



VMware

Enterprise applications

NetApp
April 25, 2024

Table of Contents

- VMware 1
 - VMware vSphere with ONTAP 1
 - Virtual Volumes (vVols) with ONTAP 39
 - VMware Site Recovery Manager with ONTAP 60
 - vSphere Metro Storage Cluster with ONTAP 79
 - Product Security 109

VMware

VMware vSphere with ONTAP

VMware vSphere with ONTAP

ONTAP has been a leading storage solution for VMware vSphere environments for almost two decades and continues to add innovative capabilities to simplify management while reducing costs. This document introduces the ONTAP solution for vSphere, including the latest product information and best practices, to streamline deployment, reduce risk, and simplify management.



This documentation replaces previously published technical reports *TR-4597: VMware vSphere for ONTAP*

Best practices supplement other documents such as guides and compatibility lists. They are developed based on lab testing and extensive field experience by NetApp engineers and customers. They might not be the only supported practices that work in every environment, but they are generally the simplest solutions that meet the needs of most customers.

This document is focused on capabilities in recent releases of ONTAP (9.x) running on vSphere 7.0 or later. See the [NetApp Interoperability Matrix Tool](#) and [VMware Compatibility Guide](#) for details related to specific releases.

Why ONTAP for vSphere?

There are many reasons why tens of thousands of customers have selected ONTAP as their storage solution for vSphere, such as a unified storage system supporting both SAN and NAS protocols, robust data protection capabilities using space-efficient snapshots and a wealth of tools to help you manage application data. Using a storage system separate from the hypervisor allows you to offload many functions and maximize your investment in vSphere host systems. This approach not only makes sure your host resources are focused on application workloads, but it also avoids random performance effects on applications from storage operations.

Using ONTAP together with vSphere is a great combination that lets you reduce host hardware and VMware software expenses. You can also protect your data at lower cost with consistent high performance. Because virtualized workloads are mobile, you can explore different approaches using Storage vMotion to move VMs across VMFS, NFS, or vVols datastores, all on the same storage system.

Here are key factors customers value today:

- **Unified storage.** Systems running ONTAP software are unified in several significant ways. Originally this approach referred to both NAS and SAN protocols, and ONTAP continues to be a leading platform for SAN along with its original strength in NAS. In the vSphere world, this approach could also mean a unified system for virtual desktop infrastructure (VDI) together with virtual server infrastructure (VSI). Systems running ONTAP software are typically less expensive for VSI than traditional enterprise arrays and yet have advanced storage efficiency capabilities to handle VDI in the same system. ONTAP also unifies a variety of storage media, from SSDs to SATA, and can extend that easily into the cloud. There's no need to buy one flash array for performance, a SATA array for archives, and separate systems for the cloud. ONTAP ties them all together.
- **Virtual volumes and storage policy based management.** NetApp was an early design partner with VMware in the development of vSphere Virtual Volumes (vVols), providing architectural input and early

support for vVols and VMware vSphere APIs for Storage Awareness (VASA). Not only did this approach bring granular VM storage management to VMFS, it also supported automation of storage provisioning through storage policy based management. This approach allows storage architects to design storage pools with different capabilities that can be easily consumed by VM administrators. ONTAP leads the storage industry in vVol scale, supporting hundreds of thousands of vVols in a single cluster, whereas enterprise array and smaller flash array vendors support as few as several thousand vVols per array. NetApp is also driving the evolution of granular VM management with upcoming capabilities in support of vVols 3.0.


- **Storage efficiency.** Although NetApp was the first to deliver deduplication for production workloads, this innovation wasn't the first or last one in this area. It started with snapshots, a space-efficient data protection mechanism with no performance effect, along with FlexClone technology to instantly make read/write copies of VMs for production and backup use. NetApp went on to deliver inline capabilities, including deduplication, compression, and zero-block deduplication, to squeeze out the most storage from expensive SSDs. Most recently, ONTAP added the ability to pack smaller I/O operations and files into a disk block using compaction. The combination of these capabilities has resulted in customers seeing savings of up to 5:1 for VSI and up to 30:1 for VDI.
- **Hybrid cloud.** Whether used for on-premises private cloud, public cloud infrastructure, or a hybrid cloud that combines the best of both, ONTAP solutions help you build your data fabric to streamline and optimize data management. Start with high-performance all-flash systems, then couple them with either disk or cloud storage systems for data protection and cloud compute. Choose from Azure, AWS, IBM, or Google clouds to optimize costs and avoid lock-in. Leverage advanced support for OpenStack and container technologies as needed. NetApp also offers cloud-based backup (SnapMirror Cloud, Cloud Backup Service, and Cloud Sync) and storage tiering and archiving tools (FabricPool) for ONTAP to help reduce operating expenses and leverage the broad reach of the cloud.
- **And more.** Take advantage of the extreme performance of NetApp AFF A-Series arrays to accelerate your virtualized infrastructure while managing costs. Enjoy completely nondisruptive operations, from maintenance to upgrades to complete replacement of your storage system, using scale-out ONTAP clusters. Protect data at rest with NetApp encryption capabilities at no additional cost. Make sure performance meets business service levels through fine-grained quality of service capabilities. They are all part of the broad range of capabilities that come with ONTAP, the industry's leading enterprise data management software.

Unified Storage

NetApp ONTAP unifies storage through a simplified, software-defined approach for secure and efficient management, improved performance, and seamless scalability. This approach enhances data protection and enables effective use of cloud resources.

Originally this unified approach referred to supporting both NAS and SAN protocols on one storage system, and ONTAP continues to be a leading platform for SAN along with its original strength in NAS. ONTAP now also provides S3 object protocol support. Though S3 isn't used for datastores, you can use it for in-guest applications. You can learn more about the S3 protocol support in ONTAP in the [S3 configuration overview](#).

A storage virtual machine (SVM) is the unit of secure multi-tenancy in ONTAP. It is a logical construct allowing client access to systems running ONTAP software. SVMs can serve data concurrently through multiple data access protocols via logical interfaces (LIFs). SVMs provide file-level data access through NAS protocols, such as CIFS and NFS, and block-level data access through SAN protocols, such as iSCSI, FC/FCoE, and NVMe. SVMs can serve data to SAN and NAS clients independently at the same time, as well as with S3.



In the vSphere world, this approach could also mean a unified system for virtual desktop infrastructure (VDI) together with virtual server infrastructure (VSI). Systems running ONTAP software are typically less expensive for VSI than traditional enterprise arrays and yet have advanced storage efficiency capabilities to handle VDI in the same system. ONTAP also unifies a variety of storage media, from SSDs to SATA, and can extend that easily into the cloud. There's no need to buy one flash array for performance, a SATA array for archives, and separate systems for the cloud. ONTAP ties them all together.

NOTE: For more information on SVMs, unified storage and client access, see [Storage Virtualization](#) in the ONTAP 9 Documentation center.

Virtualization tools for ONTAP

NetApp offers several standalone software tools that can be used together with ONTAP and vSphere to manage your virtualized environment.

The following tools are included with the ONTAP license at no additional cost. See Figure 1 for a depiction of how these tools work together in your vSphere environment.

ONTAP tools for VMware vSphere


ONTAP tools for VMware vSphere is a set of tools for using ONTAP storage together with vSphere. The vCenter plug-in, formerly known as the Virtual Storage Console (VSC), simplifies storage management and efficiency features, enhances availability, and reduces storage costs and operational overhead, whether you are using SAN or NAS. It uses best practices for provisioning datastores and optimizes ESXi host settings for NFS and block storage environments. For all these benefits, NetApp recommends using these ONTAP tools as a best practice when using vSphere with systems running ONTAP software. It includes a server appliance, user interface extensions for vCenter, VASA Provider, and Storage Replication Adapter. Nearly everything in ONTAP tools can be automated by using simple REST APIs, consumable by most modern automation tools.

- **vCenter UI extensions.** The ONTAP tools UI extensions simplify the job of operations teams and vCenter

admins by embedding easy-to-use, context-sensitive menus for managing hosts and storage, informational portlets, and native alerting capabilities directly in the vCenter UI for streamlined workflows.

- **VASA Provider for ONTAP.** The VASA Provider for ONTAP supports the VMware vStorage APIs for Storage Awareness (VASA) framework. It is supplied as part of ONTAP tools for VMware vSphere as a single virtual appliance for ease of deployment. VASA Provider connects vCenter Server with ONTAP to aid in provisioning and monitoring VM storage. It enables VMware Virtual Volumes (vVols) support, management of storage capability profiles and individual VM vVols performance, and alarms for monitoring capacity and compliance with the profiles.
- **Storage Replication Adapter.** The SRA is used together with VMware Site Recovery Manager (SRM) to manage data replication between production and disaster recovery sites and test the DR replicas nondisruptively. It helps automate the tasks of discovery, recovery, and reprotection. It includes both an SRA server appliance and SRA adapters for the Windows SRM server and SRM appliance.

The following figure depicts ONTAP tools for vSphere.



NFS Plug-In for VMware VAAI

The NetApp NFS Plug-In for VMware VAAI is a plug-in for ESXi hosts that allows them to use VAAI features with NFS datastores on ONTAP. It supports copy offload for clone operations, space reservation for thick virtual disk files, and snapshot offload. Offloading copy operations to storage is not necessarily faster to complete, but it does reduce network bandwidth requirements and offloads host resources such as CPU cycles, buffers, and queues. You can use ONTAP tools for VMware vSphere to install the plug-in on ESXi hosts or, where supported, vSphere Lifecycle Manager (vLCM).

Virtual Volumes (vVols) and Storage Policy Based Management (SPBM)

NetApp was an early design partner with VMware in the development of vSphere Virtual Volumes (vVols), providing architectural input and early support for vVols and VMware vSphere APIs for Storage Awareness (VASA). Not only did this approach bring VM granular storage management to VMFS, it also supported automation of storage provisioning through Storage Policy based Management (SPBM).

SPBM provides a framework that serves as an abstraction layer between the storage services available to your virtualization environment and the provisioned storage elements via policies. This approach allows storage architects to design storage pools with different capabilities that can be easily consumed by VM administrators. Administrators can then match virtual machine workload requirements against the provisioned storage pools, allowing for granular control of various settings on a per-VM or virtual disk level.

ONTAP leads the storage industry in vVols scale, supporting hundreds of thousands of vVols in a single cluster, whereas enterprise array and smaller flash array vendors support as few as several thousand vVols per array. NetApp is also driving the evolution of VM granular management with upcoming capabilities in support of vVols 3.0.



For more information on VMware vSphere Virtual Volumes, SPBM, and ONTAP, see [TR-4400: VMware vSphere Virtual Volumes with ONTAP](#).

Datstores and protocols

vSphere datastore and protocol features overview

Seven protocols are used to connect VMware vSphere to datastores on a system running ONTAP software:

- FCP
- FCoE
- NVMe/FC
- NVMe/TCP
- iSCSI
- NFS v3
- NFS v4.1

FCP, FCoE, NVMe/FC, NVMe/TCP, and iSCSI are block protocols that use the vSphere Virtual Machine File System (VMFS) to store VMs inside ONTAP LUNs or NVMe namespaces that are contained in an ONTAP FlexVol volume. Note that, starting from vSphere 7.0, VMware no longer supports software FCoE in production environments. NFS is a file protocol that places VMs into datastores (which are simply ONTAP volumes) without the need for VMFS. SMB (CIFS), iSCSI, NVMe/TCP, or NFS can also be used directly from a guest OS to ONTAP.

The following tables present vSphere-supported traditional datastore features with ONTAP. This information does not apply to vVols datastores, but it does generally apply to vSphere 6.x and later releases using supported ONTAP releases. You can also consult [VMware configuration maximums](#) for specific vSphere releases to confirm specific limits.

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
Format	VMFS or raw device mapping (RDM)	VMFS or RDM	VMFS	N/A
Maximum number of datastores or LUNs	1024 LUNs per host	1024 LUNs per server	256 Namespaces per server	256 mounts Default NFS. MaxVolumes is 8. Use ONTAP tools for VMware vSphere to increase to 256.
Maximum datastore size	64TB	64TB	64TB	100TB FlexVol volume or greater with FlexGroup volume
Maximum datastore file size	62TB	62TB	62TB	62TB with ONTAP 9.12.1P2 and later
Optimal queue depth per LUN or file system	64-256	64-256	Autonegotiated	Refer to NFS.MaxQueueDepth in Recommended ESXi host and other ONTAP settings .

The following table lists supported VMware storage-related functionalities.

Capacity/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
vMotion	Yes	Yes	Yes	Yes
Storage vMotion	Yes	Yes	Yes	Yes
VMware HA	Yes	Yes	Yes	Yes
Storage Distributed Resource Scheduler (SDRS)	Yes	Yes	Yes	Yes
VMware vStorage APIs for Data Protection (VADP)-enabled backup software	Yes	Yes	Yes	Yes
Microsoft Cluster Service (MSCS) or failover clustering within a VM	Yes	Yes*	Yes*	Not supported
Fault Tolerance	Yes	Yes	Yes	Yes
Site Recovery Manager	Yes	Yes	No**	V3 only**

Capacity/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
Thin-provisioned VMs (virtual disks)	Yes	Yes	Yes	Yes This setting is the default for all VMs on NFS when not using VAAI.
VMware native multipathing	Yes	Yes	Yes, using the new High Performance Plugin (HPP)	NFS v4.1 session trunking requires ONTAP 9.14.1 and later

The following table lists supported ONTAP storage management features.

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
Data deduplication	Savings in the array	Savings in the array	Savings in the array	Savings in the datastore
Thin provisioning	Datastore or RDM	Datastore or RDM	Datastore	Datastore
Resize datastore	Grow only	Grow only	Grow only	Grow, autogrow, and shrink
SnapCenter plug-ins for Windows, Linux applications (in guest)	Yes	Yes	No	Yes
Monitoring and host configuration using ONTAP tools for VMware vSphere	Yes	Yes	No	Yes
Provisioning using ONTAP tools for VMware vSphere	Yes	Yes	No	Yes

The following table lists supported backup features.

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
ONTAP Snapshots	Yes	Yes	Yes	Yes
SRM supported by replicated backups	Yes	Yes	No**	V3 only**
Volume SnapMirror	Yes	Yes	Yes	Yes
VMDK image access	VADP-enabled backup software	VADP-enabled backup software	VADP-enabled backup software	VADP-enabled backup software, vSphere Client, and vSphere Web Client datastore browser

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
VMDK file-level access	VADP-enabled backup software, Windows only	VADP-enabled backup software, Windows only	VADP-enabled backup software, Windows only	VADP-enabled backup software and third-party applications
NDMP granularity	Datastore	Datastore	Datastore	Datastore or VM

*NetApp recommends using in-guest iSCSI for Microsoft clusters rather than multiwriter-enabled VMDKs in a VMFS datastore. This approach is fully supported by Microsoft and VMware, offers great flexibility with ONTAP (SnapMirror to ONTAP systems on-premises or in the cloud), is easy to configure and automate, and can be protected with SnapCenter. vSphere 7 adds a new clustered VMDK option. This is different from multiwriter-enabled VMDKs, which requires a datastore presented via the FC protocol that has clustered VMDK support enabled. Other restrictions apply. See VMware's [Setup for Windows Server Failover Clustering](#) documentation for configuration guidelines.

**Datastores using NVMe-oF and NFS v4.1 require vSphere replication. Array-based replication is not supported by SRM.

Selecting a storage protocol

Systems running ONTAP software support all major storage protocols, so customers can choose what is best for their environment, depending on existing and planned networking infrastructure and staff skills. NetApp testing has generally shown little difference between protocols running at similar line speeds, so it is best to focus on your network infrastructure and staff capabilities over raw protocol performance.

The following factors might be useful in considering a choice of protocol:

- **Current customer environment.** Although IT teams are generally skilled at managing Ethernet IP infrastructure, not all are skilled at managing an FC SAN fabric. However, using a general-purpose IP network that's not designed for storage traffic might not work well. Consider the networking infrastructure you have in place, any planned improvements, and the skills and availability of staff to manage them.
- **Ease of setup.** Beyond initial configuration of the FC fabric (additional switches and cabling, zoning, and the interoperability verification of HBA and firmware), block protocols also require creation and mapping of LUNs and discovery and formatting by the guest OS. After the NFS volumes are created and exported, they are mounted by the ESXi host and ready to use. NFS has no special hardware qualification or firmware to manage.
- **Ease of management.** With SAN protocols, if more space is needed, several steps are necessary, including growing a LUN, rescanning to discover the new size, and then growing the file system). Although growing a LUN is possible, reducing the size of a LUN is not, and recovering unused space can require additional effort. NFS allows easy sizing up or down, and this resizing can be automated by the storage system. SAN offers space reclamation through guest OS TRIM/UNMAP commands, allowing space from deleted files to be returned to the array. This type of space reclamation is more difficult with NFS datastores.
- **Storage space transparency.** Storage utilization is typically easier to see in NFS environments because thin provisioning returns savings immediately. Likewise, deduplication and cloning savings are immediately available for other VMs in the same datastore or for other storage system volumes. VM density is also typically greater in an NFS datastore, which can improve deduplication savings as well as reduce management costs by having fewer datastores to manage.

Datstore layout

ONTAP storage systems offer great flexibility in creating datstores for VMs and virtual disks. Although many ONTAP best practices are applied when using the VSC to provision datstores for vSphere (listed in the section [Recommended ESXi host and other ONTAP settings](#)), here are some additional guidelines to consider:

- Deploying vSphere with ONTAP NFS datstores results in a high-performing, easy-to-manage implementation that provides VM-to-datstore ratios that cannot be obtained with block-based storage protocols. This architecture can result in a tenfold increase in datstore density with a correlating reduction in the number of datstores. Although a larger datstore can benefit storage efficiency and provide operational benefits, consider using at least four datstores (FlexVol volumes) to store your VMs on a single ONTAP controller to get maximum performance from the hardware resources. This approach also allows you to establish datstores with different recovery policies. Some can be backed up or replicated more frequently than others based on business needs. Multiple datstores are not required with FlexGroup volumes for performance because they scale by design.
- NetApp recommends the use of FlexVol volumes for most NFS datstores. Starting with ONTAP 9.8 FlexGroup volumes are supported for use as datstores as well, and are generally recommended for certain use cases. Other ONTAP storage containers such as qtrees are not generally recommended because these are not currently supported by either ONTAP tools for VMware vSphere or the NetApp SnapCenter plugin for VMware vSphere. That being said, deploying datstores as multiple qtrees in a single volume might be useful for highly automated environments that can benefit from datstore-level quotas or VM file clones.
- A good size for a FlexVol volume datstore is around 4TB to 8TB. This size is a good balance point for performance, ease of management, and data protection. Start small (say, 4TB) and grow the datstore as needed (up to the maximum 100TB). Smaller datstores are faster to recover from backup or after a disaster and can be moved quickly across the cluster. Consider the use of ONTAP autosize to automatically grow and shrink the volume as used space changes. The ONTAP tools for VMware vSphere Datstore Provisioning Wizard use autosize by default for new datstores. Additional customization of the grow and shrink thresholds and maximum and minimum size can be done with System Manager or the command line.
- Alternately, VMFS datstores can be configured with LUNs that are accessed by FC, iSCSI, or FCoE. VMFS allows traditional LUNs to be accessed simultaneously by every ESX server in a cluster. VMFS datstores can be up to 64TB in size and consist of up to 32 2TB LUNs (VMFS 3) or a single 64TB LUN (VMFS 5). The ONTAP maximum LUN size is 16TB on most systems, and 128TB on All-SAN-Array systems. Therefore, a maximum size VMFS 5 datstore on most ONTAP systems can be created by using four 16TB LUNs. While there can be a performance benefit for high-I/O workloads with multiple LUNs (with high-end FAS or AFF systems), this benefit is offset by added management complexity to create, manage, and protect the datstore LUNs and increased availability risk. NetApp generally recommends using a single, large LUN for each datstore and only span if there is a special need to go beyond a 16TB datstore. As with NFS, consider using multiple datstores (volumes) to maximize performance on a single ONTAP controller.
- Older guest operating systems (OSs) needed alignment with the storage system for best performance and storage efficiency. However, modern vendor-supported OSs from Microsoft and Linux distributors such as Red Hat no longer require adjustments to align the file system partition with the blocks of the underlying storage system in a virtual environment. If you are using an old OS that might require alignment, search the NetApp Support Knowledgebase for articles using “VM alignment” or request a copy of TR-3747 from a NetApp sales or partner contact.
- Avoid the use of defragmentation utilities within the guest OS, as this offers no performance benefit and affects storage efficiency and snapshot space usage. Also consider turning off search indexing in the guest OS for virtual desktops.
- ONTAP has led the industry with innovative storage efficiency features, allowing you to get the most out of your usable disk space. AFF systems take this efficiency further with default inline deduplication and

compression. Data is deduplicated across all volumes in an aggregate, so you no longer need to group similar operating systems and similar applications within a single datastore to maximize savings.

- In some cases, you might not even need a datastore. For the best performance and manageability, avoid using a datastore for high-I/O applications such as databases and some applications. Instead, consider guest-owned file systems such as NFS or iSCSI file systems managed by the guest or with RDMS. For specific application guidance, see NetApp technical reports for your application. For example, [Oracle Databases on ONTAP](#) has a section about virtualization with helpful details.
- First Class Disks (or Improved Virtual Disks) allow for vCenter-managed disks independent of a VM with vSphere 6.5 and later. While primarily managed by API, they can be useful with vVols, especially when managed by OpenStack or Kubernetes tools. They are supported by ONTAP as well as ONTAP tools for VMware vSphere.

Datastore and VM migration

When migrating VMs from an existing datastore on another storage system to ONTAP, here are some practices to keep in mind:

- Use Storage vMotion to move the bulk of your virtual machines to ONTAP. Not only is this approach nondisruptive to running VMs, it also allows ONTAP storage efficiency features such as inline deduplication and compression to process the data as it migrates. Consider using vCenter capabilities to select multiple VMs from the inventory list and then schedule the migration (use Ctrl key while clicking Actions) at an appropriate time.
- While you could carefully plan a migration to appropriate destination datastores, it is often simpler to migrate in bulk and then organize later as needed. You might want to use this approach to guide your migration to different datastores if you have specific data protection needs, such as different Snapshot schedules.
- Most VMs and their storage may be migrated while running (hot), but migrating attached (not in datastore) storage such as ISOs, LUNs, or NFS volumes from another storage system might require cold migration.
- Virtual machines that need more careful migration include databases and applications that use attached storage. In general, consider the use of the application's tools to manage migration. For Oracle, consider using Oracle tools such as RMAN or ASM to migrate the database files. See [TR-4534](#) for more information. Likewise, for SQL Server, consider using either SQL Server Management Studio or NetApp tools such as SnapManager for SQL Server or SnapCenter.

ONTAP tools for VMware vSphere

The most important best practice when using vSphere with systems running ONTAP software is to install and use the ONTAP tools for VMware vSphere plug-in (formerly known as Virtual Storage Console). This vCenter plug-in simplifies storage management, enhances availability, and reduces storage costs and operational overhead, whether using SAN or NAS. It uses best practices for provisioning datastores and optimizes ESXi host settings for multipath and HBA timeouts (these are described in Appendix B). Because it's a vCenter plug-in, it's available to all vSphere web clients that connect to the vCenter server.

The plug-in also helps you use other ONTAP tools in vSphere environments. It allows you to install the NFS Plug-In for VMware VAAI, which enables copy offload to ONTAP for VM cloning operations, space reservation for thick virtual disk files, and ONTAP snapshot offload.

The plug-in is also the management interface for many functions of the VASA Provider for ONTAP, supporting storage policy-based management with vVols. After ONTAP tools for VMware vSphere is registered, use it to create storage capability profiles, map them to storage, and make sure of datastore compliance with the profiles over time. The VASA Provider also provides an interface to create and manage vVol datastores.

In general, NetApp recommends using the ONTAP tools for VMware vSphere interface within vCenter to

provision traditional and vVols datastores to make sure best practices are followed.

General Networking

Configuring network settings when using vSphere with systems running ONTAP software is straightforward and similar to other network configuration. Here are some things to consider:

- Separate storage network traffic from other networks. A separate network can be achieved by using a dedicated VLAN or separate switches for storage. If the storage network shares physical paths such as uplinks, you might need QoS or additional uplink ports to make sure of sufficient bandwidth. Don't connect hosts directly to storage; use switches to have redundant paths and allow VMware HA to work without intervention. See [Direct connect networking](#) for additional information.
- Jumbo frames can be used if desired and supported by your network, especially when using iSCSI. If they are used, make sure they are configured identically on all network devices, VLANs, and so on in the path between storage and the ESXi host. Otherwise, you might see performance or connection problems. The MTU must also be set identically on the ESXi virtual switch, the VMkernel port, and also on the physical ports or interface groups of each ONTAP node.
- NetApp only recommends disabling network flow control on the cluster network ports within an ONTAP cluster. NetApp makes no other recommendations for best practices for the remaining network ports used for data traffic. You should enable or disable as necessary. See [TR-4182](#) for more background on flow control.
- When ESXi and ONTAP storage arrays are connected to Ethernet storage networks, NetApp recommends configuring the Ethernet ports to which these systems connect as Rapid Spanning Tree Protocol (RSTP) edge ports or by using the Cisco PortFast feature. NetApp recommends enabling the Spanning-Tree PortFast trunk feature in environments that use the Cisco PortFast feature and that have 802.1Q VLAN trunking enabled to either the ESXi server or the ONTAP storage arrays.
- NetApp recommends the following best practices for link aggregation:
 - Use switches that support link aggregation of ports on two separate switch chassis using a multi-chassis link aggregation group approach such as Cisco's Virtual PortChannel (vPC).
 - Disable LACP for switch ports connected to ESXi unless you are using dvSwitches 5.1 or later with LACP configured.
 - Use LACP to create link aggregates for ONTAP storage systems with dynamic multimode interface groups with port or IP hash. Refer to [Network Management](#) for further guidance.
 - Use an IP hash teaming policy on ESXi when using static link aggregation (e.g., EtherChannel) and standard vSwitches, or LACP-based link aggregation with vSphere Distributed Switches. If link aggregation is not used, then use "Route based on the originating virtual port ID" instead.

The following table provides a summary of network configuration items and indicates where the settings are applied.

Item	ESXi	Switch	Node	SVM
IP address	VMkernel	No**	No**	Yes
Link aggregation	Virtual switch	Yes	Yes	No*
VLAN	VMkernel and VM port groups	Yes	Yes	No*
Flow control	NIC	Yes	Yes	No*
Spanning tree	No	Yes	No	No

Item	ESXi	Switch	Node	SVM
MTU (for jumbo frames)	Virtual switch and VMkernel port (9000)	Yes (set to max)	Yes (9000)	No*
Failover groups	No	No	Yes (create)	Yes (select)

*SVM LIFs connect to ports, interface groups, or VLAN interfaces that have VLAN, MTU, and other settings. However, the settings are not managed at the SVM level.

**These devices have IP addresses of their own for management, but these addresses are not used in the context of ESXi storage networking.

SAN (FC, FCoE, NVMe/FC, iSCSI), RDM

NetApp ONTAP provides enterprise-class block storage for VMware vSphere using iSCSI, Fibre Channel Protocol (FCP, or FC for short), and NVMe over Fabrics (NVMe-oF). The following are best practices for implementing block protocols for VM storage with vSphere and ONTAP.

In vSphere, there are three ways to use block storage LUNs:

- With VMFS datastores
- With raw device mapping (RDM)
- As a LUN accessed and controlled by a software initiator from a VM guest OS

VMFS is a high-performance clustered file system that provides datastores that are shared storage pools. VMFS datastores can be configured with LUNs accessed using FC, iSCSI, FCoE, or with NVMe namespaces accessed using the NVMe/FC or NVMe/TCP protocols. VMFS allows storage to be accessed simultaneously by every ESX server in a cluster. The maximum LUN size is generally 128TB beginning with ONTAP 9.12.1P2 (and earlier with ASA systems); therefore, a maximum-size VMFS 5 or 6 datastore of 64TB can be created by using a single LUN.

vSphere includes built-in support for multiple paths to storage devices, referred to as native multipathing (NMP). NMP can detect the type of storage for supported storage systems and automatically configures the NMP stack to support the capabilities of the storage system in use.

Both NMP and ONTAP support Asymmetric Logical Unit Access (ALUA) to negotiate optimized and nonoptimized paths. In ONTAP, an ALUA-optimized path follows a direct data path, using a target port on the node that hosts the LUN being accessed. ALUA is turned on by default in both vSphere and ONTAP. The NMP recognizes the ONTAP cluster as ALUA, and it uses the ALUA storage array type plug-in (`VMW_SATP_ALUA`) and selects the round robin path selection plug-in (`VMW_PSP_RR`).

ESXi 6 supports up to 256 LUNs and up to 1,024 total paths to LUNs. ESXi does not see any LUNs or paths beyond these limits. Assuming the maximum number of LUNs, the path limit allows four paths per LUN. In a larger ONTAP cluster, it is possible to reach the path limit before the LUN limit. To address this limitation, ONTAP supports selective LUN map (SLM) in release 8.3 and later.

SLM limits the nodes that advertise paths to a given LUN. It is a NetApp best practice to have at least one LIF per node per SVM and to use SLM to limit the paths advertised to the node hosting the LUN and its HA partner. Although other paths exist, they aren't advertised by default. It is possible to modify the paths advertised with the add and remove reporting node arguments within SLM. Note that LUNs created in releases


prior to 8.3 advertise all paths and need to be modified to only advertise the paths to the hosting HA pair. For more information about SLM, review section 5.9 of [TR-4080](#). The previous method of portsets can also be used to further reduce the available paths for a LUN. Portsets help by reducing the number of visible paths through which initiators in an igroup can see LUNs.

- SLM is enabled by default. Unless you are using portsets, no additional configuration is required.
- For LUNs created prior to Data ONTAP 8.3, manually apply SLM by running the `lun mapping remove-reporting-nodes` command to remove the LUN reporting nodes and restrict LUN access to the LUN-owning node and its HA partner.

Block protocols (iSCSI, FC, and FCoE) access LUNs by using LUN IDs and serial numbers, along with unique names. FC and FCoE use worldwide names (WWNNs and WWPNs), and iSCSI uses iSCSI qualified names (IQNs). The path to LUNs inside the storage is meaningless to the block protocols and is not presented anywhere in the protocol. Therefore, a volume that contains only LUNs does not need to be internally mounted at all, and a junction path is not needed for volumes that contain LUNs used in datastores. The NVMe subsystem in ONTAP works similarly.

Other best practices to consider:

- Make sure that a logical interface (LIF) is created for each SVM on each node in the ONTAP cluster for maximum availability and mobility. ONTAP SAN best practice is to use two physical ports and LIFs per node, one for each fabric. ALUA is used to parse paths and identify active optimized (direct) paths versus active nonoptimized paths. ALUA is used for FC, FCoE, and iSCSI.
- For iSCSI networks, use multiple VMkernel network interfaces on different network subnets with NIC teaming when multiple virtual switches are present. You can also use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput. The following figure provides an example of multipath connectivity. In ONTAP, configure either a single-mode interface group for failover with two or more links that are connected to two or more switches, or use LACP or other link-aggregation technology with multimode interface groups to provide HA and the benefits of link aggregation.
- If the Challenge-Handshake Authentication Protocol (CHAP) is used in ESXi for target authentication, it must also be configured in ONTAP using the CLI (`vserver iscsi security create`) or with System Manager (edit Initiator Security under Storage > SVMs > SVM Settings > Protocols > iSCSI).
- Use ONTAP tools for VMware vSphere to create and manage LUNs and igroups. The plug-in automatically determines the WWPNs of servers and creates appropriate igroups. It also configures LUNs according to best practices and maps them to the correct igroups.
- Use RDMs with care because they can be more difficult to manage, and they also use paths, which are limited as described earlier. ONTAP LUNs support both [physical and virtual compatibility mode](#) RDMs.
- For more on using NVMe/FC with vSphere 7.0, see this [ONTAP NVMe/FC Host Configuration guide](#) and [TR-4684](#). The following figure depicts multipath connectivity from a vSphere host to an ONTAP LUN.




NFS

NetApp ONTAP is, among many other things, an enterprise-class scale-out NAS array. ONTAP empowers VMware vSphere with concurrent access to NFS-connected datastores from many ESXi hosts, far exceeding the limits imposed on VMFS file systems. Using NFS with vSphere provides some ease of use and storage efficiency visibility benefits, as mentioned in the [datastores](#) section.

The following best practices are recommended when using ONTAP NFS with vSphere:

- Use a single logical interface (LIF) for each SVM on each node in the ONTAP cluster. Past recommendations of a LIF per datastore are no longer necessary. While direct access (LIF and datastore on same node) is best, don't worry about indirect access because the performance effect is generally minimal (microseconds).
- VMware has supported NFSv3 since VMware Infrastructure 3. vSphere 6.0 added support for NFSv4.1, which enables some advanced capabilities such as Kerberos security. Where NFSv3 uses client-side locking, NFSv4.1 uses server-side locking. Although an ONTAP volume can be exported through both protocols, ESXi can only mount through one protocol. This single protocol mount does not preclude other ESXi hosts from mounting the same datastore through a different version. Make sure to specify the protocol version to use when mounting so that all hosts use the same version and, therefore, the same locking style. Do not mix NFS versions across hosts. If possible, use host profiles to check compliance.
 - Because there is no automatic datastore conversion between NFSv3 and NFSv4.1, create a new NFSv4.1 datastore and use Storage vMotion to migrate VMs to the new datastore.
 - Please refer to the NFS v4.1 Interoperability table notes in the [NetApp Interoperability Matrix tool](#) for specific ESXi patch levels required for support.
 - VMware supports nconnect with NFSv3 beginning in vSphere 8.0U2. More information on nconnect can be found at the [NFSv3 nConnect feature with NetApp and VMware](#)
- NFS export policies are used to control access by vSphere hosts. You can use one policy with multiple volumes (datastores). With NFSv3, ESXi uses the sys (UNIX) security style and requires the root mount option to execute VMs. In ONTAP, this option is referred to as superuser, and when the superuser option is used, it is not necessary to specify the anonymous user ID. Note that export policy rules with different values for `-anon` and `-allow-suid` can cause SVM discovery problems with the ONTAP tools. Here's a sample policy:
 - Access Protocol: nfs (which includes both nfs3 and nfs4)
 - Client Match Spec: 192.168.42.21

- RO Access Rule: sys
- RW Access Rule: sys
- Anonymous UID
- Superuser: sys
- If the NetApp NFS Plug-In for VMware VAAI is used, the protocol should be set as `nfs` instead of `nfs3` when the export policy rule is created or modified. The VAAI copy offload feature requires the NFSv4 protocol to function, even if the data protocol is NFSv3. Specifying the protocol as `nfs` includes both the NFSv3 and NFSv4 versions.
- NFS datastore volumes are junctioned from the root volume of the SVM; therefore, ESXi must also have access to the root volume to navigate and mount datastore volumes. The export policy for the root volume, and for any other volumes in which the datastore volume's junction is nested, must include a rule or rules for the ESXi servers granting them read-only access. Here's a sample policy for the root volume, also using the VAAI plug-in:
 - Access Protocol: `nfs` (which includes both `nfs3` and `nfs4`)
 - Client Match Spec: `192.168.42.21`
 - RO Access Rule: `sys`
 - RW Access Rule: `never` (best security for root volume)
 - Anonymous UID
 - Superuser: `sys` (also required for root volume with VAAI)
- Use ONTAP tools for VMware vSphere (the most important best practice):
 - Use ONTAP tools for VMware vSphere to provision datastores because it simplifies management of export policies automatically.
 - When creating datastores for VMware clusters with the plug-in, select the cluster rather than a single ESX server. This choice triggers it to automatically mount the datastore to all hosts in the cluster.
 - Use the plug-in mount function to apply existing datastores to new servers.
 - When not using ONTAP tools for VMware vSphere, use a single export policy for all servers or for each cluster of servers where additional access control is needed.
- Although ONTAP offers a flexible volume namespace structure to arrange volumes in a tree using junctions, this approach has no value for vSphere. It creates a directory for each VM at the root of the datastore, regardless of the namespace hierarchy of the storage. Thus, the best practice is to simply mount the junction path for volumes for vSphere at the root volume of the SVM, which is how ONTAP tools for VMware vSphere provisions datastores. Not having nested junction paths also means that no volume is dependent on any volume other than the root volume and that taking a volume offline or destroying it, even intentionally, does not affect the path to other volumes.
- A block size of 4K is fine for NTFS partitions on NFS datastores. The following figure depicts connectivity from a vSphere host to an ONTAP NFS datastore.



The following table lists NFS versions and supported features.

vSphere Features	NFSv3	NFSv4.1
vMotion and Storage vMotion	Yes	Yes
High availability	Yes	Yes
Fault tolerance	Yes	Yes
DRS	Yes	Yes
Host profiles	Yes	Yes
Storage DRS	Yes	No
Storage I/O control	Yes	No
SRM	Yes	No
Virtual volumes	Yes	No
Hardware acceleration (VAAI)	Yes	Yes
Kerberos authentication	No	Yes (enhanced with vSphere 6.5 and later to support AES, krb5i)
Multipathing support	No	Yes

FlexGroup volumes

Use ONTAP and FlexGroup volumes with VMware vSphere for simple and scalable datastores that leverage the full power of an entire ONTAP cluster.

ONTAP 9.8, along with the ONTAP tools for VMware vSphere 9.8 and SnapCenter plugin for VMware 4.4 releases added support for FlexGroup volume-backed datastores in vSphere. FlexGroup volumes simplify the creation of large datastores and automatically create the necessary distributed constituent volumes across the ONTAP cluster to get the maximum performance from an ONTAP system.

Learn more about FlexGroup volumes in [FlexCache](#) and [FlexGroup volume technical reports](#).

Use FlexGroup volumes with vSphere if you require a single, scalable vSphere datastore with the power of a full ONTAP cluster, or if you have very large cloning workloads that can benefit from the new FlexGroup cloning mechanism.

Copy offload

In addition to extensive system testing with vSphere workloads, ONTAP 9.8 added a new copy offload mechanism for FlexGroup datastores. This new system uses an improved copy engine to replicate files between constituents in the background while allowing access to both source and destination. This local cache is then used to rapidly instantiate VM clones on demand.

To enable FlexGroup optimized copy offload, refer to [How to Configure ONTAP FlexGroups to allow VAAI copy offload](#)

You may find that if you use VAAI cloning, but do not clone enough to keep the cache warm, your clones may be no faster than a host-based copy. If that is the case you may tune the cache timeout to better suit your needs.

Consider the following scenario:

- You've created a new FlexGroup with 8 constituents
- The cache timeout for the new FlexGroup is set to 160 minutes

In this scenario, the first 8 clones to complete will be full copies, not local file clones. Any additional cloning of that VM before the 160-second timeout expires will use the file clone engine inside of each constituent in a round-robin fashion to create nearly immediate copies evenly distributed across the constituent volumes.

Every new clone job a volume receives resets the timeout. If a constituent volume in the example FlexGroup does not receive a clone request before the timeout, the cache for that particular VM will be cleared and the volume will need to be populated again. Also, if the source of the original clone changes (e.g., you've updated the template) then the local cache on each constituent will be invalidated to prevent any conflict. As previously stated, the cache is tunable and can be set to match the needs of your environment.

For more information on using FlexGroups with VAAI, refer to this KB article: [VAAI: How does caching work with FlexGroup volumes?](#)

In environments where you are not able to take full advantage of the FlexGroup cache, but still require rapid cross-volume cloning, consider using vVols. Cross-volume cloning with vVols is much faster than using traditional datastores, and does not rely on a cache.

QoS settings

Configuring QoS at the FlexGroup level using ONTAP System Manager or the cluster shell is supported, however it does not provide VM awareness or vCenter integration.

QoS (max/min IOPS) can be set on individual VMs or on all VMs in a datastore at that time in the vCenter UI or via REST APIs by using ONTAP tools. Setting QoS on all VMs replaces any separate per-VM settings. Settings do not extend to new or migrated VMs in the future; either set QoS on the new VMs or re-apply QoS to all VMs in the datastore.

Note that VMware vSphere treats all IO for an NFS datastore as a single queue per host, and QoS throttling on one VM can impact performance for other VMs in the same datastore. This is in contrast with vVols which can maintain their QoS policy settings if they migrate to another datastore and do not impact IO of other VMs when throttled.

Metrics

ONTAP 9.8 also added new file-based performance metrics (IOPS, throughput, and latency) for FlexGroup files, and these metrics can be viewed in the ONTAP tools for VMware vSphere dashboard and VM reports.

The ONTAP tools for VMware vSphere plug-in also allows you to set Quality of Service (QoS) rules using a combination of maximum and/or minimum IOPS. These can be set across all VMs in a datastore or individually for specific VMs.

Best practices

- Use ONTAP tools to create FlexGroup datastores to ensure your FlexGroup is created optimally and export policies are configured to match your vSphere environment. However, after creating the FlexGroup volume with ONTAP tools, you will find that all nodes in your vSphere cluster are using a single IP address to mount the datastore. This could result in a bottleneck on the network port. To avoid this problem, unmount the datastore, and then remount it using the standard vSphere datastore wizard using a round-robin DNS name that load balancing across LIFs on the SVM. After remounting, ONTAP tools will again be able to manage the datastore. If ONTAP tools isn't available, use the FlexGroup defaults and create your export policy following the guidelines in [datastores and protocols - NFS](#).
- When sizing a FlexGroup datastore, keep in mind that the FlexGroup consists of multiple smaller FlexVol volumes that create a larger namespace. As such, size the datastore to be at least 8x (assuming the default 8 constituents) the size of your largest VMDK file plus 10-20% unused headroom to allow for flexibility in rebalancing. For example, if you have a 6TB VMDK in your environment, size the FlexGroup datastore no smaller than 52.8TB (6x8+10%).
- VMware and NetApp support NFSv4.1 session trunking beginning with ONTAP 9.14.1. Refer to the NetApp NFS 4.1 interoperability matrix notes for specific version details. NFSv3 does not support multiple physical paths to a volume but does support nconnect beginning in vSphere 8.0U2. More information on nconnect can be found at the [NFSv3 nConnect feature with NetApp and VMware](#).
- Use the NFS Plug-In for VMware VAAI for copy offload. Note that while cloning is enhanced within a FlexGroup datastore, as mentioned previously, ONTAP does not provide significant performance advantages versus ESXi host copy when copying VMs between FlexVol and/or FlexGroup volumes. Therefore consider your cloning workloads when deciding to use VAAI or FlexGroups. Modifying the number of constituent volumes is one way to optimize for FlexGroup-based cloning. As is tuning the cache timeout previously mentioned.
- Use ONTAP tools for VMware vSphere 9.8 or later to monitor the performance of FlexGroup VMs using ONTAP metrics (dashboard and VM reports), and to manage QoS on individual VMs. These metrics are not currently available through ONTAP commands or APIs.
- SnapCenter Plug-In for VMware vSphere release 4.4 and later supports backup and recovery of VMs in a FlexGroup datastore on the primary storage system. SCV 4.6 adds SnapMirror support for FlexGroup-based datastores. Using array-based snapshots and replication is the most efficient way to protect your data.

Network configuration

Configuring network settings when using vSphere with systems running ONTAP software is straightforward and similar to other network configuration.

Here are some things to consider:

- Separate storage network traffic from other networks. A separate network can be achieved by using a dedicated VLAN or separate switches for storage. If the storage network shares physical paths such as uplinks, you might need QoS or additional uplink ports to make sure of sufficient bandwidth. Don't connect hosts directly to storage; use switches to have redundant paths and allow VMware HA to work without intervention. See [Direct connect networking](#) for additional information.
- Jumbo frames can be used if desired and supported by your network, especially when using iSCSI. If they are used, make sure they are configured identically on all network devices, VLANs, and so on in the path

between storage and the ESXi host. Otherwise, you might see performance or connection problems. The MTU must also be set identically on the ESXi virtual switch, the VMkernel port, and also on the physical ports or interface groups of each ONTAP node.

- NetApp only recommends disabling network flow control on the cluster network ports within an ONTAP cluster. NetApp makes no other recommendations for best practices for the remaining network ports used for data traffic. You should enable or disable it as necessary. See [TR-4182](#) for more background on flow control.
- When ESXi and ONTAP storage arrays are connected to Ethernet storage networks, NetApp recommends configuring the Ethernet ports to which these systems connect as Rapid Spanning Tree Protocol (RSTP) edge ports or by using the Cisco PortFast feature. NetApp recommends enabling the Spanning-Tree PortFast trunk feature in environments that use the Cisco PortFast feature and that have 802.1Q VLAN trunking enabled to either the ESXi server or the ONTAP storage arrays.
- NetApp recommends the following best practices for link aggregation:
 - Use switches that support link aggregation of ports on two separate switch chassis using a multi-chassis link aggregation group approach such as Cisco's Virtual PortChannel (vPC).
 - Disable LACP for switch ports connected to ESXi unless you are using dvSwitches 5.1 or later with LACP configured.
 - Use LACP to create link aggregates for ONTAP storage systems with dynamic multimode interface groups with IP hash.
 - Use an IP hash teaming policy on ESXi.

The following table provides a summary of network configuration items and indicates where the settings are applied.

Item	ESXi	Switch	Node	SVM
IP address	VMkernel	No**	No**	Yes
Link aggregation	Virtual switch	Yes	Yes	No*
VLAN	VMkernel and VM port groups	Yes	Yes	No*
Flow control	NIC	Yes	Yes	No*
Spanning tree	No	Yes	No	No
MTU (for jumbo frames)	Virtual switch and VMkernel port (9000)	Yes (set to max)	Yes (9000)	No*
Failover groups	No	No	Yes (create)	Yes (select)

*SVM LIFs connect to ports, interface groups, or VLAN interfaces that have VLAN, MTU, and other settings. However, the settings are not managed at the SVM level.

**These devices have IP addresses of their own for management, but these addresses are not used in the context of ESXi storage networking.

SAN (FC, FCoE, NVMe/FC, iSCSI), RDM

In vSphere, there are three ways to use block storage LUNs:

- With VMFS datastores

- With raw device mapping (RDM)
- As a LUN accessed and controlled by a software initiator from a VM guest OS

VMFS is a high-performance clustered file system that provides datastores that are shared storage pools. VMFS datastores can be configured with LUNs that are accessed using FC, iSCSI, FCoE, or NVMe namespaces accessed by the NVMe/FC protocol. VMFS allows traditional LUNs to be accessed simultaneously by every ESX server in a cluster. The ONTAP maximum LUN size is generally 16TB; therefore, a maximum-size VMFS 5 datastore of 64TB (see the first table in this section) is created by using four 16TB LUNs (All SAN Array systems support the maximum VMFS LUN size of 64TB). Because the ONTAP LUN architecture does not have small individual queue depths, VMFS datastores in ONTAP can scale to a greater degree than with traditional array architectures in a relatively simple manner.

vSphere includes built-in support for multiple paths to storage devices, referred to as native multipathing (NMP). NMP can detect the type of storage for supported storage systems and automatically configures the NMP stack to support the capabilities of the storage system in use.

Both NMP and ONTAP support Asymmetric Logical Unit Access (ALUA) to negotiate optimized and nonoptimized paths. In ONTAP, an ALUA-optimized path follows a direct data path, using a target port on the node that hosts the LUN being accessed. ALUA is turned on by default in both vSphere and ONTAP. The NMP recognizes the ONTAP cluster as ALUA, and it uses the ALUA storage array type plug-in (`VMW_SATP_ALUA`) and selects the round-robin path selection plug-in (`VMW_PSP_RR`).

ESXi 6 supports up to 256 LUNs and up to 1,024 total paths to LUNs. Any LUNs or paths beyond these limits are not seen by ESXi. Assuming the maximum number of LUNs, the path limit allows four paths per LUN. In a larger ONTAP cluster, it is possible to reach the path limit before the LUN limit. To address this limitation, ONTAP supports selective LUN map (SLM) in release 8.3 and later.

SLM limits the nodes that advertise paths to a given LUN. It is a NetApp best practice to have at least one LIF per node per SVM and to use SLM to limit the paths advertised to the node hosting the LUN and its HA partner. Although other paths exist, they aren't advertised by default. It is possible to modify the paths advertised with the add and remove reporting node arguments within SLM. Note that LUNs created in releases before 8.3 advertise all paths and need to be modified to only advertise the paths to the hosting HA pair. For more information about SLM, review section 5.9 of [TR-4080](#). The previous method of portsets can also be used to further reduce the available paths for a LUN. Portsets help by reducing the number of visible paths through which initiators in an igroup can see LUNs.


- SLM is enabled by default. Unless you are using portsets, no additional configuration is required.
- For LUNs created before Data ONTAP 8.3, manually apply SLM by running the `lun mapping remove-reporting-nodes` command to remove the LUN reporting nodes and restrict LUN access to the LUN-owning node and its HA partner.

Block protocols (iSCSI, FC, and FCoE) access LUNs by using LUN IDs and serial numbers, along with unique names. FC and FCoE use worldwide names (WWNNs and WWPNs), and iSCSI uses iSCSI qualified names (IQNs). The path to LUNs inside the storage is meaningless to the block protocols and is not presented anywhere in the protocol. Therefore, a volume that contains only LUNs does not need to be internally mounted at all, and a junction path is not needed for volumes that contain LUNs used in datastores. The NVMe subsystem in ONTAP works similarly.

Other best practices to consider:

- Make sure that a logical interface (LIF) is created for each SVM on each node in the ONTAP cluster for maximum availability and mobility. ONTAP SAN best practice is to use two physical ports and LIFs per node, one for each fabric. ALUA is used to parse paths and identify active optimized (direct) paths versus active nonoptimized paths. ALUA is used for FC, FCoE, and iSCSI.

- For iSCSI networks, use multiple VMkernel network interfaces on different network subnets with NIC teaming when multiple virtual switches are present. You can also use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput. The following figure provides an example of multipath connectivity. In ONTAP, use a single-mode interface group with multiple links to different switches or LACP with multimode interface groups for high availability and link aggregation benefits.
- If the Challenge-Handshake Authentication Protocol (CHAP) is used in ESXi for target authentication, it must also be configured in ONTAP using the CLI (`vserver iscsi security create`) or with System Manager (edit Initiator Security under Storage > SVMs > SVM Settings > Protocols > iSCSI).
- Use ONTAP tools for VMware vSphere to create and manage LUNs and igroups. The plug-in automatically determines the WWPNs of servers and creates appropriate igroups. It also configures LUNs according to best practices and maps them to the correct igroups.
- Use RDMs with care because they can be more difficult to manage, and they also use paths, which are limited as described earlier. ONTAP LUNs support both [physical and virtual compatibility mode](#) RDMs.
- For more on using NVMe/FC with vSphere 7.0, see this [ONTAP NVMe/FC Host Configuration guide](#) and [TR-4684](#). The following figure depicts multipath connectivity from a vSphere host to an ONTAP LUN.



NFS

vSphere allows customers to use enterprise-class NFS arrays to provide concurrent access to datastores to all the nodes in an ESXi cluster. As mentioned in the datastore section, there are some ease of use and storage efficiency visibility benefits when using NFS with vSphere.

The following best practices are recommended when using ONTAP NFS with vSphere:

- Use a single logical interface (LIF) for each SVM on each node in the ONTAP cluster. Past recommendations of a LIF per datastore are no longer necessary. While direct access (LIF and datastore on the same node) is best, don't worry about indirect access because the performance effect is generally minimal (microseconds).
- All versions of VMware vSphere that are currently supported can use both NFS v3 and v4.1. Official support for nconnect was added to vSphere 8.0 update 2 for NFS v3. For NFS v4.1, vSphere continues to support session trunking, Kerberos authentication, and Kerberos authentication with integrity. It's important to note that session trunking requires ONTAP 9.14.1 or a later version. You can learn more about the nconnect feature and how it improves performance at [NFSv3 nConnect feature with NetApp and VMware](#).

It's worth noting that NFSv3 and NFSv4.1 use different locking mechanisms. NFSv3 uses client-side locking, while NFSv4.1 uses server-side locking. Although an ONTAP volume can be exported through both protocols, ESXi can only mount a datastore through one protocol. However, this doesn't mean that other ESXi hosts

cannot mount the same datastore through a different version. To avoid any issues, it's essential to specify the protocol version to use when mounting, ensuring that all hosts use the same version and, therefore, the same locking style. It's critical to avoid mixing NFS versions across hosts. If possible, use host profiles to check compliance.

Because there is no automatic datastore conversion between NFSv3 and NFSv4.1, create a new NFSv4.1 datastore and use Storage vMotion to migrate VMs to the new datastore.

Please refer to the NFS v4.1 Interoperability table notes in the [NetApp Interoperability Matrix tool](#) for specific ESXi patch levels required for support.

* NFS export policies are used to control access by vSphere hosts. You can use one policy with multiple volumes (datastores). With NFSv3, ESXi uses the sys (UNIX) security style and requires the root mount option to execute VMs. In ONTAP, this option is referred to as superuser, and when the superuser option is used, it is not necessary to specify the anonymous user ID. Note that export policy rules with different values for `-anon` and `-allow-suid` can cause SVM discovery problems with the ONTAP tools. Here's a sample policy:

Access Protocol: nfs3

Client Match Spec: 192.168.42.21

RO Access Rule: sys

RW Access Rule: sys

Anonymous UID

Superuser: sys

* If the NetApp NFS Plug-In for VMware VAAI is used, the protocol should be set as `nfs` when the export policy rule is created or modified. The NFSv4 protocol is required for VAAI copy offload to work, and specifying the protocol as `nfs` automatically includes both the NFSv3 and the NFSv4 versions.

* NFS datastore volumes are junctioned from the root volume of the SVM; therefore, ESXi must also have access to the root volume to navigate and mount datastore volumes. The export policy for the root volume, and for any other volumes in which the datastore volume's junction is nested, must include a rule or rules for the ESXi servers granting them read-only access. Here's a sample policy for the root volume, also using the VAAI plug-in:

Access Protocol: nfs (which includes both nfs3 and nfs4)

Client Match Spec: 192.168.42.21

RO Access Rule: sys

RW Access Rule: never (best security for root volume)

Anonymous UID

Superuser: sys (also required for root volume with VAAI)

* Use ONTAP tools for VMware vSphere (the most important best practice):

Use ONTAP tools for VMware vSphere to provision datastores because it simplifies the management of export policies automatically.


When creating datastores for VMware clusters with the plug-in, select the cluster rather than a single ESX server. This choice triggers it to automatically mount the datastore to all hosts in the cluster.

Use the plug-in mount function to apply existing datastores to new servers.

When not using ONTAP tools for VMware vSphere, use a single export policy for all servers or for each cluster of servers where additional access control is needed.

* Although ONTAP offers a flexible volume namespace structure to arrange volumes in a tree using junctions, this approach has no value for vSphere. It creates a directory for each VM at the root of the datastore, regardless of the namespace hierarchy of the storage. Thus, the best practice is to simply mount the junction path for volumes for vSphere at the root volume of the SVM, which is how ONTAP tools for VMware vSphere provisions datastores. Not having nested junction paths also means that no volume is dependent on any volume other than the root volume and that taking a volume offline or destroying it, even intentionally, does not affect the path to other volumes.

* A block size of 4K is fine for NTFS partitions on NFS datastores. The following figure depicts connectivity from a vSphere host to an ONTAP NFS datastore.



The following table lists NFS versions and supported features.

vSphere Features	NFSv3	NFSv4.1
vMotion and Storage vMotion	Yes	Yes
High availability	Yes	Yes
Fault tolerance	Yes	Yes
DRS	Yes	Yes
Host profiles	Yes	Yes
Storage DRS	Yes	No
Storage I/O control	Yes	No
SRM	Yes	No
Virtual volumes	Yes	No
Hardware acceleration (VAAI)	Yes	Yes
Kerberos authentication	No	Yes (enhanced with vSphere 6.5 and later to support AES, krb5i)
Multipathing support	No	Yes (ONTAP 9.14.1)

Direct connect networking

Storage administrators sometimes prefer to simplify their infrastructures by removing network switches from the configuration. This can be supported in some scenarios.

iSCSI and NVMe/TCP

A host using iSCSI or NVMe/TCP can be directly connected to a storage system and operate normally. The reason is pathing. Direct connections to two different storage controllers result in two independent paths for data flow. The loss of path, port, or controller does not prevent the other path from being used.

NFS

Direct-connected NFS storage can be used, but with a significant limitation - failover will not work without a significant scripting effort, which would be the responsibility of the customer.

The reason nondisruptive failover is complicated with direct-connected NFS storage is the routing that occurs

on the local OS. For example, assume a host has an IP address of 192.168.1.1/24 and is directly connected to an ONTAP controller with an IP address of 192.168.1.50/24. During failover, that 192.168.1.50 address can fail over to the other controller, and it will be available to the host, but how does the host detect its presence? The original 192.168.1.1 address still exists on the host NIC that no longer connects to an operational system. Traffic destined for 192.168.1.50 would continue to be sent to an inoperable network port.

The second OS NIC could be configured as 192.168.1.2 and would be capable of communicating with the failed over 192.168.1.50 address, but the local routing tables would have a default of using one **and only one** address to communicate with the 192.168.1.0/24 subnet. A sysadmin could create a scripting framework that would detect a failed network connection and alter the local routing tables or bring interfaces up and down. The exact procedure would depend on the OS in use.

In practice, NetApp customers do have direct-connected NFS, but normally only for workloads where IO pauses during failovers are acceptable. When hard mounts are used, there should not be any IO errors during such pauses. The IO should hang until services are restored, either by a failback or manual intervention to move IP addresses between NICs on the host.

FC Direct Connect

It is not possible to directly connect a host to an ONTAP storage system using the FC protocol. The reason is the use of NPIV. The WWN that identifies an ONTAP FC port to the FC network uses a type of virtualization called NPIV. Any device connected to an ONTAP system must be able to recognize an NPIV WWN. There are no current HBA vendors who offer an HBA that can be installed in a host that would be able to support an NPIV target.


VM and datastore cloning

Cloning a storage object allows you to quickly create copies for further use, such as provisioning additional VMs, backup/recovery operations, and so on.

In vSphere, you can clone a VM, virtual disk, vVol, or datastore. After being cloned, the object can be further customized, often through an automated process. vSphere supports both full copy clones, as well as linked clones, where it tracks changes separately from the original object.

Linked clones are great for saving space, but they increase the amount of I/O that vSphere handles for the VM, affecting performance of that VM and perhaps the host overall. That's why NetApp customers often use storage system-based clones to get the best of both worlds: efficient use of storage and increased performance.

The following figure depicts ONTAP cloning.



Cloning can be offloaded to systems running ONTAP software through several mechanisms, typically at the VM, vVol, or datastore level. These include the following:

- vVols using the NetApp vSphere APIs for Storage Awareness (VASA) Provider. ONTAP clones are used to support vVol snapshots managed by vCenter that are space-efficient with minimal I/O effect to create and delete them. VMs can also be cloned using vCenter, and these are also offloaded to ONTAP, whether within a single datastore/volume or between datastores/volumes.
- vSphere cloning and migration using vSphere APIs – Array Integration (VAAI). VM cloning operations can be offloaded to ONTAP in both SAN and NAS environments (NetApp supplies an ESXi plug-in to enable VAAI for NFS). vSphere only offloads operations on cold (powered off) VMs in a NAS datastore, whereas operations on hot VMs (cloning and storage vMotion) are also offloaded for SAN. ONTAP uses the most efficient approach based on source, destination, and installed product licenses. This capability is also used by VMware Horizon View.

- SRA (used with VMware Site Recovery Manager). Here, clones are used to test recovery of the DR replica nondisruptively.
- Backup and recovery using NetApp tools such as SnapCenter. VM clones are used to verify backup operations as well as to mount a VM backup so that individual files can be copied.

ONTAP offloaded cloning can be invoked by VMware, NetApp, and third-party tools. Clones that are offloaded to ONTAP have several advantages. They are space-efficient in most cases, needing storage only for changes to the object; there is no additional performance effect to read and write them, and in some cases performance is improved by sharing blocks in high-speed caches. They also offload CPU cycles and network I/O from the ESXi server. Copy offload within a traditional datastore using a FlexVol volume can be fast and efficient with FlexClone licensed, but copies between FlexVol volumes might be slower. If you maintain VM templates as a source of clones, consider placing them within the datastore volume (use folders or content libraries to organize them) for fast, space efficient clones.

You can also clone a volume or LUN directly within ONTAP to clone a datastore. With NFS datastores, FlexClone technology can clone an entire volume, and the clone can be exported from ONTAP and mounted by ESXi as another datastore. For VMFS datastores, ONTAP can clone a LUN within a volume or a whole volume, including one or more LUNs within it. A LUN containing a VMFS must be mapped to an ESXi initiator group (igroup) and then resignatured by ESXi to be mounted and used as a regular datastore. For some temporary use cases, a cloned VMFS can be mounted without resignaturing. After a datastore is cloned, VMs inside it can be registered, reconfigured, and customized as if they were individually cloned VMs.

In some cases, additional licensed features can be used to enhance cloning, such as SnapRestore for backup or FlexClone. These licenses are often included in license bundles at no additional cost. A FlexClone license is required for vVol cloning operations as well as to support managed snapshots of a vVol (which are offloaded from the hypervisor to ONTAP). A FlexClone license can also improve certain VAAI-based clones when used within a datastore/volume (creates instant, space-efficient copies instead of block copies). It is also used by the SRA when testing recovery of a DR replica, and SnapCenter for clone operations and to browse backup copies to restore individual files.

Data protection

Backing up your VMs and quickly recovering them are among the great strengths of ONTAP for vSphere, and it is easy to manage this ability inside vCenter with the SnapCenter Plug-In for VMware vSphere.

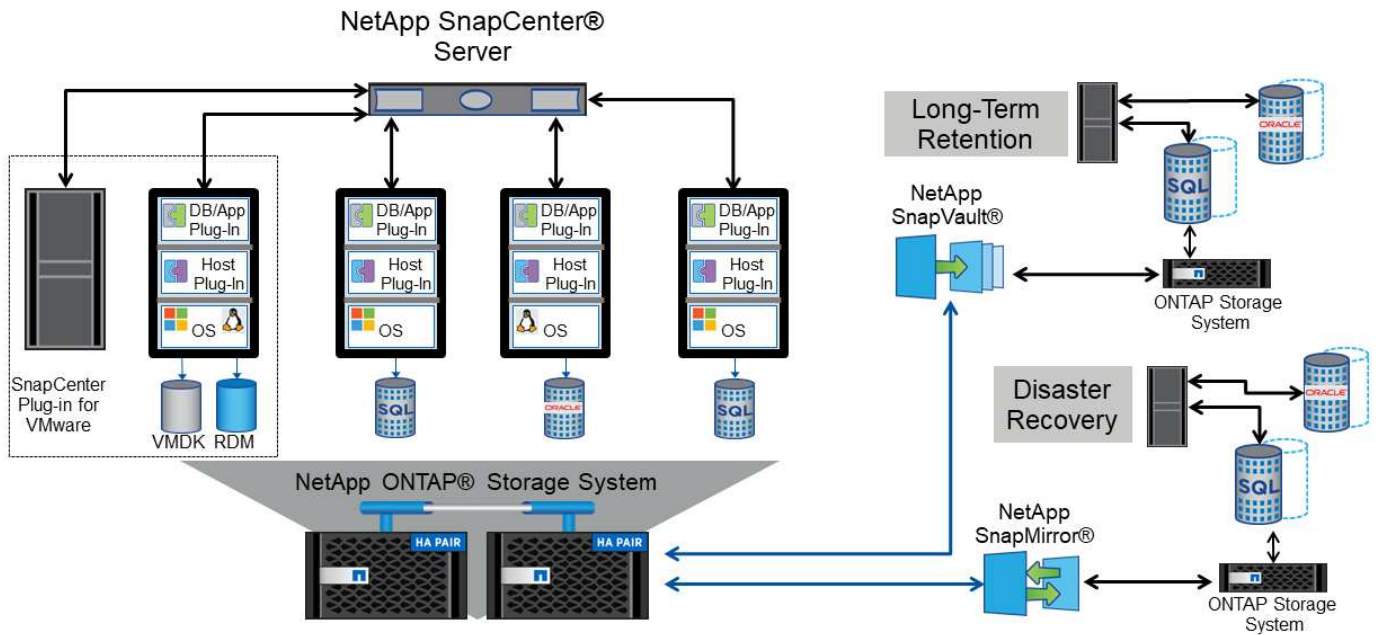
Use snapshots to make quick copies of your VM or datastore without affecting performance, and then send them to a secondary system using SnapMirror for longer-term off-site data protection. This approach minimizes storage space and network bandwidth by only storing changed information.

SnapCenter allows you to create backup policies that can be applied to multiple jobs. These policies can define schedule, retention, replication, and other capabilities. They continue to allow optional selection of VM-consistent snapshots, which leverages the hypervisor's ability to quiesce I/O before taking a VMware snapshot. However, due to the performance effect of VMware snapshots, they are generally not recommended unless you need the guest file system to be quiesced. Instead, use snapshots for general protection, and use application tools such as SnapCenter plug-ins to protect transactional data such as SQL Server or Oracle. These snapshots are different from VMware (consistency) snapshots and are suitable for longer term protection. VMware snapshots are only [recommended](#) for short term use due to performance and other effects.

These plug-ins offer extended capabilities to protect the databases in both physical and virtual environments. With vSphere, you can use them to protect SQL Server or Oracle databases where data is stored on RDM LUNs, iSCSI LUNs directly connected to the guest OS, or VMDK files on either VMFS or NFS datastores. The plug-ins allow specification of different types of database backups, supporting online or offline backup, and

protecting database files along with log files. In addition to backup and recovery, the plug-ins also support cloning of databases for development or test purposes.

The following figure depicts an example of SnapCenter deployment.



For enhanced disaster recovery capabilities, consider using the NetApp SRA for ONTAP with VMware Site Recovery Manager. In addition to support for the replication of datastores to a DR site, it also enables nondisruptive testing in the DR environment by cloning the replicated datastores. Recovery from a disaster and reprotecting production after the outage has been resolved are also made easy by automation built into SRA.

Finally, for the highest level of data protection, consider a VMware vSphere Metro Storage Cluster (vMSC) configuration using NetApp MetroCluster. vMSC is a VMware-certified solution that combines synchronous replication with array-based clustering, giving the same benefits of a high-availability cluster but distributed across separate sites to protect against site disaster. NetApp MetroCluster offers cost-effective configurations for synchronous replication with transparent recovery from any single storage component failure as well as single-command recovery in the event of a site disaster. vMSC is described in greater detail in [TR-4128](#).

Quality of service (QoS)

Systems running ONTAP software can use the ONTAP storage QoS feature to limit throughput in MBps and/or I/Os per second (IOPS) for different storage objects such as files, LUNs, volumes, or entire SVMs.

Throughput limits are useful in controlling unknown or test workloads before deployment to make sure they don't affect other workloads. They can also be used to constrain a bully workload after it is identified. Minimum levels of service based on IOPS are also supported to provide consistent performance for SAN objects in ONTAP 9.2 and for NAS objects in ONTAP 9.3.

With an NFS datastore, a QoS policy can be applied to the entire FlexVol volume or individual VMDK files within it. With VMFS datastores using ONTAP LUNs, the QoS policies can be applied to the FlexVol volume that contains the LUNs or individual LUNs, but not individual VMDK files because ONTAP has no awareness of the VMFS file system. When using vVols, minimum and/or maximum QoS can be set on individual VMs using the storage capability profile and VM storage policy.

The QoS maximum throughput limit on an object can be set in MBps and/or IOPS. If both are used, the first limit reached is enforced by ONTAP. A workload can contain multiple objects, and a QoS policy can be applied to one or more workloads. When a policy is applied to multiple workloads, the workloads share the total limit of the policy. Nested objects are not supported (for example, files within a volume cannot each have their own policy). QoS minimums can only be set in IOPS.

The following tools are currently available for managing ONTAP QoS policies and applying them to objects:

- ONTAP CLI
- ONTAP System Manager
- OnCommand Workflow Automation
- Active IQ Unified Manager
- NetApp PowerShell Toolkit for ONTAP
- ONTAP tools for VMware vSphere VASA Provider

To assign a QoS policy to a VMDK on NFS, note the following guidelines:

- The policy must be applied to the `vmname-flat.vmdk` that contains the actual virtual disk image, not the `vmname.vmdk` (virtual disk descriptor file) or `vmname.vmx` (VM descriptor file).
- Do not apply policies to other VM files such as virtual swap files (`vmname.vswp`).
- When using the vSphere web client to find file paths (Datastore > Files), be aware that it combines the information of the `-flat.vmdk` and `.vmdk` and simply shows one file with the name of the `.vmdk` but the size of the `-flat.vmdk`. Add `-flat` into the file name to get the correct path.

To assign a QoS policy to a LUN, including VMFS and RDM, the ONTAP SVM (displayed as Vserver), LUN path, and serial number can be obtained from the Storage Systems menu on the ONTAP tools for VMware vSphere home page. Select the storage system (SVM), and then Related Objects > SAN. Use this approach when specifying QoS using one of the ONTAP tools.

Maximum and minimum QoS can be easily assigned to a vVol-based VM with ONTAP tools for VMware vSphere or Virtual Storage Console 7.1 and later. When creating the storage capability profile for the vVol container, specify a max and/or min IOPS value under the performance capability and then reference this SCP with the VM's storage policy. Use this policy when creating the VM or apply the policy to an existing VM.

FlexGroup datastores offer enhanced QoS capabilities when using ONTAP tools for VMware vSphere 9.8 and later. You can easily set QoS on all VMs in a datastore or on specific VMs. See the FlexGroup section of this report for more information.

ONTAP QoS and VMware SIOC

ONTAP QoS and VMware vSphere Storage I/O Control (SIOC) are complementary technologies that vSphere and storage administrators can use together to manage performance of vSphere VMs hosted on systems running ONTAP software. Each tool has its own strengths, as shown in the following table. Because of the different scopes of VMware vCenter and ONTAP, some objects can be seen and managed by one system and not the other.

Property	ONTAP QoS	VMware SIOC
When active	Policy is always active	Active when contention exists (datastore latency over threshold)

Property	ONTAP QoS	VMware SIOC
Type of units	IOPS, MBps	IOPS, shares
vCenter or application scope	Multiple vCenter environments, other hypervisors and applications	Single vCenter server
Set QoS on VM?	VMDK on NFS only	VMDK on NFS or VMFS
Set QoS on LUN (RDM)?	Yes	No
Set QoS on LUN (VMFS)?	Yes	No
Set QoS on volume (NFS datastore)?	Yes	No
Set QoS on SVM (tenant)?	Yes	No
Policy based approach?	Yes; can be shared by all workloads in the policy or applied in full to each workload in the policy.	Yes, with vSphere 6.5 and later.
License required	Included with ONTAP	Enterprise Plus

VMware Storage Distributed Resource Scheduler

VMware Storage Distributed Resource Scheduler (SDRS) is a vSphere feature that places VMs on storage based on the current I/O latency and space usage. It then moves the VM or VMDKs nondisruptively between the datastores in a datastore cluster (also referred to as a pod), selecting the best datastore in which to place the VM or VMDKs in the datastore cluster. A datastore cluster is a collection of similar datastores that are aggregated into a single unit of consumption from the vSphere administrator's perspective.

When using SDRS with ONTAP tools for VMware vSphere, you must first create a datastore with the plug-in, use vCenter to create the datastore cluster, and then add the datastore to it. After the datastore cluster is created, additional datastores can be added to the datastore cluster directly from the provisioning wizard on the Details page.

Other ONTAP best practices for SDRS include the following:

- All datastores in the cluster should use the same type of storage (such as SAS, SATA, or SSD), be either all VMFS or NFS datastores, and have the same replication and protection settings.
- Consider using SDRS in default (manual) mode. This approach allows you to review the recommendations and decide whether to apply them or not. Be aware of these effects of VMDK migrations:
 - When SDRS moves VMDKs between datastores, any space savings from ONTAP cloning or deduplication are lost. You can rerun deduplication to regain these savings.
 - After SDRS moves VMDKs, NetApp recommends recreating the snapshots at the source datastore because space is otherwise locked by the VM that was moved.
 - Moving VMDKs between datastores on the same aggregate has little benefit, and SDRS does not have visibility into other workloads that might share the aggregate.

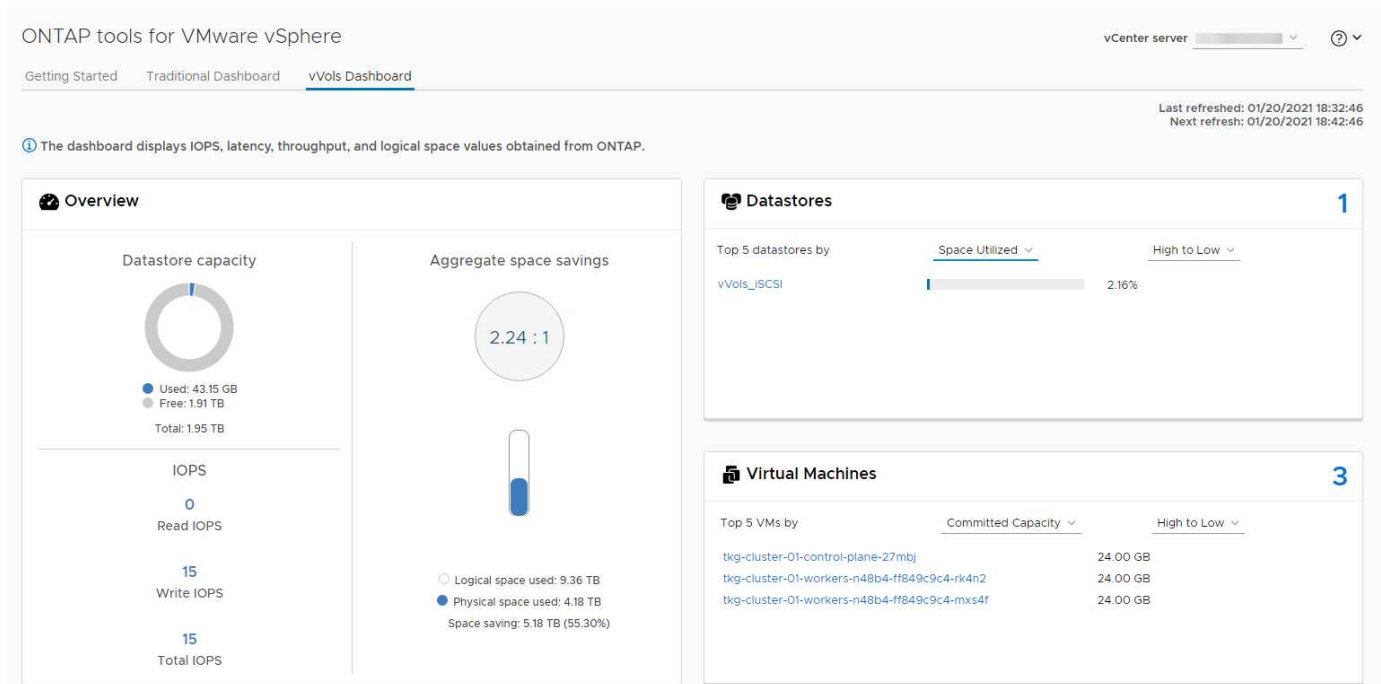
Storage policy based management and vVols

VMware vSphere APIs for Storage Awareness (VASA) make it easy for a storage administrator to configure datastores with well-defined capabilities and let the VM administrator use those whenever needed to provision VMs without having to interact with each other. It's worth taking a look at this approach to see how it can streamline your virtualization storage operations and avoid a lot of trivial work.

Prior to VASA, VM administrators could define VM storage policies, but they had to work with the storage administrator to identify appropriate datastores, often by using documentation or naming conventions. With VASA, the storage administrator can define a range of storage capabilities, including performance, tiering, encryption, and replication. A set of capabilities for a volume or a set of volumes is called a storage capability profile (SCP).

The SCP supports minimum and/or maximum QoS for a VM's data vVols. Minimum QoS is supported only on AFF systems. ONTAP tools for VMware vSphere includes a dashboard that displays VM granular performance and logical capacity for vVols on ONTAP systems.

The following figure depicts ONTAP tools for VMware vSphere 9.8 vVols dashboard.



After the storage capability profile is defined, it can be used to provision VMs using the storage policy that identifies its requirements. The mapping between the VM storage policy and the datastore storage capability profile allows vCenter to display a list of compatible datastores for selection. This approach is known as storage policy based management.

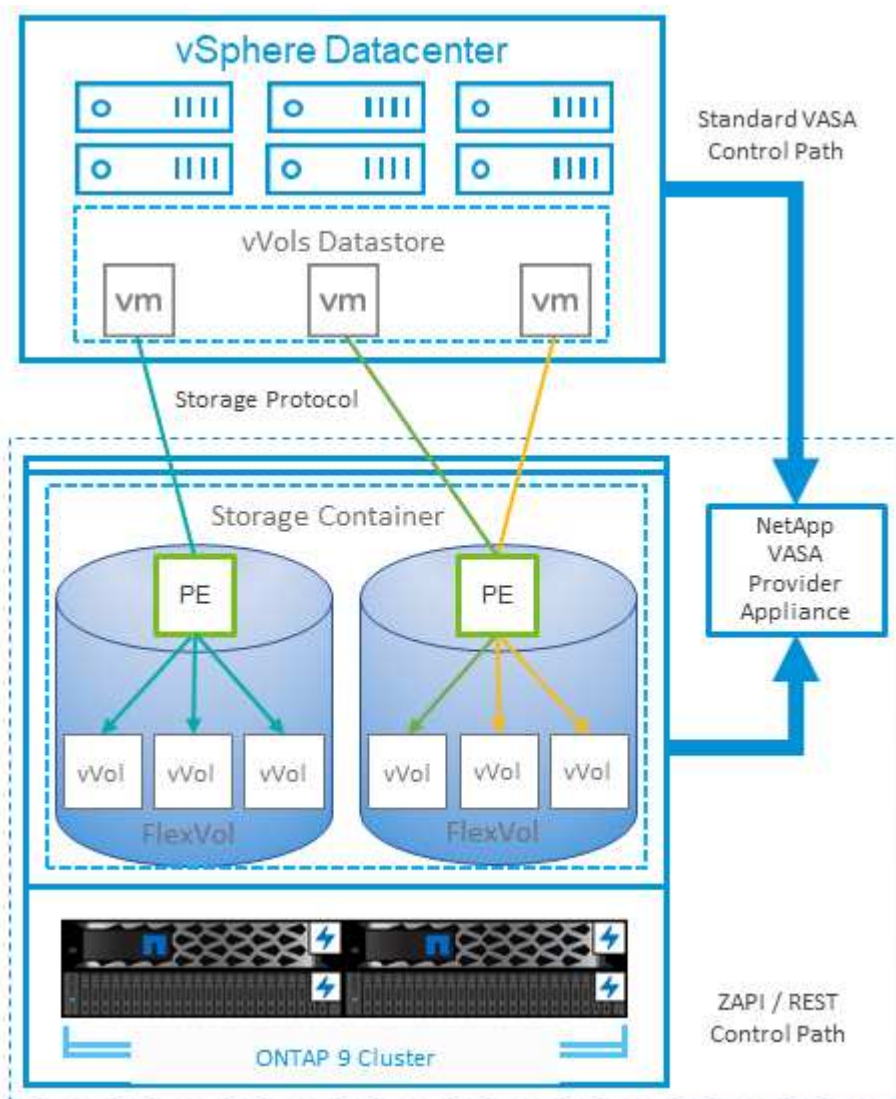
VASA provides the technology to query storage and return a set of storage capabilities to vCenter. VASA vendor providers supply the translation between the storage system APIs and constructs and the VMware APIs that are understood by vCenter. NetApp's VASA Provider for ONTAP is offered as part of the ONTAP tools for VMware vSphere appliance VM, and the vCenter plug-in provides the interface to provision and manage vVol datastores, as well as the ability to define storage capability profiles (SCPs).

ONTAP supports both VMFS and NFS vVol datastores. Using vVols with SAN datastores brings some of the benefits of NFS such as VM-level granularity. Here are some best practices to consider, and you can find additional information in [TR-4400](#):

- A vVol datastore can consist of multiple FlexVol volumes on multiple cluster nodes. The simplest approach is a single datastore, even when the volumes have different capabilities. SPBM makes sure that a compatible volume is used for the VM. However, the volumes must all be part of a single ONTAP SVM and accessed using a single protocol. One LIF per node for each protocol is sufficient. Avoid using multiple ONTAP releases within a single vVol datastore because the storage capabilities might vary across releases.

- Use the ONTAP tools for VMware vSphere plug-in to create and manage vVol datastores. In addition to managing the datastore and its profile, it automatically creates a protocol endpoint to access the vVols if needed. If LUNs are used, note that LUN PEs are mapped using LUN IDs 300 and higher. Verify that the ESXi host advanced system setting `Disk.MaxLUN` allows a LUN ID number that is higher than 300 (the default is 1,024). Do this step by selecting the ESXi host in vCenter, then the Configure tab, and find `Disk.MaxLUN` in the list of Advanced System Settings.
- Do not install or migrate VASA Provider, vCenter Server (appliance or Windows based), or ONTAP tools for VMware vSphere itself onto a vVols datastore, because they are then mutually dependent, limiting your ability to manage them in the event of a power outage or other data center disruption.
- Back up the VASA Provider VM regularly. At a minimum, create hourly snapshots of the traditional datastore that contains VASA Provider. For more about protecting and recovering the VASA Provider, see this [KB article](#).

The following figure shows vVols components.



Cloud migration and backup

Another ONTAP strength is broad support for the hybrid cloud, merging systems in your on-premises private cloud with public cloud capabilities. Here are some NetApp cloud solutions that can be used in conjunction with vSphere:

- **Cloud Volumes.** NetApp Cloud Volumes Service for Amazon Web Services or Google Cloud Platform and Azure NetApp Files for ANF provide high-performance, multi-protocol managed storage services in the leading public cloud environments. They can be used directly by VMware Cloud VM guests.
- **Cloud Volumes ONTAP.** NetApp Cloud Volumes ONTAP data management software delivers control, protection, flexibility, and efficiency to your data on your choice of cloud. Cloud Volumes ONTAP is cloud-native data management software built on ONTAP storage. Use together with Cloud Manager to deploy and manage Cloud Volumes ONTAP instances together with your on-premises ONTAP systems. Take advantage of advanced NAS and iSCSI SAN capabilities together with unified data management, including snapshots and SnapMirror replication.
- **Cloud Services.** Use Cloud Backup Service or SnapMirror Cloud to protect data from on-premises systems using public cloud storage. Cloud Sync helps migrate and keep your data in sync across NAS, object stores, and Cloud Volumes Service storage.
- **FabricPool.** FabricPool offers quick and easy tiering for ONTAP data. Cold blocks can be migrated to an object store in either public clouds or a private StorageGRID object store and are automatically recalled when the ONTAP data is accessed again. Or use the object tier as a third level of protection for data that is already managed by SnapVault. This approach can allow you to [store more snapshots of your VMs](#) on primary and/or secondary ONTAP storage systems.
- **ONTAP Select.** Use NetApp software-defined storage to extend your private cloud across the Internet to remote facilities and offices, where you can use ONTAP Select to support block and file services as well as the same vSphere data management capabilities you have in your enterprise data center.

When designing your VM-based applications, consider future cloud mobility. For example, rather than placing application and data files together use a separate LUN or NFS export for the data. This allows you to migrate the VM and data separately to cloud services.

Encryption for vSphere data

Today, there are increasing demands to protect data at rest through encryption. Although the initial focus was on financial and healthcare information, there is growing interest in protecting all information, whether it's stored in files, databases, or other data types.

Systems running ONTAP software make it easy to protect any data with at-rest encryption. NetApp Storage Encryption (NSE) uses self-encrypting disk drives with ONTAP to protect SAN and NAS data. NetApp also offers NetApp Volume Encryption and NetApp Aggregate Encryption as a simple, software-based approach to encrypt volumes on any disk drives. This software encryption doesn't require special disk drives or external key managers and is available to ONTAP customers at no additional cost. You can upgrade and start using it without any disruption to your clients or applications, and they are validated to the FIPS 140-2 level 1 standard, including the onboard key manager.

There are several approaches for protecting the data of virtualized applications running on VMware vSphere. One approach is to protect the data with software inside the VM at the guest OS level. Newer hypervisors such as vSphere 6.5 now support encryption at the VM level as another alternative. However, NetApp software encryption is simple and easy and has these benefits:

- **No effect on the virtual server CPU.** Some virtual server environments need every available CPU cycle for their applications, yet tests have shown up to 5x CPU resources are needed with hypervisor-level encryption. Even if the encryption software supports Intel's AES-NI instruction set to offload encryption workload (as NetApp software encryption does), this approach might not be feasible due to the requirement for new CPUs that are not compatible with older servers.
- **Onboard key manager included.** NetApp software encryption includes an onboard key manager at no additional cost, which makes it easy to get started without high-availability key management servers that are complex to purchase and use.

- **No effect on storage efficiency.** Storage efficiency techniques such as deduplication and compression are widely used today and are key to using flash disk media cost-effectively. However, encrypted data cannot typically be deduplicated or compressed. NetApp hardware and storage encryption operate at a lower level and allow full use of industry-leading NetApp storage efficiency features, unlike other approaches.
- **Easy datastore granular encryption.** With NetApp Volume Encryption, each volume gets its own AES 256-bit key. If you need to change it, you can do so with a single command. This approach is great if you have multiple tenants or need to prove independent encryption for different departments or apps. This encryption is managed at the datastore level, which is a lot easier than managing individual VMs.


It's simple to get started with software encryption. After the license is installed, simply configure the onboard key manager by specifying a passphrase and then either create a new volume or do a storage-side volume move to enable encryption. NetApp is working to add more integrated support for encryption capabilities in future releases of its VMware tools.

Active IQ Unified Manager

Active IQ Unified Manager provides visibility into the VMs in your virtual infrastructure and enables monitoring and troubleshooting storage and performance issues in your virtual environment.

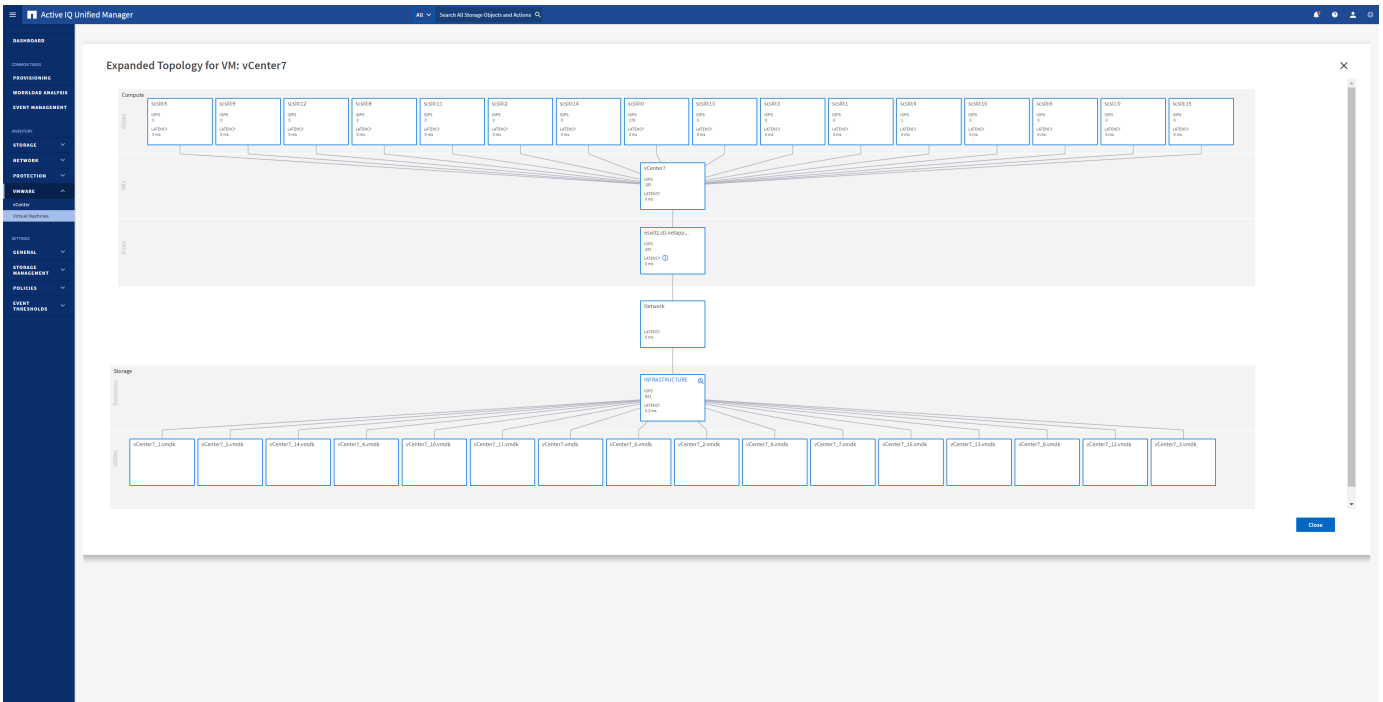
A typical virtual infrastructure deployment on ONTAP has various components that are spread across compute, network, and storage layers. Any performance lag in a VM application might occur due to a combination of latencies faced by the various components at the respective layers.

The following screenshot shows the Active IQ Unified Manager Virtual Machines view.



Unified Manager presents the underlying sub-system of a virtual environment in a topological view for determining whether a latency issue has occurred in the compute node, network, or storage. The view also highlights the specific object that causes the performance lag for taking remedial steps and addressing the underlying issue.

The following screenshot shows the AIQUM expanded topology.



Storage policy based management and vVols

VMware vSphere APIs for Storage Awareness (VASA) make it easy for a storage administrator to configure datastores with well-defined capabilities and let the VM administrator use those whenever needed to provision VMs without having to interact with each other.


It's worth taking a look at this approach to see how it can streamline your virtualization storage operations and avoid a lot of trivial work.

Prior to VASA, VM administrators could define VM storage policies, but they had to work with the storage administrator to identify appropriate datastores, often by using documentation or naming conventions. With VASA, the storage administrator can define a range of storage capabilities, including performance, tiering, encryption, and replication. A set of capabilities for a volume or a set of volumes is called a storage capability profile (SCP).

The SCP supports minimum and/or maximum QoS for a VM's data vVols. Minimum QoS is supported only on AFF systems. ONTAP tools for VMware vSphere includes a dashboard that displays VM granular performance and logical capacity for vVols on ONTAP systems.

The following figure depicts ONTAP tools for VMware vSphere 9.8 vVols dashboard.

! The dashboard displays IOPS, latency, throughput, and logical space values obtained from ONTAP.



After the storage capability profile is defined, it can be used to provision VMs using the storage policy that identifies its requirements. The mapping between the VM storage policy and the datastore storage capability profile allows vCenter to display a list of compatible datastores for selection. This approach is known as storage policy based management.

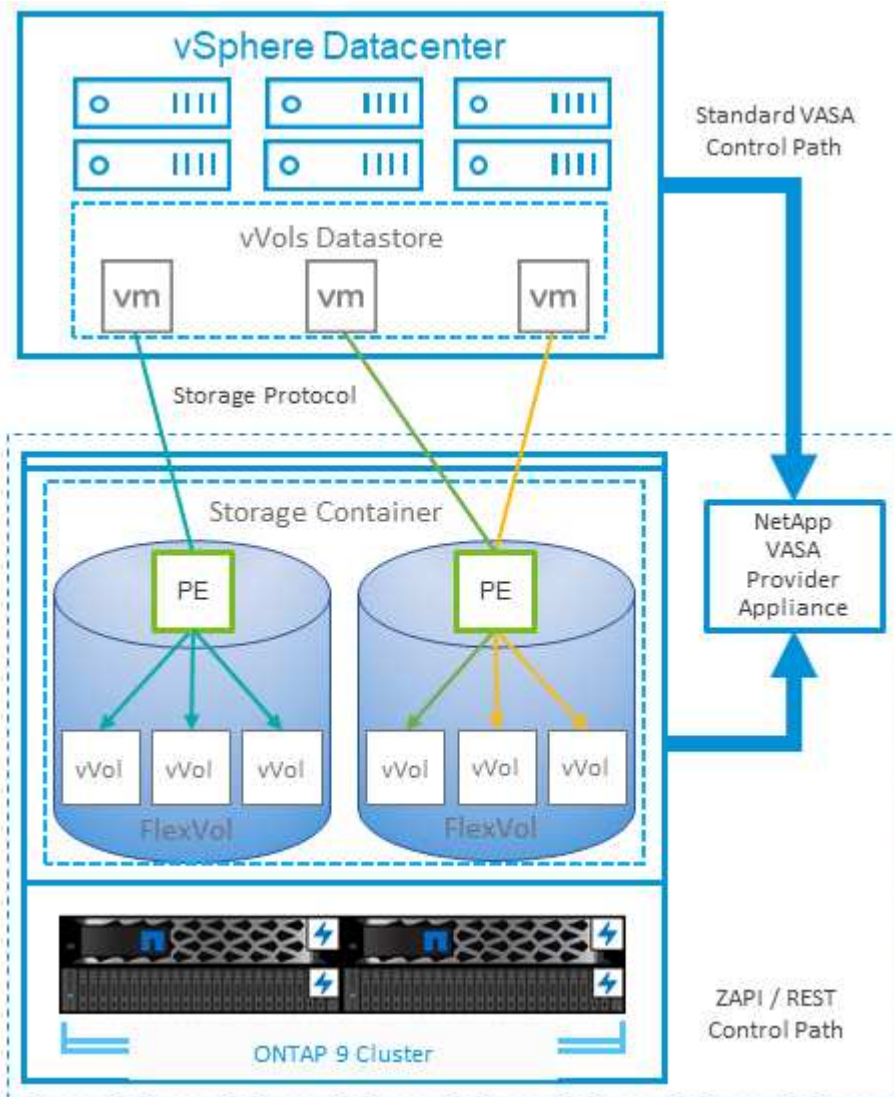
VASA provides the technology to query storage and return a set of storage capabilities to vCenter. VASA vendor providers supply the translation between the storage system APIs and constructs and the VMware APIs that are understood by vCenter. NetApp's VASA Provider for ONTAP is offered as part of the ONTAP tools for VMware vSphere appliance VM, and the vCenter plug-in provides the interface to provision and manage vVol datastores, as well as the ability to define storage capability profiles (SCPs).

ONTAP supports both VMFS and NFS vVol datastores. Using vVols with SAN datastores brings some of the benefits of NFS such as VM-level granularity. Here are some best practices to consider, and you can find additional information in [TR-4400](#):

- A vVol datastore can consist of multiple FlexVol volumes on multiple cluster nodes. The simplest approach is a single datastore, even when the volumes have different capabilities. SPBM makes sure that a compatible volume is used for the VM. However, the volumes must all be part of a single ONTAP SVM and accessed using a single protocol. One LIF per node for each protocol is sufficient. Avoid using multiple ONTAP releases within a single vVol datastore because the storage capabilities might vary across releases.
- Use the ONTAP tools for VMware vSphere plug-in to create and manage vVol datastores. In addition to managing the datastore and its profile, it automatically creates a protocol endpoint to access the vVols if needed. If LUNs are used, note that LUN PEs are mapped using LUN IDs 300 and higher. Verify that the ESXi host advanced system setting `Disk.MaxLUN` allows a LUN ID number that is higher than 300 (the default is 1,024). Do this step by selecting the ESXi host in vCenter, then the Configure tab, and find `Disk.MaxLUN` in the list of Advanced System Settings.
- Do not install or migrate VASA Provider, vCenter Server (appliance or Windows based), or ONTAP tools for VMware vSphere itself onto a vVols datastore, because they are then mutually dependent, limiting your ability to manage them in the event of a power outage or other data center disruption.
- Back up the VASA Provider VM regularly. At a minimum, create hourly snapshots of the traditional datastore that contains VASA Provider. For more about protecting and recovering the VASA Provider, see

this [KB article](#).

The following figure shows vVols components.



VMware Storage Distributed Resource Scheduler

VMware Storage Distributed Resource Scheduler (SDRS) is a vSphere feature that places VMs on storage based on the current I/O latency and space usage.

It then moves the VM or VMDKs nondisruptively between the datastores in a datastore cluster (also referred to as a pod), selecting the best datastore in which to place the VM or VMDKs in the datastore cluster. A datastore cluster is a collection of similar datastores that are aggregated into a single unit of consumption from the vSphere administrator's perspective.

When using SDRS with ONTAP tools for VMware vSphere, you must first create a datastore with the plug-in, use vCenter to create the datastore cluster, and then add the datastore to it. After the datastore cluster is created, additional datastores can be added to the datastore cluster directly from the provisioning wizard on the Details page.

Other ONTAP best practices for SDRS include the following:

- All datastores in the cluster should use the same type of storage (such as SAS, SATA, or SSD), be either all VMFS or NFS datastores, and have the same replication and protection settings.
- Consider using SDRS in default (manual) mode. This approach allows you to review the recommendations and decide whether to apply them or not. Be aware of these effects of VMDK migrations:
 - When SDRS moves VMDKs between datastores, any space savings from ONTAP cloning or deduplication are lost. You can rerun deduplication to regain these savings.
 - After SDRS moves VMDKs, NetApp recommends recreating the snapshots at the source datastore because space is otherwise locked by the VM that was moved.
 - Moving VMDKs between datastores on the same aggregate has little benefit, and SDRS does not have visibility into other workloads that might share the aggregate.

Recommended ESXi host and other ONTAP settings

NetApp has developed a set of optimal ESXi host settings for both NFS and block protocols. Specific guidance is also provided for multipathing and HBA timeout settings for proper behavior with ONTAP based on NetApp and VMware internal testing.

These values are easily set using ONTAP tools for VMware vSphere: From the Summary dashboard, click Edit Settings in the Host Systems portlet or right-click the host in vCenter, then navigate to ONTAP tools > Set Recommended Values.

Here are the currently recommended host settings with the 9.8-9.13 releases.

Host Setting	NetApp Recommended Value	Reboot Required
ESXi Advanced Configuration		
VMFS3.HardwareAcceleratedLocking	Keep default (1)	No
VMFS3.EnableBlockDelete	Keep default (0), but can be changed if needed. For more information, see VMware KB 2007427	No
VMFS3.EnableVMFS6Unmap	Keep default (1) For more information, see VMware vSphere APIs: Array Integration (VAAI)	No
NFS Settings		
Net.TcpipHeapSize	vSphere 6.0 or later, set to 32. All other NFS configurations, set to 30	Yes
Net.TcpipHeapMax	Set to 512MB for most vSphere 6.X releases. Set to 1024MB for 6.5U3, 6.7U3, and 7.0 or later.	Yes
NFS.MaxVolumes	vSphere 6.0 or later, set to 256 All other NFS configurations set to 64.	No
NFS41.MaxVolumes	vSphere 6.0 or later, set to 256.	No
NFS.MaxQueueDepth ¹	vSphere 6.0 or later, set to 128	Yes
NFS.HeartbeatMaxFailures	Set to 10 for all NFS configurations	No

NFS.HeartbeatFrequency	Set to 12 for all NFS configurations	No
NFS.HeartbeatTimeout	Set to 5 for all NFS configurations.	No
SunRPC.MaxConnPerLP	vSphere 7.0 or later, set to 128.	No
FC/FCoE Settings		
Path selection policy	Set to RR (round robin) when FC paths with ALUA are used. Set to FIXED for all other configurations. Setting this value to RR helps provide load balancing across all active/optimized paths. The value FIXED is for older, non-ALUA configurations and helps prevent proxy I/O. In other words, it helps keep I/O from going to the other node of a high-availability (HA) pair in an environment that has Data ONTAP operating in 7-Mode	No
Disk.QFullSampleSize	Set to 32 for all configurations. Setting this value helps prevent I/O errors.	No
Disk.QFullThreshold	Set to 8 for all configurations. Setting this value helps prevent I/O errors.	No
Emulex FC HBA timeouts	Use the default value.	No
QLogic FC HBA timeouts	Use the default value.	No
iSCSI Settings		
Path selection policy	Set to RR (round robin) for all iSCSI paths. Setting this value to RR helps provide load balancing across all active/optimized paths.	No
Disk.QFullSampleSize	Set to 32 for all configurations. Setting this value helps prevent I/O errors	No
Disk.QFullThreshold	Set to 8 for all configurations. Setting this value helps prevent I/O errors.	No



1 - NFS advanced configuration option MaxQueueDepth may not work as intended when using VMware vSphere ESXi 7.0.1 and VMware vSphere ESXi 7.0.2. Please reference [VMware KB 86331](#) for more information.

ONTAP tools also specify certain default settings when creating ONTAP FlexVol volumes and LUNs:

ONTAP Tool	Default Setting
Snapshot reserve (-percent-snapshot-space)	0
Fractional reserve (-fractional-reserve)	0
Access time update (-atime-update)	False
Minimum readahead (-min-readahead)	False

Scheduled snapshots	None
Storage efficiency	Enabled
Volume guarantee	None (thin provisioned)
Volume Autosize	grow_shrink
LUN space reservation	Disabled
LUN space allocation	Enabled

Multipath settings for performance

While not currently configured by available ONTAP tools, NetApp suggests these configuration options:

- In high-performance environments or when testing performance with a single LUN datastore, consider changing the load balance setting of the round-robin (VMW_PSP_RR) path selection policy (PSP) from the default IOPS setting of 1000 to a value of 1. See VMware KB [2069356](#) for more info.
- In vSphere 6.7 Update 1, VMware introduced a new latency load balance mechanism for the Round Robin PSP. The new option considers I/O bandwidth and path latency when selecting the optimal path for I/O. You might benefit from using it in environments with non-equivalent path connectivity, such as cases with more network hops on one path than another, or when using a NetApp All SAN Array system. See [Path Selection Plug-Ins and Policies](#) for more information.

Additional documentation

For FCP and iSCSI with vSphere 7, more details can be found at [Use VMware vSphere 7.x with ONTAP](#)

For FCP and iSCSI with vSphere 8, more details can be found at [Use VMware vSphere 8.x with ONTAP](#)

For NVMe-oF with vSphere 7, more details can be found at [For NVMe-oF, more details can be found at NVMe-oF Host Configuration for ESXi 7.x with ONTAP](#)

For NVMe-oF with vSphere 8, more details can be found at [For NVMe-oF, more details can be found at NVMe-oF Host Configuration for ESXi 8.x with ONTAP](#)

Virtual Volumes (vVols) with ONTAP

Overview

ONTAP has been a leading storage solution for VMware vSphere environments for over two decades and continues to add innovative capabilities to simplify management while reducing costs.

This document covers ONTAP capabilities for VMware vSphere Virtual Volumes (vVols), including the latest product information and use cases along with best practices and other information to streamline deployment and reduce errors.



This documentation replaces previously published technical reports *TR-4400: VMware vSphere Virtual Volumes (vVols) with ONTAP*

Best practices supplement other documents such as guides and compatibility lists. They are developed based on lab testing and extensive field experience by NetApp engineers and customers. They might not be the only practices that work or are supported but are generally the simplest solutions that meet the needs of most customers.



This document has been updated to include new vVols features found in vSphere 8.0 update 1 which are supported with the ONTAP tools 9.12 release.

Virtual Volumes (vVols) overview

NetApp began working with VMware to support vSphere APIs for Storage Awareness (VASA) for vSphere 5 in 2012. This early VASA Provider allowed for the definition of storage capabilities in a profile that could be used to filter datastores when provisioning and for checking compliance with the policy afterwards. Over time this evolved to add new capabilities to enable more automation in provisioning, as well as adding Virtual Volumes or vVols, where individual storage objects are used for virtual machine files and virtual disks. These objects could be LUNs, files, and now with vSphere 8 - NVMe namespaces. NetApp worked closely with VMware as a reference partner for vVols released with vSphere 6 in 2015, and again as a design partner for vVols using NVMe over fabrics in vSphere 8. NetApp continues to enhance vVols to take advantage of the latest capabilities in ONTAP.

There are several components to be aware of:

VASA Provider

This is the software component that handles communication between VMware vSphere and the storage system. For ONTAP, the VASA Provider runs in an appliance known as ONTAP tools for VMware vSphere (ONTAP tools for short). ONTAP tools also includes a vCenter plugin, a storage replication adapter (SRA) for VMware Site Recovery Manager, and REST API server for building your own automation. Once ONTAP tools is configured and registered with vCenter, there is little need to directly interact with the ONTAP system anymore, since nearly all of your storage needs can be managed from directly within the vCenter UI, or through REST API automation.

Protocol Endpoint (PE)

The protocol endpoint is a proxy for I/O between the ESXi hosts and the vVols datastore. The ONTAP VASA Provider creates these automatically, either one protocol endpoint LUN (4MB in size) per FlexVol volume of the vVols datastore, or one NFS mount point per NFS interface (LIF) on the storage node hosting a FlexVol volume in the datastore. The ESXi host mounts these protocol endpoints directly rather than individual vVol LUNs and virtual disk files. There is no need to manage the protocol endpoints as they are created, mounted, unmounted, and deleted automatically by the VASA Provider, along with any necessary interface groups or export policies.

Virtual Protocol Endpoint (vPE)

New in vSphere 8, when using NVMe over Fabrics (NVMe-oF) with vVols, the concept of a protocol endpoint is no longer relevant in ONTAP. Instead, a virtual PE is instantiated automatically by the ESXi host for each ANA group as soon as the first VM is powered on. ONTAP automatically creates ANA groups for each FlexVol volume used by the datastore.

An additional advantage to using NVMe-oF for vVols is that there are no bind requests required of the VASA Provider. Instead, the ESXi host handles vVol binding functionality internally based on the vPE. This reduces the opportunity for a vVol bind storm to impact service.

For more information, see [NVMe and Virtual Volumes](#) on [VMware.com](#)

Virtual Volume datastore

The Virtual Volume datastore is a logical datastore representation of a vVols container which is created and maintained by a VASA Provider. The container represents a pool of storage capacity provisioned from storage systems managed by the VASA Provider. ONTAP tools supports allocating multiple FlexVol volumes (referred to as backing volumes) to a single vVols datastore, and these vVols datastores can span multiple nodes in an ONTAP cluster, combining flash and hybrid systems with different capabilities. The administrator may create new FlexVol volumes using the provisioning wizard or REST API, or select pre-created FlexVol volumes for backing storage if they are available.

Virtual Volumes (vVols)

vVols are the actual virtual machine files and disks stored in the vVols datastore. Using the term vVol (singular) is referring to a single specific file, LUN, or namespace. ONTAP creates NVMe namespaces, LUNs or files depending on what protocol the datastore uses. There are several distinct types of vVols; most common are Config (metadata files), Data (virtual disk or VMDK), and Swap (created when VM is powered on). vVols protected by VMware VM encryption will be of type Other. VMware VM encryption should not be confused with ONTAP volume or aggregate encryption.

Policy based management

VMware vSphere APIs for Storage Awareness (VASA) make it easy for a VM administrator to use whatever storage capabilities are needed to provision VMs without having to interact with their storage team. Prior to VASA, VM administrators could define VM storage policies, but had to work with their storage administrators to identify appropriate datastores, often by using documentation or naming conventions. With VASA, vCenter administrators with the appropriate permissions can define a range of storage capabilities which vCenter users can then use to provision VMs. The mapping between VM storage policy and datastore storage capability profile allows vCenter to display a list of compatible datastores for selection, as well as enabling other technologies like Aria (formerly known as vRealize) Automation or Tanzu Kubernetes Grid to automatically select storage from an assigned policy. This approach is known as storage policy based management. While storage capability profiles and policies may also be used with traditional datastores, our focus here is on vVols datastores.

There are two elements:

Storage Capability Profile (SCP)

A storage capability profile (SCP) is a form of storage template that allows the vCenter admin to define what storage features they require without actually needing to understand how to manage those features in ONTAP. By taking a template style approach, it allows the admin to easily deliver storage services in a consistent and predictable way. Capabilities described in an SCP include performance, protocol, storage efficiency, and other features. Specific features vary by version. They are created using the ONTAP tools for VMware vSphere menu within the vCenter UI. You can also use REST APIs to create SCPs. They may be manually created by selecting individual capabilities, or automatically generated from existing (traditional) datastores.

VM Storage Policy

VM Storage Policies are created in vCenter under Policies and Profiles. For vVols, create a ruleset using rules from the NetApp vVols storage type provider. ONTAP tools provides a simplified approach by allowing you to simply select an SCP rather than forcing you to specify individual rules.

As mentioned above, using policies can help streamline the task of provisioning a volume. Simply select an appropriate policy, and the VASA Provider will show vVols datastores that support that policy and place the vVol into an individual FlexVol volume that is compliant (Figure 1).

Deploy VM using Storage Policy

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Select storage**
- 5 Select compatibility
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy

Disable Storage DRS for this virtual machine

	Name	Storage Compatibility	Capacity	Provisioned	Free	Type	Clu
<input checked="" type="radio"/>	vVolsiSCSI	Compatible	100 GB	40.74 GB	64.88 GB	vVol	
<input type="radio"/>	vVolsNFS2202...	Compatible	2 TB	36.88 GB	1.96 TB	vVol	
<input type="radio"/>	local-esx01	Incompatible	3.63 TB	1.46 GB	3.63 TB	VMFS 6	
<input type="radio"/>	local-esx07	Incompatible	1.81 TB	3.85 GB	1.81 TB	VMFS 6	
<input type="radio"/>	local-esx08	Incompatible	1.69 TB	1.43 GB	1.69 TB	VMFS 6	
<input type="radio"/>	local-esx09	Incompatible	1.81 TB	3.85 GB	1.81 TB	VMFS 6	
<input type="radio"/>	local-esx15	Incompatible	3.63 TB	1.46 GB	3.63 TB	VMFS 6	
<input type="radio"/>	tier001_ds	Incompatible	22 TB	23.73 TB	18.09 TB	NFS v3	

CANCEL

BACK

NEXT

Once a VM is provisioned, the VASA Provider will continue to check compliance, and alert the VM administrator with an alarm in vCenter when the backing volume is no longer compliant with the policy (Figure 2).

VM Storage Policy Compliance

Storage Policies



VM Storage Policies

AFF_VASA10

VM Storage Policy Compliance

Noncompliant

Last Checked Date

5/20/2022, 12:59:35 PM

VM Replication Groups

[CHECK COMPLIANCE](#)

NetApp vVols support

ONTAP has supported the VASA specification since its initial release in 2012. While other NetApp storage systems may support VASA, this document focuses on currently supported releases of ONTAP 9.

ONTAP

In addition to ONTAP 9 on AFF, ASA, and FAS systems, NetApp supports VMware workloads on ONTAP Select, Amazon FSx for NetApp with VMware Cloud on AWS, Azure NetApp Files with Azure VMware Solution, Cloud Volumes Service with Google Cloud VMware Engine, and NetApp Private Storage in Equinix, but specific functionality may vary based on service provider and available network connectivity. Access from vSphere guests to data stored in those configurations as well as Cloud Volumes ONTAP is also available.

At the time of publication, hyperscaler environments are limited to traditional NFS v3 datastores only, therefore, vVols are only available with on-premises ONTAP systems, or cloud connected systems that offer the full functionality of an on-premises systems such as those hosted by NetApp partners and services providers around the world.

For more information about ONTAP, see [ONTAP product documentation](#)

For more information about ONTAP and VMware vSphere best practices, see [TR-4597](#)

Benefits of using vVols with ONTAP

When VMware introduced vVols support with VASA 2.0 in 2015 they described it as “an integration and management framework delivering a new operational model for external storage (SAN/NAS).” This operational model offers several benefits together with ONTAP storage.

Policy based management

As covered in section 1.2, policy based management allows VMs to be provisioned and subsequently managed using pre-defined policies. This can help IT operations in several ways:

- **Increase velocity.** ONTAP tools eliminates the requirement for the vCenter administrator to open tickets with the storage team for storage provisioning activities. However, ONTAP tools RBAC roles in vCenter and on the ONTAP system still allow for independent teams (such as storage teams), or independent activities by the same team by restricting access to specific functions if desired.
- **Smarter provisioning.** Storage system capabilities can be exposed through the VASA APIs, allowing provisioning workflows to take advantage of advanced capabilities without the VM administrator needing to understand how to manage the storage system.
- **Faster provisioning.** Different storage capabilities can be supported in a single datastore and automatically selected as appropriate for a VM based on the VM policy.
- **Avoid mistakes.** Storage and VM policies are developed in advance and applied as needed without having to customize storage each time a VM is provisioned. Compliance alarms are raised when storage capabilities drift from the defined policies. As previously mentioned, SCPs make the initial provisioning predictable and repeatable, while basing VM storage policies on the SCPs guarantees accurate placement.
- **Better capacity management.** VASA and ONTAP tools make it possible to view storage capacity down to the individual aggregate level if needed and provide multiple layers of alerting in the event capacity starts to run low.

VM granular management on the modern SAN

SAN storage systems using Fibre Channel and iSCSI were the first to be supported by VMware for ESX, but they have lacked the ability to manage individual VM files and disks from the storage system. Instead, LUNs are provisioned and VMFS manages the individual files. This makes it difficult for the storage system to directly manage individual VM storage performance, cloning, and protection. vVols bring storage granularity that customers using NFS storage already enjoy, with the robust, high performance SAN capabilities of ONTAP.

Now, with vSphere 8 and ONTAP tools for VMware vSphere 9.12 and later, those same granular controls used by vVols for legacy SCSI based protocols are now available in the modern Fibre Channel SAN using NVMe over Fabrics for even greater performance at scale. With vSphere 8.0 update 1, it is now possible to deploy a complete end-to-end NVMe solution using vVols without any I/O translation in the hypervisor storage stack.

Greater storage offload capabilities

While VAAI offers a variety of operations that are offloaded to storage, there are some gaps that are addressed by the VASA Provider. SAN VAAI is not able to offload VMware managed snapshots to the storage system. NFS VAAI can offload VM managed snapshots, but there are limitations placed a VM with storage native snapshots. Since vVols use individual LUNs, namespaces, or files for virtual machine disks, ONTAP can quickly and efficiently clone the files or LUNs to create VM-granular snapshots that no longer require delta files. NFS VAAI also does not support offloading clone operations for hot (powered on) Storage vMotion migrations. The VM must be powered off to allow offload of the migration when using VAAI with traditional NFS datastores. The VASA Provider in ONTAP tools allows for near instant, storage efficient clones for hot and cold migrations, and it also supports near instant copies for cross-volume migrations of vVols. Because of these significant storage efficiency benefits, you may be able to take full advantage of vVols workloads under the [Efficiency Guarantee](#) program. Likewise, if cross volume clones using VAAI don't meet your requirements, you will likely be able to solve your business challenge thanks to the improvements in the copy experience with vVols.

Common use cases for vVols

In addition to these benefits, we also see these common use cases for vVol storage:

- **On-Demand provisioning of VMs**
 - Private cloud or service provider IaaS.
 - Leverage automation and orchestration via the Aria (formerly vRealize) suite, OpenStack, etc.
- **First Class Disks (FCDs)**
 - VMware Tanzu Kubernetes Grid [TKG] persistent volumes.
 - Provide Amazon EBS-like services through independent VMDK lifecycle management.
- **On-Demand Provisioning of Temporary VMs**
 - Test/dev labs
 - Training environments

Common benefits with vVols

When used to their full advantage, such as in the above use cases, vVols provide the following specific improvements:

- Clones are quickly created within a single volume, or across multiple volumes in an ONTAP cluster, which is an advantage when compared to traditional VAAI enabled clones. They are also storage efficient. Clones within a volume use ONTAP file clone, which are like FlexClone volumes and only store changes from the source vVol file/LUN/namespace. So long-term VMs for production or other application purposes are created quickly, take minimal space, and can benefit from VM level protection (using NetApp SnapCenter plugin for VMware vSphere, VMware managed snapshots or VADP backup) and performance management (with ONTAP QoS).
- vVols are the ideal storage technology when using TKG with the vSphere CSI, providing discrete storage classes and capacities managed by the vCenter administrator.
- Amazon EBS-like services can be delivered through FCDs because an FCD VMDK, as the name suggests, is a first-class citizen in vSphere and has a lifecycle which can be independently managed separate from VMs that it might be attached to.

Using vVols with ONTAP

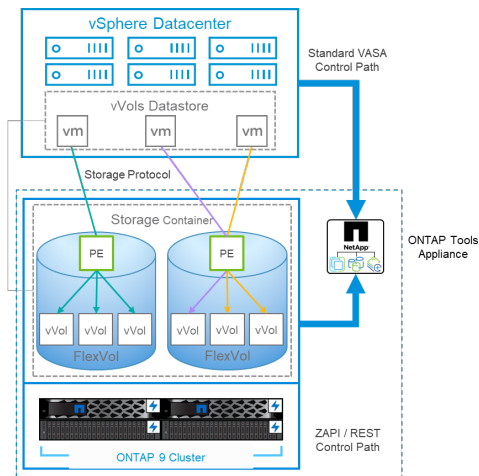
The key to using vVols with ONTAP is the VASA Provider software included as part of the ONTAP tools for VMware vSphere virtual appliance.

ONTAP tools also includes the vCenter UI extensions, REST API server, Storage Replication Adapter for VMware Site Recovery Manager, Monitoring and Host configuration tools, and an array of reports which help you better manage your VMware environment.

Products and Documentation

The ONTAP FlexClone license (included with ONTAP One) and the ONTAP tools appliance are the only additional products required to use vVols with ONTAP. Recent releases of ONTAP tools are supplied as a single unified appliance that runs on ESXi, providing the functionality of what formerly were three different appliances and servers. For vVols, it is important to use the ONTAP tools vCenter UI extensions or REST APIs as general management tools and user interfaces for ONTAP functions with vSphere, together with the VASA Provider which provides specific vVols functionality. The SRA component is included for traditional datastores, but VMware Site Recovery Manager does not use SRA for vVols, instead implementing new services in SRM 8.3 and later which leverage the VASA provider for vVols replication.

ONTAP tools VASA Provider architecture when using iSCSI or FCP



Product Installation

For new installations, deploy the virtual appliance into your vSphere environment. Current releases of ONTAP tools will automatically register themselves with your vCenter and enable the VASA Provider by default. In addition to ESXi host and vCenter Server information, you will also need the IP address configuration details for the appliance. As previously stated, the VASA Provider requires the ONTAP FlexClone license be already installed onto any ONTAP clusters you plan to use for vVols. The appliance has a built-in watchdog to ensure availability, and as a best practice should be configured with VMware High Availability and optionally Fault Tolerance features. See section 4.1 for additional details. Do not install or move the ONTAP tools appliance or vCenter Server appliance (VCSA) to vVols storage as this can prevent the appliances from restarting.

In-place upgrades of ONTAP tools are supported by using the upgrade ISO file available for download on the NetApp Support Site (NSS). Follow the Deployment and Setup Guide instructions to upgrade the appliance.

For sizing your virtual appliance, and understanding the configuration limits, refer to this knowledge base article: [Sizing Guide for ONTAP tools for VMware vSphere](#)

Product Documentation

The following documentation is available to help you deploy ONTAP tools.

For the complete documentation repository, visit [this link to docs.netapp.com](https://docs.netapp.com)

Get started

- [Release notes](#)
- [Learn about ONTAP tools for VMware vSphere](#)
- [ONTAP tools Quick start](#)
- [Deploy ONTAP tools](#)
- [Upgrade ONTAP tools](#)

Use ONTAP tools

- [Provision traditional datastores](#)
- [Provision vVols datastores](#)

- [Configure role-based access control](#)
- [Configure remote diagnostics](#)
- [Configure high availability](#)

Protect and manage datastores

- [Protect traditional datastores with SRM](#)
- [Protect vVols based virtual machines with SRM](#)
- [Monitor traditional datastores and virtual machines](#)
- [Monitor vVols datastores and virtual machines](#)

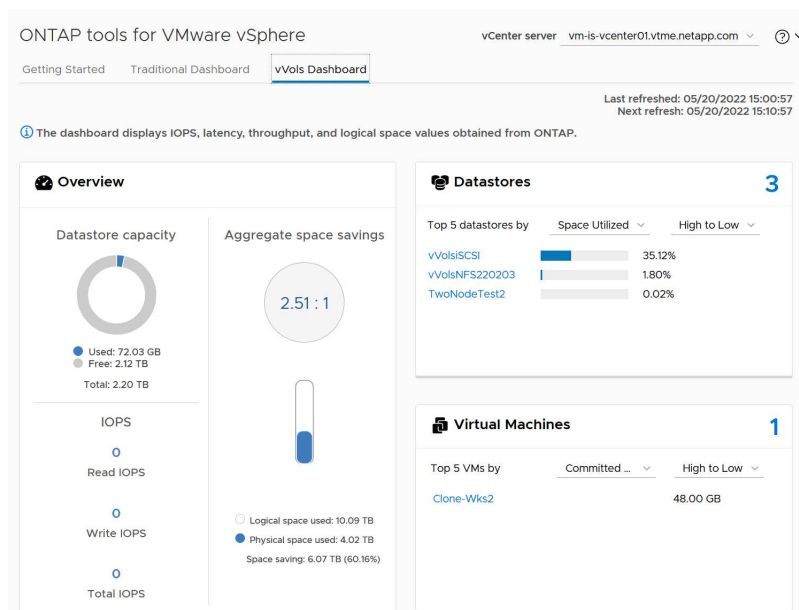
In addition to product documentation, there are Support Knowledgebase articles that may be useful.

- [How to perform a VASA Provider Disaster Recovery - Resolution Guide](#)

VASA Provider Dashboard

The VASA Provider includes a dashboard with performance and capacity information for individual vVols VMs. This information comes directly from ONTAP for the vVol files and LUNs, including latency, IOPS, throughput, and uptime for the top 5 VMs, and latency and IOPS for the top 5 datastores. It is enabled by default when using ONTAP 9.7 or later. It can take up to 30 minutes for initial data to be retrieved and displayed in the dashboard.

ONTAP tools vVols dashboard



Best Practices

Using ONTAP vVols with vSphere is simple and follows published vSphere methods (see Working with Virtual Volumes under vSphere Storage in VMware documentation for your version of ESXi). Here are a few additional practices to consider in conjunction with ONTAP.

Limits

In general, ONTAP supports vVols limits as defined by VMware (see published [Configuration Maximums](#)). The following table summarizes specific ONTAP limits in size and number of vVols. Always check the [NetApp Hardware Universe](#) for updated limits on numbers and sizes of LUNs and files.

ONTAP vVols Limits

Capacity/Feature	SAN (SCSI or NVMe-oF)	NFS
Maximum vVols size	62 TiB*	62 TiB*
Maximum number of vVols per FlexVol volume	1024	2 billion
Maximum number of vVols per ONTAP node	Up to 12,288**	50 billion
Maximum number of vVols per ONTAP pair	Up to 24,576**	50 billion
Maximum number of vVols per ONTAP cluster	Up to 98,304**	No specific cluster limit
Maximum QoS objects (shared policy group and individual vVols service level)	12,000 through ONTAP 9.3; 40,000 with ONTAP 9.4 and later	

- Size limit based on ASA systems or AFF and FAS systems running ONTAP 9.12.1P2 and later.
 - Number of SAN vVols (NVMe namespaces or LUNs) varies based on platform. Always check the [NetApp Hardware Universe](#) for updated limits on numbers and sizes of LUNs and files.

Use ONTAP tools for VMware vSphere’s UI extensions or REST APIs to provision vVols datastores and Protocol Endpoints.

While it’s possible to create vVols datastores with the general vSphere interface, using ONTAP tools will automatically create protocol endpoints as needed, and creates FlexVol volumes using ONTAP best practices and in compliance with your defined storage capability profiles. Simply right click on the host/cluster/datacenter, then select *ONTAP tools* and *Provision datastore*. From there simply choose the desired vVols options in the wizard.

Never store the ONTAP tools appliance or vCenter Server Appliance (VCSA) on a vVols datastore that they are managing.

This can result in a “chicken and egg situation” if you need to reboot the appliances because they won’t be able to rebind their own vVols while they are rebooting. You may store them on a vVols datastore managed by a different ONTAP tools and vCenter deployment.


Avoid vVols operations across different ONTAP releases.

Supported storage capabilities such as QoS, personality and more have changed in various releases of the VASA Provider, and some are dependent on ONTAP release. Using different releases in an ONTAP cluster or moving vVols between clusters with different releases can result in unexpected behavior or compliance alarms.

Zone your Fibre Channel fabric before using NVMe/FC or FCP for vVols.

The ONTAP tools VASA provider takes care of managing FCP and iSCSI igroups as well as NVMe subsystems in ONTAP based on discovered initiators of managed ESXi hosts. However, it does not integrate with Fibre Channel switches to manage zoning. Zoning must be done according to best practices before any provisioning can take place. The following is an example of single initiator zoning to four ONTAP systems:

Single initiator zoning:



Refer to the following documents for more best practices:

[TR-4080 Best practices for modern SAN ONTAP 9](#)

[TR-4684 Implementing and configuring modern SANs with NVMe-oF](#)

Plan your backing FlexVols according to your needs.

It can be desirable to add several backing volumes to your vVols datastore to distribute workload across the ONTAP cluster, to support different policy options, or to increase the number of allowed LUNs or files. However, if maximum storage efficiency is required, then place all your backing volumes on a single aggregate. Or if maximum cloning performance is required, then consider using a single FlexVol volume and keeping your templates or content library in the same volume. The VASA Provider offloads many vVols storage operations to ONTAP, including migration, cloning and snapshots. When this is done within a single FlexVol volume, space efficient file clones are used and are almost instantly available. When this is done across FlexVol volumes, the copies are quickly available and use inline deduplication and compression, but maximum storage efficiency may not be recovered until background jobs run on volumes using background deduplication and compression. Depending on the source and destination, some efficiency may be degraded.

Keep Storage Capability Profiles (SCPs) simple.

Avoid specifying capabilities that aren't required by setting them to Any. This will minimize problems when selecting or creating FlexVol volumes. For example, with VASA Provider 7.1 and earlier, if compression is left at the default SCP setting of No, it will attempt to disable compression, even on an AFF system.

Use the default SCPs as example templates to create your own.

The included SCPs are suitable for most general-purpose uses, but your requirements may be different.

Consider using Max IOPS to control unknown or test VMs.

First available in VASA Provider 7.1, Max IOPS can be used to limit IOPS to a specific vVol for an unknown workload to avoid impact on other, more critical workloads. See Table 4 for more on performance management.

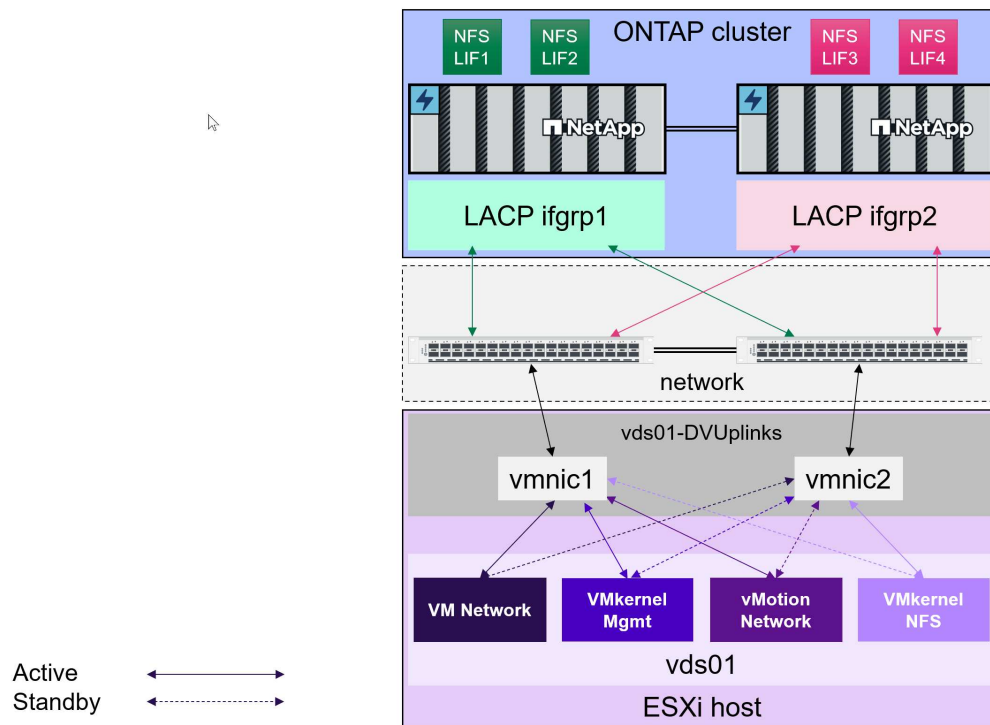
Ensure you have sufficient data LIFs.

Create at least two LIFs per node per HA pair. More may be required based on your workload.

Follow all protocol best practices.

Refer to NetApp and VMware's other best practice guides specific to the protocol you've selected. In general, there are not any changes other than those already mentioned.

Example network configuration using vVols over NFS v3



Deploying vVols Storage

There are several steps to creating vVols storage for your VMs.

The first two steps may not be needed for an existing vSphere environment that uses ONTAP for traditional datastores. You may already be using ONTAP tools for managing, automating, and reporting with your VMFS or traditional NFS based storage. These steps are covered in more detail in the following section.

1. Create the Storage Virtual Machine (SVM) and its protocol configuration. You will select NVMe/FC, NFSv3, NFSv4.1, iSCSI, FCP, or a mix of those options. You may use either ONTAP System Manager wizards, or the cluster shell command line.
 - At least one LIF per node for each switch/fabric connection. As a best practice, create two or more per node for FCP, iSCSI, or NVMe-based protocols.
 - Volumes may be created at this time, but it is simpler to let the *Provision Datastore* wizard create them. The only exception to this rule is if you plan to use vVols replication with VMware Site Recovery Manager. This is easier to set up with pre-existing FlexVol volumes with existing SnapMirror relationships. Be mindful to not enable QoS on any volumes to be used for vVols as this is intended to be managed by SPBM and ONTAP tools.
2. Deploy ONTAP tools for VMware vSphere using the OVA downloaded from the NetApp Support Site.
3. Configure ONTAP tools for your environment.
 - Add the ONTAP cluster to ONTAP tools under *Storage Systems*
 - While ONTAP tools and SRA support both cluster-level and SVM-level credentials, the VASA Provider supports only cluster-level credentials for storage systems. This is because many of the APIs used for vVols are only available at the cluster level. Therefore, if you plan to use vVols, you must add your ONTAP clusters using cluster-scoped credentials.
 - If your ONTAP data LIFs are on different subnets from your VMkernel adapters, then you must add the

VMkernel adapter subnets to the selected subnets list in the settings menu of ONTAP tools. By default, ONTAP tools will secure your storage traffic by only allowing local subnet access.

- The ONTAP tools comes with several pre-defined policies that may be used or see [Managing VMs with policies](#) for guidance on creating SCPs.

4. Use the *ONTAP tools* menu in vCenter to start the *Provision datastore* wizard.
5. Provide a meaningful name and select the desired protocol. You may provide a description of the datastore as well.
6. Select one or more SCPs to be supported by the vVols datastore. This will filter out any ONTAP systems which are unable to match the profile. From the resulting list, select your desired cluster and SVM.
7. Use the wizard to create new FlexVol volumes for each of the specified SCPs or use existing volumes by selecting the appropriate radio button.
8. Create VM policies for each SCP that will be used in the datastore from the *Policies and Profiles* menu in the vCenter UI.
9. Choose the "NetApp.clustered.Data.ONTAP.VP.vvol" storage rule set. The "NetApp.clustered.Data.ONTAP.VP.VASA10" storage rule set is for SPBM support with non-vVols datastores
10. You will specify the Storage Capability Profile by name when creating a VM Storage Policy. While at this step, you may also configure SnapMirror policy matching by using the replication tab, and tag-based matching using the tags tab. Note that tags must already be created in order to be selectable.
11. Create your VMs, selecting the VM Storage Policy and compatible datastore under Select storage.

Migrating VMs from traditional datastores to vVols

Migration of VMs from traditional datastores to a vVols datastore is as simple as moving VMs between traditional datastores. Simply select the VM(s), then select Migrate from the list of Actions, and select a migration type of *change storage only*. Migration copy operations will be offloaded with vSphere 6.0 and later for SAN VMFS to vVols migrations, but not from NAS VMDKs to vVols.

Managing VMs with policies

To automate storage provisioning with policy based management, we need to:

- Define the capabilities of the storage (ONTAP node and FlexVol volume) with Storage Capability Profiles (SCPs).
- Create VM storage policies that map to the defined SCPs.

NetApp has simplified the capabilities and mapping beginning with VASA Provider 7.2 with continuing improvements throughout later versions. This section focuses on this new approach. Earlier releases supported a greater number of capabilities and allowed them to be mapped individually to storage policies, but this approach is no longer supported.

Storage Capability Profile capabilities by ONTAP tools release

SCP Capability	Capability Values	Release Supported	Notes
Compression	Yes, No, Any	All	Mandatory for AFF in 7.2 and later.
Deduplication	Yes, No, Any	All	Mandatory for AFF in 7.2 and later.

SCP Capability	Capability Values	Release Supported	Notes
Encryption	Yes, No, Any	7.2 and later	Selects/creates encrypted FlexVol volume.. ONTAP license required.
Max IOPS	<number>	7.1 and later, but differences	Listed under QoS Policy Group for 7.2 and later. See Performance management with ONTAP tools 9.10 and later for more information.
Personality	A FF, FAS	7.2 and later	FAS also includes other non-AFF systems, such as ONTAP Select. AFF includes ASA.
Protocol	NFS, NFS 4.1, iSCSI, FCP, NVMe/FC, Any	7.1 and earlier, 9.10 and later	7.2-9.8 is effectively "Any". Beginning again in 9.10 where NFS 4.1 and NVMe/FC were added to the original list.
Space Reserve (Thin Provisioning)	Thin, Thick, (Any)	All, but differences	Called Thin Provisioning in 7.1 and earlier, which also allowed value of Any. Called Space Reserve in 7.2. All releases default to Thin.
Tiering Policy	Any, None, Snapshot, Auto	7.2 and later	Used for FabricPool - requires AFF or ASA with ONTAP 9.4 or later. Only Snapshot is recommended unless using an on-premise S3 solution like NetApp StorageGRID.

Creating Storage Capability Profiles

The NetApp VASA Provider comes with several pre-defined SCPs. New SCPs may be created manually, using the vCenter UI, or via automation using REST APIs. By specifying capabilities in a new profile, cloning an existing profile, or by auto-generating profile(s) from existing traditional datastores. This is done using the menus under ONTAP tools. Use *Storage Capability Profiles* to create or clone a profile, and *Storage Mapping* to auto-generate a profile.

Storage Capabilities for ONTAP tools 9.10 and later

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary

General

Specify a name and description for the storage capability profile.

Name:

Description:

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary

Platform

Platform:

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary

Protocol

Protocol:

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary

Performance

None

QoS policy group

Min IOPS:

Max IOPS:

Unlimited

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary

Storage attributes

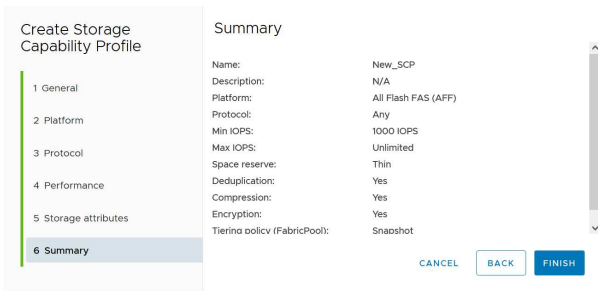
Deduplication:

Compression:

Space reserve:

Encryption:

Tiering policy (FabricPool):



Creating vVols Datastores

Once the necessary SCPs have been created, they may be used to create the vVols datastore (and optionally, FlexVol volumes for the datastore). Right-click on the host, cluster, or datacenter on which you want to create the vVols datastore, then select *ONTAP tools > Provision Datastore*. Select one or more SCPs to be supported by the datastore, then select from existing FlexVol volumes and/or provision new FlexVol volumes for the datastore. Finally, specify the default SCP for the datastore, which will be used for VMs that do not have an SCP specified by policy, as well as for swap vVols (these do not require high performance storage).


Creating VM Storage Policies

VM Storage Policies are used in vSphere to manage optional features such as Storage I/O Control or vSphere Encryption. They are also used with vVols to apply specific storage capabilities to the VM. Use the “NetApp.clustered.Data.ONTAP.VP.vvol” storage type and “ProfileName” rule to apply a specific SCP to VMs through use of the Policy. See [example network configuration using vVols over NFS v3](#) for an example of this with the ONTAP tools VASA Provider. Rules for “NetApp.clustered.Data.ONTAP.VP.VASA10” storage are to be used with non-vVols based datastores.

Earlier releases are similar, but as mentioned in [Storage Capability Profile capabilities by ONTAP tools release](#), your options will vary.

Once the storage policy has been created, it can be used when provisioning new VMs as shown in [Deploy VM using Storage Policy](#). Guidelines for using performance management capabilities with VASA Provider 7.2 are covered in [Performance management with ONTAP tools 9.10 and later](#).

VM storage policy creation with ONTAP tools VASA Provider 9.10

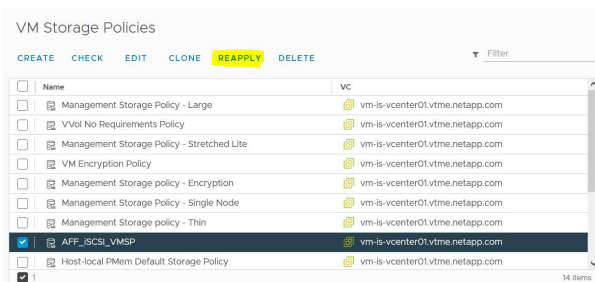


Performance management with ONTAP tools 9.10 and later

- ONTAP tools 9.10 uses its own balanced placement algorithm to place a new vVol in the best FlexVol volume within a vVols datastore. Placement is based on the specified SCP and matching FlexVol volumes. This makes sure that the datastore and backing storage can meet the specified performance requirements.
- Changing Performance capabilities such as Min and Max IOPS requires some attention to the specific configuration.
 - **Min and Max IOPS** may be specified in an SCP and used in a VM Policy.
 - Changing the IOPS in the SCP will not change QoS on the vVols until the VM Policy is edited, and then reapplied to the VMs that use it (see [Storage Capabilities for ONTAP tools 9.10 and later](#)). Or create a new SCP with the desired IOPS and change the policy to use it (and reapply to VMs). Generally it is recommended to simply define separate SCPs and VM storage policies for different tiers of service and simply change the VM storage policy on the VM.

- AFF and FAS personalities have different IOPs settings. Both Min and Max are available on AFF. However non-AFF systems can only use Max IOPs settings.
- In some cases, a vVol may need to be migrated after a policy change (either manually, or automatically by VASA Provider and ONTAP):
 - Some changes require no migration (such as changing Max IOPS, which can be applied immediately to the VM as outlined above).
 - If the policy change cannot be supported by the current FlexVol volume that stores the vVol (for example, the platform does not support the encryption or tiering policy requested), you will need to manually migrate the VM in vCenter.
- ONTAP tools creates individual non-shared QoS policies with currently supported versions of ONTAP. Therefore, each individual VMDK will receive its own allocation of IOPs.

Reapplying VM Storage Policy



Protecting vVols

The following sections outline the procedures and best practices for using VMware vVols with ONTAP storage.

VASA Provider High Availability

The NetApp VASA Provider runs as part of the virtual appliance together with the vCenter plugin and REST API server (formerly known as the Virtual Storage Console [VSC]) and Storage Replication Adapter. If the VASA Provider is not available, VMs using vVols will continue to run. However, new vVols datastores cannot be created, and vVols cannot be created or bound by vSphere. This means that VMs using vVols cannot be powered on as vCenter will not be able to request creation of the swap vVol. And running VMs cannot use vMotion for migration to another host because the vVols cannot be bound to the new host.

VASA Provider 7.1 and later support new capabilities to make sure the services are available when needed. It includes new watchdog processes that monitor VASA Provider and integrated database services. If it detects a failure, it updates the log files and then restarts the services automatically.

Further protection must be configured by the vSphere administrator using the same availability features used to protect other mission critical VMs from faults in software, host hardware and network. No additional configuration is required on the virtual appliance to use these features; simply configure them using standard vSphere approaches. They have been tested and are supported by NetApp.

vSphere High Availability is easily configured to restart a VM on another host in the host cluster in the event of failure. vSphere Fault Tolerance provides higher availability by creating a secondary VM that is continuously replicated and can take over at any point. Additional information on these features is available in the [ONTAP tools for VMware vSphere documentation \(Configure high availability for ONTAP tools\)](#), as well as VMware vSphere documentation (look for vSphere Availability under ESXi and vCenter Server).

configuration, it is not currently supported with vVols.

Refer to these guides for more information on NetApp MetroCluster:

[TR-4689 MetroCluster IP Solution architecture and design](#)

[TR-4705 NetApp MetroCluster Solution architecture and design](#)

[VMware KB 2031038 VMware vSphere Support with NetApp MetroCluster](#)

vVols Backup Overview

There are several approaches to protecting VMs such as using in-guest backup agents, attaching VM data files to a backup proxy, or using defined APIs such as VMware VADP. vVols may be protected using the same mechanisms and many NetApp partners support VM backups, including vVols.

As mentioned earlier, VMware vCenter managed snapshots are offloaded to space efficient and fast ONTAP file/LUN clones. These may be used for quick, manual backups, but are limited by vCenter to a maximum of 32 snapshots. You may use vCenter to take snapshots and revert as needed.

Beginning with SnapCenter Plugin for VMware vSphere (SCV) 4.6 when used in conjunction with ONTAP tools 9.10 and later adds support for crash consistent backup and recovery of vVols based VMs leveraging ONTAP FlexVol volume snapshots with support for SnapMirror and SnapVault replication. Up to 1023 snapshots are supported per volume. SCV can also store more snapshots with longer retention on secondary volumes using SnapMirror with a mirror-vault policy.

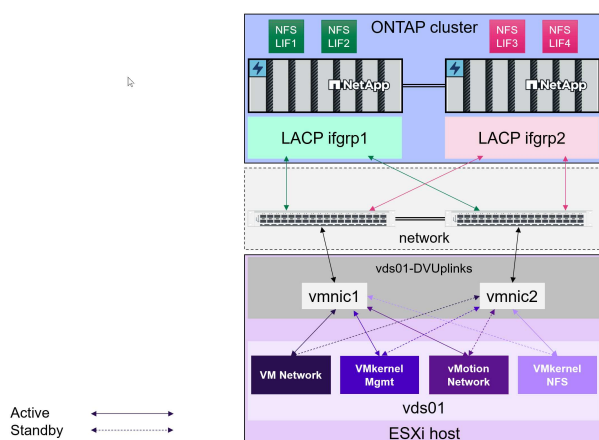
vSphere 8.0 support was introduced with SCV 4.7, which used an isolated local plugin architecture. vSphere 8.0U1 support was added to SCV 4.8 which fully transitioned to the new remote plugin architecture.

vVols Backup with SnapCenter plugin for VMware vSphere

With NetApp SnapCenter you can now create resource groups for vVols based on tags and/or folders to automatically take advantage of ONTAP's FlexVol based snapshots for vVols based VMs. This allows you to define backup and recovery services which will protect VMs automatically as they get dynamically provisioned within your environment.

SnapCenter plugin for VMware vSphere is deployed as a standalone appliance registered as a vCenter extension, managed through the vCenter UI or via REST APIs for backup and recovery service automation.

SnapCenter architecture



Since the other SnapCenter plugins don't yet support vVols at the time of this writing, we will focus on the standalone deployment model in this document.

Because SnapCenter uses ONTAP FlexVol snapshots there is no overhead placed on vSphere, nor is there any performance penalty as one might see with traditional VMs using vCenter managed snapshots. Furthermore, because SCV's functionality is exposed via REST APIs, it makes it easy to create automated workflows using tools like VMware Aria Automation, Ansible, Terraform, and virtually any other automation tool that is capable of using standard REST APIs.

For information on SnapCenter REST APIs, see [Overview of REST APIs](#)

For information on SnapCenter Plug-in for VMware vSphere REST APIs, see [SnapCenter Plug-in for VMware vSphere REST APIs](#)

Best Practices

The following best practices can help you get the most out of your SnapCenter deployment.

- SCV supports both vCenter Server RBAC and ONTAP RBAC and includes predefined vCenter roles which are automatically created for you when the plugin is registered. You can read more about the supported types of RBAC [here](#).
 - Use the vCenter UI to assign least privileged account access using the predefined roles described [here](#).
 - If you use SCV with SnapCenter Server, you must assign the *SnapCenterAdmin* role.
 - ONTAP RBAC refers to the user account used to add and manage the storage systems used by SCV. ONTAP RBAC doesn't apply to vVols based backups. Read more about ONTAP RBAC and SCV [here](#).
- Replicate your backup datasets to a second system using SnapMirror for complete replicas of source volumes. As previously mentioned, you may also use mirror-vault policies for longer term retention of backup data independent of source volume snapshot retention settings. Both mechanisms are supported with vVols.
- Because SCV also requires ONTAP tools for VMware vSphere for vVols functionality, always check the NetApp Interoperability Matrix Tool (IMT) for specific version compatibility
- If you are using vVols replication with VMware SRM, be mindful of your policy RPO and backup schedule
- Design your backup policies with retention settings that meet your organizations defined recovery point objectives (RPOs)
- Configure notification settings on your resource groups to be notified of the status when backups run (see figure 10 below)

Resource group notification options

Edit Resource Group

1. General info & notification

- 2. Resource
- 3. Spanning disks
- 4. Policies
- 5. Schedules
- 6. Summary

vCenter Server:

Name:

Description:

Notification:

Email send from:

Email send to:

Email subject:

Latest Snapshot name: Enable _recent suffix for latest Snapshot Copy

Custom snapshot format: Use custom name format for Snapshot copy

Note that the Plug-in for VMware vSphere cannot do the following:

Get started with SCV using these documents

[Learn about SnapCenter Plug-in for VMware vSphere](#)

[Deploy SnapCenter Plug-in for VMware vSphere](#)

Troubleshooting

There are several troubleshooting resources available with additional information.

NetApp Support Site

In addition to a variety of Knowledgebase articles for NetApp virtualization products, the NetApp Support site also offers a convenient landing page for the [ONTAP tools for VMware vSphere](#) product. This portal provides links to articles, downloads, technical reports, and VMware Solutions Discussions on NetApp Community. It is available at:

[NetApp Support Site](#)

Additional solution documentation is available here:

[NetApp Solutions for Virtualization](#)

Product Troubleshooting

The various components of ONTAP tools, such as the vCenter plugin, VASA Provider, and Storage Replication Adapter are all documented together in the NetApp documents repository. However, each has a separate subsection of the Knowledge Base and may have specific troubleshooting procedures. These address the most common issues that may be encountered with the VASA Provider.

VASA Provider UI Problems

Occasionally the vCenter vSphere Web Client encounters problems with the Serenity components, causing the VASA Provider for ONTAP menu items not to display. See [Resolving VASA Provider registration issues](#) in the Deployment Guide, or this Knowledgebase [article](#).

vVols Datastore Provisioning Fails

Occasionally vCenter services may time out when creating the vVols datastore. To correct it, restart the vmware-sps service, and re-mount the vVols datastore using the vCenter menus (Storage > New Datastore). This is covered under vVols datastore provisioning fails with vCenter Server 6.5 in the Administration Guide.

Upgrading Unified Appliance Fails to Mount ISO

Due to a bug in vCenter, the ISO used to upgrade the Unified Appliance from one release to the next may fail to mount. If the ISO is able to be attached to the appliance in vCenter, follow the process in this Knowledgebase [article](#) to resolve.

VMware Site Recovery Manager with ONTAP

VMware Site Recovery Manager with ONTAP

ONTAP has been a leading storage solution for VMware vSphere environments since its introduction into the modern datacenter in 2002, and it continues to add innovative capabilities to simplify management while reducing costs.

This document introduces the ONTAP solution for VMware Site Recovery Manager (SRM), VMware's industry leading disaster recovery (DR) software, including the latest product information and best practices to streamline deployment, reduce risk, and simplify ongoing management.



This documentation replaces previously published technical report *TR-4900: VMware Site Recovery Manager with ONTAP*

Best practices supplement other documents such as guides and compatibility tools. They are developed based on lab testing and extensive field experience by NetApp engineers and customers. In some cases, recommended best practices might not be the right fit for your environment; however, they are generally the simplest solutions that meet the needs of the most customers.

This document is focused on capabilities in recent releases of ONTAP 9 when used in conjunction with ONTAP tools for VMware vSphere 9.12 (which includes the NetApp Storage Replication Adapter [SRA] and VASA Provider [VP]), as well as VMware Site Recovery Manager 8.7.

Why use ONTAP with SRM?

NetApp data management platforms powered by ONTAP software are some of the most widely adopted storage solutions for SRM. The reasons are plentiful: A secure, high performance, unified protocol (NAS and SAN together) data management platform that provides industry defining storage efficiency, multitenancy, quality of service controls, data protection with space-efficient snapshots and replication with SnapMirror. All leveraging native hybrid multi-cloud integration for the protection of VMware workloads and a plethora of automation and orchestration tools at your fingertips.

When you use SnapMirror for array-based replication, you take advantage of one of ONTAP's most proven and mature technologies. SnapMirror gives you the advantage of secure and highly efficient data transfers, copying only changed file system blocks, not entire VMs or datastores. Even those blocks take advantage of space savings, such as deduplication, compression, and compaction. Modern ONTAP systems now use version-independent SnapMirror, allowing you flexibility in selecting your source and destination clusters. SnapMirror has truly become one of the most powerful tools available for disaster recovery.

Whether you are using traditional NFS, iSCSI, or Fibre Channel- attached datastores (now with support for vVols datastores), SRM provides a robust first party offering that leverages the best of ONTAP capabilities for disaster recovery or datacenter migration planning and orchestration.

How SRM leverages ONTAP 9

SRM leverages the advanced data management technologies of ONTAP systems by integrating with ONTAP tools for VMware vSphere, a virtual appliance that includes three primary components:

- The vCenter plug-in, formerly known as Virtual Storage Console (VSC), simplifies storage management and efficiency features, enhances availability, and reduces storage costs and operational overhead, whether you are using SAN or NAS. It uses best practices for provisioning datastores and optimizes ESXi host settings for NFS and block storage environments. For all these benefits, NetApp recommends this plug-in when using vSphere with systems running ONTAP software.
- The VASA Provider for ONTAP supports the VMware vStorage APIs for Storage Awareness (VASA) framework. VASA Provider connects vCenter Server with ONTAP to aid in provisioning and monitoring VM storage. It enables VMware Virtual Volumes (vVols) support and the management of storage capability profiles (including vVols replication capabilities) and individual VM vVols performance. It also provides alarms for monitoring capacity and compliance with the profiles. When used in conjunction with SRM, the VASA Provider for ONTAP enables support for vVols-based virtual machines without requiring the installation of an SRA adapter on the SRM server.
- The SRA is used together with SRM to manage the replication of VM data between production and disaster recovery sites for traditional VMFS and NFS datastores and also for the nondisruptive testing of DR replicas. It helps automate the tasks of discovery, recovery, and reprotection. It includes both an SRA server appliance and SRA adapters for the Windows SRM server and the SRM appliance.

After you have installed and configured the SRA adapters on the SRM server for protecting non-vVols datastores and/or enabled vVols replication in the VASA Provider settings, you can begin the task of configuring your vSphere environment for disaster recovery.

The SRA and VASA Provider deliver a command-and-control interface for the SRM server to manage the ONTAP FlexVols that contain your VMware Virtual Machines (VMs), as well as the SnapMirror replication protecting them.

Starting with SRM 8.3, a new SRM vVols Provider control path was introduced into the SRM server, allowing it to communicate with the vCenter server and, through it, to the VASA Provider without needing an SRA. This enabled the SRM server to leverage much deeper control over the ONTAP cluster than was possible before, because VASA provides a complete API for closely coupled integration.

SRM can test your DR plan nondisruptively using NetApp's proprietary FlexClone technology to make nearly instantaneous clones of your protected datastores at your DR site. SRM creates a sandbox to safely test so that your organization, and your customers, are protected in the event of a true disaster, giving you confidence in your organization's ability to execute a failover during a disaster.

In the event of a true disaster or even a planned migration, SRM allows you to send any last-minute changes to the dataset via a final SnapMirror update (if you choose to do so). It then breaks the mirror and mounts the datastore to your DR hosts. At that point, your VMs can be automatically powered up in any order according to your pre-planned strategy.

SRM with ONTAP and other use cases: hybrid cloud and migration

Integrating your SRM deployment with ONTAP advanced data management capabilities allows for vastly improved scale and performance when compared with local storage options. But more than that, it brings the flexibility of the hybrid cloud. The hybrid cloud enables you to save money by tiering unused data blocks from your high-performance array to your preferred hyperscaler using FabricPool, which could be an on-premises S3 store such as NetApp StorageGRID. You can also use SnapMirror for edge-based systems with software-defined ONTAP Select or cloud-based DR using Cloud Volumes ONTAP (CVO) or [NetApp Private Storage in Equinix](#) for Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) to create a fully

integrated storage, networking, and compute- services stack in the cloud.

You could then perform test failover inside a cloud service provider's datacenter with near-zero storage footprint thanks to FlexClone. Protecting your organization can now cost less than ever before.

SRM can also be used to execute planned migrations by leveraging SnapMirror to efficiently transfer your VMs from one datacenter to another or even within the same datacenter, whether your own, or via any number of NetApp partner service providers.

Deployment best practices

The following sections outline the deployment best practices with ONTAP and VMware SRM.

SVM layout and segmentation for SMT

With ONTAP, the concept of the storage virtual machine (SVM) provides strict segmentation in secure multitenant environments. SVM users on one SVM cannot access or manage resources from another. In this way, you can leverage ONTAP technology by creating separate SVMs for different business units who manage their own SRM workflows on the same cluster for greater overall storage efficiency.

Consider managing ONTAP using SVM-scoped accounts and SVM management LIFs to not only improve security controls, but also improve performance. Performance is inherently greater when using SVM-scoped connections because the SRA is not required to process all the resources in an entire cluster, including physical resources. Instead, it only needs to understand the logical assets that are abstracted to the particular SVM.

When using NAS protocols only (no SAN access), you can even leverage the new NAS optimized mode by setting the following parameter (note that the name is such because SRA and VASA use the same backend services in the appliance):

1. Log into the control panel at `https://<IP address>:9083` and click Web based CLI interface.
2. Run the command `vp updateconfig -key=enable.qtree.discovery -value=true`.
3. Run the command `vp updateconfig -key=enable.optimised.sra -value=true`.
4. Run the command `vp reloadconfig`.

Deploy ONTAP tools and considerations for vVols

If you intend to use SRM with vVols, you must manage the storage using cluster- scoped credentials and a cluster management LIF. This is because the VASA Provider must understand the underlying physical architecture to satisfy the policy requires for VM storage policies. For example, if you have a policy that requires all- flash storage, the VASA Provider must be able to see which systems are all flash.

Another deployment best practice is to never store your ONTAP tools appliance on a vVols datastore that it is managing. This could lead to a situation whereby you cannot power on the VASA Provider because you cannot create the swap vVol for the appliance because the appliance is offline.

Best practices for managing ONTAP 9 systems

As previously mentioned, you can manage ONTAP clusters using either cluster or SVM scoped credentials and management LIFs. For optimum performance, you may want to consider using SVM- scoped credentials whenever you aren't using vVols. However, in doing so, you should be aware of some requirements, and that

you do lose some functionality.

- The default vsadmin SVM account does not have the required access level to perform ONTAP tools tasks. Therefore, you need to create a new SVM account.
- If you are using ONTAP 9.8 or later, NetApp recommends creating an RBAC least privileged user account using ONTAP System Manager's users menu together with the JSON file available on your ONTAP tools appliance at <https://<IP address>:9083/vsc/config/>. Use your administrator password to download the JSON file. This can be used for SVM or cluster scoped accounts.

If you are using ONTAP 9.6 or earlier, you should use the RBAC User Creator (RUC) tool available in the [NetApp Support Site Toolchest](#).

- Because the vCenter UI plugin, VASA Provider, and SRA server are all fully integrated services, you must add storage to the SRA adapter in SRM the same way you add storage in the vCenter UI for ONTAP tools. Otherwise, the SRA server might not recognize the requests being sent from SRM via the SRA adapter.
- NFS path checking is not performed when using SVM-scoped credentials. This is because the physical location is logically abstracted from the SVM. This is not a cause for concern though, as modern ONTAP systems no longer suffer any noticeable performance decline when using indirect paths.
- Aggregate space savings due to storage efficiency might not be reported.
- Where supported, load-sharing mirrors cannot be updated.
- EMS logging might not be performed on ONTAP systems managed with SVM scoped credentials.

Operational best practices

The following sections outline the operational best practices for VMware SRM and ONTAP storage.

Datstores and protocols

- If possible, always use ONTAP tools to provision datstores and volumes. This makes sure that volumes, junction paths, LUNs, igroups, export policies, and other settings are configured in a compatible manner.
- SRM supports iSCSI, Fibre Channel, and NFS version 3 with ONTAP 9 when using array-based replication through SRA. SRM does not support array-based replication for NFS version 4.1 with either traditional or vVols datstores.
- To confirm connectivity, always verify that you can mount and unmount a new test datstore at the DR site from the destination ONTAP cluster. Test each protocol you intend to use for datstore connectivity. A best practice is to use ONTAP tools to create your test datstore, since it is doing all the datstore automation as directed by SRM.
- SAN protocols should be homogeneous for each site. You can mix NFS and SAN, but the SAN protocols should not be mixed within a site. For example, you can use FCP in site A, and iSCSI in site B. You should not use both FCP and iSCSI at site A. The reason for this is that the SRA does not create mixed igroups at the recovery site and SRM does not filter the initiator list given to the SRA.
- Previous guides advised creating LIF to data locality. That is to say, always mount a datstore using a LIF located on the node that physically owns the volume. That is no longer a requirement in modern versions of ONTAP 9. Whenever possible, and if given cluster-scoped credentials, ONTAP tools will still choose to load balance across LIFs local to the data, but it is not a requirement for high availability or performance.
- ONTAP 9 can be configured to automatically remove snapshots to preserve uptime in the event of an out-of-space condition when autosize is not able to supply sufficient emergency capacity. The default setting for this capability does not automatically delete the snapshots that are created by SnapMirror. If SnapMirror

snapshots are deleted, then the NetApp SRA cannot reverse and resynchronize replication for the affected volume. To prevent ONTAP from deleting SnapMirror snapshots, configure the Snapshot autodelete capability to try.

```
snap autodelete modify -volume -commitment try
```

- Volume autosize should be set to `grow` for volumes containing SAN datastores and `grow_shrink` for NFS datastores. Learn more about [configuring volumes to automatically grow or shrink](#).
- SRM performs best when the number of datastores and thus protection groups is minimized in your recovery plans. Therefore you should consider optimizing for VM density in SRM-protected environments where RTO is of key importance.
- Use Distributed Resource Scheduler (DRS) to help balance the load on your protected and recovery ESXi clusters. Remember that if you plan to failback, when you run a reprotect the previously protected clusters will become the new recovery clusters. DRS will help balance placement going in both directions.
- Where possible, avoid using IP customization with SRM as this can increase your RTO.

Storage Policy Based Management (SPBM) and vVols

Starting with SRM 8.3, the protection of VMs using vVols datastores is supported. SnapMirror schedules are exposed to VM storage policies by the VASA Provider when vVols replication is enabled in the ONTAP tools settings menu, as shown in the following screenshots.

The following example shows the enablement of vVols replication.


Manage Capabilities

- Enable VASA Provider**
vStorage APIs for Storage Awareness (VASA) is a set of application program interfaces (APIs) that enables vSphere vCenter to recognize the capabilities of storage arrays.
- Enable vVols replication**
Enables replication of vVols when used with VMware Site Recovery Manager 8.3 or later.
- Enable Storage Replication Adapter (SRA)**
Storage Replication Adapter (SRA) allows VMware Site Recovery Manager (SRM) to integrate with third party storage array technology.

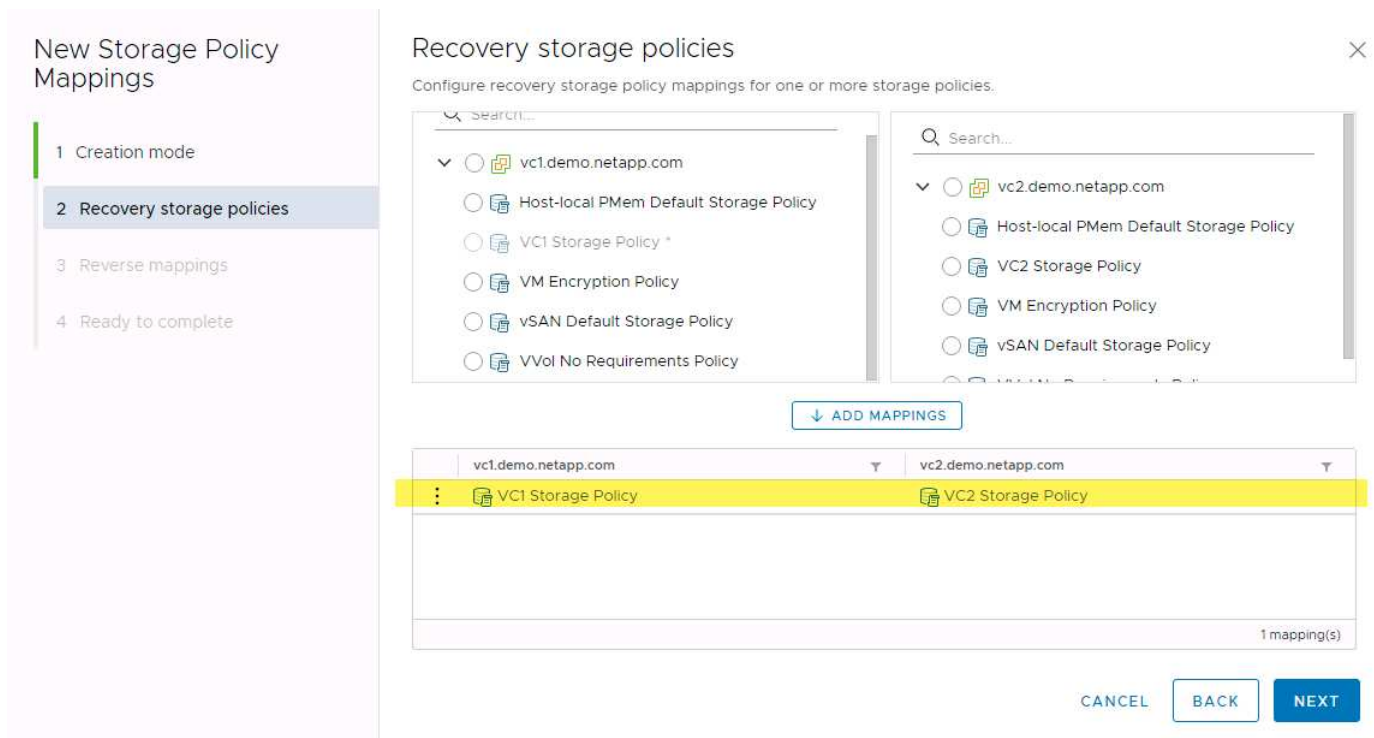
Enter authentication details for VASA Provider and SRA server:

IP address or hostname: 192.168.64.7
Username: Administrator
Password: _____

The following screenshot provides an example of SnapMirror schedules displayed in the Create VM Storage Policy wizard.



The ONTAP VASA Provider supports failover to dissimilar storage. For example, the system can fail over from ONTAP Select at an edge location to an AFF system in the core datacenter. Regardless of storage similarity, you must always configure storage policy mappings and reverse mappings for replication-enabled VM storage policies to make sure that services provided at the recovery site meet expectations and requirements. The following screenshot highlights a sample policy mapping.



Create replicated volumes for vVols datastores

Unlike previous vVols datastores, replicated vVols datastores must be created from the start with replication

enabled, and they must use volumes that were pre-created on the ONTAP systems with SnapMirror relationships. This requires pre-configuring things like cluster peering and SVM peering. These activities should be performed by your ONTAP administrator because this facilitates a strict separation of responsibilities between those who manage the ONTAP systems across multiple sites and those who are primarily responsible for vSphere operations.

This does come with a new requirement on behalf of the vSphere administrator. Because volumes are being created outside the scope of ONTAP tools, it is unaware of the changes your ONTAP administrator has made until the regularly scheduled rediscovery period. For that reason, it is a best practice to always run rediscovery whenever you create a volume or SnapMirror relationship to be used with vVols. Simply right-click on the host or cluster and select ONTAP tools > Update Host and Storage Data, as shown in the following screenshot.



One caution should be taken when it comes to vVols and SRM. Never mix protected and unprotected VMs in the same vVols datastore. The reason for this is that when you use SRM to failover to your DR site, only those VMs that are part of the protection group are brought online in DR. Therefore, when you reprotect (reverse the SnapMirror from DR back to production again), you may overwrite the VMs that were not failed over and could contain valuable data.

About array pairs

An array manager is created for each array pair. With SRM and ONTAP tools, each array pairing is done with the scope of an SVM, even if you are using cluster credentials. This allows you to segment DR workflows between tenants based on which SVMs they have been assigned to manage. You can create multiple array managers for a given cluster, and they can be asymmetric. You can fan out or fan in between different ONTAP 9 clusters. For example, you can have SVM-A and SVM-B on Cluster-1 replicating to SVM-C on Cluster-2, SVM-D on Cluster-3, or vice-versa.

When configuring array pairs in SRM, you should always add them in SRM the same way as you added them to ONTAP Tools, meaning, they must use the same username, password, and management LIF. This requirement ensures that SRA communicates properly with the array. The following screenshot illustrates how a cluster might appear in ONTAP Tools and how it might be added to an array manager.

vm vSphere Client Menu Search in all environments

ONTAP tools

- Overview
- Storage Systems**
- Storage Capability Profiles
- Storage Mapping
- Settings
- Reports

Storage Systems

ADD REDISCOVER ALL

Name	Type	IP Address
cluster2	Cluster	cluster2.demo.netapp.com

Edit Local Array Manager

Enter a name for the array manager on "vc2.demo.netapp.com":

Storage Array Parameters

Storage Management IP Address or Hostname

Enter the cluster management IP address/hostname. To connect directly to a Storage Virtual Machine(SVM), enter the SVM management IP address/hostname.

About replication groups

Replication groups contain logical collections of virtual machines that are recovered together. The ONTAP tools VASA Provider automatically creates replication groups for you. Because ONTAP SnapMirror replication occurs at the volume level, all VMs in a volume are in the same replication group.

There are several factors to consider with replication groups and how you distribute VMs across FlexVol volumes. Grouping similar VMs in the same volume can increase storage efficiency with older ONTAP systems that lack aggregate-level deduplication, but grouping increases the size of the volume and reduces volume I/O concurrency. The best balance of performance and storage efficiency can be achieved in modern ONTAP systems by distributing VMs across FlexVol volumes in the same aggregate, thereby leveraging aggregate-level deduplication and gaining greater I/O parallelization across multiple volumes. You can recover VMs in the volumes together because a protection group (discussed below) can contain multiple replication groups. The downside to this layout is that blocks might be transmitted over the wire multiple times because volume SnapMirror doesn't take aggregate deduplication into account.

One final consideration for replication groups is that each one is by its nature a logical consistency group (not to be confused with SRM consistency groups). This is because all VMs in the volume are transferred together using the same snapshot. So if you have VMs that must be consistent with each other, consider storing them in the same FlexVol.

About protection groups

Protection groups define VMs and datastores in groups that are recovered together from the protected site. The protected site is where the VMs that are configured in a protection group exist during normal steady-state operations. It is important to note that even though SRM might display multiple array managers for a protection group, a protection group cannot span multiple array managers. For this reason, you should not span VM files across datastores on different SVMs.

About recovery plans

Recovery plans define which protection groups are recovered in the same process. Multiple protection groups

can be configured in the same recovery plan. Also, to enable more options for the execution of recovery plans, a single protection group can be included in multiple recovery plans.

Recovery plans allow SRM administrators to define recovery workflows by assigning VMs to a priority group from 1 (highest) to 5 (lowest), with 3 (medium) being the default. Within a priority group, VMs can be configured for dependencies.

For example, your company could have a tier-1 business-critical application that relies on a Microsoft SQL server for its database. So, you decide to place your VMs in priority group 1. Within priority group 1, you begin planning the order to bring up services. You probably want your Microsoft Windows domain controller to boot up before your Microsoft SQL server, which would need to be online before your application server, and so on. You would add all these VMs to the priority group and then set the dependencies because dependencies only apply within a given priority group.

NetApp strongly recommends working with your application teams to understand the order of operations required in a failover scenario and to construct your recovery plans accordingly.

Test failover

As a best practice, always perform a test failover whenever a change is made to the configuration of a protected VM storage. This ensures that, in the event of a disaster, you can trust that Site Recovery Manager can restore services within the expected RTO target.

NetApp also recommends confirming in-guest application functionality occasionally, especially after reconfiguring VM storage.

When a test recovery operation is performed, a private test bubble network is created on the ESXi host for the VMs. However, this network is not automatically connected to any physical network adapters and therefore does not provide connectivity between the ESXi hosts. To allow communication among VMs that are running on different ESXi hosts during DR testing, a physical private network is created between the ESXi hosts at the DR site. To verify that the test network is private, the test bubble network can be separated physically or by using VLANs or VLAN tagging. This network must be segregated from the production network because as the VMs are recovered, they cannot be placed on the production network with IP addresses that could conflict with actual production systems. When a recovery plan is created in SRM, the test network that was created can be selected as the private network to connect the VMs to during the test.

After the test has been validated and is no longer required, perform a cleanup operation. Running cleanup returns the protected VMs to their initial state and resets the recovery plan to the Ready state.

Failover considerations

There are several other considerations when it comes to failing over a site in addition to the order of operations mentioned in this guide.

One issue you might have to contend with is networking differences between sites. Some environments might be able to use the same network IP addresses at both the primary site and the DR site. This ability is referred to as a stretched virtual LAN (VLAN) or stretched network setup. Other environments might have a requirement to use different network IP addresses (for example, in different VLANs) at the primary site relative to the DR site.

VMware offers several ways to solve this problem. For one, network virtualization technologies like VMware NSX-T Data Center abstract the entire networking stack from layers 2 through 7 from the operating environment, allowing for more portable solutions. Learn more about [NSX-T options with SRM](#).

SRM also gives you the ability to change the network configuration of a VM as it is recovered. This

reconfiguration includes settings such as IP addresses, gateway addresses, and DNS server settings. Different network settings, which are applied to individual VMs as they are recovered, can be specified in the property's settings of a VM in the recovery plan.

To configure SRM to apply different network settings to multiple VMs without having to edit the properties of each one in the recovery plan, VMware provides a tool called the `dr-ip-customizer`. Learn how to use this utility, refer to [VMware's documentation](#).

Reprotect

After a recovery, the recovery site becomes the new production site. Because the recovery operation broke the SnapMirror replication, the new production site is not protected from any future disaster. A best practice is to protect the new production site to another site immediately after a recovery. If the original production site is operational, the VMware administrator can use the original production site as a new recovery site to protect the new production site, effectively reversing the direction of protection. Reprotection is available only in non-catastrophic failures. Therefore, the original vCenter Servers, ESXi servers, SRM servers, and corresponding databases must be eventually recoverable. If they are not available, a new protection group and a new recovery plan must be created.

Failback

A failback operation is fundamentally a failover in a different direction than before. As a best practice, you verify that the original site is back to acceptable levels of functionality before attempting to failback, or, in other words, failover to the original site. If the original site is still compromised, you should delay failback until the failure is sufficiently remediated.

Another failback best practice is to always perform a test failover after completing reprotect and before doing your final failback. This verifies that the systems in place at the original site can complete the operation.

Reprotecting the original site

After failback, you should confirm with all stakeholders that their services have been returned to normal before running reprotect again,

Running reprotect after failback essentially puts the environment back in the state it was in at the beginning, with SnapMirror replication again running from the production site to the recovery site.

Replication topologies

In ONTAP 9, the physical components of a cluster are visible to cluster administrators, but they are not directly visible to the applications and hosts that use the cluster. The physical components provide a pool of shared resources from which the logical cluster resources are constructed. Applications and hosts access data only through SVMs that contain volumes and LIFs.

Each NetApp SVM is treated as an array in VMware vCenter Site Recovery Manager. SRM supports certain array-to-array (or SVM-to-SVM) replication layouts.

A single VM cannot own data—Virtual Machine Disk (VMDK) or RDM—on more than one SRM array for the following reasons:


- SRM sees only the SVM, not an individual physical controller.
- An SVM can control LUNs and volumes that span multiple nodes in a cluster.

Best Practice

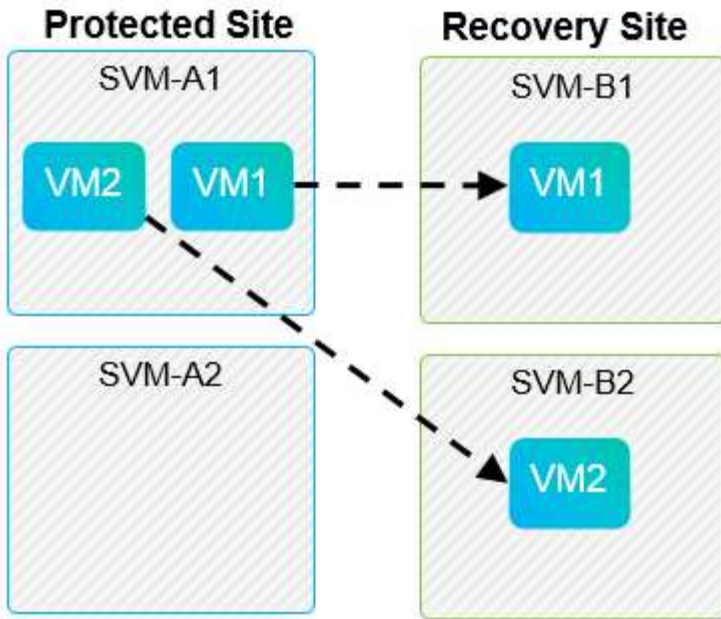
To determine supportability, keep this rule in mind: to protect a VM by using SRM and the NetApp SRA, all parts of the VM must exist on only one SVM. This rule applies at both the protected site and the recovery site.

Supported SnapMirror layouts


The following figures show the SnapMirror relationship layout scenarios that SRM and SRA support. Each VM in the replicated volumes owns data on only one SRM array (SVM) at each site.




SnapMirror Replication
----->



SnapMirror Replication
----->





Supported Array Manager layouts

When you use array-based replication (ABR) in SRM, protection groups are isolated to a single array pair, as shown in the following screenshot. In this scenario, SVM1 and SVM2 are peered with SVM3 and SVM4 at the recovery site. However, you can select only one of the two array pairs when you create a protection group.

New Protection Group

- 1 Name and direction
- 2 Type
- 3 Datastore groups
- 4 Recovery plan
- 5 Ready to complete

Type ✕

Select the type of protection group you want to create:

- Datastore groups (array-based replication)**
Protect all virtual machines which are on specific datastores.
- Individual VMs (vSphere Replication)**
Protect specific virtual machines, regardless of the datastores.
- Virtual Volumes (vVol replication)**
Protect virtual machines which are on replicated vVol storage.
- Storage policies (array-based replication)**
Protect virtual machines with specific storage policies.



Select array pair

Array Pair	Array Manager Pair
<input type="radio"/> ✓ cluster1:svm1 ↔ cluster2:svm2	vc1 array manager ↔ vc2 array manager
<input type="radio"/> ✓ cluster1:svm3 ↔ cluster2:svm4	vc1 trad datastores ↔ vc2 trad datastores


Unsupported layouts

Unsupported configurations have data (VMDK or RDM) on multiple SVMs that is owned by an individual VM. In

the examples shown in the following figures, VM1 cannot be configured for protection with SRM because VM1 has data on two SVMs.




Any replication relationship in which an individual NetApp volume is replicated from one source SVM to multiple destinations in the same SVM or in different SVMs is referred to as SnapMirror fan-out. Fan-out is not supported with SRM. In the example shown in the following figure, VM1 cannot be configured for protection in SRM because it is replicated with SnapMirror to two different locations.



SnapMirror cascade

SRM does not support cascading of SnapMirror relationships, in which a source volume is replicated to a destination volume and that destination volume is also replicated with SnapMirror to another destination volume. In the scenario shown in the following figure, SRM cannot be used for failover between any sites.



SnapMirror and SnapVault

NetApp SnapVault software enables disk-based backup of enterprise data between NetApp storage systems. SnapVault and SnapMirror can coexist in the same environment; however, SRM supports the failover of only

the SnapMirror relationships.



The NetApp SRA supports the `mirror-vault` policy type.

SnapVault was rebuilt from the ground up for ONTAP 8.2. Although former Data ONTAP 7-Mode users should find similarities, major enhancements have been made in this version of SnapVault. One major advance is the ability to preserve storage efficiencies on primary data during SnapVault transfers.

An important architectural change is that SnapVault in ONTAP 9 replicates at the volume level as opposed to at the qtrees level, as is the case in 7-Mode SnapVault. This setup means that the source of a SnapVault relationship must be a volume, and that volume must replicate to its own volume on the SnapVault secondary system.

In an environment in which SnapVault is used, specifically named snapshots are created on the primary storage system. Depending on the configuration implemented, the named snapshots can be created on the primary system by a SnapVault schedule or by an application such as NetApp Active IQ Unified Manager. The named snapshots that are created on the primary system are then replicated to the SnapMirror destination, and from there they are vaulted to the SnapVault destination.

A source volume can be created in a cascade configuration in which a volume is replicated to a SnapMirror destination in the DR site, and from there it is vaulted to a SnapVault destination. A source volume can also be created in a fan-out relationship in which one destination is a SnapMirror destination and the other destination is a SnapVault destination. However, SRA does not automatically reconfigure the SnapVault relationship to use the SnapMirror destination volume as the source for the vault when SRM failover or replication reversal occurs.

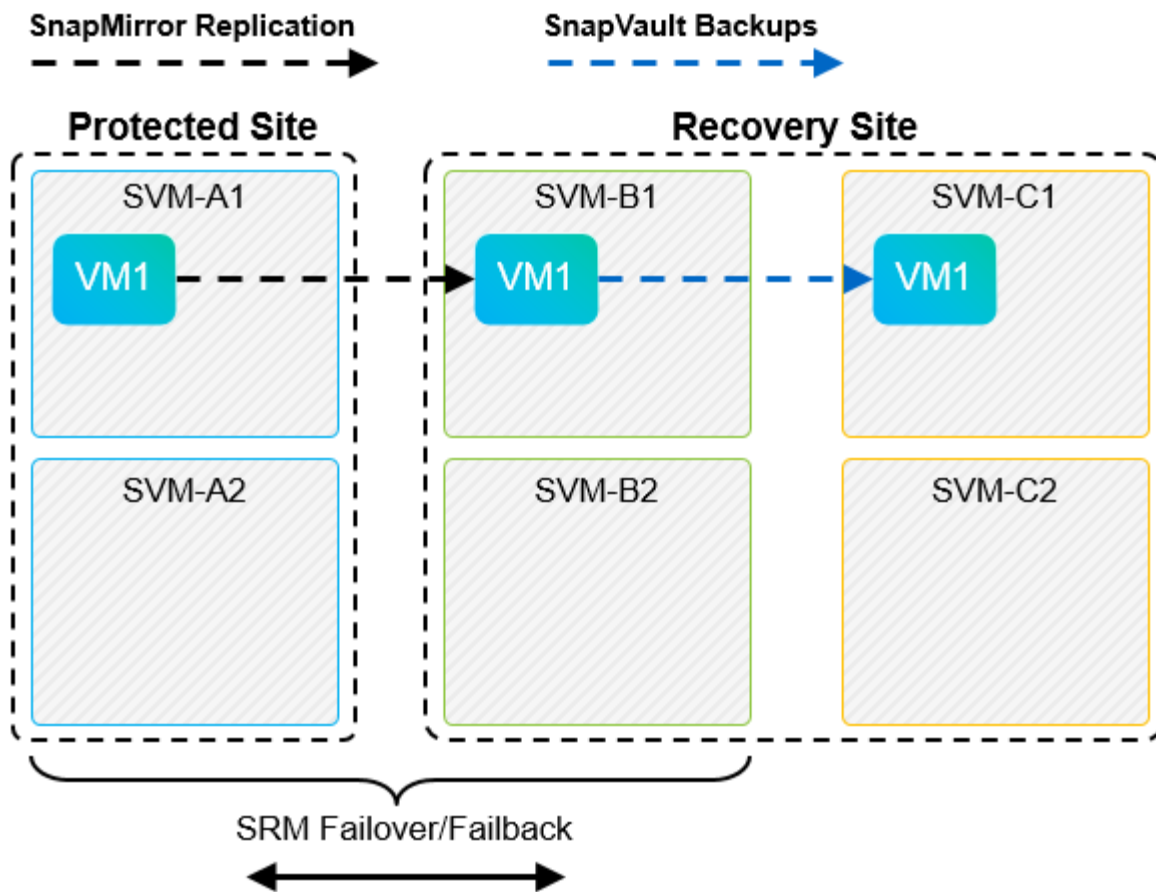
For the latest information about SnapMirror and SnapVault for ONTAP 9, see [TR-4015 SnapMirror Configuration Best Practice Guide for ONTAP 9](#).

Best Practice


If SnapVault and SRM are used in the same environment, NetApp recommends using a SnapMirror to SnapVault cascade configuration in which SnapVault backups are normally performed from the SnapMirror destination at the DR site. In the event of a disaster, this configuration makes the primary site inaccessible. Keeping the SnapVault destination at the recovery site allows SnapVault backups to be reconfigured after failover so that SnapVault backups can continue while operating at the recovery site.

In a VMware environment, each datastore has a universal unique identifier (UUID), and each VM has a unique managed object ID (MOID). These IDs are not maintained by SRM during failover or failback. Because datastore UUIDs and VM MOIDs are not maintained during failover by SRM, any applications that depend on these IDs must be reconfigured after SRM failover. An example application is NetApp Active IQ Unified Manager, which coordinates SnapVault replication with the vSphere environment.

The following figure depicts a SnapMirror to SnapVault cascade configuration. If the SnapVault destination is at the DR site or at a tertiary site that is not affected by an outage at the primary site, the environment can be reconfigured to allow backups to continue after failover.



The following figure depicts the configuration after SRM has been used to reverse SnapMirror replication back to the primary site. The environment has also been reconfigured such that SnapVault backups are occurring from what is now the SnapMirror source. This setup is a SnapMirror SnapVault fan-out configuration.



After SRM performs failback and a second reversal of the SnapMirror relationships, the production data is back at the primary site. This data is now protected in the same way that it was before the failover to the DR site—through SnapMirror and SnapVault backups.

Use of Qtrees in Site Recovery Manager environments

Qtrees are special directories that allow the application of file system quotas for NAS. ONTAP 9 allows the creation of qtrees, and qtrees can exist in volumes that are replicated with SnapMirror. However, SnapMirror does not allow replication of individual qtrees or qtree-level replication. All SnapMirror replication is at the volume level only. For this reason, NetApp does not recommend the use of qtrees with SRM.

Mixed FC and iSCSI environments

With the supported SAN protocols (FC, FCoE, and iSCSI), ONTAP 9 provides LUN services—that is, the ability to create and map LUNs to attached hosts. Because the cluster consists of multiple controllers, there are multiple logical paths that are managed by multipath I/O to any individual LUN. Asymmetric logical unit access (ALUA) is used on the hosts so that the optimized path to a LUN is selected and is made active for data transfer. If the optimized path to any LUN changes (for example, because the containing volume is moved), ONTAP 9 automatically recognizes and nondisruptively adjusts for this change. If the optimized path becomes unavailable, ONTAP can nondisruptively switch to any other available path.

VMware SRM and NetApp SRA support the use of the FC protocol at one site and the iSCSI protocol at the other site. It does not support having a mix of FC-attached datastores and iSCSI-attached datastores in the same ESXi host or in different hosts in the same cluster, however. This configuration is not supported with SRM because, during the SRM failover or test failover, SRM includes all FC and iSCSI initiators in the ESXi hosts in the request.

Best Practice

SRM and SRA support mixed FC and iSCSI protocols between the protected and recovery sites. However, each site should be configured with only one protocol, either FC or iSCSI, not both protocols at the same site. If a requirement exists to have both FC and iSCSI protocols configured at the same site, NetApp recommends that some hosts use iSCSI and other hosts use FC. NetApp also recommends in this case that SRM resource mappings be set up so that the VMs are configured to fail over into one group of hosts or the other.

Troubleshooting SRM when using vVols replication

The workflow within SRM is significantly different when using vVols replication from what is used with SRA and traditional datastores. For example, there is no array manager concept. As such, `discoverarrays` and `discoverdevices` commands are never seen.

When troubleshooting, it is beneficial to understand the new workflows, which are listed below:

1. `queryReplicationPeer`: Discovers the replication agreements between two fault domains.
2. `queryFaultDomain`: Discovers fault domain hierarchy.
3. `queryReplicationGroup`: Discovers the replication groups present in the source or target domains.
4. `syncReplicationGroup`: Synchronizes the data between source and target.
5. `queryPointInTimeReplica`: Discovers the point in time replicas on a target.
6. `testFailoverReplicationGroupStart`: Begins test failover.
7. `testFailoverReplicationGroupStop`: Ends test failover.
8. `promoteReplicationGroup`: Promotes a group currently in test to production.
9. `prepareFailoverReplicationGroup`: Prepares for a disaster recovery.
10. `failoverReplicationGroup`: Executes disaster recovery.
11. `reverseReplicateGroup`: Initiates reverse replication.
12. `queryMatchingContainer`: Finds containers (along with Hosts or Replication Groups) that might satisfy a provisioning request with a given policy.
13. `queryResourceMetadata`: Discovers the metadata of all resources from the VASA provider, the resource utilization can be returned as an answer to the `queryMatchingContainer` function.

The most common error seen when configuring vVols replication is a failure to discover the SnapMirror relationships. This occurs because the volumes and SnapMirror relationships are created outside of the purview of ONTAP Tools. Therefore, it is a best practice to always make sure your SnapMirror relationship is fully initialized and that you have run a rediscovery in ONTAP Tools at both sites before attempting to create a replicated vVols datastore.

Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- TR-4597: VMware vSphere for ONTAP
<https://docs.netapp.com/us-en/ontap-apps-dbs/vmware/vmware-vsphere-overview.html>

- TR-4400: VMware vSphere Virtual Volumes with ONTAP
<https://docs.netapp.com/us-en/ontap-apps-dbs/vmware/vmware-vvols-overview.html>
- TR-4015 SnapMirror Configuration Best Practice Guide for ONTAP 9
<https://www.netapp.com/media/17229-tr4015.pdf?v=127202175503P>
- RBAC User Creator for ONTAP
<https://mysupport.netapp.com/site/tools/tool-eula/rbac>
- ONTAP tools for VMware vSphere Resources
<https://mysupport.netapp.com/site/products/all/details/otv/docsandkb-tab>
- VMware Site Recovery Manager Documentation
<https://docs.vmware.com/en/Site-Recovery-Manager/index.html>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

vSphere Metro Storage Cluster with ONTAP

vSphere Metro Storage Cluster with ONTAP

VMware's industry-leading vSphere hypervisor can be deployed as a stretched cluster referred to as a vSphere Metro Storage Cluster (vMSC).

vMSC solutions are supported with both NetApp® MetroCluster™ and SnapMirror active sync (formerly known as SnapMirror Business Continuity, or SMBC) and provide advanced business continuity if one or more failure domains suffer a total outage. The resilience to different modes of failure depends on which configuration options you choose.

Continuous Availability Solutions for vSphere Environments

ONTAP architecture is a flexible and scalable storage platform that provides SAN (FCP, iSCSI, and NVMe-oF) and NAS (NFS v3 and v4.1) services for datastores. The NetApp AFF, ASA, and FAS storage systems use the ONTAP operating system to offer additional protocols for guest storage access like S3 and SMB/CIFS.

NetApp MetroCluster uses NetApp's HA (controller failover or CFO) function to protect against controller failures. It also includes local SyncMirror technology, cluster failover on disaster (controller failover on demand or CFOD), hardware redundancy, and geographical separation to achieve high levels of availability. SyncMirror synchronously mirrors data across the two halves of the MetroCluster configuration by writing data to two plexes: the local plex (on the local shelf) actively serving data and the remote plex (on the remote shelf) normally not serving data. Hardware redundancy is put in place for all MetroCluster components such as controllers, storage, cables, switches (used with fabric MetroCluster), and adapters.

NetApp SnapMirror active sync provides datastore-granular protection with FCP and iSCSI SAN protocols, allowing you to selectively protect only high-priority workloads. It offers active-active access to both local and remote sites, unlike NetApp MetroCluster which is an active-standby solution. At present, active sync is an asymmetric solution where one side is preferred over the other, providing better performance. This is achieved using ALUA (Asymmetric Logical Unit Access) functionality which automatically informs the ESXi host which controllers to prefer. However, NetApp has announced that active sync will soon enable fully symmetric access.

To create a VMware HA/DRS cluster across two sites, ESXi hosts are used and managed by a vCenter Server Appliance (VCSA). The vSphere management, vMotion®, and virtual machine networks are connected through a redundant network between the two sites. The vCenter Server managing the HA/DRS cluster can connect to the ESXi hosts at both sites and should be configured using vCenter HA.

Refer to [How Do You Create and Configure Clusters in the vSphere Client](#) to configure vCenter HA.

You should also refer to [VMware vSphere Metro Storage Cluster Recommended Practices](#).

What is vSphere Metro Storage Cluster?

vSphere Metro Storage Cluster (vMSC) is a certified configuration that protects virtual machines (VMs) and containers against failures. This is achieved by using stretched storage concepts along with clusters of ESXi hosts, which are distributed across different failure domains such as racks, buildings, campuses, or even cities. The NetApp MetroCluster and SnapMirror active sync storage technologies are used to provide RPO=0 or near RPO=0 protection respectively to the host clusters. The vMSC configuration is designed to ensure that data is always available even if a complete physical or logical “site” fails. A storage device that is part of the vMSC configuration must be certified after undergoing a successful vMSC certification process. All the supported storage devices can be found in the [VMware Storage Compatibility Guide](#).

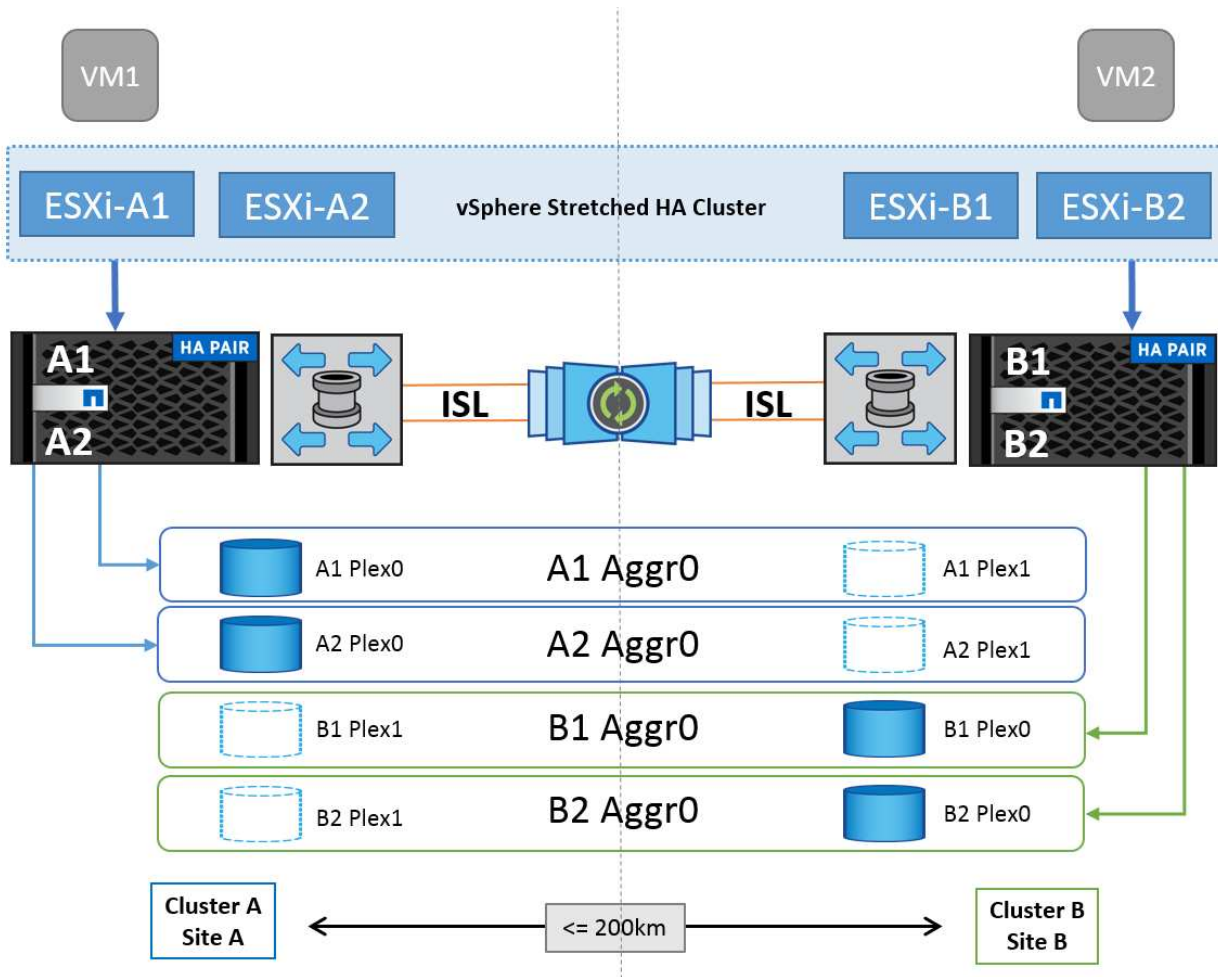
If you want more information about the design guidelines for vSphere Metro Storage Cluster, you can refer to the following documentation:

- [VMware vSphere support with NetApp MetroCluster](#)
- [VMware vSphere support with NetApp SnapMirror Business Continuity](#) (now known as SnapMirror active sync)

Depending on the latency considerations, NetApp MetroCluster can be deployed in two different configurations for use with vSphere:

- Stretch MetroCluster
- Fabric MetroCluster


The following illustrates a high-level topology diagram of stretch MetroCluster.



Refer to [MetroCluster documentation](#) for specific design and deployment information for MetroCluster.

SnapMirror active sync can also be deployed in two different ways.

- Asymmetric
- Symmetric (private preview in ONTAP 9.14.1)



Refer to [NetApp Docs](#) for specific design and deployment information for SnapMirror active sync.


VMware vSphere Solution Overview

The vCenter Server Appliance (VCSA) is the powerful centralized management system and single pane of glass for vSphere that enables administrators to effectively operate ESXi clusters. It facilitates key functions such as VM provisioning, vMotion operation, High Availability (HA), Distributed Resource Scheduler (DRS), Tanzu Kubernetes Grid, and more. It is an essential component in VMware cloud environments and should be designed with service availability in mind.

vSphere High Availability

VMware's cluster technology groups ESXi servers into pools of shared resources for virtual machines and provides vSphere High Availability (HA). vSphere HA provides easy-to-use, high availability for applications running in virtual machines. When the HA feature is enabled on the cluster, each ESXi server maintains communication with other hosts so that if any ESXi host becomes unresponsive or isolated, the HA cluster can negotiate the recovery of the virtual machines that were running on that ESXi host among surviving hosts in the cluster. In the event of a guest operating system failure, vSphere HA restarts the affected virtual machine on the same physical server. vSphere HA makes it possible to reduce planned downtime, prevent unplanned downtime, and rapidly recover from outages.

vSphere HA cluster recovering VMs from failed server.



It's important to understand that VMware vSphere has no knowledge of NetApp MetroCluster or SnapMirror active sync and sees all ESXi hosts in the vSphere cluster as eligible hosts for HA cluster operations depending on host and VM group affinity configurations.

Host Failure Detection

As soon as the HA cluster is created, all hosts in the cluster participate in election, and one of the hosts becomes a master. Each slave performs network heartbeat to the master, and the master in turn performs network heartbeat on all slave hosts. The master host of a vSphere HA cluster is responsible for detecting the failure of slave hosts.

Depending on the type of failure detected, the virtual machines running on the hosts might need to be failed over.

In a vSphere HA cluster, three types of host failure are detected:

- Failure - A host stops functioning.
- Isolation - A host becomes network isolated.
- Partition - A host loses network connectivity with the master host.

The master host monitors the slave hosts in the cluster. This communication is done through the exchange of network heartbeats every second. When the master host stops receiving these heartbeats from a slave host, it checks for host liveness before declaring the host to have failed. The liveness check that the master host performs is to determine whether the slave host is exchanging heartbeats with one of the datastores. Also, the master host checks whether the host responds to ICMP pings sent to its management IP addresses to detect whether it is merely isolated from its master node or completely isolated from the network. It does this by pinging the default gateway. One or more isolation addresses can be specified manually to enhance the reliability of isolation validation.

Best Practice

NetApp recommends specifying a minimum of two additional isolation addresses, and that each of these addresses be site-local. This will enhance the reliability of isolation validation.

Host Isolation Response

Isolation Response is a setting in vSphere HA that determines the action triggered on Virtual Machines when a host in a vSphere HA cluster loses its management network connections but continues to run. There are three options for this setting, “Disabled”, “Shut Down and Restart VMs,” and “Power Off and Restart VMs.”

“Shut Down” is better than “Power Off”, which does not flush most recent changes to disk or commit transactions. If virtual machines have not shut down in 300 seconds they are powered off. To change the wait time, use the advanced option `das.isolationshutdowntimeout`.

Before HA initiates the isolation response, it first checks to see if the vSphere HA master agent owns the datastore that contains the VM config files. If not, then the host will not trigger the isolation response, because there is no master to restart the VMs. The host will periodically check the datastore state to determine if it is claimed by a vSphere HA agent that holds the master role.

Best Practice

NetApp recommends setting the “Host Isolation Response” to Disabled.

A split-brain condition can occur if a host becomes isolated or partitioned from the vSphere HA master host and the master is unable to communicate via heartbeat datastores or by ping. The master declares the isolated host dead and restarts the VMs on other hosts in the cluster. A split-brain condition now exists because there are two instances of the virtual machine running, only one of which can read or write the virtual disks. Split-brain conditions can now be avoided by configuring VM Component Protection (VMCP).

VM Component Protection (VMCP)

One of the feature enhancements in vSphere 6, relevant to HA, is VMCP. VMCP provides enhanced protection from All Paths Down (APD) and Permanent Device Loss (PDL) conditions for block (FC, iSCSI, FCoE) and file storage (NFS).

Permanent Device Loss (PDL)

PDL is a condition that occurs when a storage device permanently fails or is administratively removed and is not expected to return. The NetApp storage array issues a SCSI Sense code to ESXi declaring that the device is permanently lost. In the Failure Conditions and VM Response section of vSphere HA, you can configure what the response should be after a PDL condition is detected.

Best Practice

NetApp recommends setting the “Response for Datastore with PDL” to “**Power off and restart VMs**”. When this condition is detected a VM will be restarted instantly on a healthy host within the vSphere HA cluster.

All Paths Down (APD)

APD is a condition that occurs when a storage device becomes inaccessible to the host and no paths to the array are available. ESXi considers this a temporary problem with the device and is expecting it to become available again.

When an APD condition is detected, a timer is started. After 140 seconds, the APD condition is officially declared, and the device is marked as APD time out. When the 140 seconds have passed, HA will start counting the number of minutes specified in the Delay for VM Failover APD. When the specified time has passed, HA will restart the impacted virtual machines. You can configure VMCP to respond differently if desired (Disabled, Issue Events, or Power Off and Restart VMs).

Best Practice

NetApp recommends configuring the “Response for Datastore with APD” to “**Power off and restart VMs (conservative)**”.

Conservative refers to the likelihood of HA being able to restart VMs. When set to Conservative, HA will only restart the VM that is impacted by the APD if it knows another host can restart it. In the case of Aggressive, HA will try to restart the VM even if it doesn't know the state of the other hosts. This can result in VMs not being restarted if there is no host with access to the datastore it is located on.

If the APD status is resolved and access to the storage is restored before the time-out has passed, HA will not unnecessarily restart the virtual machine unless you explicitly configure it to do so. If a response is desired even when the environment has recovered from the APD condition, then Response for APD Recovery After APD Timeout should be configured to Reset VMs.

Best Practice

NetApp recommends configuring Response for APD Recovery After APD Timeout to Disabled.

VMware DRS Implementation for NetApp MetroCluster

VMware DRS is a feature that aggregates the host resources in a cluster and is primarily used to load balance within a cluster in a virtual infrastructure. VMware DRS primarily calculates the CPU and memory resources to perform load balancing in a cluster. Because vSphere is unaware of stretched clustering, it considers all hosts in both sites when load balancing. To avoid cross-site traffic, NetApp recommends configuring DRS affinity rules to manage a logical separation of VMs. This will ensure that unless there is a complete site failure, HA and DRS will only use local hosts.


If you create a DRS affinity rule for your cluster, you can specify how vSphere applies that rule during a virtual machine failover.

There are two types of rules you can specify vSphere HA failover behavior:

- VM anti-affinity rules force specified virtual machines to remain apart during failover actions.
- VM host affinity rules place specified virtual machines on a particular host or a member of a defined group of hosts during failover actions.

Using VM host affinity rules in VMware DRS, one can have a logical separation between site A and site B so that the VM runs on the host at the same site as the array that is configured as the primary read/write controller for a given datastore. Also, VM host affinity rules enable virtual machines to stay local to the storage, which in turn ascertains the virtual machine connection in case of network failures between the sites.

The following is an example of VM host groups and affinity rules.



Best Practice

NetApp recommends implementing “should” rules instead of “must” rules because they are violated by vSphere HA in the case of a failure. Using “must” rules could potentially lead to service outages.

Availability of services should always prevail over performance. In the scenario where a full data center fails, “must” rules must choose hosts from the VM host affinity group, and when the data center is unavailable, the virtual machines will not restart.

VMware Storage DRS Implementation with NetApp MetroCluster

The VMware Storage DRS feature enables the aggregation of datastores into a single unit and balances virtual machine disks when storage I/O control thresholds are exceeded.

Storage I/O control is enabled by default on Storage DRS-enabled DRS clusters. Storage I/O control allows an administrator to control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which enables more important virtual machines to have preference over less important virtual machines for I/O resource allocation.

Storage DRS uses Storage vMotion to migrate the virtual machines to different datastores within a datastore cluster. In a NetApp MetroCluster environment, a virtual machine migration needs to be controlled within the datastores of that site. For example, virtual machine A, running on a host at site A, should ideally migrate within the datastores of the SVM at site A. If it fails to do so, the virtual machine will continue to operate but with degraded performance, since the virtual disk read/write will be from site B through inter-site links.

Best Practice

NetApp recommends creating datastore clusters with respect to storage site affinity; that is, datastores with site affinity for site A should not be mixed with datastores with site affinity for site B.

Whenever a virtual machine is newly provisioned or migrated using Storage vMotion, NetApp recommends that all the VMware DRS rules specific to those virtual machines be manually updated, accordingly. This will

ascertain the virtual machine affinity at the site level for both host and datastore and thus reduce the network and storage overhead.

vMSC Design and Implementation Guidelines

This document outlines the design and implementation guidelines for vMSC with ONTAP storage systems.

NetApp Storage Configuration

Setup instructions for NetApp MetroCluster (referred to as an MCC configuration) are available at [MetroCluster Documentation](#). Instructions for SnapMirror active sync are also available at [SnapMirror Business Continuity overview](#).

Once you have configured MetroCluster, administering it is like managing a traditional ONTAP environment. You can set up Storage Virtual Machines (SVMs) using various tools like the Command Line Interface (CLI), System Manager, or Ansible. Once the SVMs are configured, create Logical Interfaces (LIFs), volumes, and Logical Unit Numbers (LUNs) on the cluster that will be used for normal operations. These objects will automatically be replicated to the other cluster using the cluster peering network.

If not using MetroCluster, you can use SnapMirror active sync which provides datastore-granular protection and active-active access across multiple ONTAP clusters in different failure domains. SnapMirror active sync uses consistency groups to ensure write-order consistency among one or more datastores and you can create multiple consistency groups depending on your application and datastore requirements. Consistency groups are especially useful for applications that require data synchronization between multiple datastores. SnapMirror active sync also supports Raw Device Mappings (RDMs) and guest-connected storage with in-guest iSCSI initiators. You can learn more about consistency groups at [Consistency groups overview](#).

There is some difference in managing a vMSC configuration with SnapMirror active sync when compared to a MetroCluster. First, this is a SAN-only configuration, no NFS datastores can be protected with SnapMirror active sync. Second, you must map both copies of the LUNs to your ESXi hosts for them to access the replicated datastores in both failure domains.

VMware vSphere HA

Create a vSphere HA Cluster

Creating a vSphere HA cluster is a multi-step process that is fully documented at [How Do You Create and Configure Clusters in the vSphere Client on docs.vmware.com](#). In short, you must first create an empty cluster, and then, using vCenter, you must add hosts and specify the cluster's vSphere HA and other settings.

Note: Nothing in this document supersedes [VMware vSphere Metro Storage Cluster Recommended Practices](#)

To configure an HA cluster, complete the following steps:

1. Connect to the vCenter UI.
2. In Hosts and Clusters, browse to the data center where you want to create your HA cluster.
3. Right-click the data center object and select New Cluster. Under basics ensure you have enabled vSphere DRS and vSphere HA. Complete the wizard.

New Cluster

- 1 Basics
- 2 Image
- 3 Review

Basics ✕

Name	MCC Cluster
Location	Raleigh
vSphere DRS	<input checked="" type="checkbox"/>
vSphere HA	<input checked="" type="checkbox"/>
vSAN	<input type="checkbox"/> Enable vSAN ESA

Manage all hosts in the cluster with a single image

Choose how to set up the cluster's image

- Compose a new image
- Import image from an existing host in the vCenter inventory
- Import image from a new host

Manage configuration at a cluster level

4. Select the cluster and go to the configure tab. Select vSphere HA and click edit.
5. Under Host Monitoring, select the Enable Host Monitoring option.

Edit Cluster Settings | MCC Cluster ✕

vSphere HA

Failures and responses | Admission Control | Heartbeat Datastores | Advanced Options

You can configure how vSphere HA responds to the failure conditions on this cluster. The following failure conditions are supported: host, host isolation, VM component protection (datastore with PDL and APD), VM and application.

Enable Host Monitoring

> Host Failure Response	Restart VMs ▼
> Response for Host Isolation	Disabled ▼
> Datastore with PDL	Power off and restart VMs ▼
> Datastore with APD	Power off and restart VMs - Conservative restart policy ▼
> VM Monitoring	Disabled ▼

CANCEL
OK

6. While still on the Failures and Responses tab, Under VM Monitoring, select the VM Monitoring Only option or VM and Application Monitoring option.

> Response for Host Isolation Disabled

> Datastore with PDL Power off and restart VMs

> Datastore with APD Power off and restart VMs - Conservative restart policy

▼ VM Monitoring

Enable heartbeat monitoring

VM monitoring resets individual VMs if their VMware tools heartbeats are not received within a set time. Application monitoring resets individual VMs if their in-guest heartbeats are not received within a set time.

Disabled

VM Monitoring Only

Turns on VMware tools heartbeats. When heartbeats are not received within a set time, the VM is reset.

VM and Application Monitoring

Turns on application heartbeats. When heartbeats are not received within a set time, the VM is reset.

7. Under Admission Control, set the HA admission control option to cluster resource reserve; use 50% CPU/MEM.

vSphere HA

Failures and responses | Admission Control | Heartbeat Datastores | Advanced Options

Admission control is a policy used by vSphere HA to ensure failover capacity within a cluster. Raising the number of potential host failures will increase the availability constraints and capacity reserved.

Host failures cluster tolerates: 1
Maximum is one less than number of hosts in cluster.

Define host failover capacity by: Cluster resource Percentage

Override calculated failover capacity.

Reserved failover CPU capacity: 50 % CPU

Reserved failover Memory capacity: 50 % Memory

Reserve Persistent Memory failover capacity

Override calculated Persistent Memory failover capacity

CANCEL OK

- Click "OK".
- Select DRS and click EDIT.
- Set the automation level to manual unless required by your applications.

vSphere DRS

Automation | Additional Options | Power Management | Advanced Options

Automation Level: Manual
DRS generates both power-on placement recommendations, and migration recommendations for virtual machines. Recommendations need to be manually applied or ignored.

Migration Threshold: Conservative (Less Frequent vMotions) to Aggressive (More Frequent vMotions)

Predictive DRS: Enable

Virtual Machine Automation: Enable

11. Enable VM Component Protection, refer to docs.vmware.com.

12. The following additional vSphere HA settings are recommended for vMSC with MCC:

Failure	Response
Host failure	Restart VMs
Host isolation	Disabled
Datastore with Permanent Device Loss (PDL)	Power off and restart VMs
Datastore with All paths Down (APD)	Power off and restart VMs
Guest not heartbeating	Reset VMs
VM restart policy	Determined by the importance of the VM
Response for host isolation	Shut down and restart VMs
Response for datastore with PDL	Power off and restart VMs
Response for datastore with APD	Power off and restart VMs (conservative)
Delay for VM failover for APD	3 minutes
Response for APD recovery with APD timeout	Disabled
VM monitoring sensitivity	Preset high

Configure Datastores for Heartbeating

vSphere HA uses datastores to monitor hosts and virtual machines when the management network has failed. You can configure how vCenter selects heartbeat datastores. To configure datastores for heartbeating, complete the following steps:

1. In the Datastore Heartbeating section, select Use Datastores from the Specified List and Compliment Automatically if Needed.
2. Select the datastores you want vCenter to use from both sites and press OK.

vSphere HA









Failures and responses Admission Control **Heartbeat Datastores** Advanced Options

vSphere HA uses datastores to monitor hosts and virtual machines when the HA network has failed. vCenter Server selects 4 datastores for each host using the policy and datastore preferences specified below.

Heartbeat datastore selection policy:

- Automatically select datastores accessible from the hosts
- Use datastores only from the specified list
- Use datastores from the specified list and complement automatically if needed

Available heartbeat datastores

	Name ↑	Datastore Cluster	Hosts Mounting Datastore
<input checked="" type="checkbox"/>	 d11	N/A	2
<input checked="" type="checkbox"/>	 d12	N/A	2
<input checked="" type="checkbox"/>	 d21	N/A	2
<input checked="" type="checkbox"/>	 d22	N/A	2
<input type="checkbox"/>	 d31	N/A	2
<input type="checkbox"/>	 d32	N/A	2
<input type="checkbox"/>	 d41	N/A	2
<input type="checkbox"/>	 d42	N/A	2

11 items

Configure Advanced Options

Host Failure Detection

Isolation events occur when hosts within an HA cluster lose connectivity to either the network or other hosts in the cluster. By default, vSphere HA will use the default gateway for its management network as the default isolation address. However, you can specify additional isolation addresses for the host to ping to determine whether an isolation response should be triggered. Add two isolation IPs that can ping, one per site. Do not use the gateway IP. The vSphere HA advanced setting used is `das.isolationaddress`. You can use ONTAP or Mediator IP addresses for this purpose.

Refer to core.vmware.com for more information.

vSphere HA

Failures and responses Admission Control Heartbeat Datastores **Advanced Options**

You can set advanced options that affect the behavior of your vSphere HA cluster.

+ Add ✕ Delete

Option	Value
das.IgnoreRedundantNetWarning	true
das.Isolationaddress0	10.61.99.100
das.Isolationaddress1	10.61.99.110
das.heartbeatDsPerHost	4

4 items

CANCEL OK

Adding an advanced setting called `das.heartbeatDsPerHost` can increase the number of heartbeat datastores. Use four heartbeat datastores (HB DSs)—two per site. Use the “Select from List but Compliment” option. This is needed because if one site fails, you still need two HB DSs. However, those don’t have to be protected with MCC or SnapMirror active sync.

Refer to core.vmware.com for more information.

VMware DRS Affinity for NetApp MetroCluster

In this section, we create DRS groups for VMs and hosts for each site\cluster in the MetroCluster environment. Then we configure VM\Host rules to align VM host affinity with local storage resources. For example, site A VMs belong to VM group `sitea_vms` and site A hosts belong to host group `sitea_hosts`. Next, in VM\Host Rules, we state that `sitea_vms` should run on hosts in `sitea_hosts`.

Best Practice

- NetApp highly recommends the specification **Should Run on Hosts in Group** rather than the specification **Must Run on Hosts in Group**. In the event of a site A host failure, the VMs of site A need to be restarted on hosts at site B through vSphere HA, but the latter specification does not allow HA to restart VMs on site B because it’s a hard rule. The former specification is a soft rule and will be violated in the event of HA, thus enabling availability rather than performance.

Note: You can create an event-based alarm that is triggered when a virtual machine violates a VM-Host affinity rule. In the vSphere Client, add a new alarm for the virtual machine and select “VM is violating VM-Host Affinity Rule” as the event trigger. For more information about creating and editing alarms, refer to [vSphere Monitoring and Performance](#) documentation.

Create DRS Host Groups

To create DRS host groups specific to site A and site B, complete the following steps:

1. In the vSphere web client, right-click the cluster in the inventory and select Settings.
2. Click VM\Host Groups.
3. Click Add.
4. Type the name for the group (for instance, sitea_hosts).
5. From the Type menu, select Host Group.
6. Click Add and select the desired hosts from site A and click OK.
7. Repeat these steps to add another host group for site B.
8. Click OK.

Create DRS VM Groups

To create DRS VM groups specific to site A and site B, complete the following steps:

1. In the vSphere web client, right-click the cluster in the inventory and select Settings.
9. Click VM\Host Groups.
10. Click Add.
11. Type the name for the group (for instance, sitea_vms).
12. From the Type menu, select VM Group.
13. Click Add and select the desired VMs from site A and click OK.
14. Repeat these steps to add another host group for site B.
15. Click OK.

Create VM Host Rules

To create DRS affinity rules specific to site A and site B, complete the following steps:

1. In the vSphere web client, right-click the cluster in the inventory and select Settings.
1. Click VM\Host Rules.
2. Click Add.
3. Type the name for the rule (for instance, sitea_affinity).
4. Verify the Enable Rule option is checked.
5. From the Type menu, select Virtual Machines to Hosts.
6. Select the VM group (for instance, sitea_vms).
7. Select the Host group (for instance, sitea_hosts).
8. Repeat these steps to add another VM\Host Rule for site B.

9. Click OK.

Create VM/Host Rule | Cluster-01 ×

Name	sitea_affinity <input checked="" type="checkbox"/> Enable rule.
Type	Virtual Machines to Hosts ▼

Virtual machines that are members of the Cluster VM Group sitea_vms should run on host group sitea_hosts.

VM Group:

sitea_vms ▼
Should run on hosts in group ▼

Host Group:

sitea_hosts ▼

VMWare vSphere Storage DRS for NetApp MetroCluster

Create Datastore Clusters

To configure a datastore cluster for each site, complete the following steps:


1. Using the vSphere web client, browse to the data center where the HA cluster resides under Storage.
2. Right-click the data center object and select Storage > New Datastore Cluster.
3. Select the Turn ON Storage DRS option and click Next.
4. Set all options to No Automation (Manual Mode) and click Next.

Best Practice

- NetApp recommends that Storage DRS be configured in manual mode, so that the administrator gets to decide and control when migrations need to happen.

Storage DRS automation	
Cluster automation level	<input checked="" type="radio"/> No Automation (Manual Mode) vCenter Server will make migration recommendations for virtual machine storage, but will not perform automatic migrations.
	<input type="radio"/> Fully Automated Files will be migrated automatically to optimize resource usage.

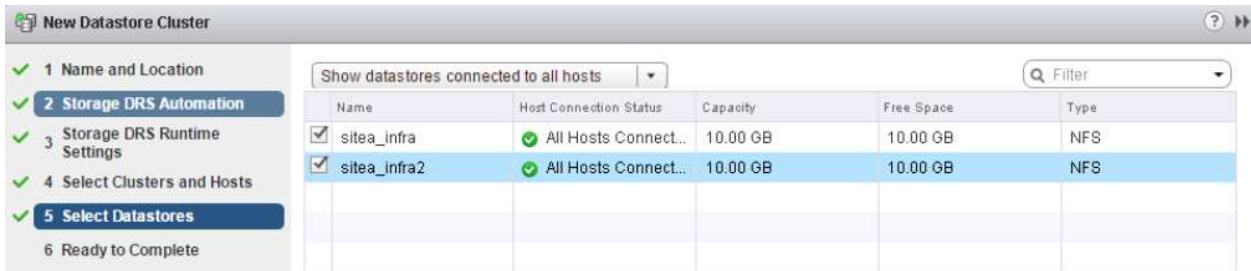
5. Verify that the Enable I/O Metric for SDRS Recommendations checkbox is checked; metric settings can be left with default values.



6. Select the HA cluster and click Next.



7. Select the datastores belonging to site A and click Next.



8. Review options and click Finish.

9. Repeat these steps to create the site B datastore cluster and verify that only datastores of site B are selected.

vCenter Server Availability

Your vCenter Server Appliances (VCSAs) should be protected with vCenter HA. vCenter HA allows you to deploy two VCSAs in an active-passive HA pair. One in each failure domain. You can read more about vCenter HA on docs.vmware.com.

Resiliency for Planned and Unplanned Events

NetApp MetroCluster and SnapMirror active sync are powerful tools that enhance the high availability and non-disruptive operations of NetApp hardware and ONTAP® software.

These tools provide site-wide protection for the entire storage environment, ensuring that your data is always available. Whether you are using standalone servers, high-availability server clusters, Docker containers, or virtualized servers, NetApp technology seamlessly maintains storage availability in the event of a total outage due to loss of power, cooling, or network connectivity, storage array shutdown, or operational error.

MetroCluster and SnapMirror active sync provide three basic methods for data continuity in the event of planned or unplanned events:

- Redundant components for protection against single-component failure
- Local HA takeover for events affecting a single controller
- Complete site protection – rapid resumption of service by moving storage and client access from the source cluster to the destination cluster

This means operations continue seamlessly in case of a single component failure and return automatically to redundant operation when the failed component is replaced.

All ONTAP clusters, except single-node clusters (typically software-defined versions, such as ONTAP Select for example), have built-in HA features called takeover and giveback. Each controller in the cluster is paired with another controller, forming an HA pair. These pairs ensure that each node is locally connected to the storage.

Takeover is an automated process where one node takes over the other's storage to maintain data services. Giveback is the reverse process that restores normal operation. Takeover can be planned, such as when performing hardware maintenance or ONTAP upgrades, or unplanned, resulting from a node panic or hardware failure.

During a takeover, Network Attached Storage Logical Interfaces (NAS LIFs) in MetroCluster configurations automatically failover. However, Storage Area Network LIFs (SAN LIFs) do not fail over; they will continue to use the direct path to the Logical Unit Numbers (LUNs).

For more information on HA takeover and giveback, please refer to the [HA pair management overview](#). It's worth noting that this functionality is not specific to MetroCluster or SnapMirror active sync.

Site switchover with MetroCluster occurs when one site is offline or as a planned activity for site-wide maintenance. The remaining site assumes ownership of the storage resources (disks and aggregates) of the offline cluster, and the SVMs on the failed site are brought online and restarted on the disaster site, preserving their full identity for client and host access.

With SnapMirror active sync, since both copies are actively used simultaneously, your existing hosts will continue to operate. The NetApp Mediator is required to ensure site failover occurs correctly.

Failure Scenarios for vMSC with MCC

The following sections outline the expected results from various failure scenarios with vMSC and NetApp MetroCluster systems.

Single Storage Path Failure

In this scenario, if components such as the HBA port, the network port, the front-end data switch port, or an FC or Ethernet cable fails, that particular path to the storage device is marked as dead by the ESXi host. If several paths are configured for the storage device by providing resiliency at the HBA/network/switch port, ESXi ideally performs a path switchover. During this period, virtual machines remain running without getting affected, because availability to the storage is taken care of by providing multiple paths to the storage device.

Note: There is no change in MetroCluster behavior in this scenario, and all the datastores continue to be intact from their respective sites.

Best Practice

In environments in which NFS/iSCSI volumes are used, NetApp recommends having at least two network uplinks configured for the NFS vmkernel port in the standard vSwitch and the same at the port group where the NFS vmkernel interface is mapped for the distributed vSwitch. NIC teaming can be configured in either active-active or active-standby.


Also, for iSCSI LUNs, multipathing must be configured by binding the vmkernel interfaces to the iSCSI network adapters. For more information, refer to the vSphere storage documentation.

Best Practice

In environments in which Fibre Channel LUNs are used, NetApp recommends having at least two HBAs, which guarantees resiliency at the HBA/port level. NetApp also recommends single initiator to single target zoning as the best practice to configure zoning.

Virtual Storage Console (VSC) should be used to set multipathing policies because it sets policies for all new and existing NetApp storage devices.

Single ESXi Host Failure




In this scenario, if there is an ESXi host failure, the master node in the VMware HA cluster detects the host failure since it no longer receives network heartbeats. To determine whether the host is really down or only a network partition, the master node monitors the datastore heartbeats and, if they are absent, it performs a final check by pinging the management IP addresses of the failed host. If all these checks are negative, then the master node declares this host a failed host and all the virtual machines that were running on this failed host are rebooted on the surviving host in the cluster.

If DRS VM and host affinity rules have been configured (VMs in VM group `sita_vms` should run hosts in host group `sita_hosts`), then the HA master first checks for available resources at site A. If there are no available hosts at site A, the master attempts to restart the VMs on hosts at site B.

It is possible that the virtual machines will be started on the ESXi hosts at the other site if there is a resource constraint in the local site. However, the defined DRS VM and host affinity rules will correct if any rules are violated by migrating the virtual machines back to any surviving ESXi hosts in the local site. In cases in which DRS is set to manual, NetApp recommends invoking DRS and applying the recommendations to correct the virtual machine placement.

There is no change in the MetroCluster behavior in this scenario and all the datastores continue to be intact from their respective sites.

ESXi Host Isolation




In this scenario, if the management network of the ESXi host is down, the master node in the HA cluster will not receive any heartbeats, and thus this host becomes isolated in the network. To determine whether it has failed or is only isolated, the master node starts monitoring the datastore heartbeat. If it is present then the host is declared isolated by the master node. Depending on the isolation response configured, the host may choose to power off, shut down the virtual machines, or even leave the virtual machines powered on. The default interval for the isolation response is 30 seconds.

There is no change in the MetroCluster behavior in this scenario and all the datastores continue to be intact from their respective sites.

Disk Shelf Failure


In this scenario, there is a failure of more than two disks or an entire shelf. Data is served from the surviving plex with no interruption to data services. The disk failure could affect either a local or remote plex. The aggregates will show as degraded mode because only one plex is active. Once the failed disks are replaced, the affected aggregates will automatically resync to rebuild the data. After resync, the aggregates will return automatically to normal mirrored mode. If more than two disks within a single RAID group have failed, then the plex has to be rebuilt from scratch.



Note: During this period, there is no impact on the virtual machine I/O operations, but there is degraded performance because the data is being accessed from the remote disk shelf through ISL links.

Single Storage Controller Failure


In this scenario, one of the two storage controllers fails at one site. Because there is an HA pair at each site, a failure of one node transparently and automatically triggers failover to the other node. For example, if node A1 fails, its storage and workloads are automatically transferred to node A2. Virtual machines will not be affected because all plexes remain available. The second site nodes (B1 and B2) are unaffected. In addition, vSphere HA will not take any action because the master node in the cluster will still be receiving the network heartbeats.



If the failover is part of a rolling disaster (node A1 fails over to A2), and there is a subsequent failure of A2, or the complete failure of site A, switchover following a disaster can occur at site B.

Interswitch Link Failures


Interswitch Link Failure at Management Network



In this scenario, if the ISL links at the front-end host management network fail, the ESXi hosts at site A will not be able to communicate with ESXi hosts at site B. This will lead to a network partition because ESXi hosts at a particular site will be unable to send the network heartbeats to the master node in the HA cluster. As such, there will be two network segments because of partition and there will be a master node in each segment that will protect VMs from host failures within the particular site.

Note: During this period, the virtual machines remain running and there is no change in the MetroCluster behavior in this scenario. All the datastores continue to be intact from their respective sites.


Interswitch Link Failure at Storage Network



In this scenario, if the ISL links at the backend storage network fail, the hosts at site A will lose access to the storage volumes or LUNs of cluster B at site B and vice versa. The VMware DRS rules are defined so that host-storage site affinity facilitates the virtual machines to run without impact within the site.

During this period, the virtual machines remain running in their respective sites and there is no change in the MetroCluster behavior in this scenario. All the datastores continue to be intact from their respective sites.


If for some reason the affinity rule was violated (for example, VM1, which was supposed to run from site A where its disks reside on local cluster A nodes, is running on a host at site B), the virtual machine's disk will be remotely accessed via ISL links. Because of ISL link failure, VM1 running at site B would not be able to write to its disks because the paths to the storage volume are down and that particular virtual machine is down. In these situations, VMware HA does not take any action since the hosts are actively sending heartbeats. Those virtual machines need to be manually powered off and powered on in their respective sites. The following figure illustrates a VM violating a DRS affinity rule.



All Interswitch Failure or Complete Data Center Partition

In this scenario, all the ISL links between the sites are down and both the sites are isolated from each other. As discussed in earlier scenarios, such as ISL failure at the management network and at the storage network, the virtual machines are not affected in complete ISL failure.


After ESXi hosts are partitioned between sites, the vSphere HA agent will check for datastore heartbeats and, in each site, the local ESXi hosts will be able to update the datastore heartbeats to their respective read-write volume/LUN. Hosts in site A will assume that the other ESXi hosts at site B have failed because there are no network/datastore heartbeats. vSphere HA at site A will try to restart the virtual machines of site B, which will eventually fail because the datastores of site B will not be accessible due to storage ISL failure. A similar situation is repeated in site B.



Interswitch Link Failure on Both Fabrics in NetApp MetroCluster


In a scenario of one or more ISLs failing, traffic continues through the remaining links. If all ISLs on both fabrics fail, such that there is no link between the sites for storage and NVRAM replication, each controller will continue to serve its local data. On restoration of a minimum of one ISL, resynchronization of all the plexes will happen automatically.

Any writes occurring after all ISLs are down will not be mirrored to the other site. A switchover on disaster, while the configuration is in this state, would therefore incur loss of the data that had not been synchronized. In this case, manual intervention is required for recovery after the switchover. If it is likely that no ISLs will be available for an extended period, an administrator can choose to shut down all data services to avoid the risk of data loss if a switchover on disaster is necessary. Performing this action should be weighed against the likelihood of a disaster requiring switchover before at least one ISL becomes available. Alternatively, if ISLs are failing in a cascading scenario, an administrator could trigger a planned switchover to one of the sites before all the links have failed.



Peered Cluster Link Failure

In a peered cluster link failure scenario, because the fabric ISLs are still active, data services (reads and writes) continue at both sites to both plexes. Any cluster configuration changes (for example, adding a new SVM, provisioning a volume or LUN in an existing SVM) cannot be propagated to the other site. These are kept in the local CRS metadata volumes and automatically propagated to the other cluster upon restoration of the peered cluster link. If a forced switchover is necessary before the peered cluster link can be restored, outstanding cluster configuration changes will be replayed automatically from the remote replicated copy of the metadata volumes at the surviving site as part of the switchover process.



Complete Site Failure

In a complete site A failure scenario, the ESXi hosts at site B will not get the network heartbeat from the ESXi hosts at site A because they are down. The HA master at site B will verify that the datastore heartbeats are not present, declare the hosts at site A to be failed, and try to restart the site A virtual machines in site B. During this period, the storage administrator performs a switchover to resume services of the failed nodes on the surviving site which will restore all the storage services of site A at site B. After the site A volumes or LUNs are available at site B, the HA master agent will attempt to restart the site A virtual machines in site B.

If the vSphere HA master agent's attempt to restart a VM (which involves registering it and powering it on) fails, the restart is retried after a delay. The delay between restarts can be configured to up to a maximum of 30 minutes. vSphere HA attempts these restarts for a maximum number of attempts (six attempts by default).

Note: The HA master does not begin the restart attempts until the placement manager finds suitable storage, so in the case of a complete site failure, that would be after the switchover has been performed.

If site A has been switched over, a subsequent failure of one of the surviving site B nodes can be seamlessly handled by failover to the surviving node. In this case, the work of four nodes is now being performed by only one node. Recovery in this case would consist of performing a giveback to the local node. Then, when site A is restored, a switchover operation is performed to restore steady state operation of the configuration.

any time.

- **Vulnerability scanning.** The purpose of vulnerability scanning is to detect common and known security vulnerabilities in NetApp products before they are released to customers.
- **Penetration testing.** Penetration testing is the process of evaluating a system, web application, or network to find security vulnerabilities that could be exploited by an attacker. Penetration tests (pen tests) at NetApp are conducted by a group of approved and trusted third-party companies. Their testing scope includes the launching of attacks against an application or software similar to hostile intruders or hackers using sophisticated exploitation methods or tools.

Product security features

ONTAP tools for VMware vSphere includes the following security features in each release.

- **Login banner.** SSH is disabled by default and only allows one-time logins if enabled from the VM console. The following login banner is shown after the user enters a username in the login prompt:

WARNING: Unauthorized access to this system is forbidden and will be prosecuted by law. By accessing this system, you agree that your actions may be monitored if unauthorized usage is suspected.

After the user completes login through the SSH channel, the following text is displayed:

```
Linux vscl 4.19.0-12-amd64 #1 SMP Debian 4.19.152-1 (2020-10-18) x86_64
The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.
Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
```

- **Role-based access control (RBAC).** Two kinds of RBAC controls are associated with ONTAP tools:
 - Native vCenter Server privileges
 - vCenter plug-in specific privileges. For details, see [this link](#).
- **Encrypted communications channels.** All external communication happens over HTTPS using version 1.2 of TLS.
- **Minimal port exposure.** Only the necessary ports are open on the firewall.

The following table describes the open port details.

TCP v4/v6 port #	Direction	Function
8143	inbound	HTTPS connections for REST API
8043	inbound	HTTPS connections
9060	inbound	HTTPS connections Used for SOAP over https connections This port must be opened to allow a client to connect to the ONTAP tools API server.

TCP v4/v6 port #	Direction	Function
22	inbound	SSH (Disabled by default)
9080	inbound	HTTPS connections - VP and SRA - Internal connections from loopback only
9083	inbound	HTTPS connections - VP and SRA Used for SOAP over https connections
1162	inbound	VP SNMP trap packets
1527	internal only	Derby database port, only between this computer and itself, external connections not accepted — Internal connections only
443	bi-directional	Used for connections to ONTAP clusters

- **Support for certificate authority (CA) signed certificates.** ONTAP tools for VMware vSphere supports CA signed certificates. See this [kb article](#) for more information.
- **Audit logging.** Support bundles can be downloaded and are extremely detailed. ONTAP tools logs all user login and logout activity in a separate log file. VASA API calls are logged in a dedicated VASA audit log (local cxf.log).
- **Password policies.** The following password policies are followed:
 - Passwords are not logged in any log files.
 - Passwords are not communicated in plain text.
 - Passwords are configured during the installation process itself.
 - Password history is a configurable parameter.
 - Minimum password age is set to 24 hours.
 - Auto complete for the password fields are disabled.
 - ONTAP tools encrypts all stored credential information using SHA256 hashing.

SnapCenter Plug-in VMware vSphere

NetApp SnapCenter Plug-in for VMware vSphere software engineering uses the following secure development activities:

- **Threat modeling.** The purpose of threat modelling is to discover security flaws in a feature, component, or product early in the software development life cycle. A threat model is a structured representation of all the information that affects the security of an application. In essence, it is a view of the application and its environment through the lens of security.
- **Dynamic application security testing (DAST).** Technologies that are designed to detect vulnerable conditions on applications in their running state. DAST tests the exposed HTTP and HTML interfaces of web-enabled applications.
- **Third-party code currency.** As part of developing software and using open-source software (OSS), it is

important to address security vulnerabilities that might be associated with OSS that has been incorporated into your product. This is a continuous effort as the version of the OSS component may have a newly discovered vulnerability reported at any time.

- **Vulnerability scanning.** The purpose of vulnerability scanning is to detect common and known security vulnerabilities in NetApp products before they are released to customers.
- **Penetration testing.** Penetration testing is the process of evaluating a system, web application or network to find security vulnerabilities that could be exploited by an attacker. Penetration tests (pen tests) at NetApp are conducted by a group of approved and trusted third-party companies. Their testing scope includes the launching of attacks against an application or software like hostile intruders or hackers using sophisticated exploitation methods or tools.
- **Product Security Incident Response activity.** Security vulnerabilities are discovered both internally and externally to the company and can pose a serious risk to NetApp's reputation if they are not addressed in a timely manner. To facilitate this process, a Product Security Incident Response Team (PSIRT) reports and tracks the vulnerabilities.

Product security features

NetApp SnapCenter Plug-in for VMware vSphere includes the following security features in each release:

- **Restricted shell access.** SSH is disabled by default, and one-time logins are only allowed if they are enabled from the VM console.
- **Access warning in login banner.** The following login banner is shown after the user enters a user name in the login prompt:

WARNING: Unauthorized access to this system is forbidden and will be prosecuted by law. By accessing this system, you agree that your actions may be monitored if unauthorized usage is suspected.

After the user completes login through the SSH channel, the following output displays:

```
Linux vsc1 4.19.0-12-amd64 #1 SMP Debian 4.19.152-1 (2020-10-18) x86_64
The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.
Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
```

- **Role-based access control (RBAC).** Two kinds of RBAC controls are associated with ONTAP tools:
 - Native vCenter Server privileges.
 - VMware vCenter plug-in specific privileges. For more information, see [Role-Based Access Control \(RBAC\)](#).
- **Encrypted communications channels.** All external communication happens over HTTPS by using TLS.
- **Minimal port exposure.** Only the necessary ports are open on the firewall.

The following table provides the open port details.

TCP v4/v6 port number	Function
8144	HTTPS connections for REST API

TCP v4/v6 port number	Function
8080	HTTPS connections for OVA GUI
22	SSH (disabled by default)
3306	MySQL (internal connections only; external connections disabled by default)
443	Nginx (data protection services)

- **Support for Certificate Authority (CA) signed certificates.** SnapCenter Plug-in for VMware vSphere supports the feature of CA signed certificates. See [How to create and/or import an SSL certificate to SnapCenter Plug-in for VMware vSphere \(SCV\)](#).
- **Password policies.** The following password policies are in effect:
 - Passwords are not logged in any log files.
 - Passwords are not communicated in plain text.
 - Passwords are configured during the installation process itself.
 - All credential information is stored using SHA256 hashing.
- **Base operating system image.** The product ships with Debian Base OS for OVA with restricted access and shell access disabled. This reduces the attack footprint. Every SnapCenter release base operating system is updated with latest security patches available for maximum security coverage.

NetApp develops software features and security patches with regards to SnapCenter Plug-in for VMware vSphere appliance and then releases them to customers as a bundled software platform. Because these appliances include specific Linux sub-operating system dependencies as well as our proprietary software, NetApp recommends that you do not make changes to the sub-operating system because this has a high potential to affect the NetApp appliance. This could affect the ability of NetApp to support the appliance. NetApp recommends testing and deploying our latest code version for appliances because they are released to patch any security-related issues.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.