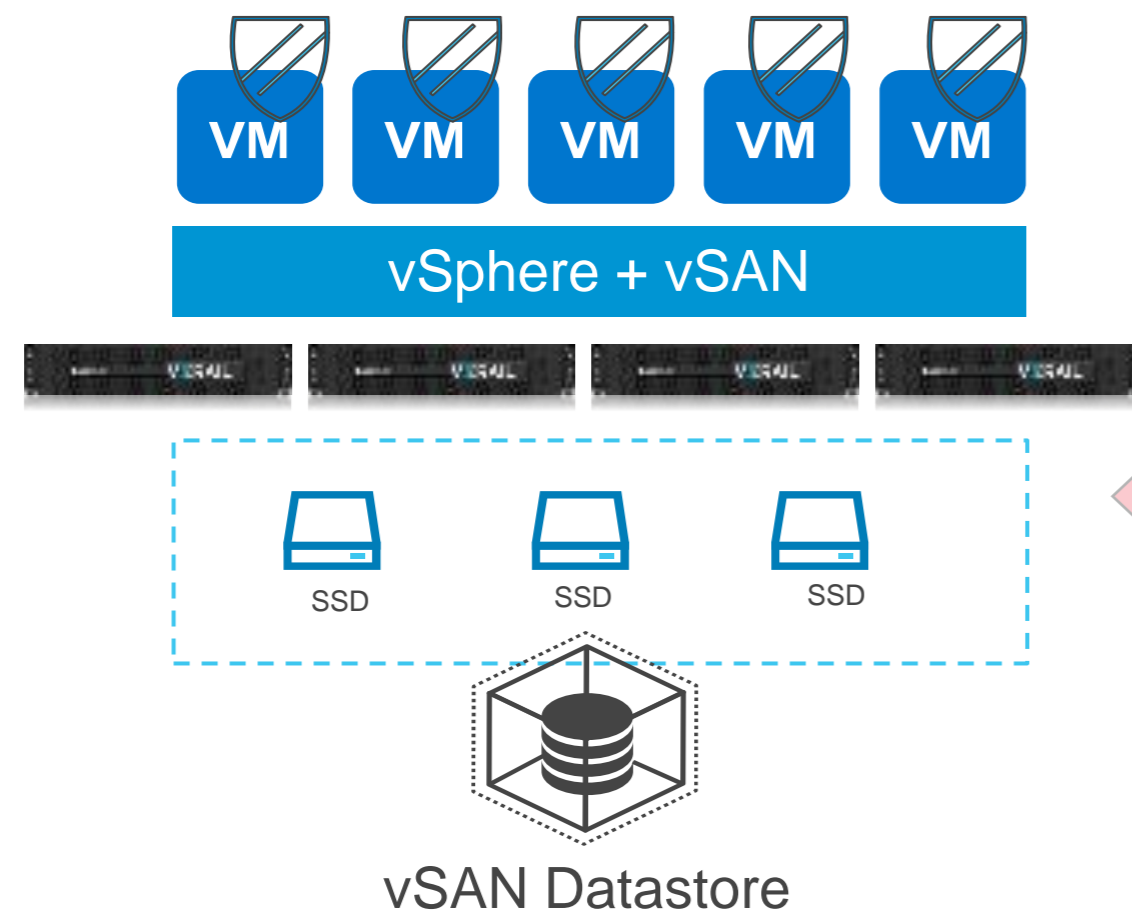
VMware VMkernel Cryptographic Module v1.0 has achieved FIPS 140-2

vSAN uses this FIPS 140-2 validated cryptographic module implemented in vSphere

Which Encryption Should I Use?

Each solves a unique challenges with trade-offs


vSphere Encryption



Addresses fears of rogue admin

- Per VM level setting
- No benefit from dedupe & compression
- In-flight (over the wire)

vSAN Encryption



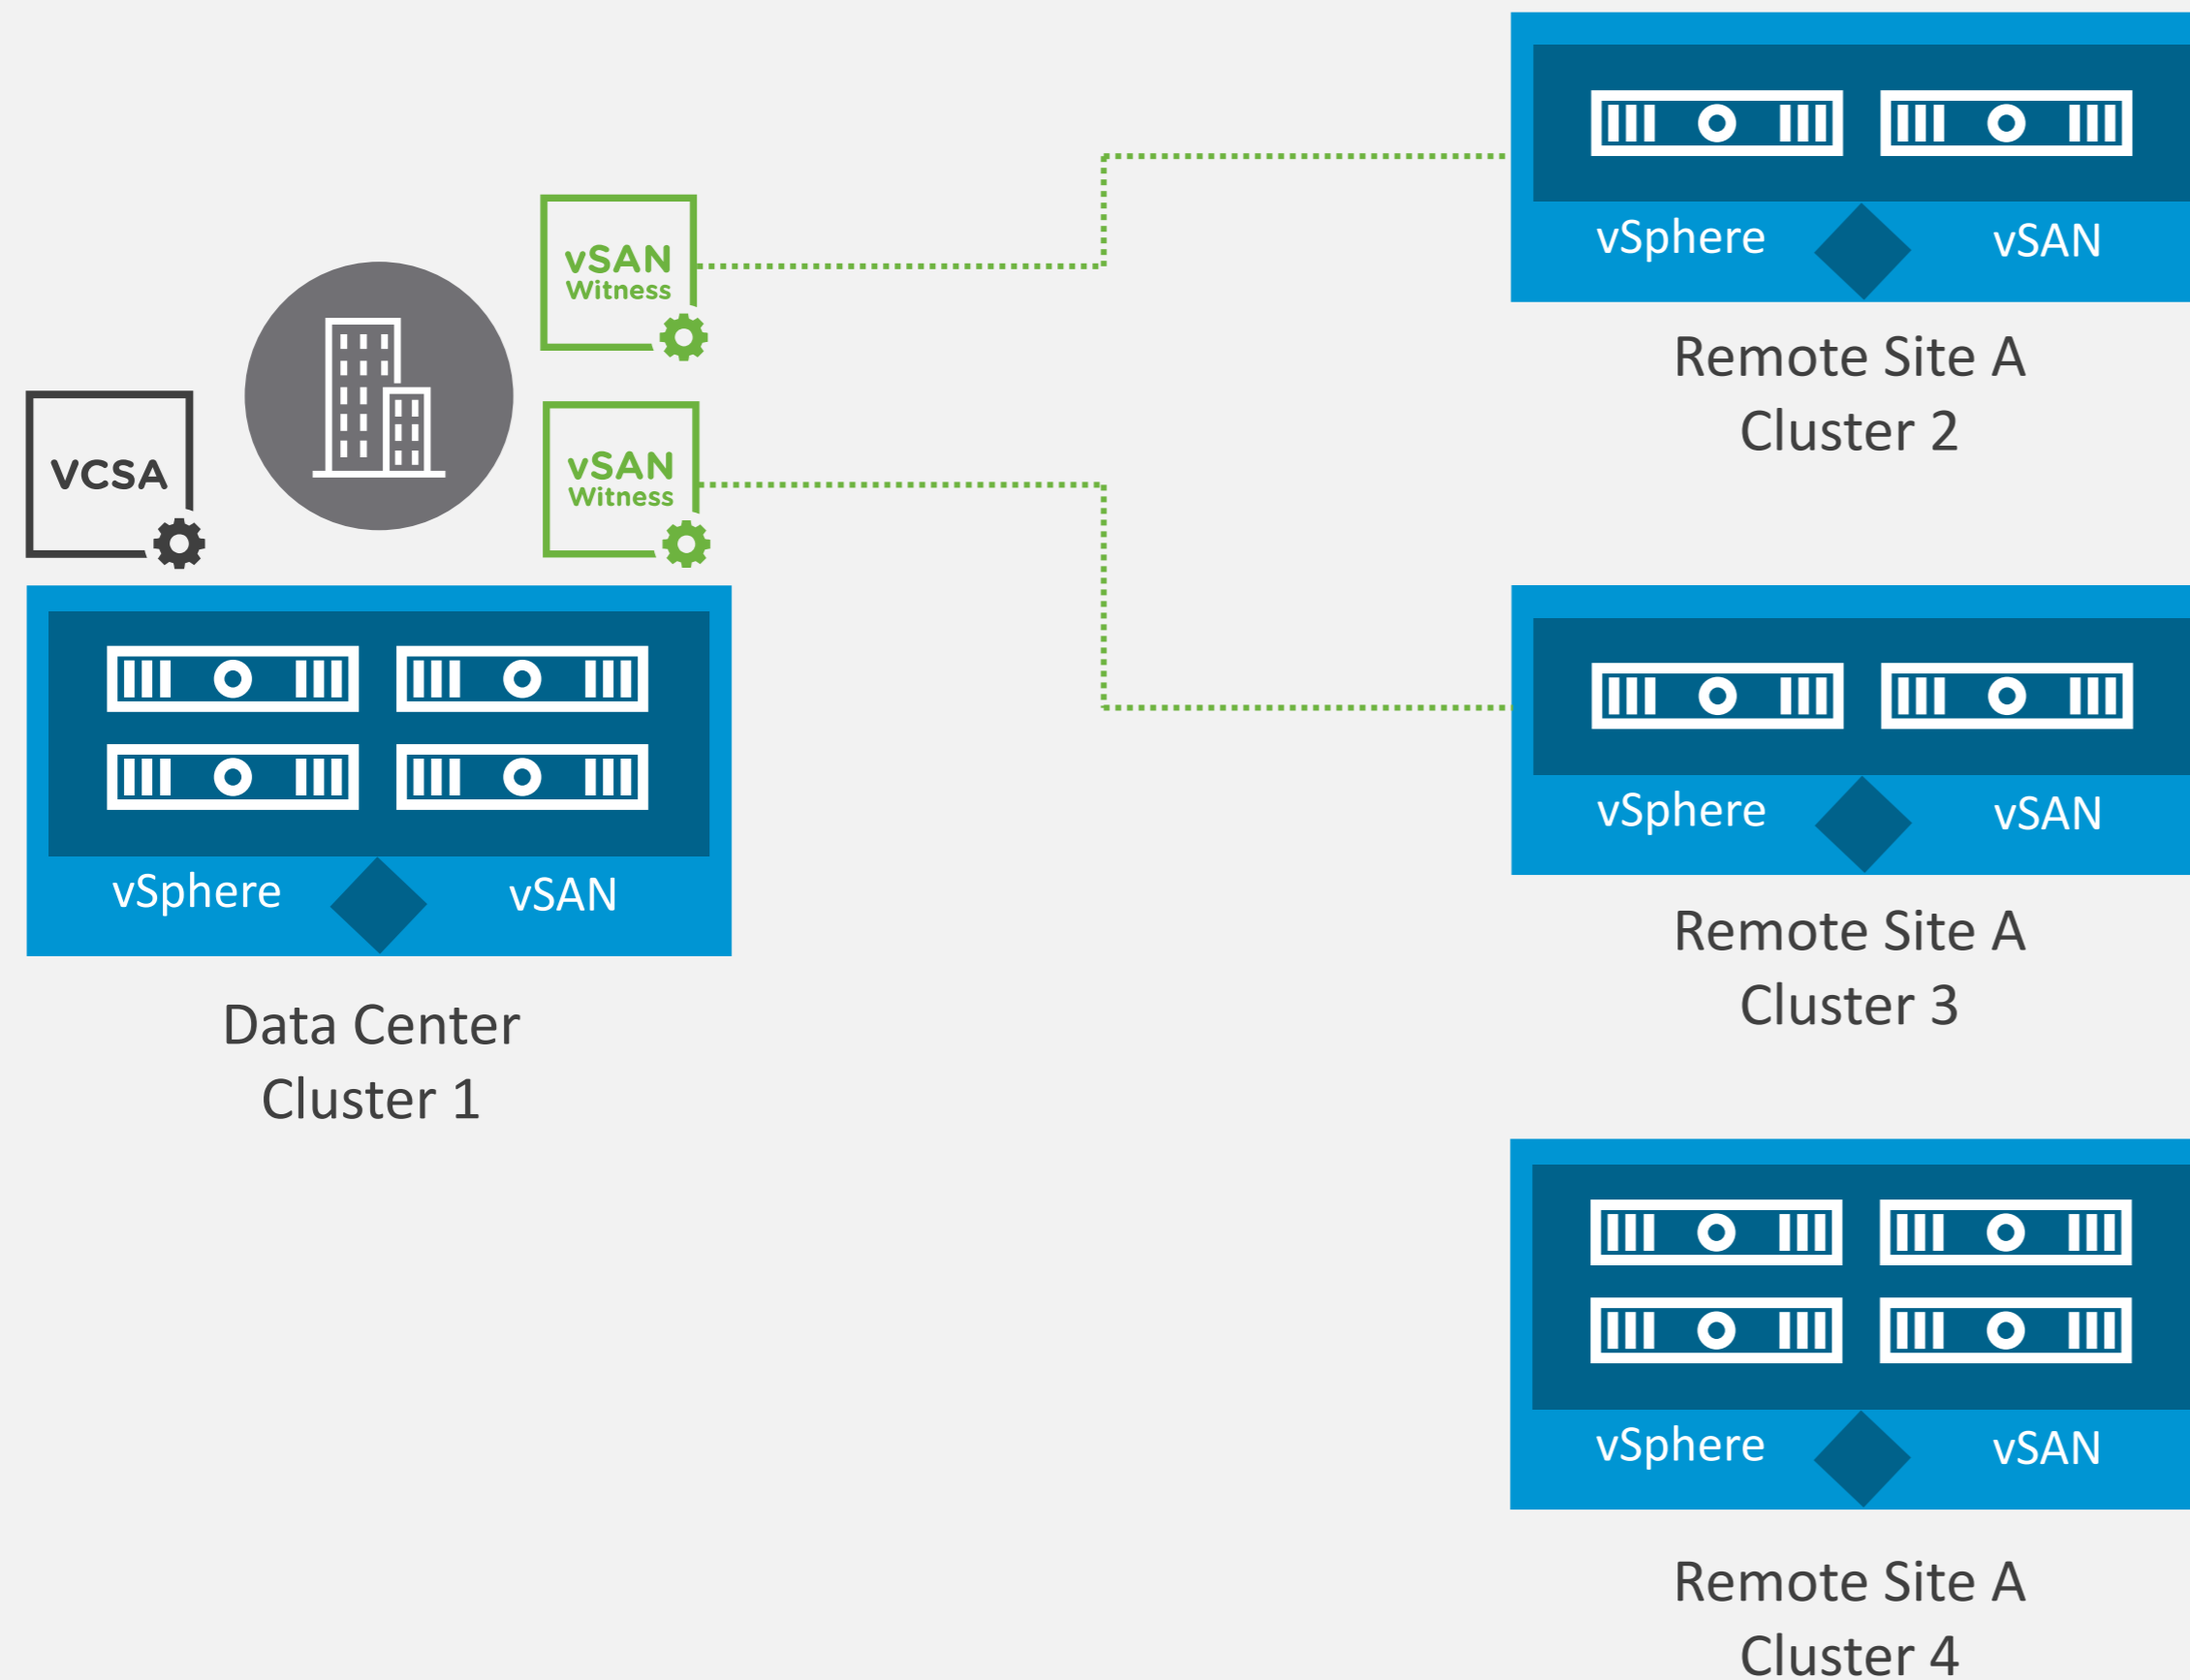
Addresses fears of media theft

- Applied to entire datastore (all or none)
- Allows dedupe & compression
- At rest encryption

Both require use of a supported external key management server (KMS) that cannot be in VxRail cluster being encrypted

Easily Accommodate Branch Office Environments

vSAN for Remote Locations



Up to **25 VM's** per site with **ROBO license**

Can scale to more nodes if needed

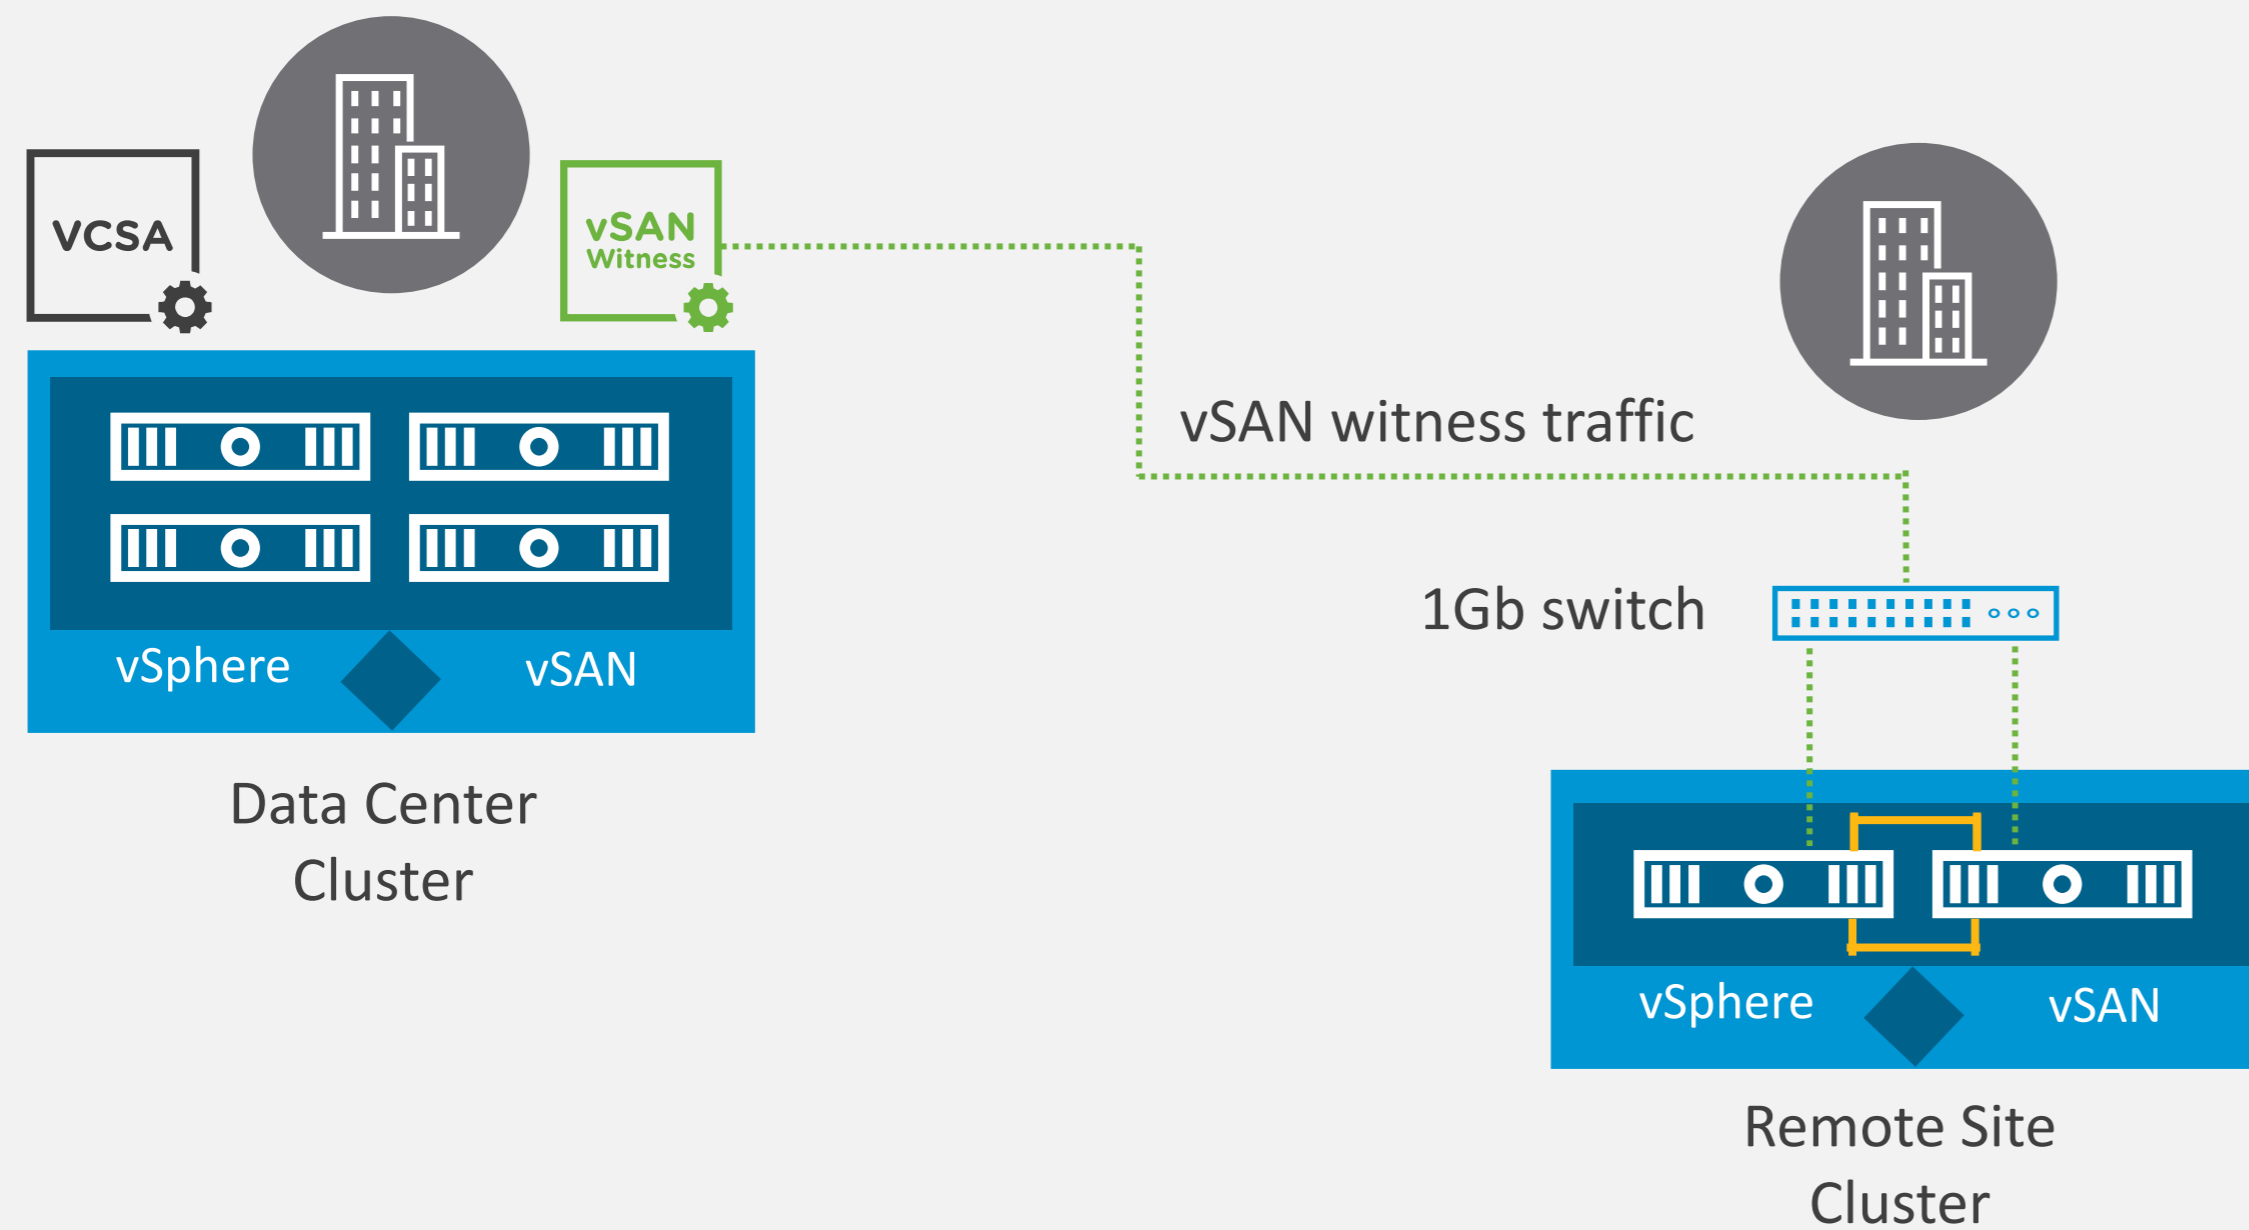
When using 2 Node

- Both hosts in remote office store data
- Witness in central office or other site
- FTT=1 Mirroring & DD+C Supported

Managed by **single vCenter Server**

Affordable Branch Office Configurations

vSAN 2 Node – Direct Connect



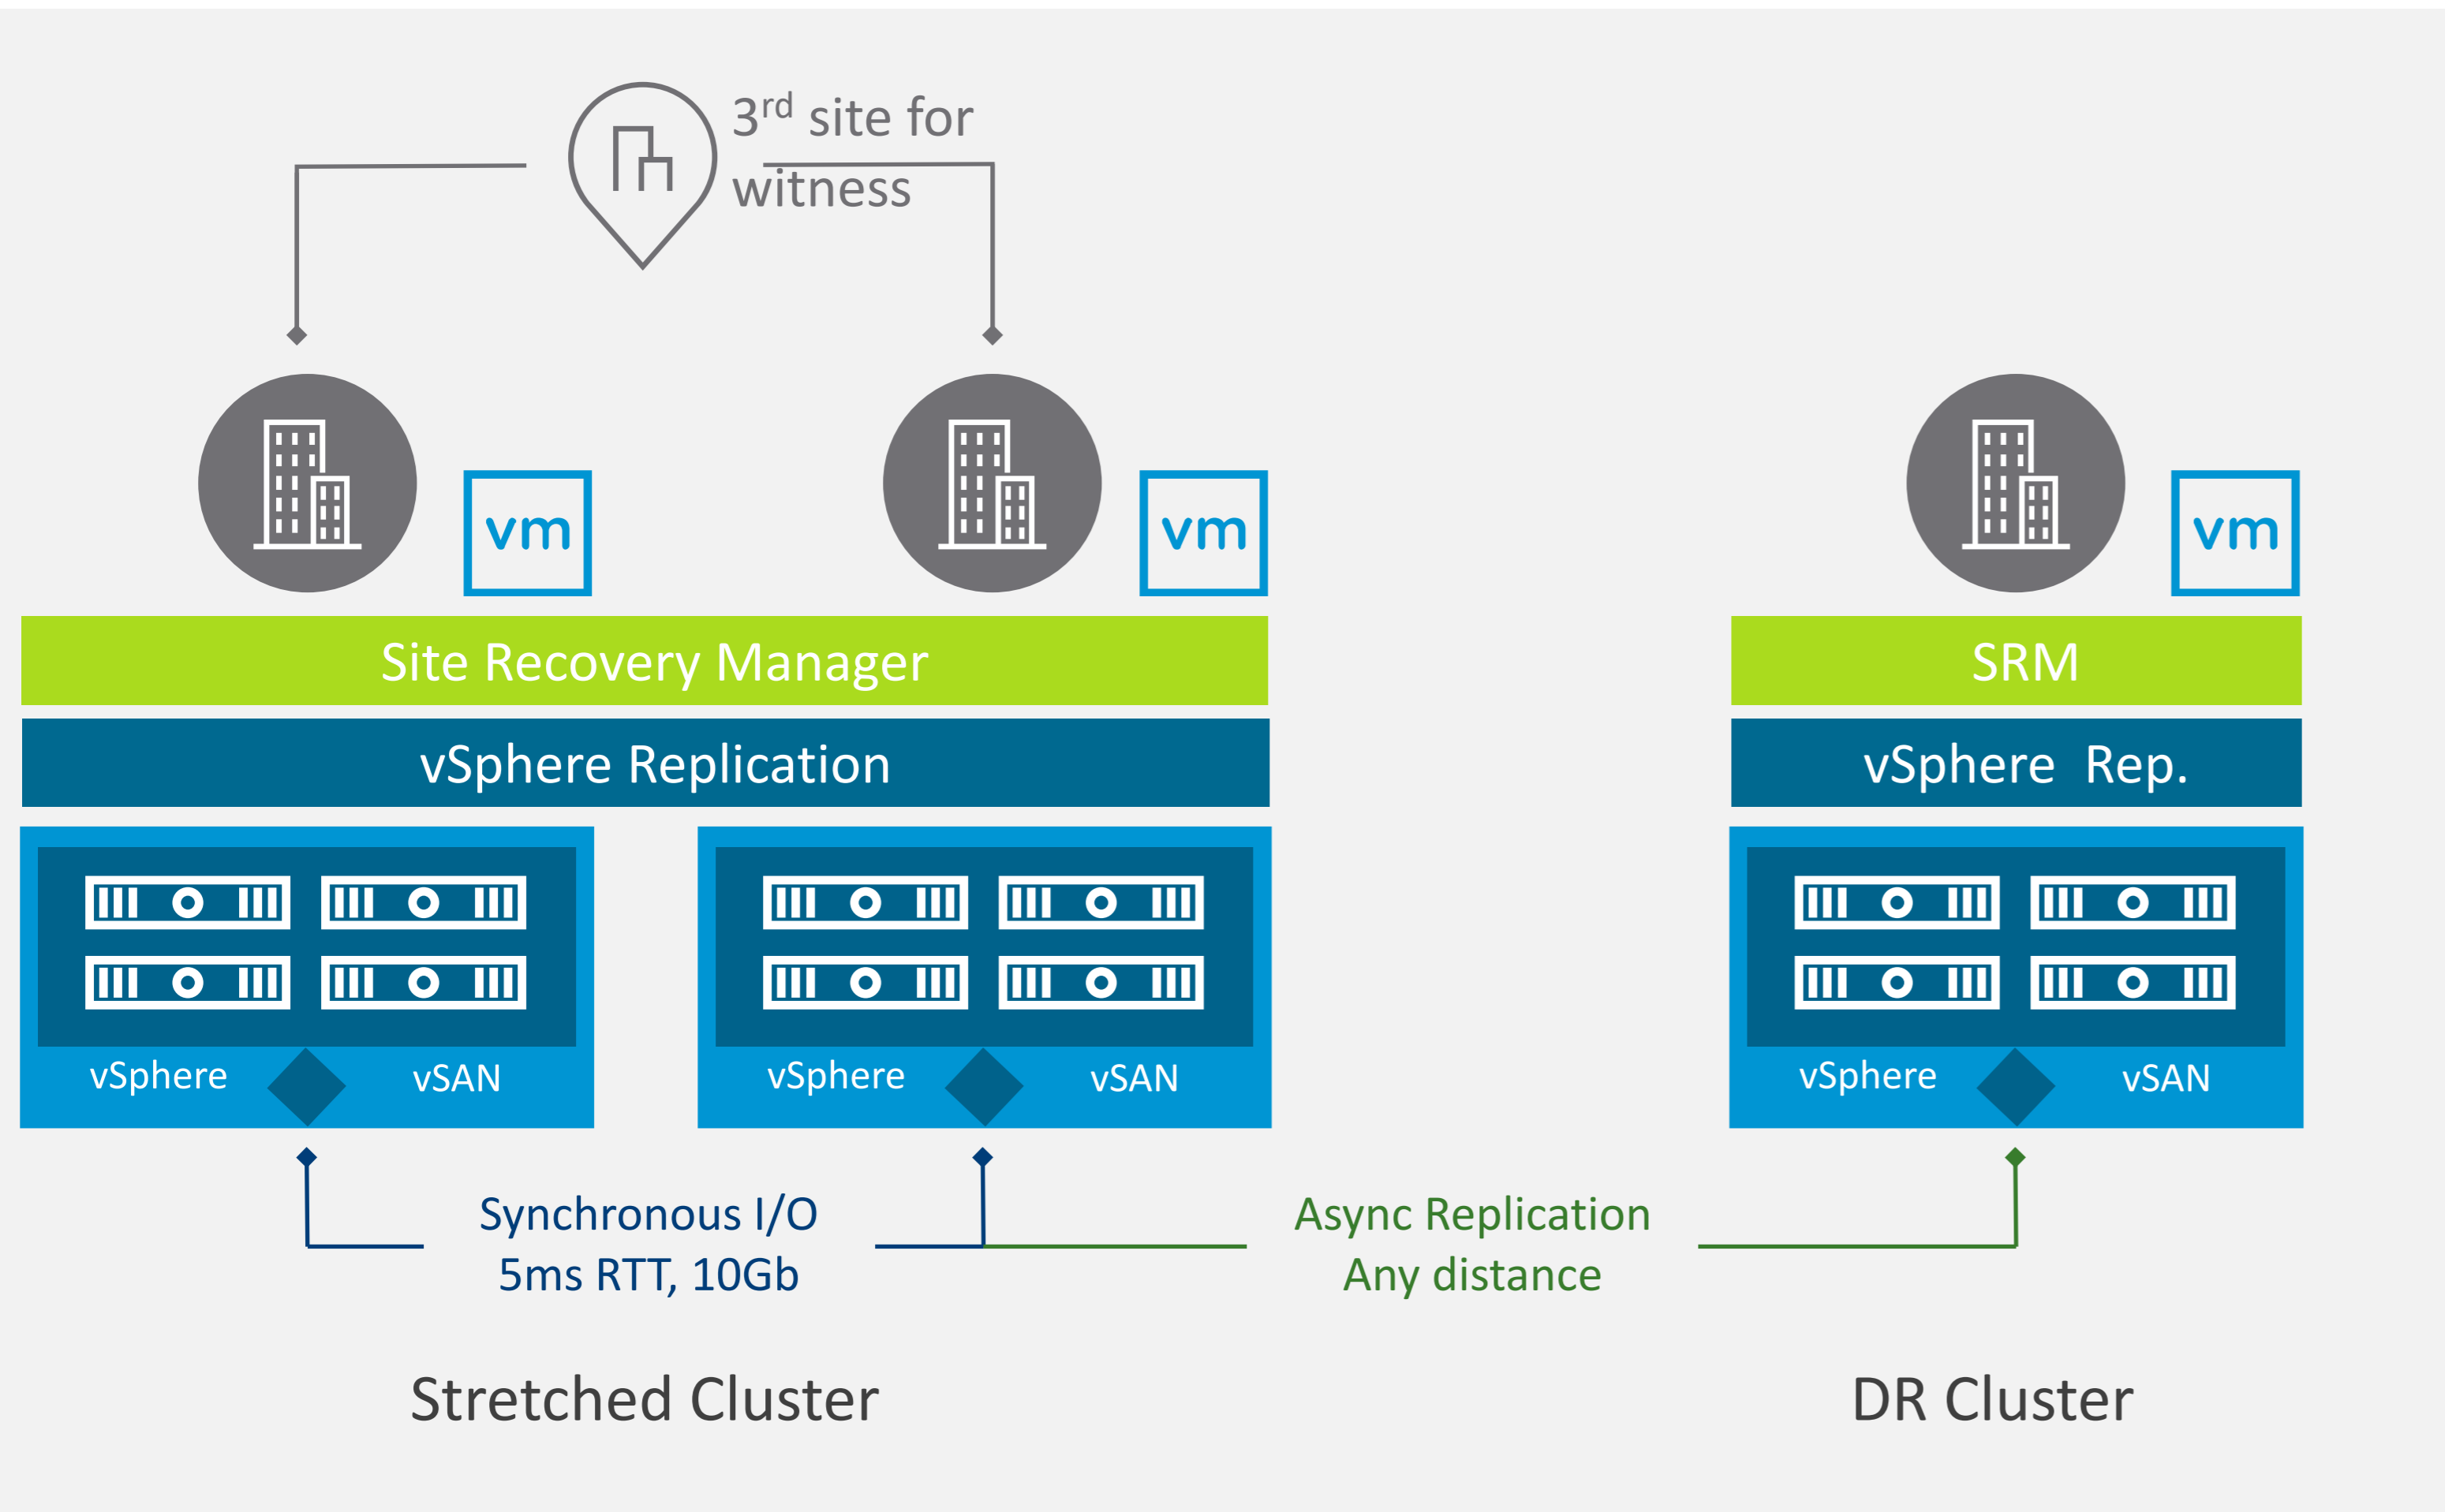
Connect **two nodes directly** between 10GbE NICs

Witness Traffic Separation **separates vSAN data** traffic from **witness traffic**

Accommodates **Layer-2** and **Layer-3** topologies

Active-Active Data Center with Additional DR Protection

vSAN stretched clusters and SRM



Use **vSphere Replication** to third site, and enable RPO as low as 5 minutes

Use **Site Recovery Manager** for disaster recovery orchestration

Stretched across metro distance, **replicated** across geo


VxRail Failure Scenarios

(mini presentation)

VxRail Failure Scenarios


Choose a scenario:

- [Single site, using mirroring](#)
- [Single site, using erasure coding](#)
- [Stretched cluster](#)
- [Absent / Degraded](#)



VxRail - Single Site

Single Site, FTT = 1, FTM = Mirror




Click on a black arrow to simulate a failure

*example of component layout

VxRail - Single Site

Single Site, FTT = 1, FTM = Erasure Coding

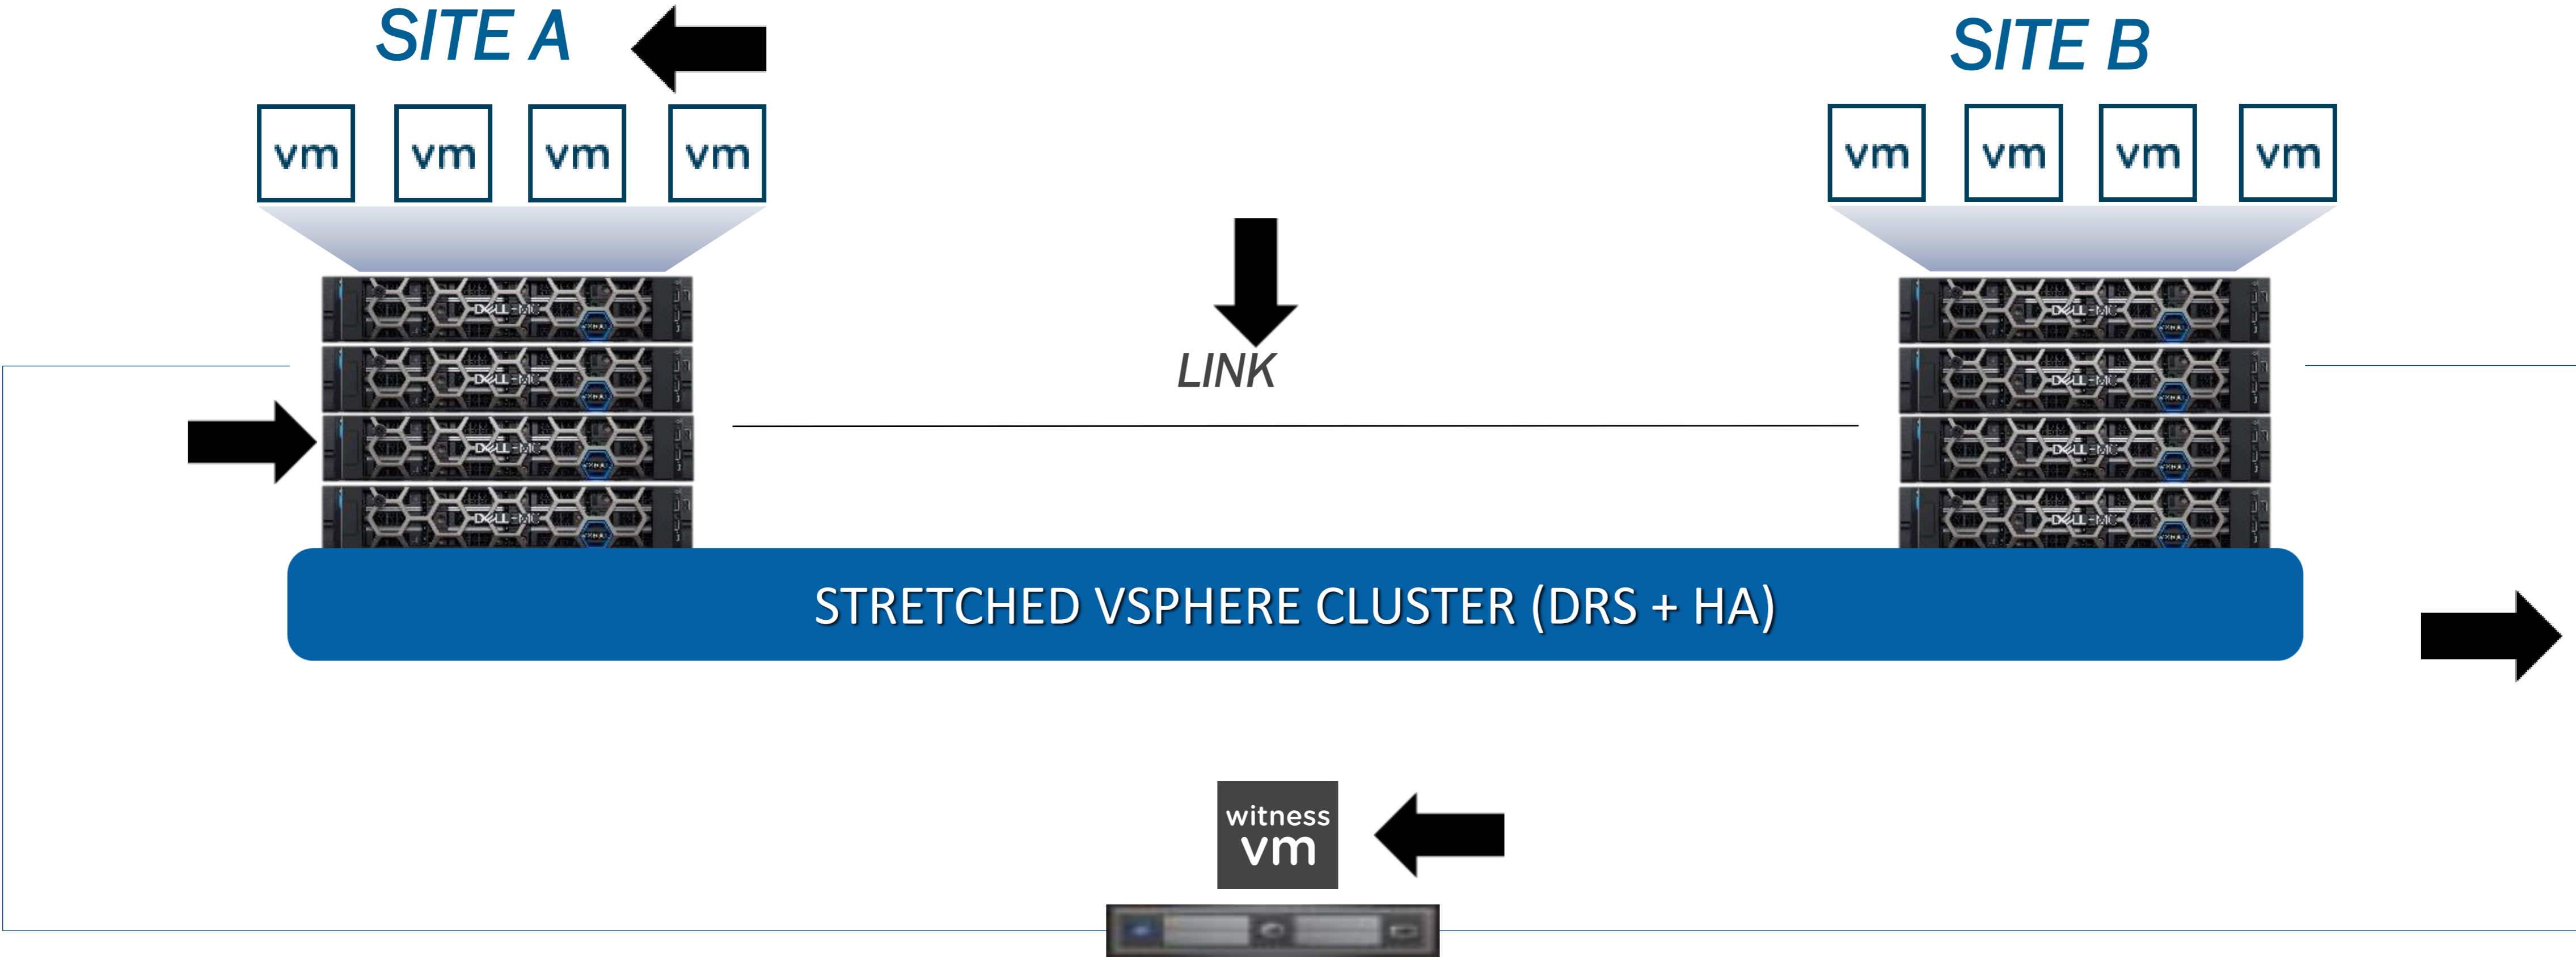


Click on a black arrow to simulate a failure

*example of component layout

VxRail - Stretched Cluster


Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror



Click on a black arrow to simulate a failure

VxRail – Absent / Degraded

Single Site, FTT =1, FTM = Mirror




Click one of the buttons above to shutdown the host or unplug the disk.

*example of component layout

VxRail - Single Site

Single Site, FTT = 1, FTM = Mirror




On the following Slides you will see the scenarios of a Single Site with FTT = 1 and FTM = Mirror

*example of component layout

Cache Drive Failure

Single Site, FTT =1, FTM = Mirror




If a cache drive fails, then the whole disk group fails.
vSAN will copy the missing components to a different disk group.

*example of component layout

vSAN status: Not Compliant

Capacity Drive Failure

Single Site, FTT =1, FTM = Mirror



If a capacity drive fails, vSAN will copy the missing components to a different drive.

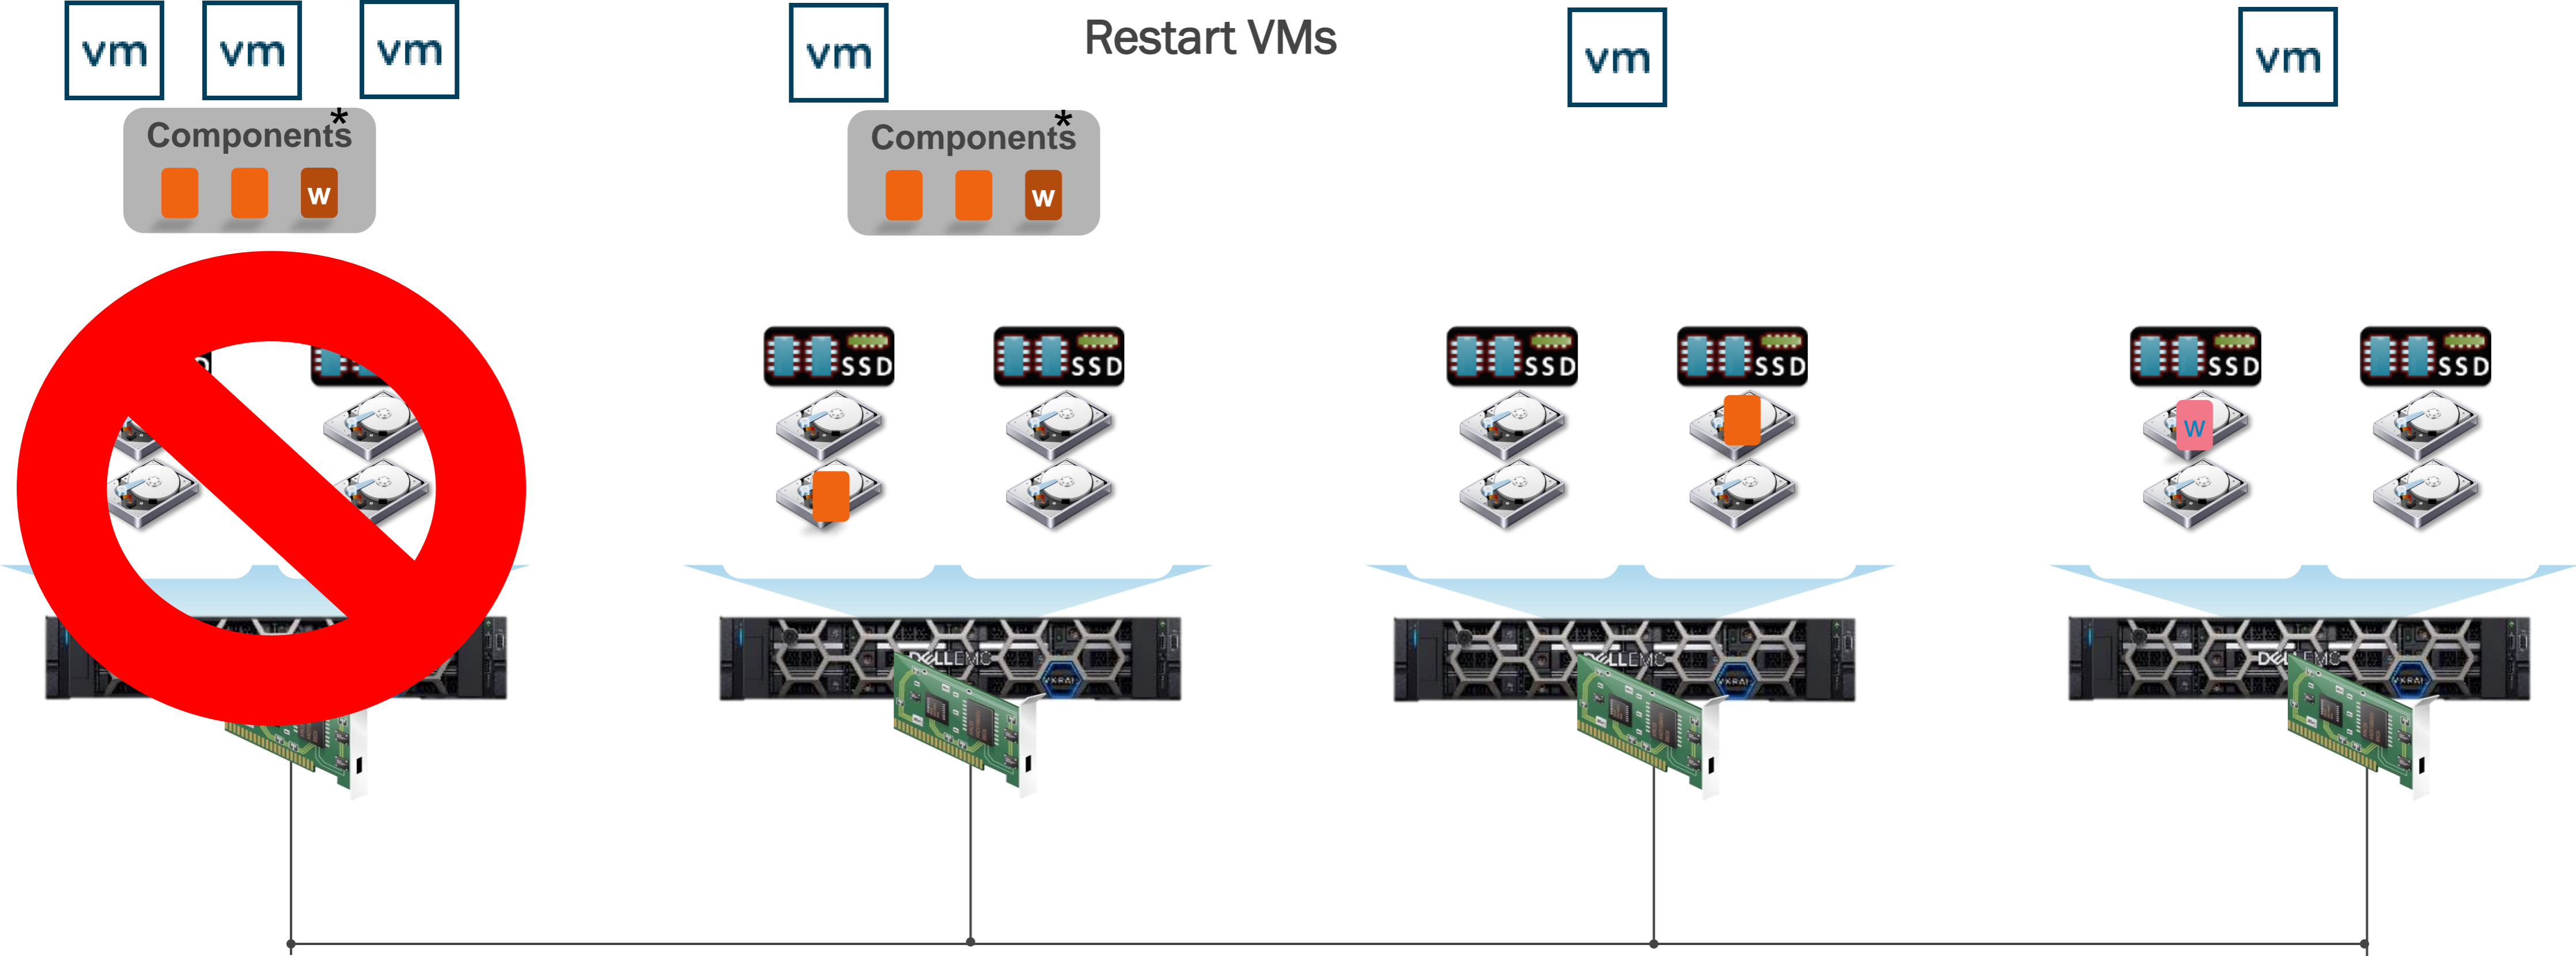
*example of component layout

vSAN status:  Not Compliant

Node Failure with VM / No Component

Back

Single Site, FTT = 1, FTM = Mirror



If a host fails, vSphere HA will restart the VMs on a different host.


*example of component layout

vSAN status: Not Compliant



Node Failure with Component

Single Site, FTT = 1, FTM = Mirror




If a host, that doesn't run any VMs, but contains vSAN components fails, then vSAN will copy the missing components to a new host.

*example of component layout

vSAN status:  Not Compliant

Onboard NIC Failure

Single Site, FTT =1, FTM = Mirror



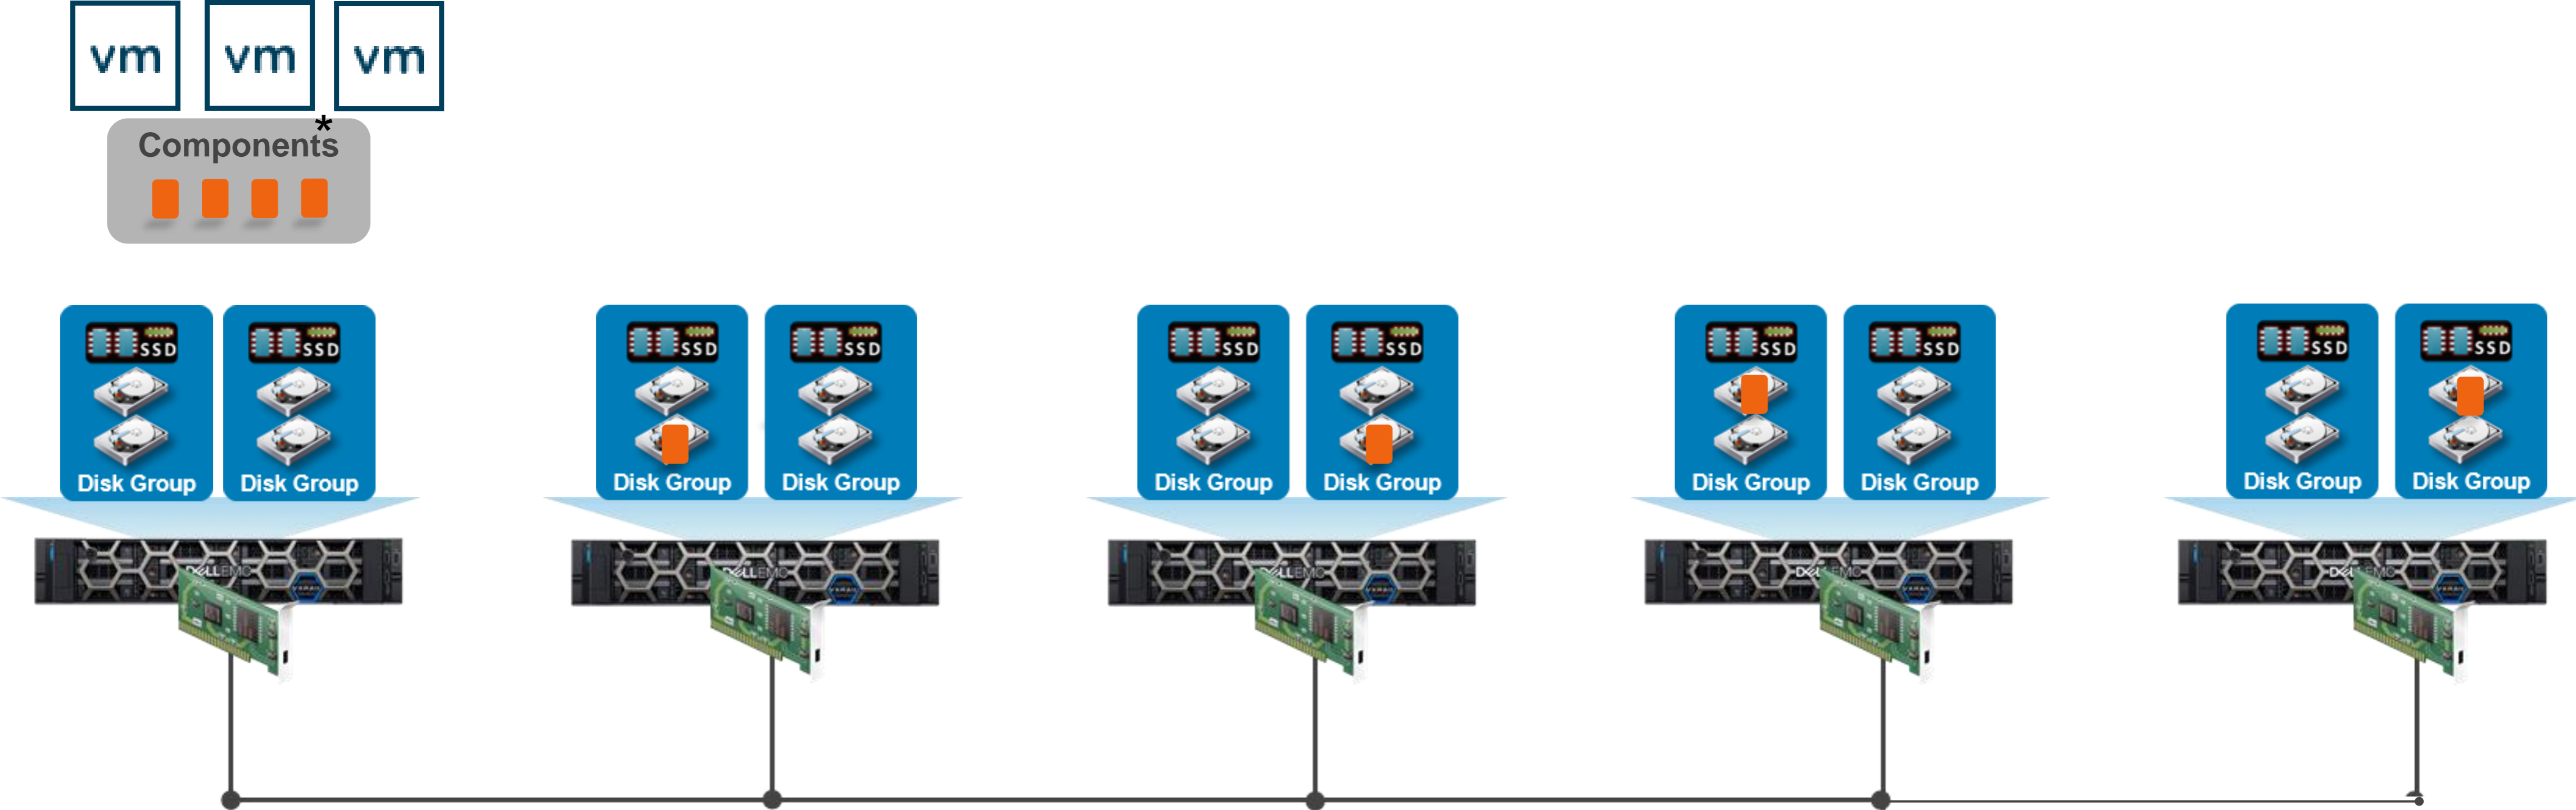
If the onboard NIC fails, then the host fails and vSAN will copy the missing components to a new host.

*example of component layout

vSAN status: Not Compliant

VxRail - Single Site

Single Site, FTT = 1, FTM = Erasure Coding

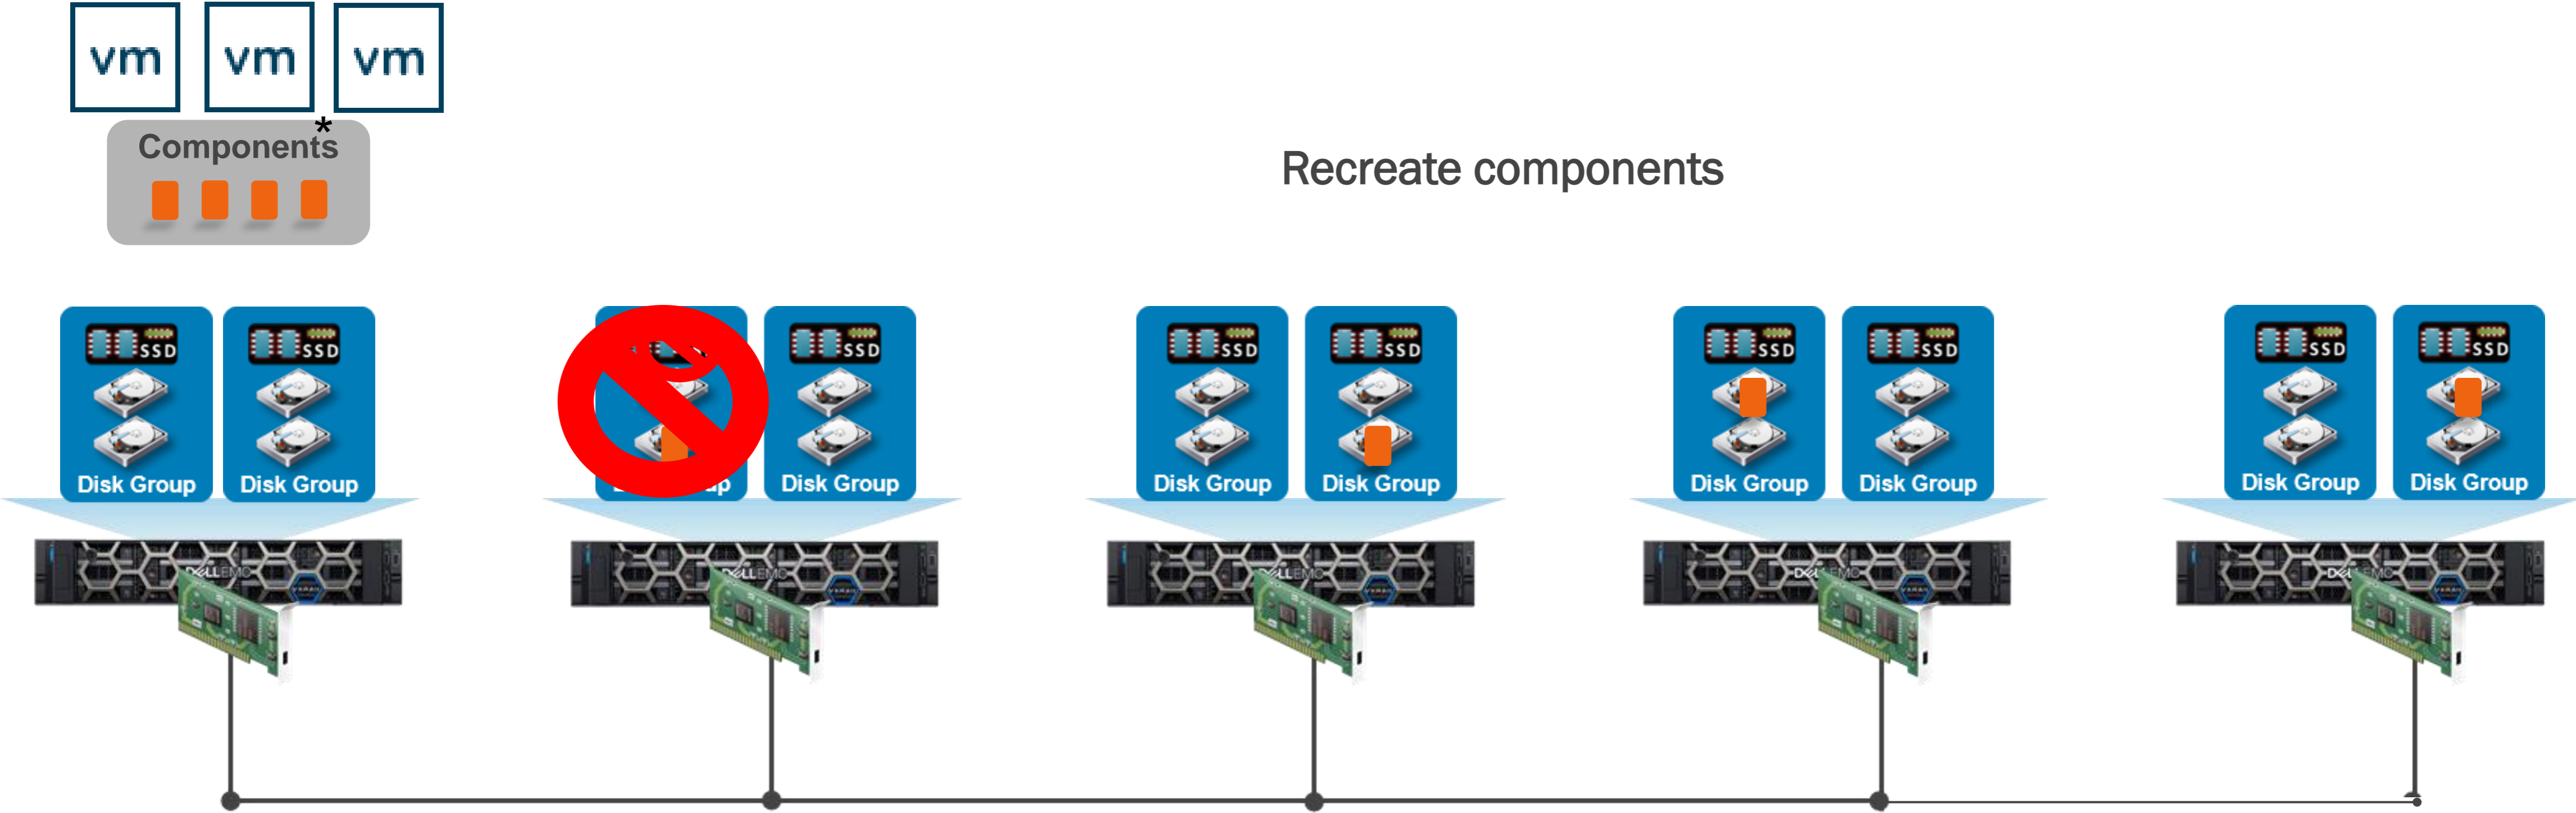


On the following Slides you will see the scenarios of a Single Site with FTT = 1 and FTM = EC

*example of component layout

Cache Drive Failure

Single Site, FTT =1, FTM = Erasure Coding



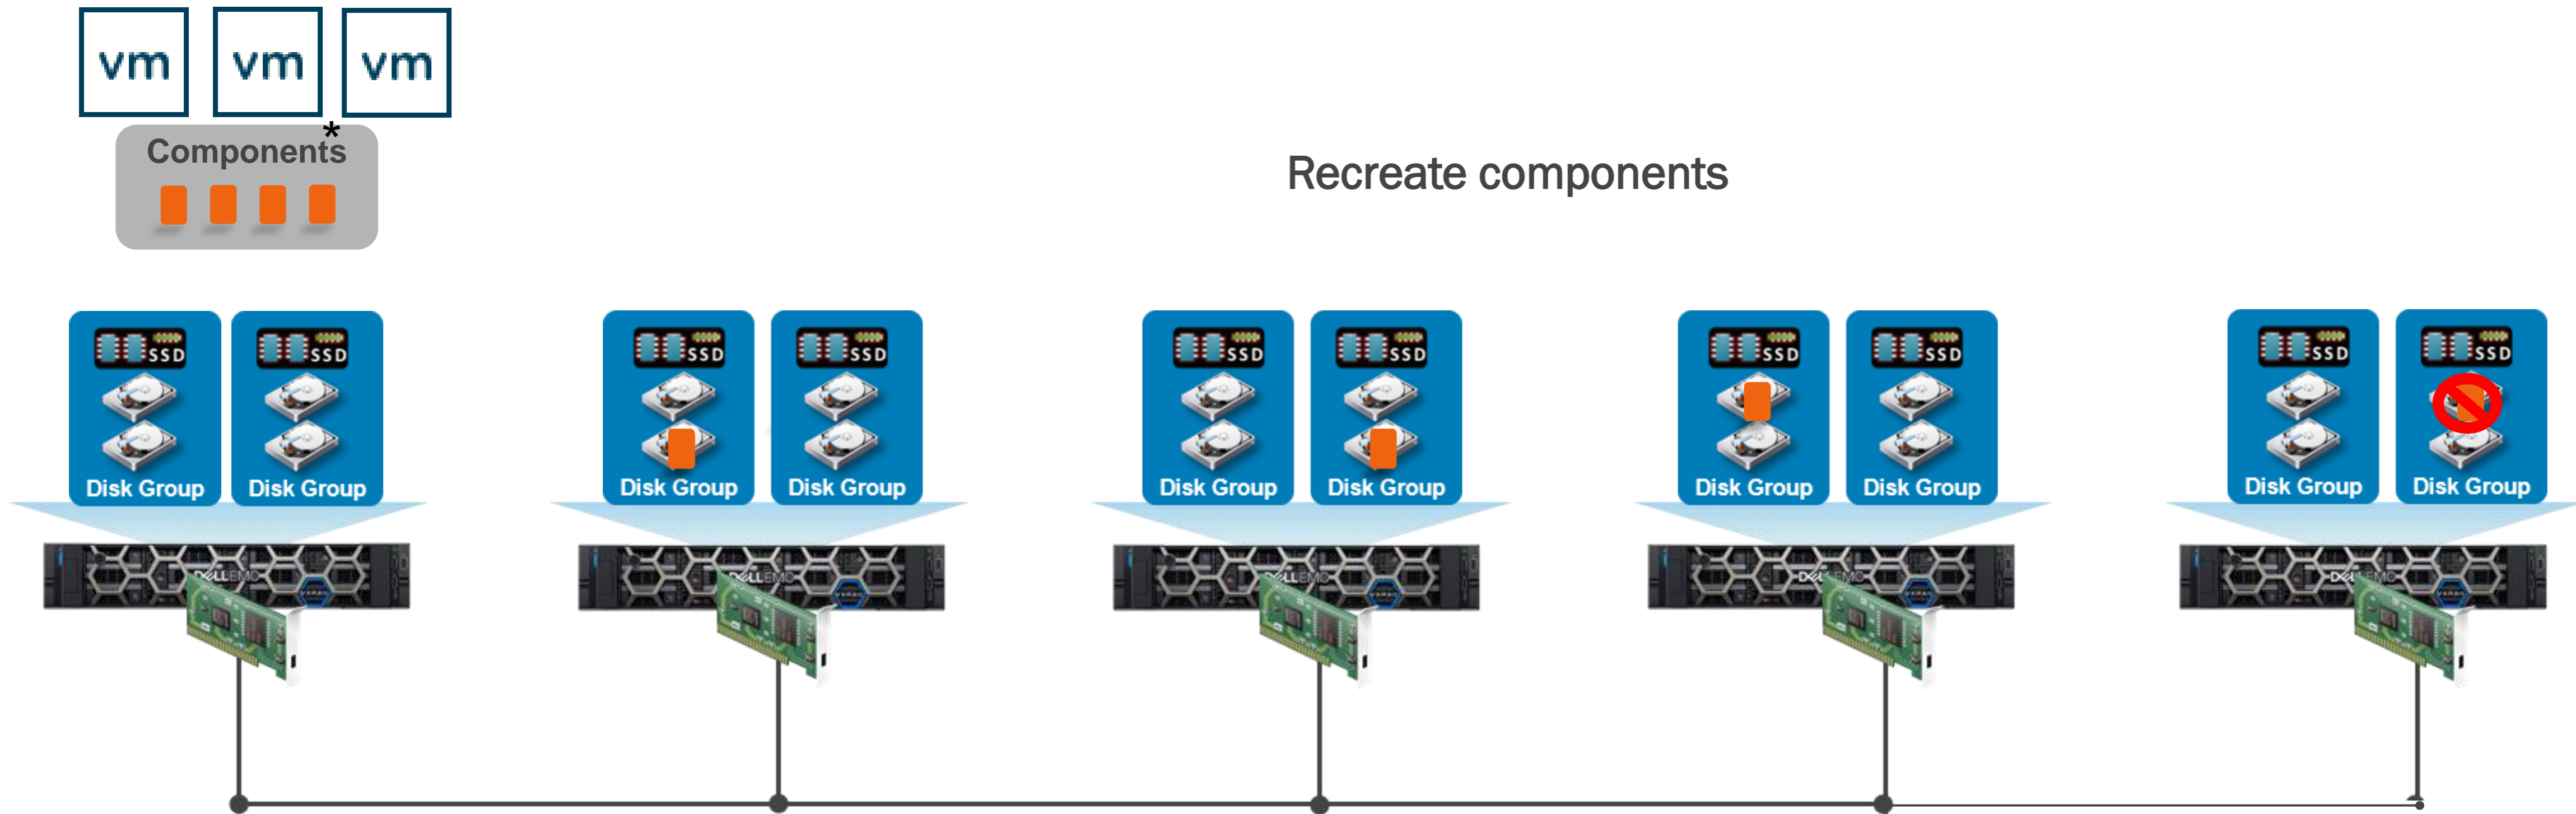
If a cache drive fails, then the whole disk group fails. vSAN will recreate the missing components at a different disk group.

*example of component layout

vSAN status: Not Compliant

Capacity Drive Failure

Single Site, FTT =1, FTM = Erasure Coding



If a capacity drive fails - vSAN will immediately recreate the missing components to a different drive.

*example of component layout

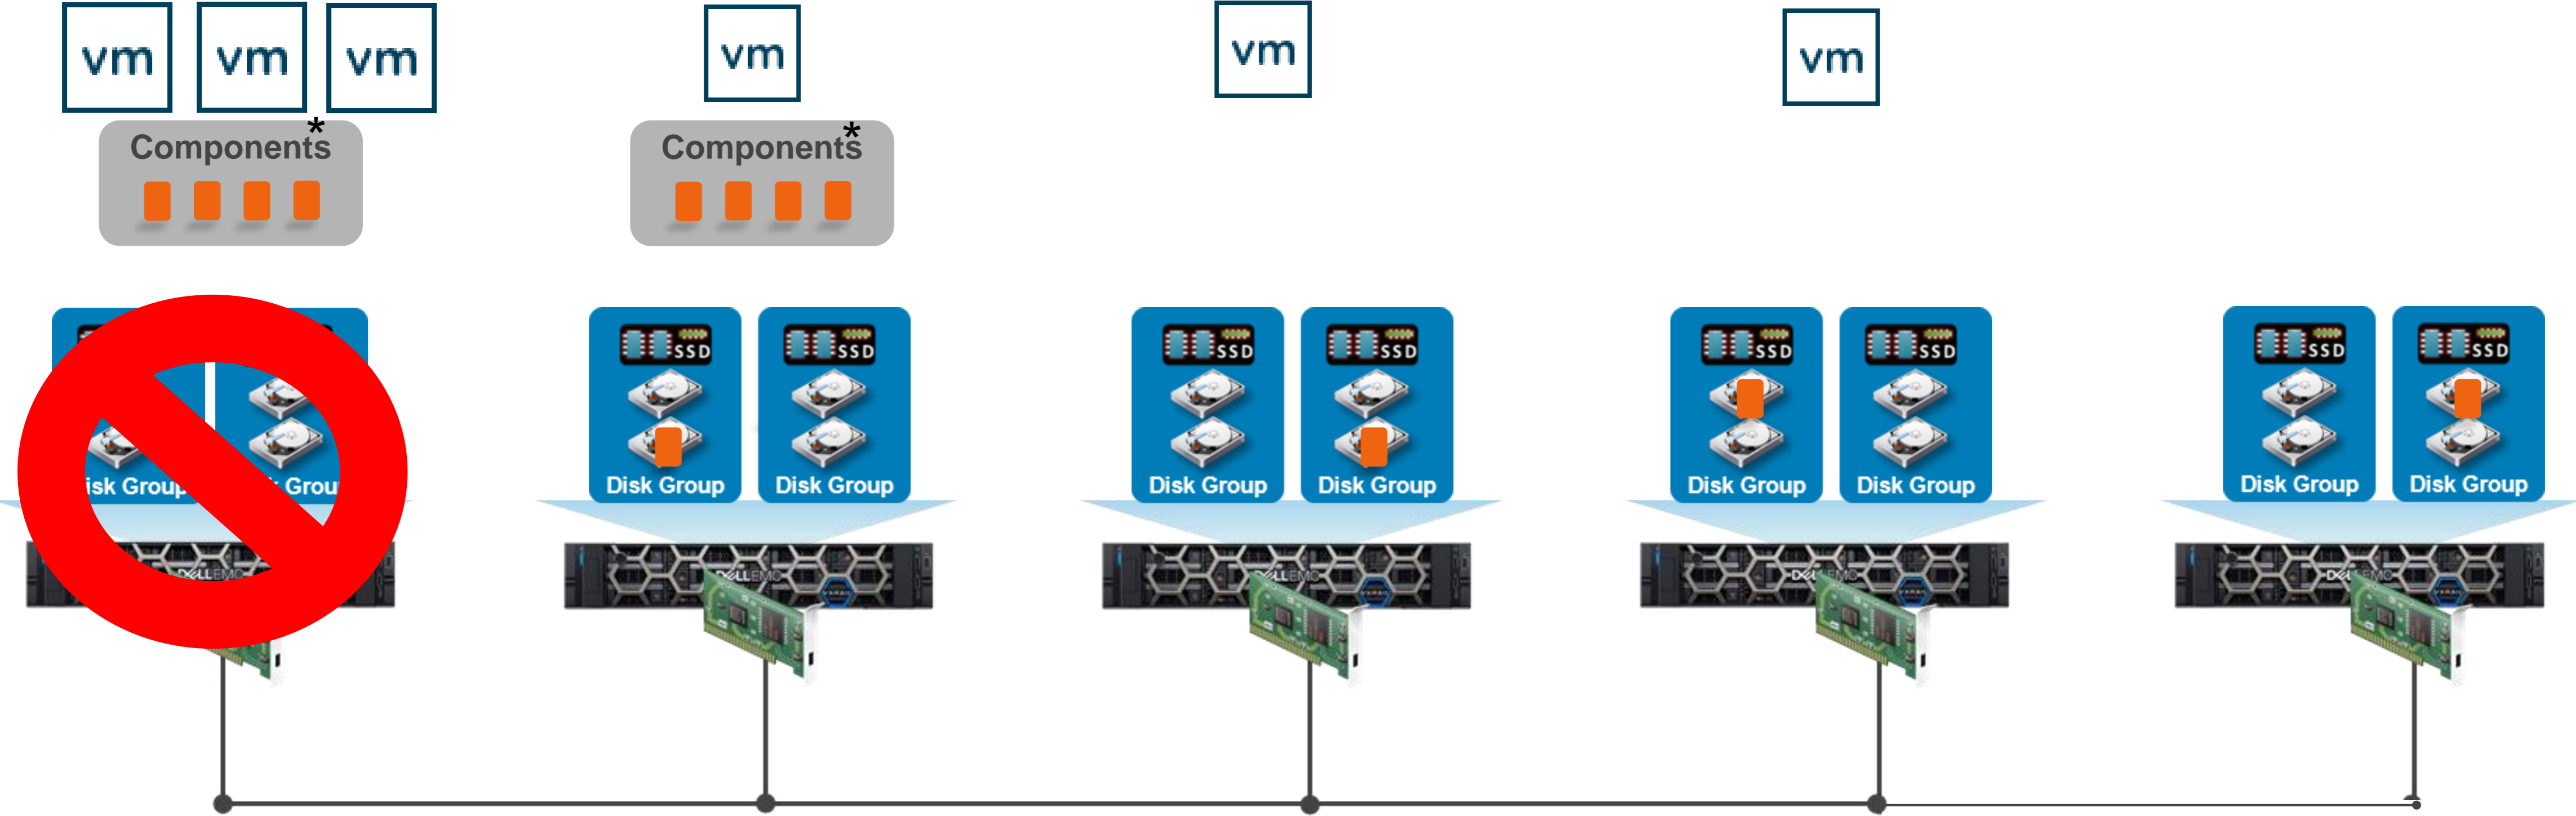
vSAN status:  Not Compliant

Node Failure with VM / No Component

Back

Single Site, FTT =1, FTM = Erasure Coding

Restart VMs



If a host fails, vSphere HA will restart the VMs on a different host.

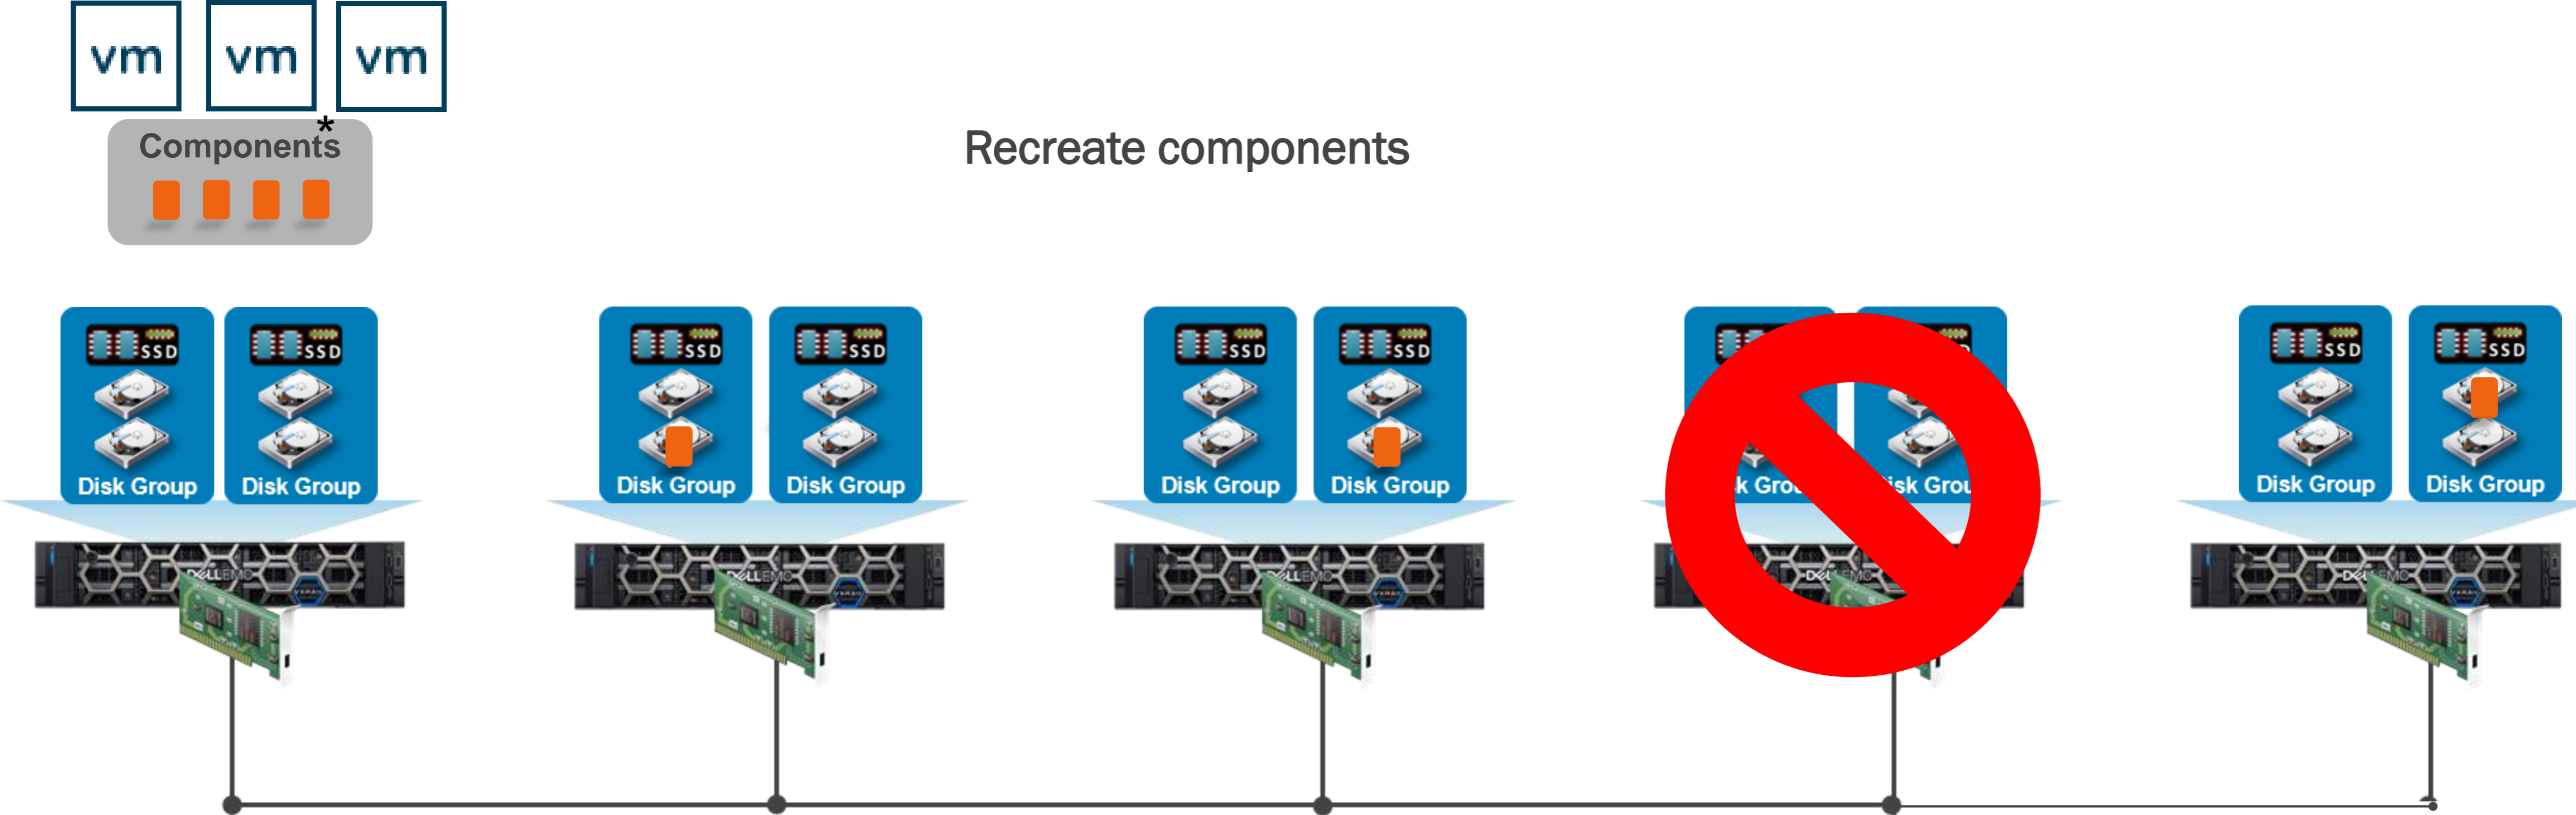
*example of component layout

vSAN status: Not Compliant



Node Failure with Component

Single Site, FTT =1, FTM = Erasure Coding




If a host, which contains vSAN components, but doesn't run any VMs fails, vSAN will recreate the missing components to a different host.

*example of component layout

vSAN status:  Not Compliant

Onboard NIC Failure

Single Site, FTT =1, FTM = Erasure Coding



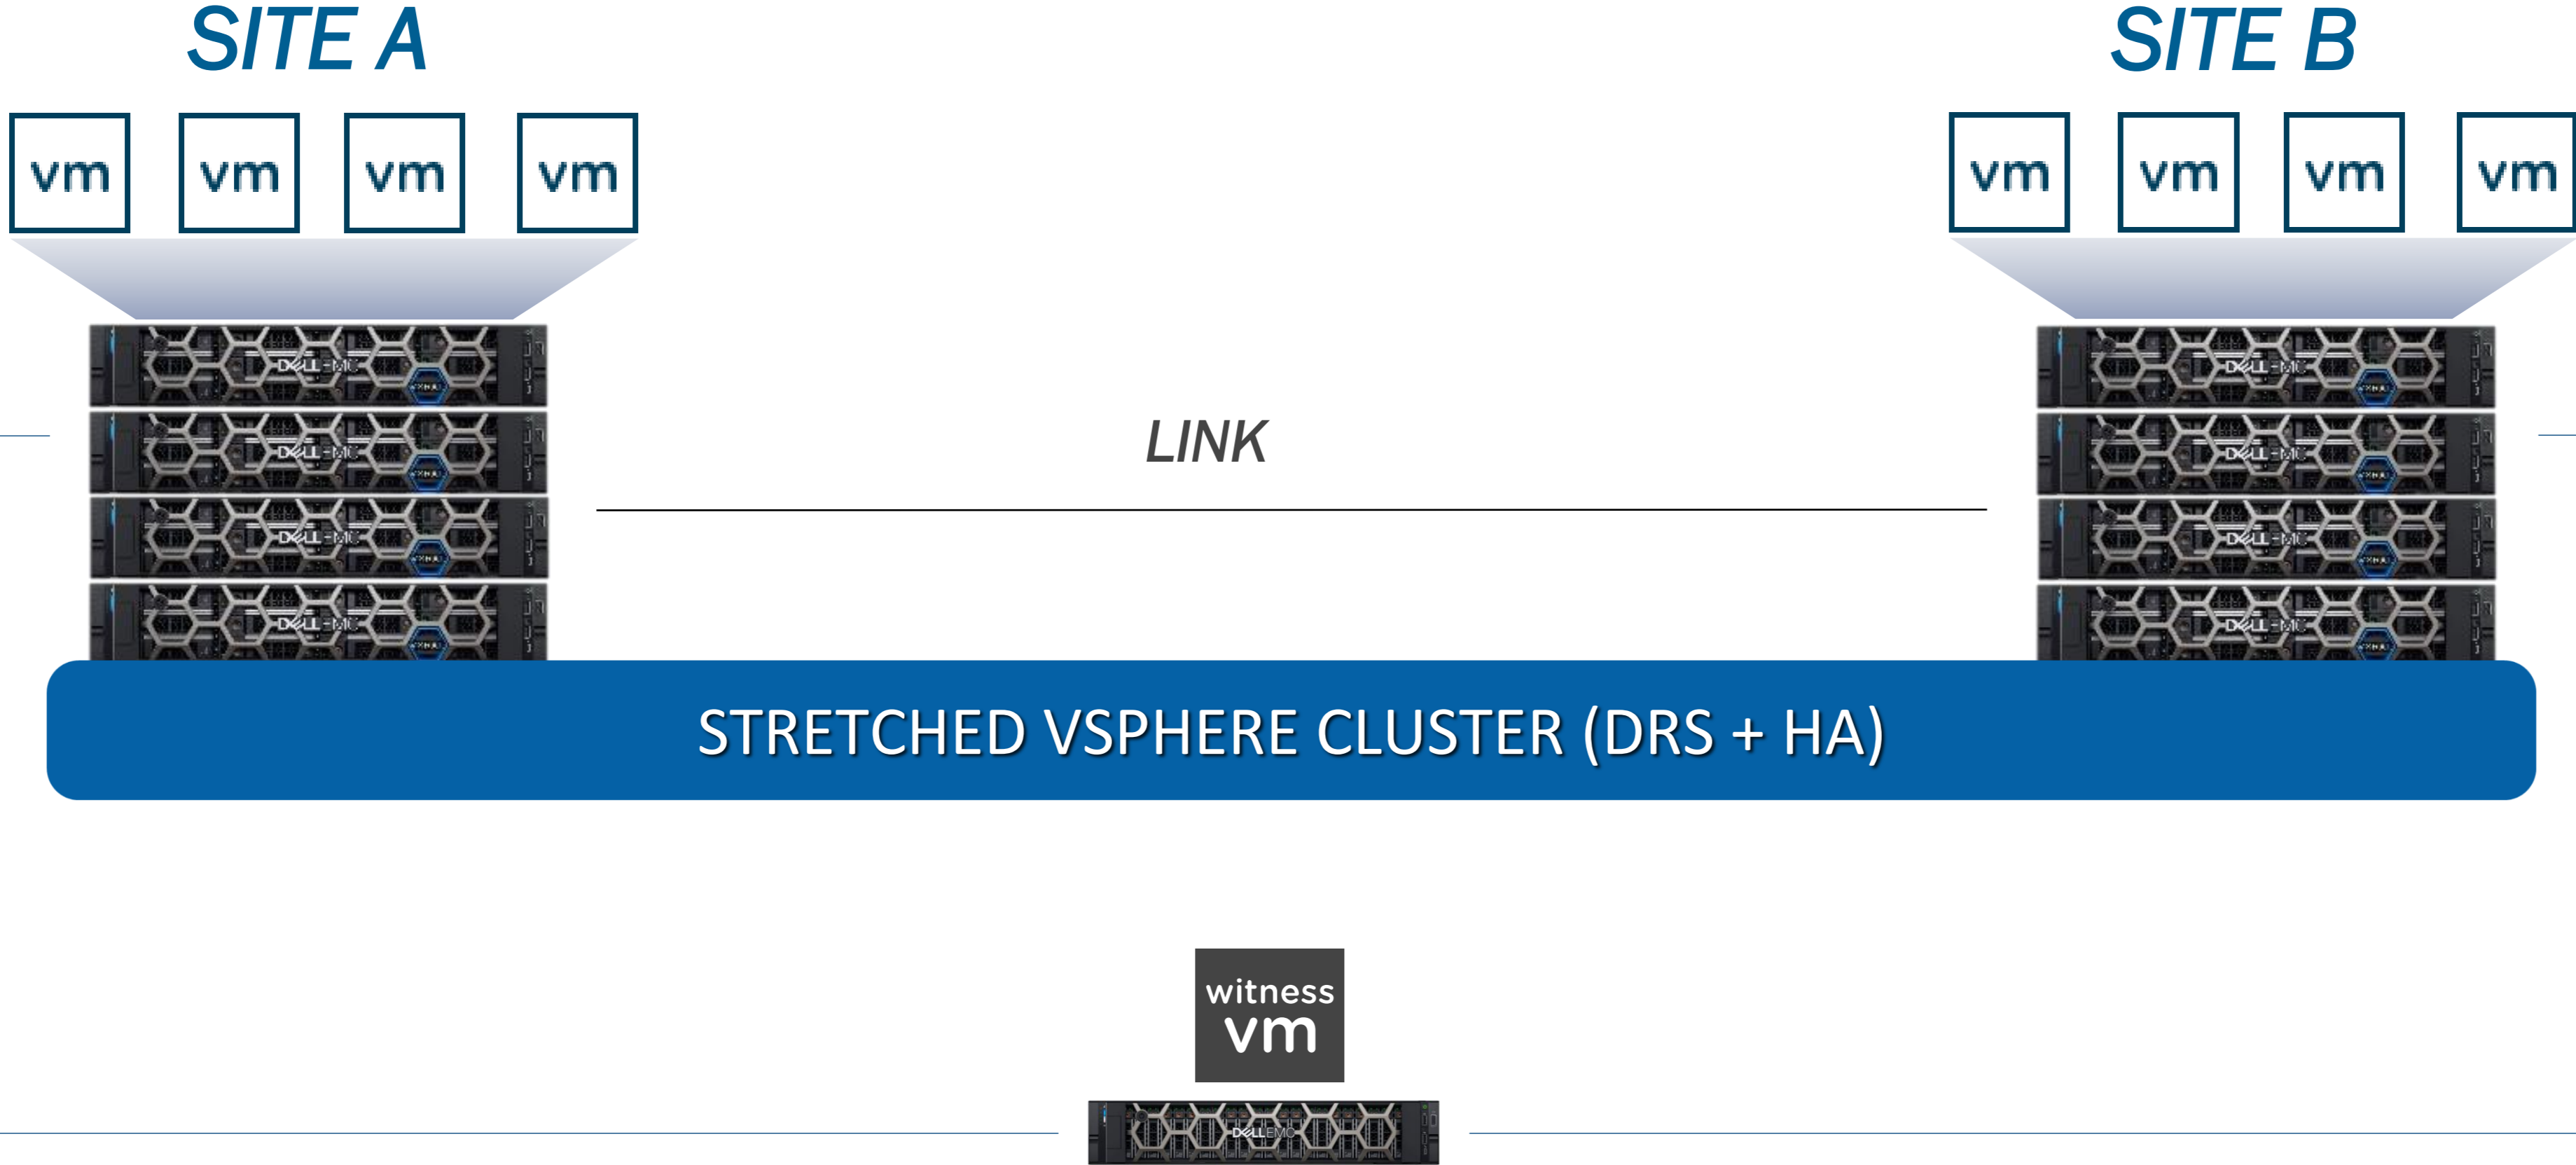
If the Onboard NIC fails, the Host fails - vSAN will recreate the missing components to a new host.

*example of component layout

vSAN status:  Not Compliant

VxRail – Stretched Cluster

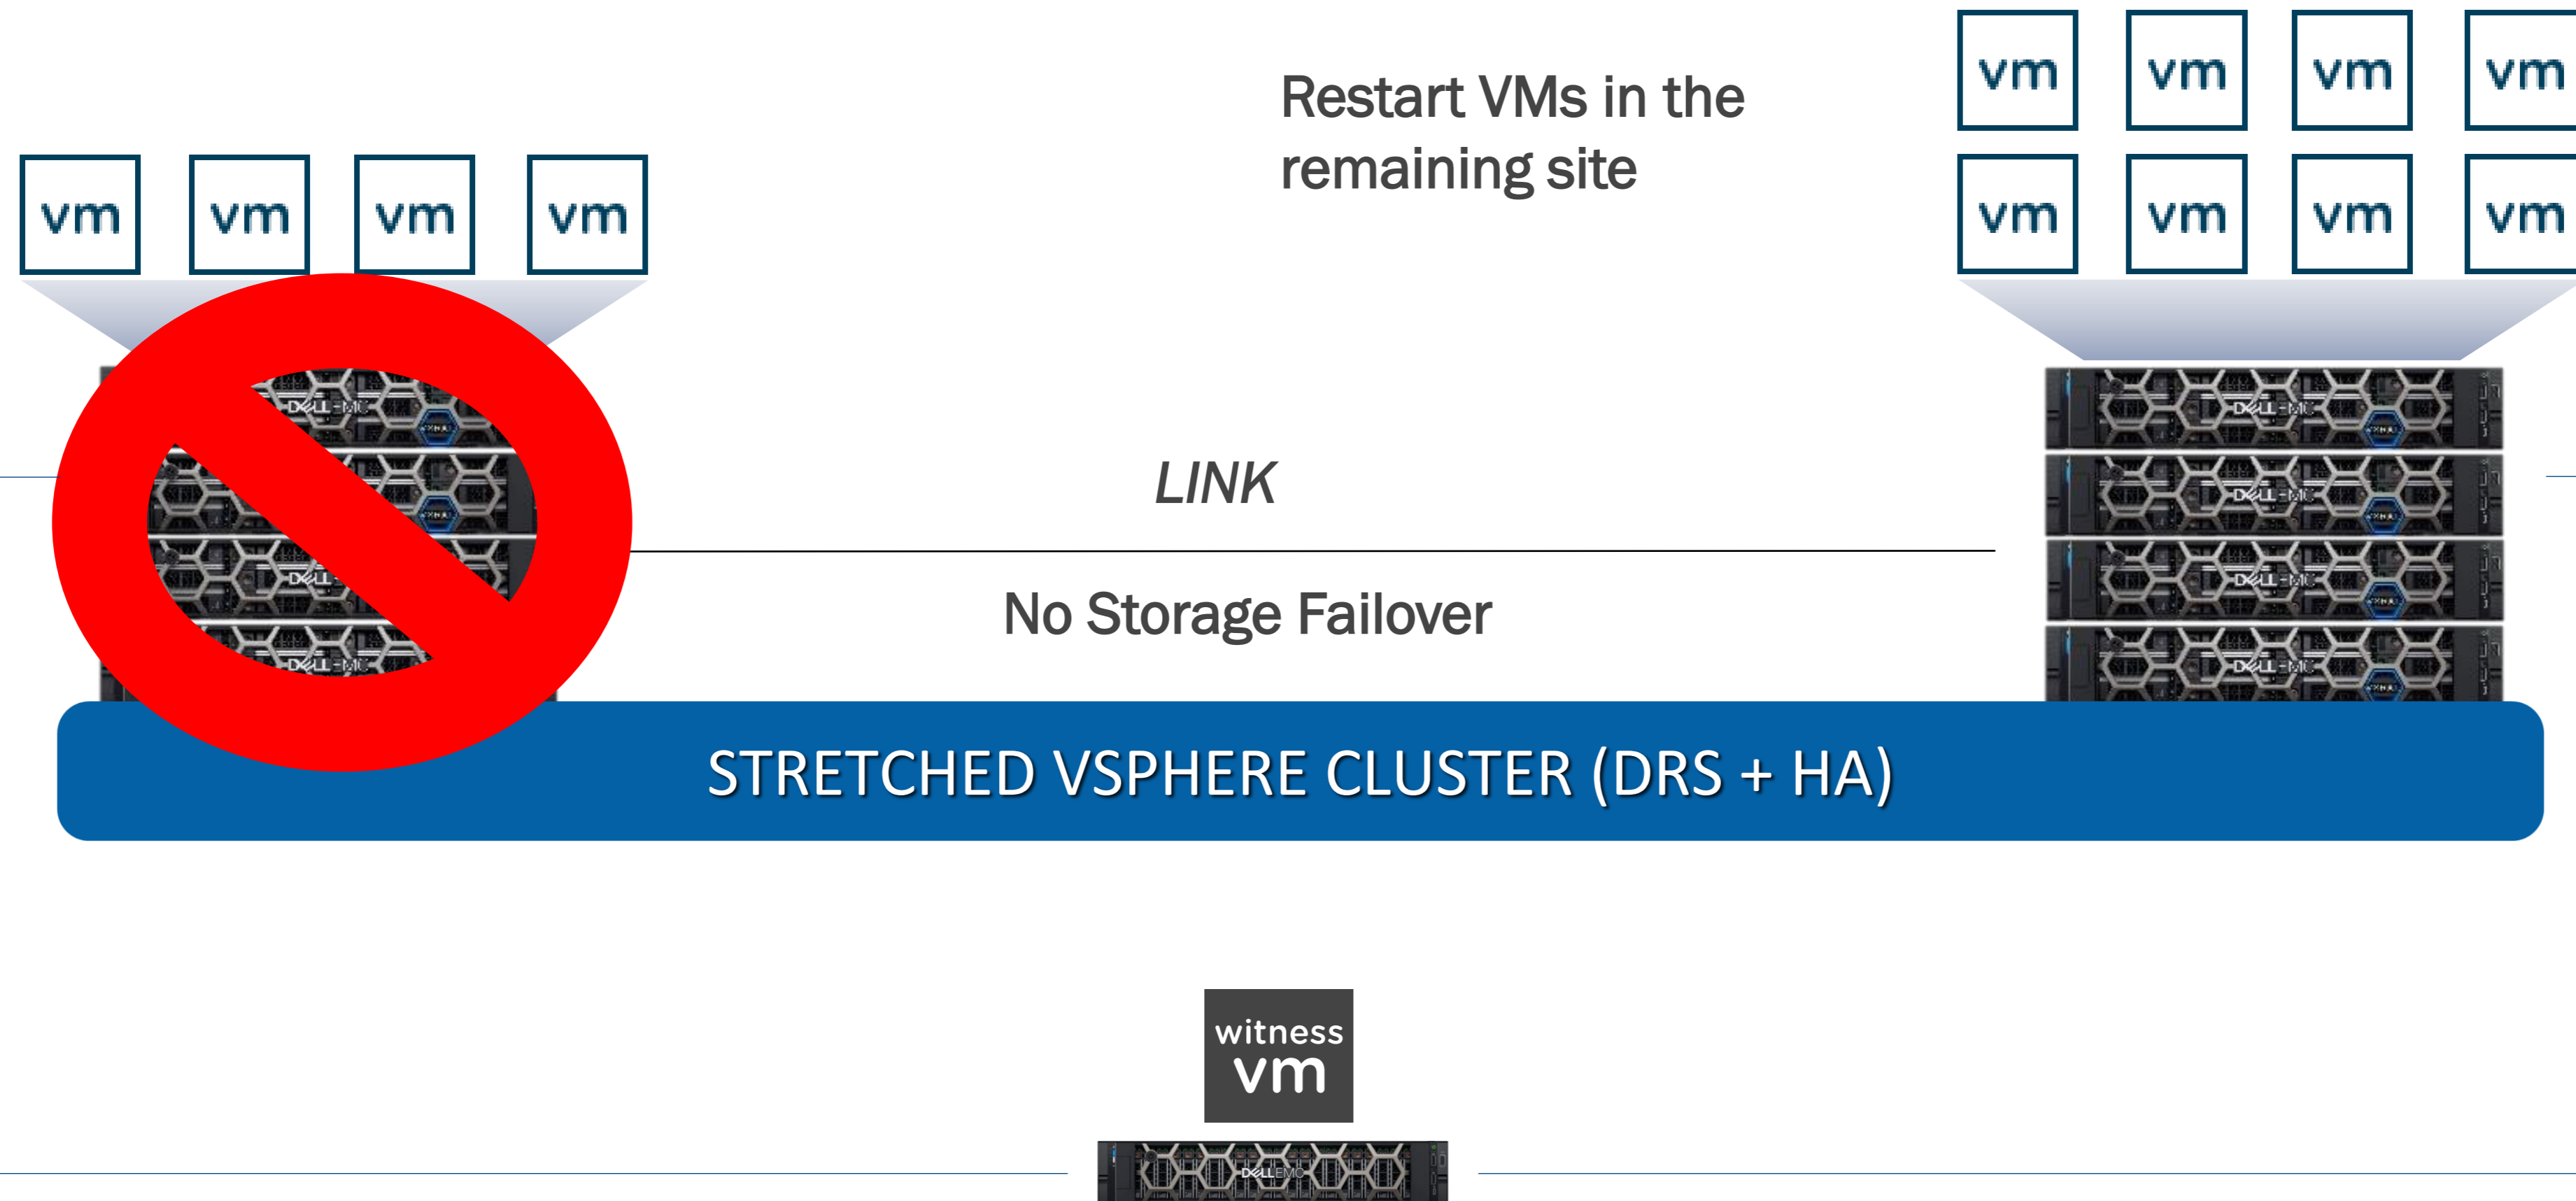
Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror



On the following Slides you will see the scenarios of a Stretched Cluster with PFTT = 1, SFTT = 1 and FTM = Mirror

Site Failure

Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror



If a whole site fails all VMs are restarted in the remaining site using vSphere HA. vSAN Data is still 100% protected due to SFTT=1.

vSAN:  Not Compliant

Node Failure

Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror

Restart VMs on the remaining hosts



Copy missing components



LINK

STRETCHED VSPHERE CLUSTER (DRS + HA)

witness
vm

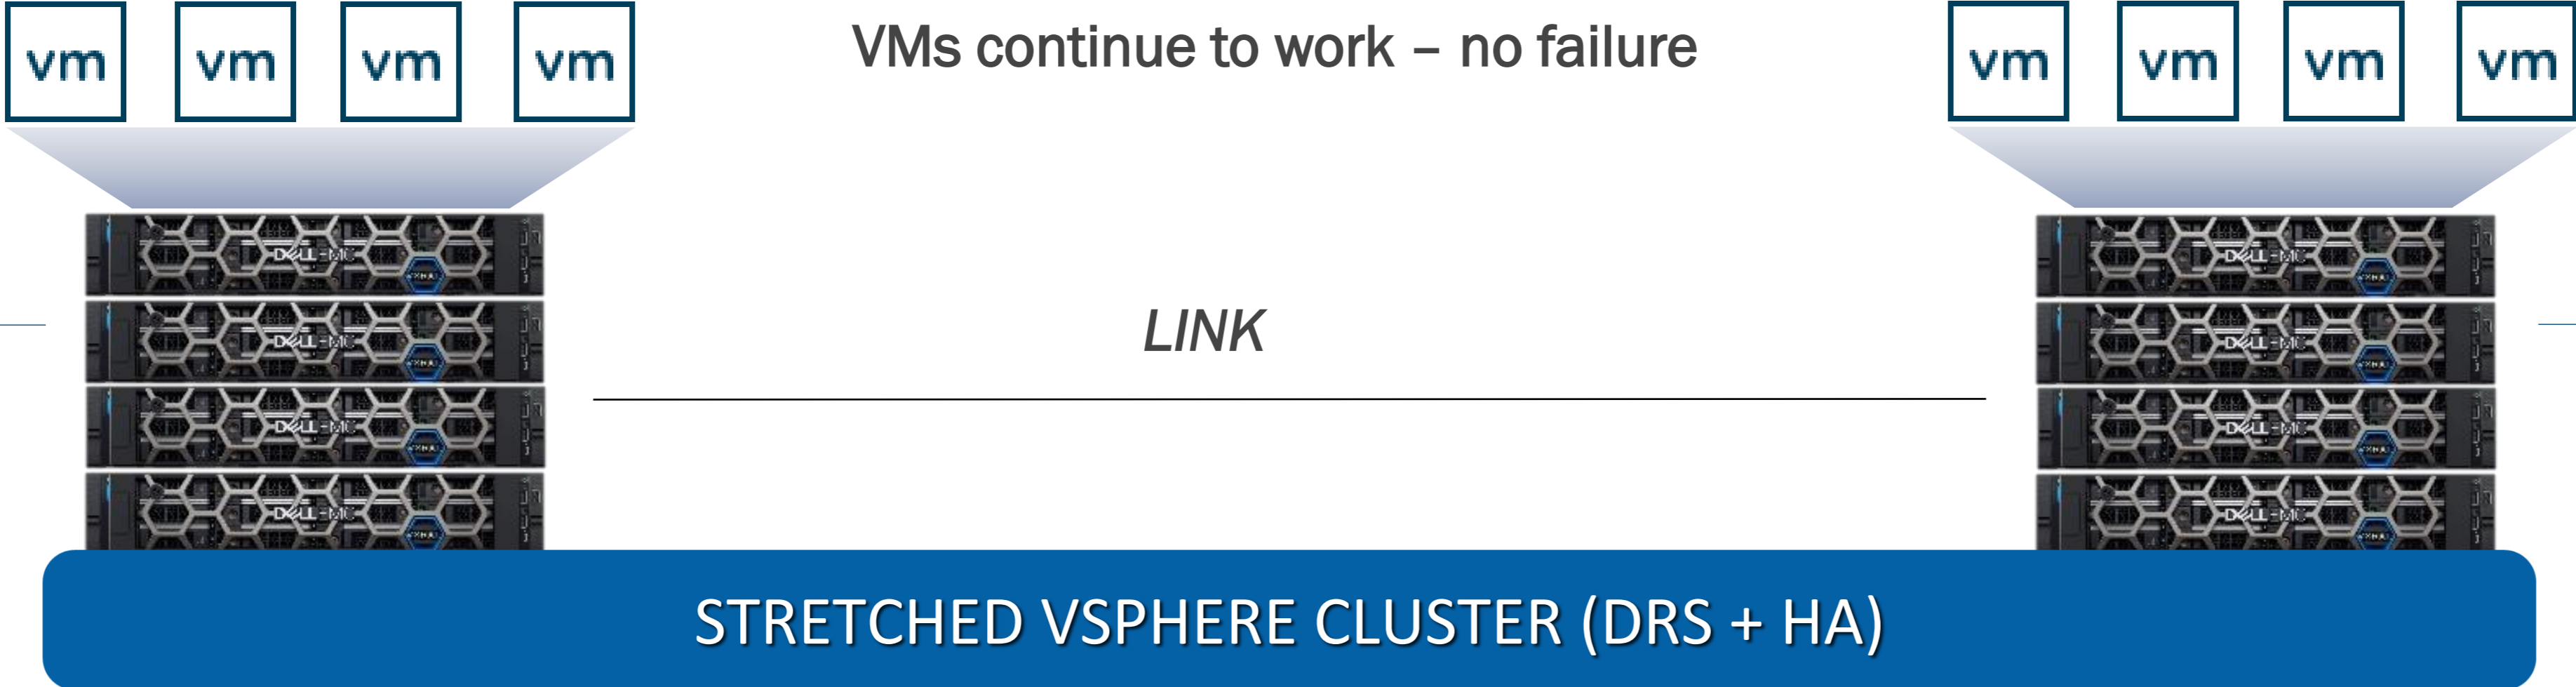


If a Node fails all VMs on this are restarted on the remaining hosts using vSphere HA. vSAN will copy missing components to remaining hosts.

vSAN status:  Not Compliant

Witness Failure

Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror



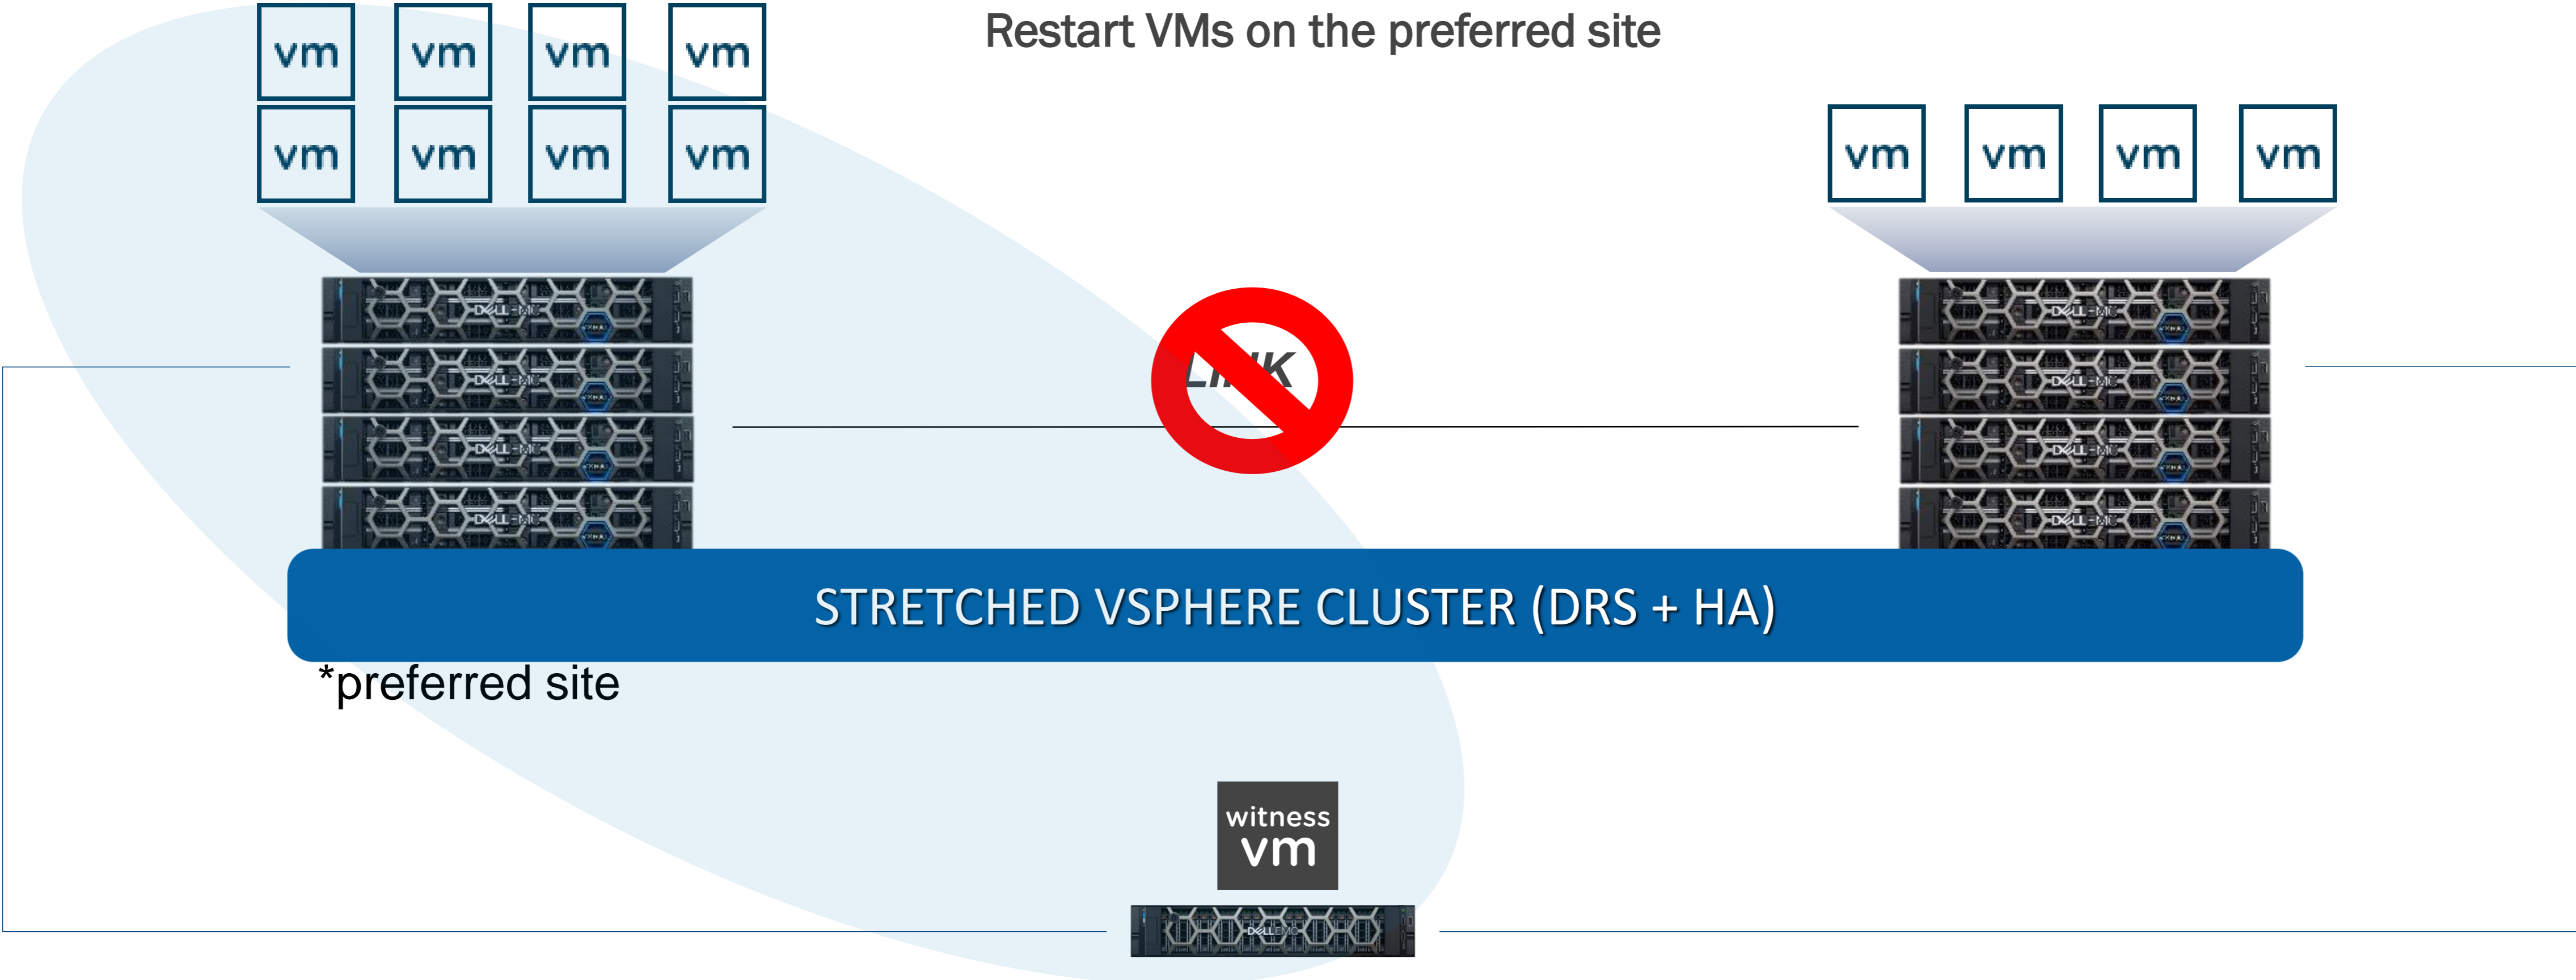
If the witness fails vCenter will issue an error / warning but VMs will continue to work normally.

vSAN status:  Not Compliant

VxRail Link Failure

Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror

Restart VMs on the preferred site

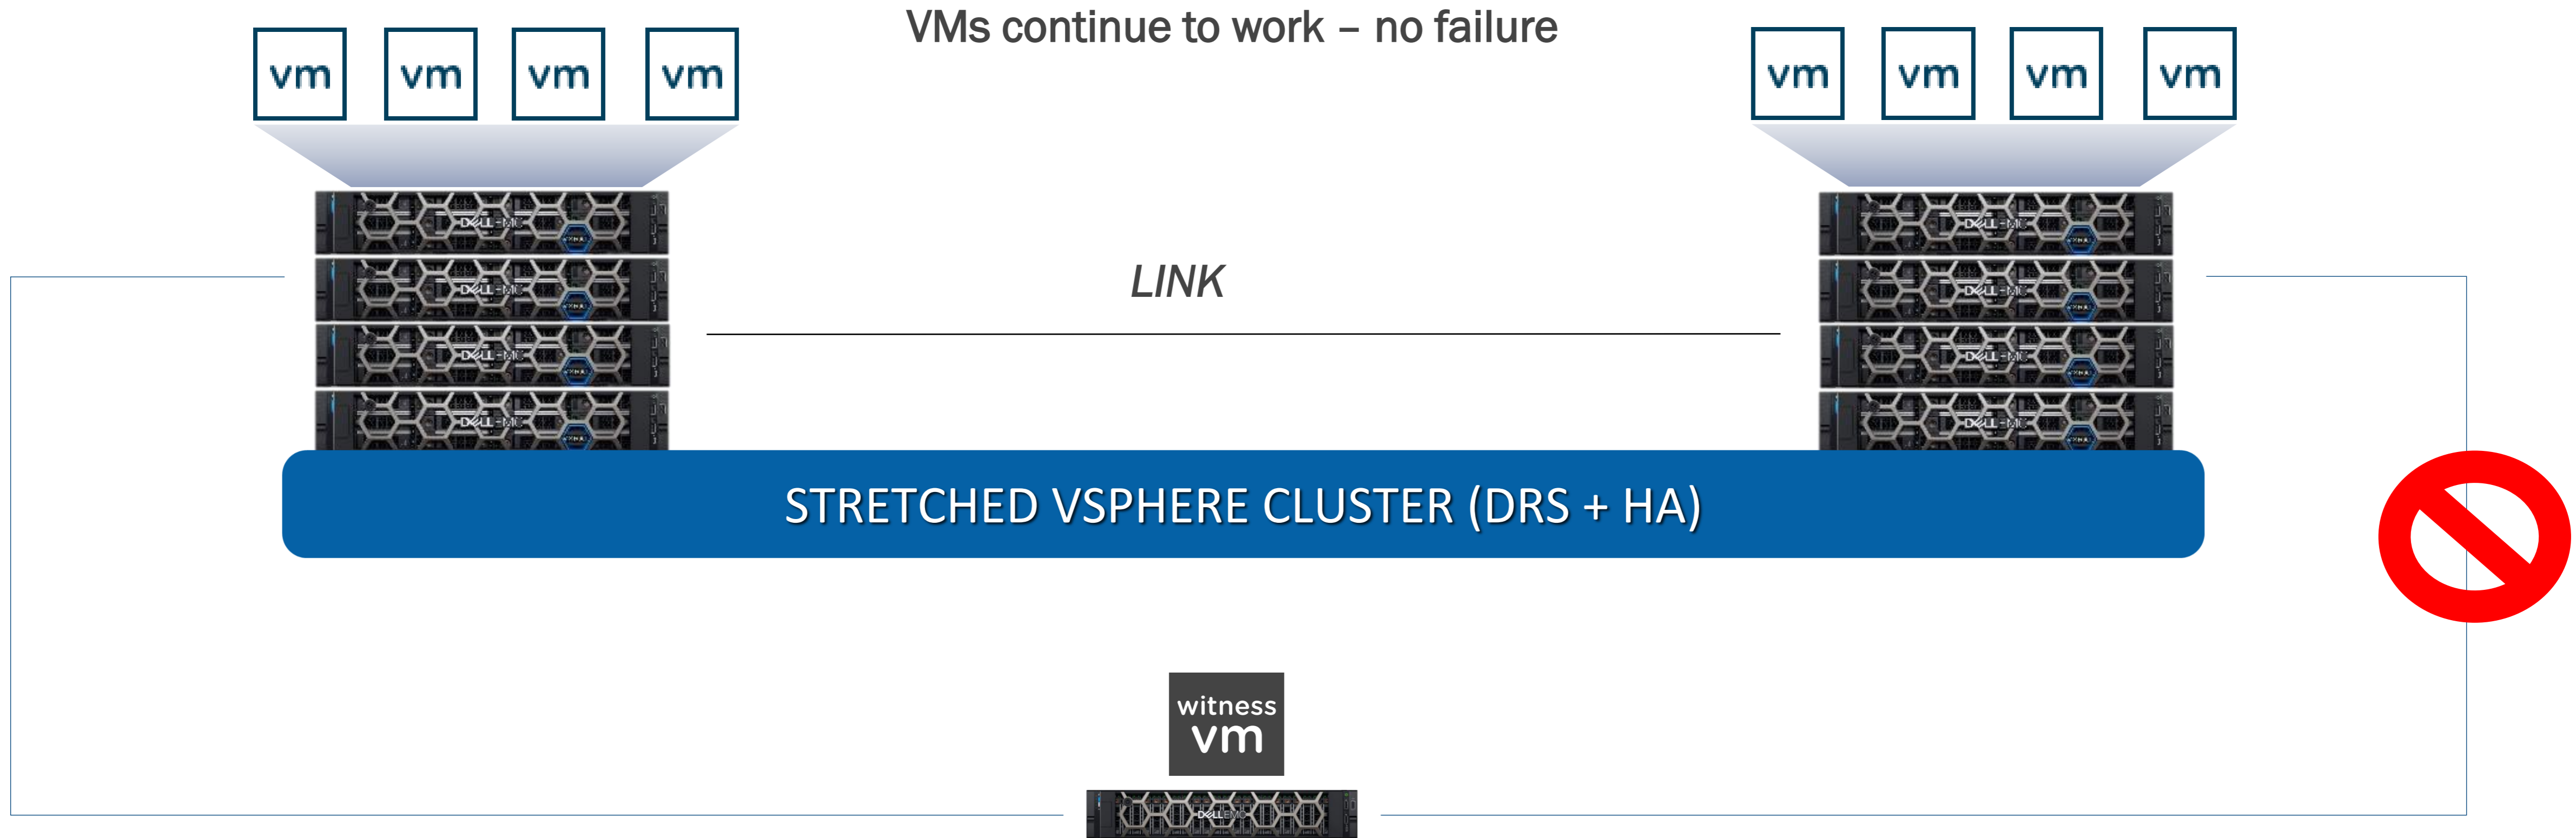


If the link between the data sites fails, the Witness will create a cluster with the preferred site and continue working. VMs from unpreferred site will restart on the preferred site.

vSAN status:  Not Compliant

Witness Link Failure

Stretched Cluster, PFTT = 1, SFTT = 1, FTM = Mirror




If a link to the witness fails, vCenter will issue an error / warning but VMs will continue to work normally.

vSAN status:  Not Compliant

VxRail – Absent / Degraded

Back

Single Site, FTT=1, FTM = Mirror




On the following slides you will see an explanation of Absent & Degraded Status in a Single Site scenario with FTT = 1 and FTM = Mirror

*example of component layout

VxRail - Reboot Host

Single Site, FTT=1, FTM = Mirror

Back




Choose between:



Node needs more than 60min

Nodes successfully reboots

Host unavailable



The host will be unavailable. Running VMs will restart on another host using HA. Data will be declared as "Absent" (see unplug drive).


*example of component layout

vSAN status: Not Compliant



VxRail – Reboot Host Successful

Single Site, FTT=1, FTM = Mirror




Host will become available again, vSAN will copy the delta of the data and will continue working. If DRS is enabled – VMs will be balanced across the nodes.

*example of component layout

VxRail – Reboot Host Wait 60+ Mins

Back

Single Site, FTT=1, FTM = Mirror



If the host is not coming back within 60min vSAN will declare the host as degraded and will start to copy data.

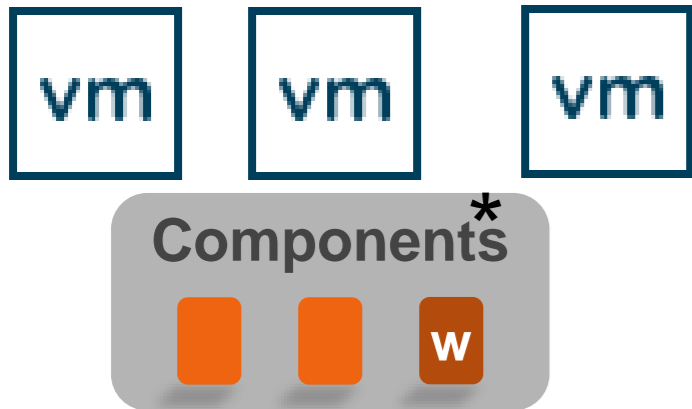
*example of component layout

vSAN status:  Not Compliant



VxRail - Hot Unplug Disk

Single Site, FTT=1, FTM = Mirror




Wait for 60+ min

Choose between:



reinsert drive within 60min

Component	Absent
-----------	--------



If you hot unplug a disk, the components on this disk will be marked as "Absent".
 vSAN will wait for 60 mins** for the disk to come back. If not, vSAN will start to rebuild the component.

*example of component layout
 **default value / possible to change


vSAN status: Not Compliant



VxRail – Plug Back in within 60 Mins

Back

Single Site, FTT=1, FTM = Mirror



vSAN will copy the delta of the data and will continue working.

*example of component layout


vSAN status: Not Compliant



VxRail – Does Not Plug Back in within 60 Mins

Back

Single Site, FTT=1, FTM = Mirror



vSAN will mark the drive as failed and will copy the missing components to a different drive.


*example of component layout

vSAN status:  Not Compliant



How do we get there?

Sterling provides end-to-end solutions.



ASSESSMENTS

- HCI Assessment
- vRNI Network Assessment
- vOA Assessment
- Tech Profile
- Product Demonstrations



DESIGN

- Gather Business Requirements
- Properly size based on assessment
- Provide options that align with business goals.



DEPLOY & IMPLEMENT

- Sterling Services
- VMware PSO
- VMware VVD
- VMC on AWS SDDCaaS



MANAGE & TRAINING

- Consulting
- Health Checks
- Staff Augmentation
- Project Management
- Training

Thank You!