



VMware vSphere with ONTAP

Enterprise applications

NetApp
February 10, 2026

Table of Contents

VMware vSphere with ONTAP	1
VMware vSphere with ONTAP	1
Why ONTAP for VMware vSphere?	1
The advantages of using ONTAP for vSphere	1
Unified Storage	3
Virtualization tools for ONTAP	4
ONTAP tools for VMware vSphere	4
SnapCenter Plug-In for VMware vSphere	5
NFS Plug-In for VMware VAAI	6
Premium software options	6
Virtual Volumes (vVols) and Storage Policy Based Management (SPBM)	6
Virtual Volumes (vVols)	6
Storage Policy Based Management (SPBM)	6
NetApp ONTAP and vVols	7
Datastores and protocols	7
vSphere datastore and protocol features overview	7
SAN (FC, FCoE, NVMe/FC, iSCSI), RDM	13
NFS	16
FlexGroup volumes	19
Network configuration	21
SAN (FC, NVMe/FC, iSCSI, NVMe/TCP), RDM	22
NFS	22
Direct connect networking	22
VM and datastore cloning	23
Data protection	25
NetApp ONTAP volume snapshots	25
SnapCenter Plug-In for VMware vSphere	26
ONTAP tools for VMware vSphere with VMware Live Site Recovery	27
NetApp Disaster Recovery	27
vSphere Metro Storage Cluster (vMSC) with NetApp MetroCluster and SnapMirror active sync	27
Quality of service (QoS)	28
ONTAP QoS policy support	28
Non-vVols NFS datastores	28
VMFS datastores	29
vVols datastores	29
ONTAP QoS and VMware SIOC	29
VMware Storage Distributed Resource Scheduler	30
Storage policy based management and vVols	31
Cloud migration and backup	33
Encryption for vSphere data	33
Active IQ Unified Manager	34
Storage policy based management and vVols	36
VMware Storage Distributed Resource Scheduler	38

Recommended ESXi host and other ONTAP settings	39
Multipath settings for performance	41
Additional documentation	42

VMware vSphere with ONTAP

VMware vSphere with ONTAP

ONTAP has served as a premier storage solution for VMware vSphere and, more recently, Cloud Foundation environments since its introduction into the modern datacenter in 2002. It continues to introduce innovative features that simplify management and lower costs.

This document presents the ONTAP solution for vSphere, highlighting the latest product information and best practices to streamline deployment, mitigate risks, 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 [Interoperability Matrix Tool \(IMT\)](#) and [VMware Compatibility Guide](#) for details related to specific releases.

Why ONTAP for VMware vSphere?

Customers confidently select ONTAP for vSphere for both SAN and NAS storage solutions. The new simplified disaggregated storage architecture, which is featured in the latest All SAN Arrays, delivers a simplified experience familiar to SAN storage administrators while keeping most of the integrations and feature set of traditional ONTAP systems. ONTAP systems provide exceptional snapshot protection and robust management tools. By offloading functions to dedicated storage, ONTAP maximizes host resources, reduces costs, and maintains optimal performance. Additionally, workloads can be easily migrated using Storage vMotion across VMFS, NFS, or vVols.

The advantages of using 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 a 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 the key factors customers value today:

- **Unified storage.** Systems running ONTAP 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 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 storage operating system for performance, another for archives, and yet another one for the cloud. ONTAP ties them all together.
- **All SAN Array (ASA).** The latest ONTAP ASA systems (beginning with the A1K, A90, A70, A50, A30, and A20) are built on a new storage architecture that eliminates the traditional ONTAP storage paradigm of managing aggregates and volumes. Since there are no file system shares, there's no need for volumes! All storage attached to an HA pair is treated as a common Storage Availability Zone (SAZ) within which LUNs and NVMe namespaces are provisioned as "Storage Units" (SUs). The latest ASA systems are designed to be simple to manage, with a familiar experience for SAN storage administrators. This new architecture is ideal for vSphere environments, as it allows for easy management of storage resources and provides a simplified experience for SAN storage administrators. The ASA architecture also supports the latest NVMe over Fabrics (NVMe-oF) technology, which provides even greater performance and scalability for vSphere workloads.
- **Snapshot technology.** ONTAP was the first to deliver snapshot technology for data protection, and it remains the most advanced in the industry. This space-efficient approach to data protection has been extended to support VMware vSphere APIs for Array Integration (VAAI). This integration allows you to take advantage of ONTAP's snapshot capabilities for backup and restore operations, reducing the impact on your production environment. This approach also allows you to use snapshots for rapid recovery of VMs, reducing the time and effort required to restore data. In addition, ONTAP's snapshot technology is integrated with VMware's Live Site Recovery (VLSR, formerly Site Recovery Manager [SRM]) solutions, providing a comprehensive data protection strategy for your virtualized environment.
- **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.
- **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. ONTAP also 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 commonly seeing savings of up to 5:1 for VSI and up to 30:1 for VDI. The newest generation of ONTAP systems also includes hardware-accelerated compression and deduplication, which can further improve storage efficiency and reduce costs. This approach allows you to store more data in less space, reducing the overall cost of storage and improving performance. NetApp is so confident in its storage efficiency capabilities that it offers a link: <https://www.netapp.com/pdf.html?item=/media/79014-ng-937-Efficiency-Guarantee-Customer-Flyer.pdf> [Efficiency Guarantee^].

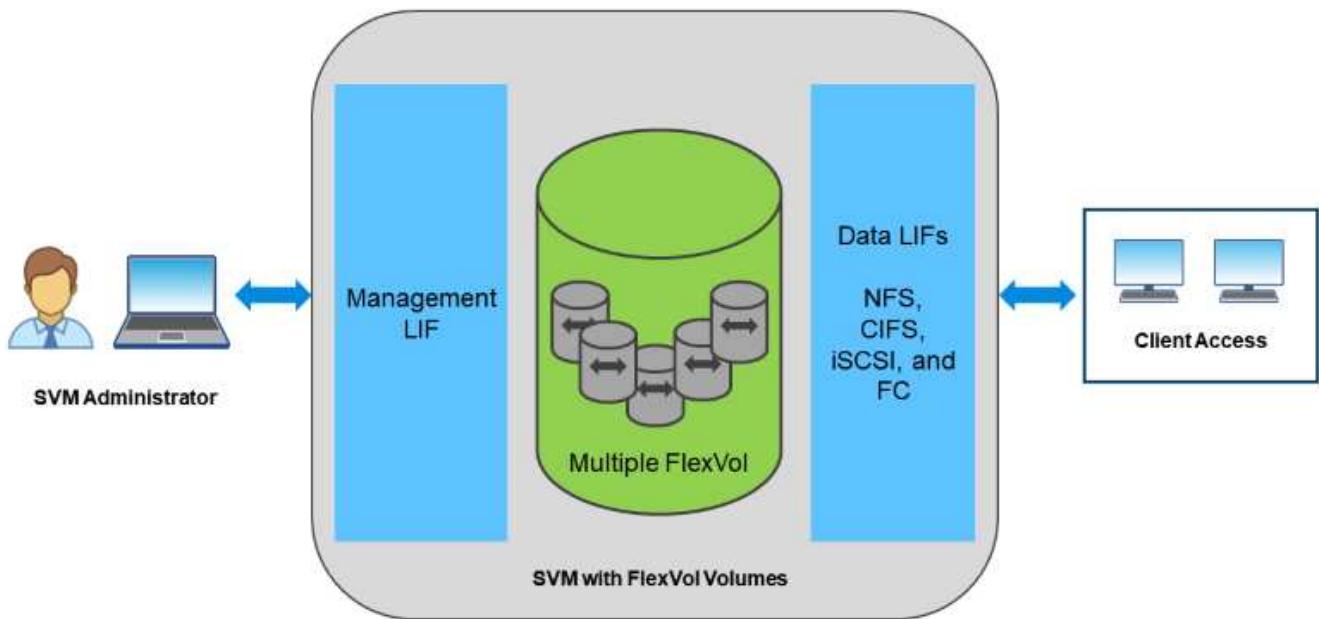
- **Multitenancy.** ONTAP has long been a leader in multitenancy, allowing you to create multiple storage virtual machines (SVMs) on a single cluster. This approach allows you to isolate workloads and provide different levels of service to different tenants, making it ideal for service providers and large enterprises. The latest generation of ONTAP systems also includes support for tenant capacity management. This feature allows you to set capacity limits for each tenant, ensuring that no single tenant can consume all the available resources. This approach helps to ensure that all tenants receive the level of service they expect, while also providing a high level of security and isolation between tenants. In addition, ONTAP's multitenancy capabilities are integrated with VMware's vSphere platform, allowing you to easily manage and monitor your virtualized environment through [ONTAP tools for VMware vSphere](#) and [Data Infrastructure Insights](#).
- **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 Cloud 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

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](#). The term unified storage has evolved to mean a unified approach to storage management, including the ability to manage all of your storage resources from a single interface. This includes the ability to manage both on-premises and cloud storage resources, the latest All SAN Array (ASA) systems, and the ability to manage multiple storage systems from a single interface.

A storage virtual machine (SVM) is the unit of secure multitenancy in ONTAP. It is a logical construct allowing client access to systems running ONTAP. 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 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 provides several standalone software tools compatible with both traditional ONTAP and ASA systems, integrating vSphere to effectively manage your virtualized environment.

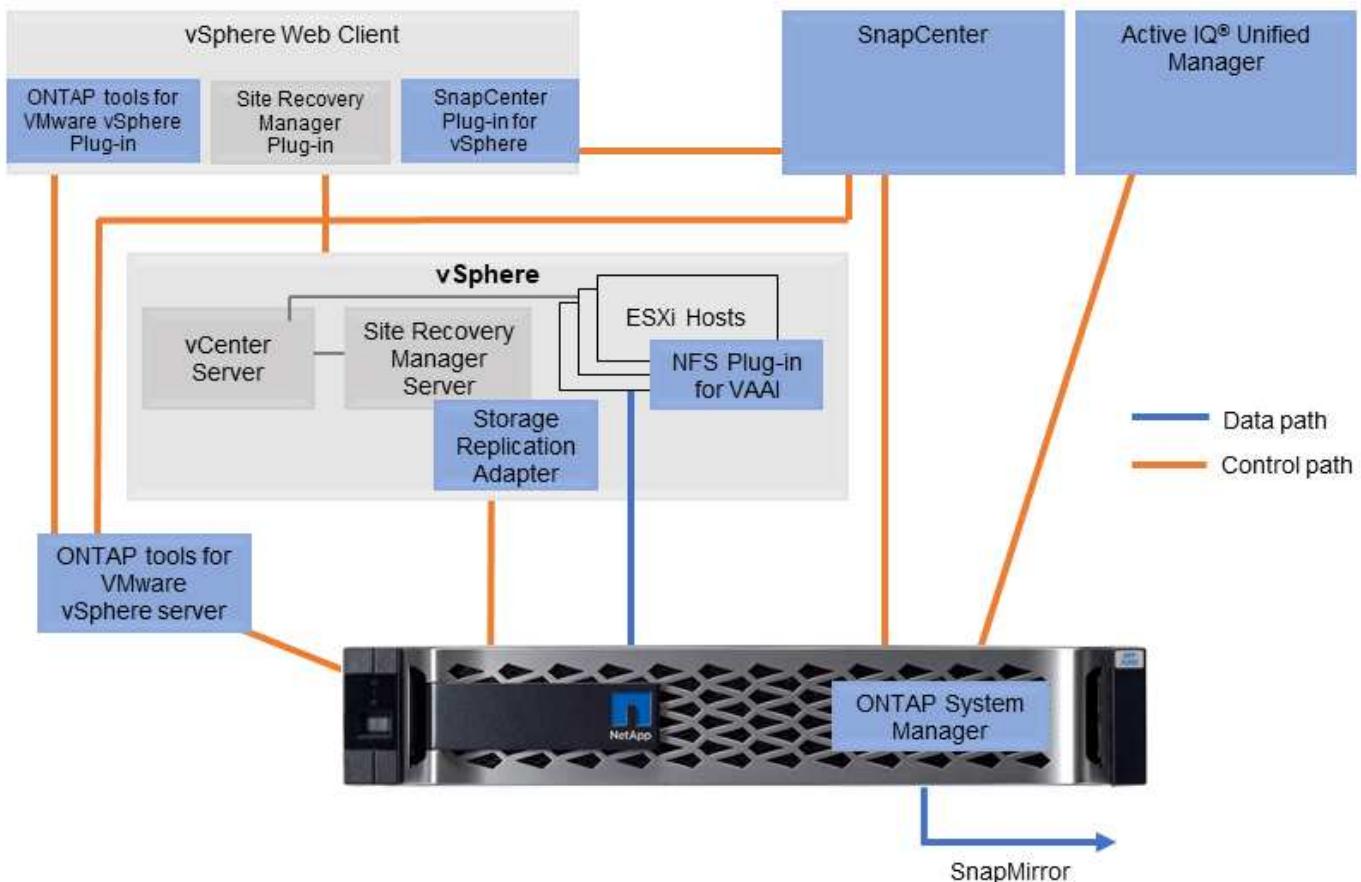
The following tools are included with the ONTAP One license at no additional cost. See Figure 1 for a depiction of how these tools work together in your vSphere environment.

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. It includes a server appliance, UI 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 Live Site Recovery (VLSR)/Site Recovery Manager (SRM) to manage data replication between production and disaster recovery sites using SnapMirror for array-based replication. It can automate the task of failover in the event of disaster, and can help test the DR replicas nondisruptively to ensure confidence in your DR solution.

The following figure depicts ONTAP tools for vSphere.



SnapCenter Plug-In for VMware vSphere

The [SnapCenter Plug-In for VMware vSphere](#) is a plug-in for vCenter Server that enables you to manage backups and restores of virtual machines (VMs) and datastores. It provides a single interface for managing backups, restores, and clones of VMs and datastores across multiple ONTAP systems. SnapCenter supports replication to and recovery from secondary sites using SnapMirror. The latest versions also support SnapMirror to cloud (S3), Tamperproof snapshots, SnapLock, and SnapMirror active sync. The SnapCenter Plug-In for VMware vSphere can be integrated with SnapCenter application plugins to provide application-consistent backups.

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).

Premium software options

The following premium software products are available from NetApp. They are not included with the ONTAP One license and must be purchased separately.

- [NetApp Disaster Recovery \(DR\)](#) for VMware vSphere. This is a cloud-based service that provides disaster recovery and backup for VMware environments. It can be used with or without SnapCenter and supports on-prem to on-prem DR using SAN or NAS, and on-prem to/from cloud using NFS, where supported.
- [Data Infrastructure Insights \(DII\)](#). This is a cloud-based service that provides monitoring and analytics for VMware environments. It supports other storage vendors in a heterogenous storage environments, as well as multiple switch vendors and other hypervisors. DII provides complete end-to-end insights into the performance, capacity, and health of your VMware environment.

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

First announced in 2012, NetApp was an early design partner with VMware in the development of VMware vSphere APIs for Storage Awareness (VASA), the foundation of Storage Policy Based Management (SPBM) with enterprise storage arrays. This approach brought limited VM granular storage management to VMFS and NFS storage.

As a technology design partner, NetApp provided architectural input and in 2015 announced support for vVols. This new technology now enabled the automation of VM-granular and truly array-native storage provisioning through SPBM.

Virtual Volumes (vVols)

vVols are a revolutionary storage architecture that enables VM granular storage management, allowing storage to be managed on not just a per-VM basis (including VM metadata), but even on a per VMDK basis. vVols are a key component of the Software Defined Data Center (SDDC) strategy that forms the basis of VMware Cloud Foundation (VCF), providing a more efficient and scalable storage architecture for virtualized environments.

vVols enable VMs to consume storage on a per-VM basis because each VM storage object is a unique entity in NetApp ONTAP. With ASA r2 systems which no longer require volume management this means that each VM storage object is a unique storage unit (SU) on the array and can be independently controlled. This allows for the creation of storage policies that can be applied to individual VMs or VMDKs (and thus individual SUs), providing granular control over storage services such as performance, availability, and data protection.

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. These pools can be easily consumed by VM administrators. Administrators can then match virtual machine workload requirements against the provisioned storage pools. This approach simplifies storage management and allows for more efficient use of storage resources.

SPBM is a key component of vVols, providing a policy-based framework for managing storage services. Policies are created by vSphere administrators using rules and capabilities exposed by the vendor's VASA Provider(VP). Policies can be created for different storage services such as performance, availability, and data protection. Policies can be assigned to individual VMs or VMDKs, providing granular control over storage services.

NetApp ONTAP and vVols

NetApp ONTAP leads the storage industry in vVols scale, supporting hundreds of thousands of vVols in a single cluster*. In contrast, enterprise array and smaller flash array vendors support as few as several thousand vVols per array. ONTAP provides a scalable and efficient storage solution for VMware vSphere environments, supporting vVols with a rich set of storage services, including data deduplication, compression, thin provisioning, and data protection. SPBM allows for seamless integration with VMware vSphere environments.

Previously we mentioned that VM administrators can consume capacity as storage pools. This is done through the use of storage containers that are represented in vSphere as logical datastores.

Storage containers are created by storage administrators and are used to group storage resources that can be consumed by VM administrators. Storage containers can be created differently depending on what type of ONTAP system you are using. With traditional ONTAP 9 clusters, containers are assigned one or more backing FlexVol volumes that together form the storage pool. With ASA r2 systems, the entire cluster is the storage pool.



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

*Depending on platform and protocol

Datastores and protocols

vSphere datastore and protocol features overview

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

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

FCP, 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. 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 the [VMware Configuration Maximums tool](#) for specific vSphere releases to confirm specific limits.

Capability/Feature	FC	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 ESXi host, up to 32 paths per LUN, up to 4096 total paths per host, up to 128 hosts per datastore	1024 LUNs per ESXi host, up to 32 paths per LUN, up to 4096 total paths per host, up to 128 hosts per datastore	256 Namespaces per ESXi host, up to 32 paths per namespace per host, 2048 total paths per host, up to 16 hosts per datastore	256 NFS connections per host (impacted by nconnect and session trunking) Default NFS. MaxVolumes is 8. Use ONTAP tools for VMware vSphere to increase to 256.
Maximum datastore size	64TB	64TB	64TB	300TB 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	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

Capacity/Feature	FC	iSCSI	NVMe-oF	NFS
Microsoft Cluster Service (MSCS) or failover clustering within a VM	Yes	Yes ¹	Yes ¹	Not supported
Fault Tolerance	Yes	Yes	Yes	Yes
Live Site Recovery/Site Recovery Manager	Yes	Yes	No ²	V3 only ²
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	NFS v4.1 session trunking requires ONTAP 9.14.1 and later

The following table lists supported ONTAP storage management features.

Capability/Feature	FC	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	Yes	Yes
Monitoring and host configuration using ONTAP tools for VMware vSphere	Yes	Yes	Yes	Yes
Provisioning using ONTAP tools for VMware vSphere	Yes	Yes	Yes	Yes

The following table lists supported backup features.

Capability/Feature	FC	iSCSI	NVMe-oF	NFS
ONTAP snapshots	Yes	Yes	Yes	Yes
SRM supported by replicated backups	Yes	Yes	No ²	V3 only ²

Capability/Feature	FC	iSCSI	NVMe-oF	NFS
Volume SnapMirror	Yes	Yes	Yes	Yes
VMDK image access	SnapCenter and VADP-enabled backup software	SnapCenter and VADP-enabled backup software	SnapCenter and VADP-enabled backup software	SnapCenter and VADP-enabled backup software, vSphere Client, and vSphere Web Client datastore browser
VMDK file-level access	SnapCenter and VADP-enabled backup software, Windows only	SnapCenter and VADP-enabled backup software, Windows only	SnapCenter and VADP-enabled backup software, Windows only	SnapCenter and VADP-enabled backup software and third-party applications
NDMP granularity	Datastore	Datastore	Datastore	Datastore or VM

¹ **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 require a VMFS 6 datastore that has clustered VMDK support enabled. Other restrictions apply. See VMware's [Setup for Windows Server Failover Clustering](#) documentation for configuration guidelines.

² Datastores using NVMe-oF and NFS v4.1 require vSphere replication. Array-based replication for NFS v4.1 is not currently supported by SRM. Array-based replication with NVMe-oF is not currently supported by the ONTAP tools for VMware vSphere Storage Replication Adapter (SRA).

Selecting a storage protocol

Systems running ONTAP 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. Historically, NetApp testing has generally shown little difference between protocols running at similar line speeds and numbers of connections. However, NVMe-oF (NVMe/TCP and NVMe/FC) shows remarkable gains in IOPS, reduction in latency, and up to 50% or more reduction in host CPU consumption by storage IO. On the other end of the spectrum, NFS provides the greatest flexibility and ease of management, especially for large numbers of VMs. All of these protocols can be used and managed with ONTAP tools for VMware vSphere, which provides a simple interface to create and manage datastores.

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

- **Current operating 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. NFS allows easy sizing up or down, and this resizing can be automated by the storage system. SAN offers space reclamation through guest OS DEALLOCATE/TRIM/UNMAP commands, allowing space from deleted files to be returned to the array. This type of space reclamation is not possible 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.

Datastore layout

ONTAP storage systems offer great flexibility in creating datastores for VMs and virtual disks. Although many ONTAP best practices are applied when using the ONTAP tools to provision datastores 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 datastores results in a high-performing, easy-to-manage implementation that provides VM-to-datastore ratios that cannot be obtained with block-based storage protocols. This architecture can result in a tenfold increase in datastore density with a corresponding reduction in the number of datastores. Although a larger datastore can benefit storage efficiency and provide operational benefits, consider using at least four datastores (FlexVol volumes) per node to store your VMs on a single ONTAP controller to get maximum performance from the hardware resources. This approach also allows you to establish datastores with different recovery policies. Some can be backed up or replicated more frequently than others based on business needs. Multiple datastores are not required with FlexGroup volumes for performance because they scale by design.
- **NetApp recommends** the use of FlexVol volumes for most NFS datastores. Starting with ONTAP 9.8 FlexGroup volumes are supported for use as datastores 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.
- A good size for a FlexVol volume datastore 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 datastore as needed (up to the maximum 300TB). Smaller datastores 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 Datastore Provisioning Wizard uses autosize by default for new datastores. 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 datastores can be configured with LUNs or NVMe namespaces (referred to as storage units in new ASA systems) that are accessed by FC, iSCSI, NVMe/FC, or NVMe/TCP. VMFS allows datastores to be accessed simultaneously by every ESX server in a cluster. VMFS datastores 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 128TB on AFF, ASA, and FAS systems. NetApp always recommends using a single, large LUN for each datastore, rather than trying to use extents. As with NFS, consider using multiple datastores (volumes or storage units) 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. Consider guest-owned file systems such as NFS, SMB, NVMe/TCP or iSCSI file systems managed by the guest. 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. Further, once the VMs are on the NetApp cluster, storage vMotion can use VAAI offloads to move VMs between datastores on the cluster without requiring a host-based copy. Note that NFS does not offload storage vMotion of powered-on VMs; however, VMFS does.
- 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 [Migration of Oracle databases to ONTAP storage systems](#) 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 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, on ASA, AFF, FAS, or even ONTAP Select (a software-defined version of ONTAP running in a VMware or KVM VM). 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.



On image-based vSphere clusters, you will still want to add the NFS Plug-In to your image so they don't go out of compliance when you install it with ONTAP tools.

ONTAP tools is also the management interface for many functions of the VASA Provider for ONTAP, supporting storage policy-based management with vVols.

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 is straightforward and similar to other network configurations. 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 interconnect 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.

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

In vSphere, there are four ways to use block storage devices:

- With VMFS datastores

- With raw device mapping (RDM)
- As an iSCSI-connected LUN or NVMe/TCP-connected namespace accessed and controlled by a software initiator from a VM guest OS
- As a vVols datastore

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.

 Extents are a vSphere storage concept whereby you can "stitch" multiple LUNs together to create a single larger datastore. You should never use extents to reach your desired datastore size. A single LUN is the best practice for a VMFS datastore.

vSphere includes built-in support for multiple paths to storage devices. vSphere can detect the type of storage device for supported storage systems and automatically configures the multipathing stack to support the capabilities of the storage system in use, regardless of the protocol used, or if using ASA, AFF, FAS, or software defined ONTAP.

Both vSphere and ONTAP support Asymmetric Logical Unit Access (ALUA) to establish active/optimized and active/non-optimized paths for Fibre Channel and iSCSI, and Asymmetric Namespace Access (ANA) for NVMe namespaces using NVMe/FC and NVMe/TCP. In ONTAP, an ALUA or ANA-optimized path follows a direct data path, using a target port on the node that hosts the LUN or namespace being accessed. ALUA/ANA is turned on by default in both vSphere and ONTAP. The multipathing software in vSphere recognizes the ONTAP cluster as ALUA or ANA, and it uses the appropriate native plug-in with the round robin load balance policy.

With NetApp's ASA systems, the LUNs and namespaces are presented to the ESXi hosts with symmetric pathing. Meaning that all paths are active and optimized. The multipathing software in vSphere recognizes the ASA system as symmetric, and it uses the appropriate native plug-in with the round robin load balance policy.

 Refer to [Recommended ESXi host and other ONTAP settings](#) for optimized multipathing settings.

ESXi does not see any LUNs, namespaces, or paths beyond its limits. 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.

 Refer to the [VMware Configuration Maximums tool](#) for the most up to date supported limits in ESXi.

SLM limits the nodes that advertise paths to a given LUN. It is a NetApp best practice to have at least two LIFs 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.

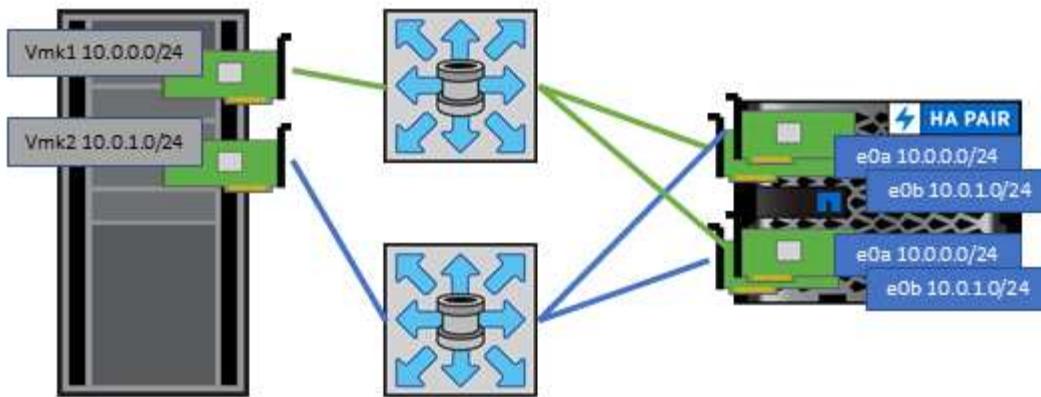
- SLM is enabled by default. Unless you are using portsets, no additional configuration is required.

- 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.

SCSI-based 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) to establish paths based on LUN to igroup mappings filtered by portsets and SLM. NVMe-based block protocols are managed by assigning the namespace with an automatically generated namespace ID to an NVMe subsystem and mapping that subsystem to the NVMe Qualified Name (NQN) of the host(s). Regardless of FC or TCP, NVMe namespaces are mapped using the NQN and not the WWPN or WWNN. The host then creates a software-defined controller for the mapped subsystem to access its namespaces. The path to LUNs and namespaces inside of ONTAP 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.

Other best practices to consider:

- Check [Recommended ESXi host and other ONTAP settings](#) for settings recommended by NetApp in collaboration with VMware.
- 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

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 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.
- 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).
- If you use fpolicy, be sure to exclude .lck files as these are used by vSphere for locking whenever a VM is powered on.
- 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, and update 3 for NFS v4.1. 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](#).

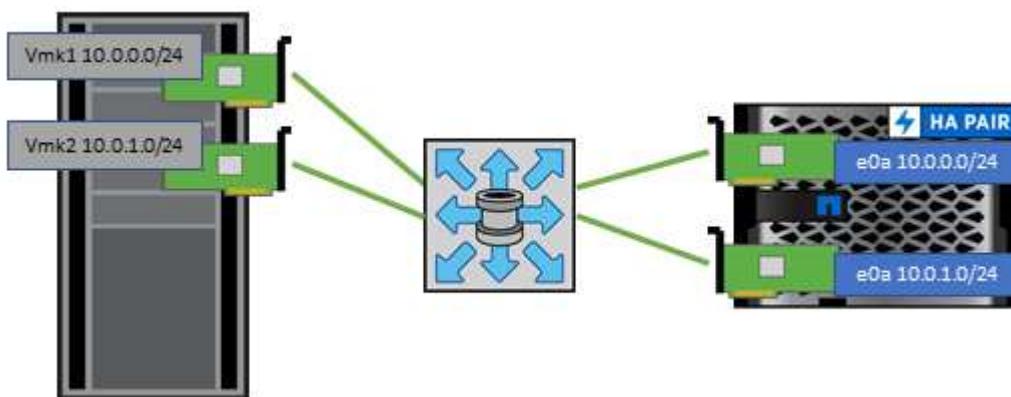
- The maximum value for nconnect in vSphere 8 is 4 and the default value is 1. The maximum value limit in vSphere can be raised on a per-host basis through advanced settings, however it is generally not needed.
- A value of 4 is recommended for environments requiring more performance than a single TCP connection can deliver.
- Be aware that ESXi has a limit of 256 NFS connections and each nconnect connection counts towards that total. E.g. two datastores with nconnect=4 would count as eight total connections.
- It is important to test the performance impact of nconnect on your environment before implementing large scale changes in production environments.



- 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.
 - Refer to the NFS v4.1 Interoperability table notes in the [NetApp Interoperability Matrix Tool](#) for specific ESXi patch levels required for support.
- As mentioned in [settings](#), if you are not using the vSphere CSI for Kubernetes, you should set the newSyncInterval per [VMware KB 386364](#)
- NFS export policy rules are used to control access by vSphere hosts. You can use one policy with multiple volumes (datastores). With NFS, 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 ONTAP tools. The IP addresses should be a comma-separated list without spaces of the vmkernel port addresses mounting the datastores. Here's a sample policy rule:
 - Access Protocol: nfs (which includes both nfs3 and nfs4)
 - List of Client Match Hostnames, IP Addresses, Netgroups, or Domains: 192.168.42.21,192.168.42.22
 - RO Access Rule: any
 - RW Access Rule: any
 - User ID To Which Anonymous Users Are Mapped: 65534
 - Superuser Security Types: any
 - Honor SetUID Bits in SETATTR: true
 - Allow Creation of Devices: true
- 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. This is required even if the datastore type is created as NFS v3.
- 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
- Client Match Spec: 192.168.42.21,192.168.42.22
- RO Access Rule: sys
- RW Access Rule: never (best security for root volume)
- Anonymous UID
- Superuser: sys (also required for root volume with VAAI)
- 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

vSphere Features	NFSv3	NFSv4.1
Kerberos authentication	No	Yes (enhanced with vSphere 6.5 and later to support AES, krb5i)
Multipathing support	No	Yes (ONTAP 9.14.1)

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 9.8-9.13 and SnapCenter plugin for VMware 4.4 and newer 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.

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 FlexGroup cloning mechanism by constantly keeping the clone cache warm.

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 constituent-local cache is then used to rapidly instantiate VM clones on demand.

To enable FlexGroup optimized copy offload, refer to [How to Configure ONTAP FlexGroup volumes 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 FlexGroup volumes 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 for that host. 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 Tool (IMT) 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 FlexGroup volumes. 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 9.8-9.13 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 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 unless your solution guide specifically calls for it; use switches to have redundant paths and allow VMware HA to work without intervention.
- Jumbo frames should be used if supported by your network. 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-interconnect ports within an ONTAP cluster. NetApp makes no other recommendations for best practices regarding flow control 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*

Item	ESXi	Switch	Node	SVM
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, NVMe/FC, iSCSI, NVMe/TCP), RDM

ONTAP offers enterprise-class block storage for VMware vSphere using traditional iSCSI and Fibre Channel Protocol (FCP) as well as the highly efficient and performant next-generation block protocol, NVMe over Fabrics (NVMe-oF), with support for both NVMe/FC and NVMe/TCP.

For detailed best practices for implementing block protocols for VM storage with vSphere and ONTAP refer to [Datastores and Protocols - SAN](#)

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 [datastores](#) section, there are some ease of use and storage efficiency visibility benefits when using NFS with vSphere.

For recommended best practices refer to [Datastores and Protocols - NFS](#)

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. However, there are some limitations and caveats to be aware of.

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 freeze 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.

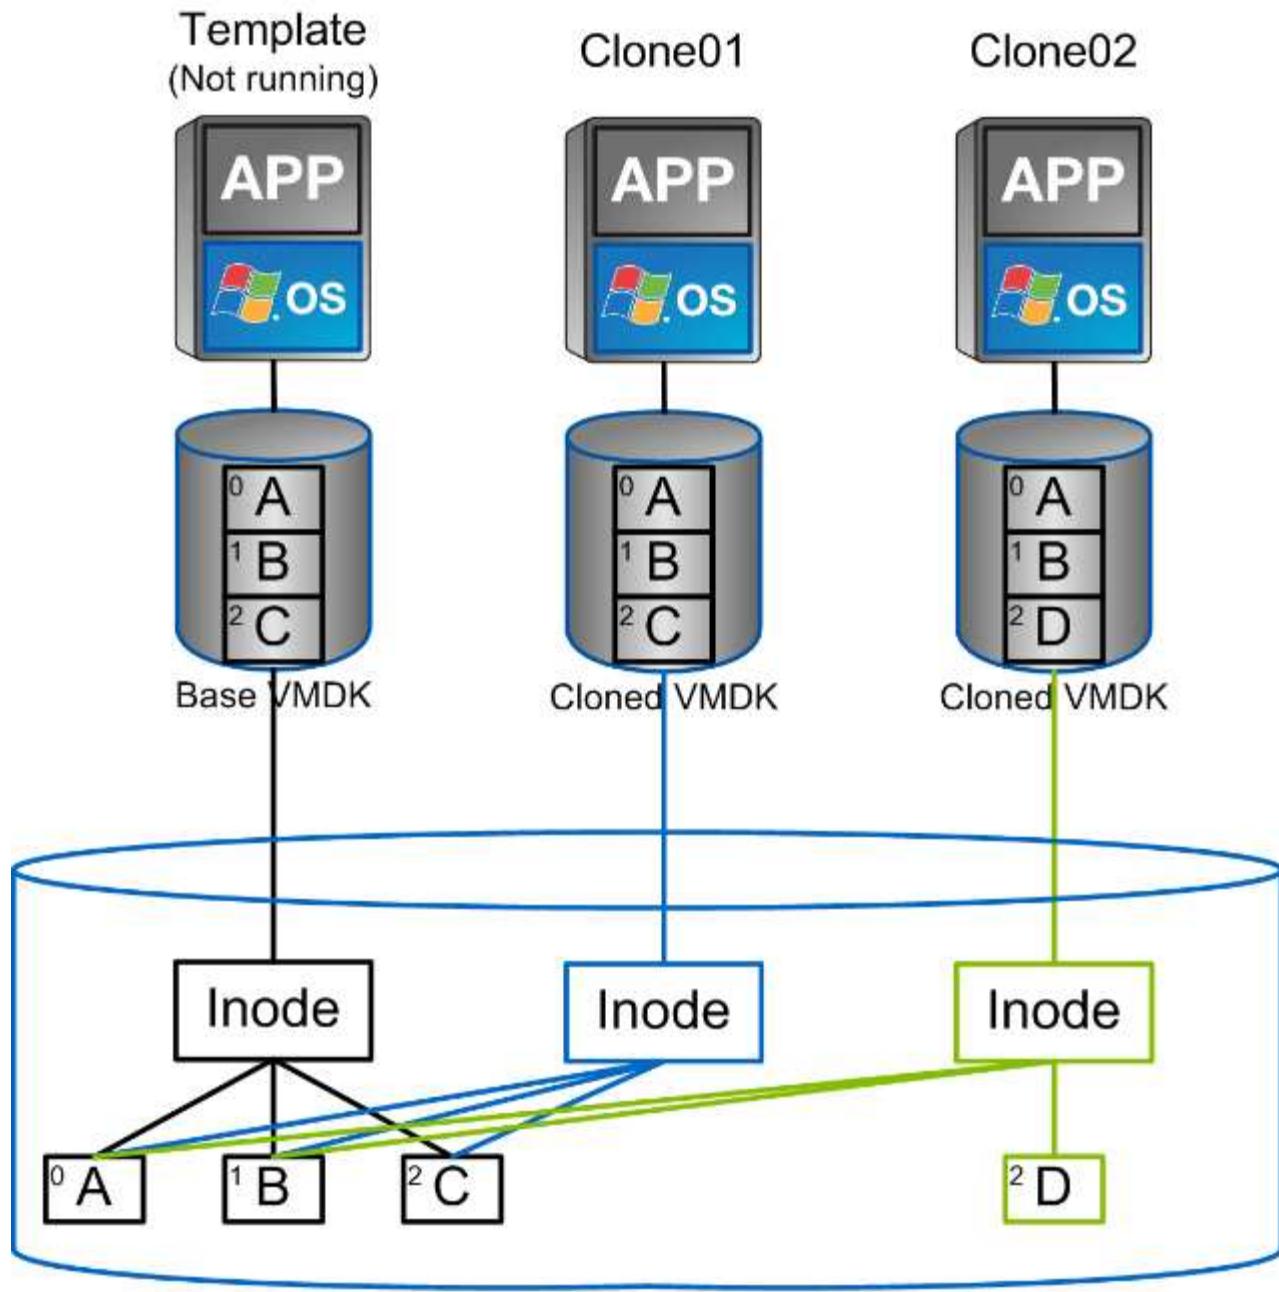
VM and datastore cloning

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 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 and destination. This capability is also used by [Omnissa Horizon View](#).

- SRA (used with VMware Live Site Recovery/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 restored.

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 (included in the ONTAP One license), 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 and quickly recovering your virtual machines (VMs) are key advantages of using ONTAP for vSphere. This functionality can be easily managed within vCenter through the SnapCenter Plug-In for VMware vSphere. Many customers enhance their third-party backup solutions with SnapCenter to leverage ONTAP's snapshot technology, as it offers the fastest and most straightforward way to recover a VM with ONTAP. SnapCenter is available for free to customers who have the ONTAP One license, and other license bundles may also be available.

Additionally, the SnapCenter Plug-In for VMware can integrate with [NetApp Backup and Recovery for virtual machines](#), enabling effective 3-2-1 backup solutions for most ONTAP systems. Note that some fees may apply if using Backup and Recovery for virtual machines with premium services, such as object stores for additional backup storage. This section outlines the various options available for protecting your VMs and datastores.

NetApp ONTAP volume snapshots

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.

Snapshots are a key feature of ONTAP, allowing you to create point-in-time copies of your data. They are

space-efficient and can be created quickly, making them ideal for protecting VMs and datastores. Snapshots can be used for various purposes, including backup, recovery, and testing. These snapshots are different from VMware (consistency) snapshots and are suitable for longer-term protection. VMware's vCenter-managed snapshots are only recommended for short-term use due to performance and other effects. Refer to [Snapshot Limitations](#) for more details.

Snapshots are created at the volume level, and they can be used to protect all the VMs and datastores within that volume. This means that you can create a snapshot of an entire datastore, which includes all the VMs within that datastore.

For NFS datastores, you can easily view VM files in snapshots by browsing the .snapshots directory. This allows you to quickly access and restore files from a snapshot without needing to use a specific backup solution.

For VMFS datastores, you can create a FlexClone of the datastore based on the desired snapshot. This allows you to create a new datastore that is based on the snapshot, which can be used for testing or development purposes. The FlexClone will only consume space for the changes made after the snapshot was taken, making it a space-efficient way to create a copy of the datastore. Once the FlexClone is created, you can map the LUN or namespace to an ESXi host just like a regular datastore. Not only does this allow you to restore specific VM files, but it allows you to quickly create test or development environments based on production data without impacting the performance of the production environment.

For more information on snapshots, refer to the ONTAP documentation. The following links provide additional details:

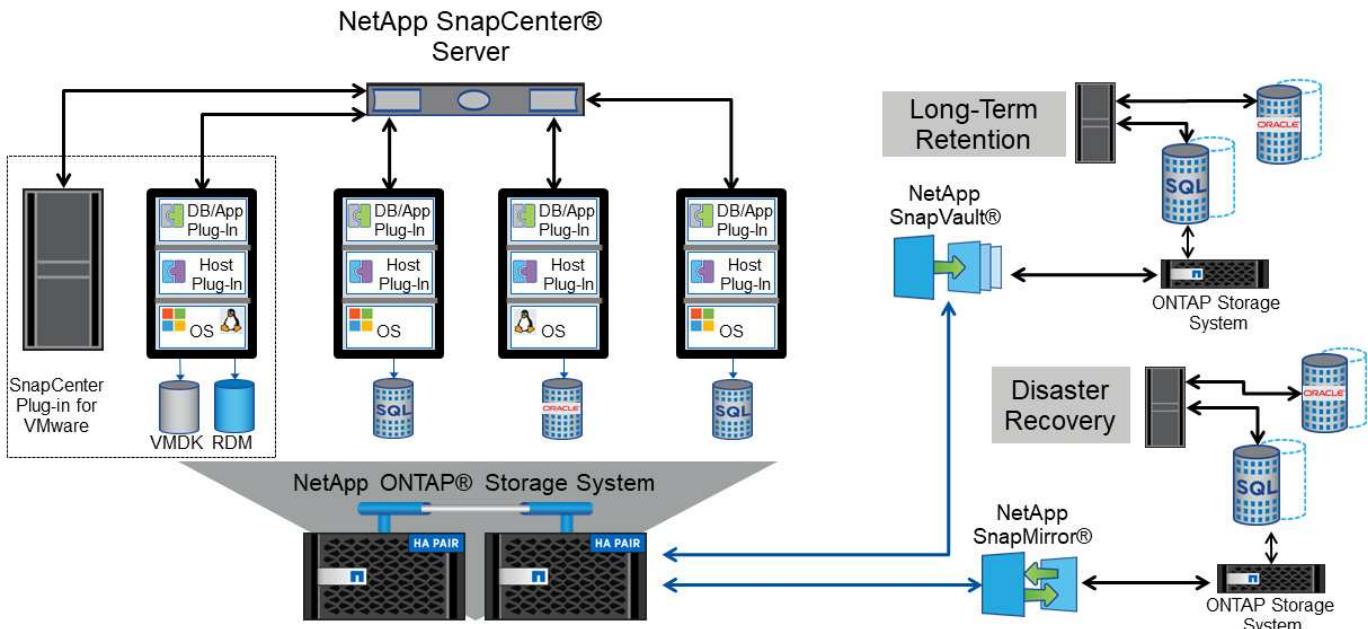
[ONTAP local snapshot copies](#)
[ONTAP SnapMirror replication workflow](#)

SnapCenter Plug-In for VMware vSphere

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 an 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 application plug-ins to protect transactional data such as SQL Server or Oracle.

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, vVols, or NVMe/TCP namespaces and iSCSI LUNs directly connected to the guest OS, or VMDK files on either VMFS or NFS datastores. The plug-ins allow the 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 the cloning of databases for development or test purposes.

The following figure depicts an example of SnapCenter deployment.



For sizing information, refer to the [Sizing Guide for SnapCenter Plugin for VMware vSphere](#)

ONTAP tools for VMware vSphere with VMware Live Site Recovery

The ONTAP tools for VMware vSphere (OT4VS) is a free plug-in that provides a seamless integration between VMware vSphere and NetApp ONTAP. It allows you to manage your ONTAP storage directly from the vSphere Web Client, making it easier to perform tasks such as provisioning storage, managing replication, and monitoring performance.

For improved disaster recovery capabilities, consider utilizing the NetApp SRA for ONTAP, which is part of ONTAP tools for VMware vSphere, alongside VMware Live Site Recovery (formerly known as Site Recovery Manager). This tool not only supports the replication of datastores to a disaster recovery site using SnapMirror, but it also allows for nondisruptive testing in the DR environment by cloning the replicated datastores. Additionally, recovery from a disaster and reprotecting production after resolving an outage is streamlined thanks to the built-in automation features.

NetApp Disaster Recovery

Disaster Recovery (DR) is a cloud-based service that provides a comprehensive solution for protecting your data and applications in the event of a disaster. It offers a range of features, including automated failover and failback, multiple point-in-time recovery points, application-consistent disaster recovery, and support for both on-prem and cloud-based ONTAP systems. NetApp Disaster Recovery is designed to work seamlessly with ONTAP and your VMware vSphere environment, providing a unified solution for disaster recovery.

vSphere Metro Storage Cluster (vMSC) with NetApp MetroCluster and SnapMirror active sync

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, NetApp supported solution that uses synchronous replication, giving the same benefits of a high-availability cluster but distributed across separate sites to protect against site disaster. NetApp SnapMirror active sync, with ASA and AFF, and MetroCluster with AFF, offers cost-effective configurations for synchronous replication with transparent recovery from any single storage component failure as well as transparent recovery in the case of SnapMirror active sync, or single-

command recovery in the event of a site disaster with MetroCluster. vMSC is described in greater detail in [TR-4128](#).

Quality of service (QoS)

Throughput limits are useful in controlling service levels, managing unknown workloads, or to test applications before deployment to make sure they don't affect other workloads in production. They can also be used to constrain a bully workload after it is identified.

ONTAP QoS policy support

Systems running ONTAP can use the 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.

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.

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 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.

Refer to [Performance monitoring and management overview](#) for more information.

Non-vVols NFS datastores

An ONTAP QoS policy can be applied to the entire datastore or individual VMDK files within it. However, it is important to understand that all VMs on a traditional (non-vVols) NFS datastore share a common I/O queue from a given host. If any VM is throttled by an ONTAP QoS policy then this will in practice result in all I/O for that datastore appearing to be throttled for that host.

Example:

- * You configure a QoS limit on vm1.vmdk for a volume that is mounted as a traditional NFS datastore by host esxi-01.
- * The same host (esxi-01) is using vm2.vmdk and it is on the same volume.
- * If vm1.vmdk gets throttled, then vm2.vmdk will also appear to be throttled since it shares the same IO queue

with vm1.vmdk.



This does not apply to vVols.

Beginning in vSphere 6.5 you can manage file-granular limits on non-vVols datastores by leveraging Storage Policy-Based Management (SPBM) with Storage I/O Control (SIOC) v2.

Refer to the following links for more information on managing performance with SIOC and SPBM policies.

[SPBM Host-Based Rules: SIOC v2](#)

[Manage Storage I/O Resources with vSphere](#)

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.

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. Be aware that the previously mentioned limitations of QoS with traditional NFS datastores still apply.

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.

vVols datastores

Minimum and/or maximum QoS can be easily set on individual VMs or VMDKs without impacting any other VM or VMDK using the Storage Policy-Based Management and vVols.

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.



vVols requires the use ONTAP tools for VMware vSphere which functions as the VASA Provider for ONTAP. Refer to [VMware vSphere Virtual Volumes \(vVols\) with ONTAP](#) for vVols best practices.

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. 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)
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	Yes (the datastore can be throttled)
Set QoS on volume (NFS datastore)?	Yes	Yes (the datastore can be throttled)
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.

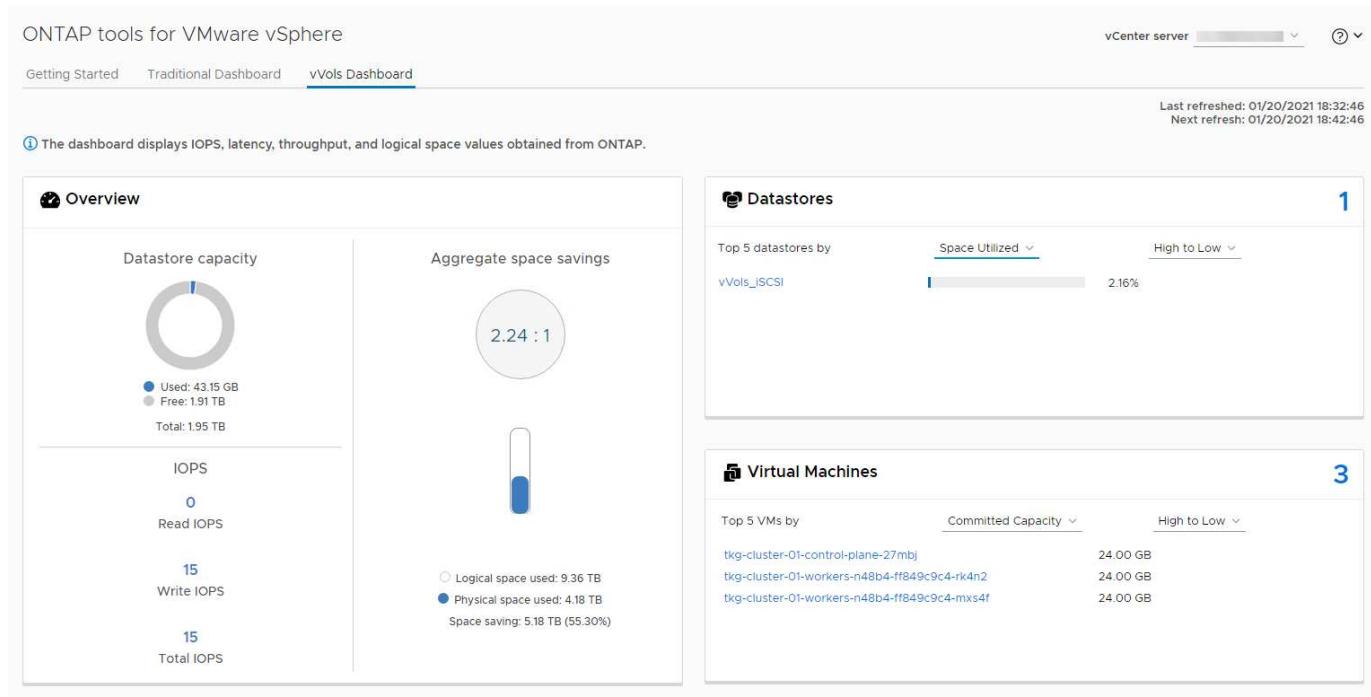
Storage policy based management and vVols

VMware vSphere APIs for Storage Awareness (VASA) make it easy for a storage administrator to configure datastores with well-defined capabilities and enable the VM administrator to 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.

Before 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.



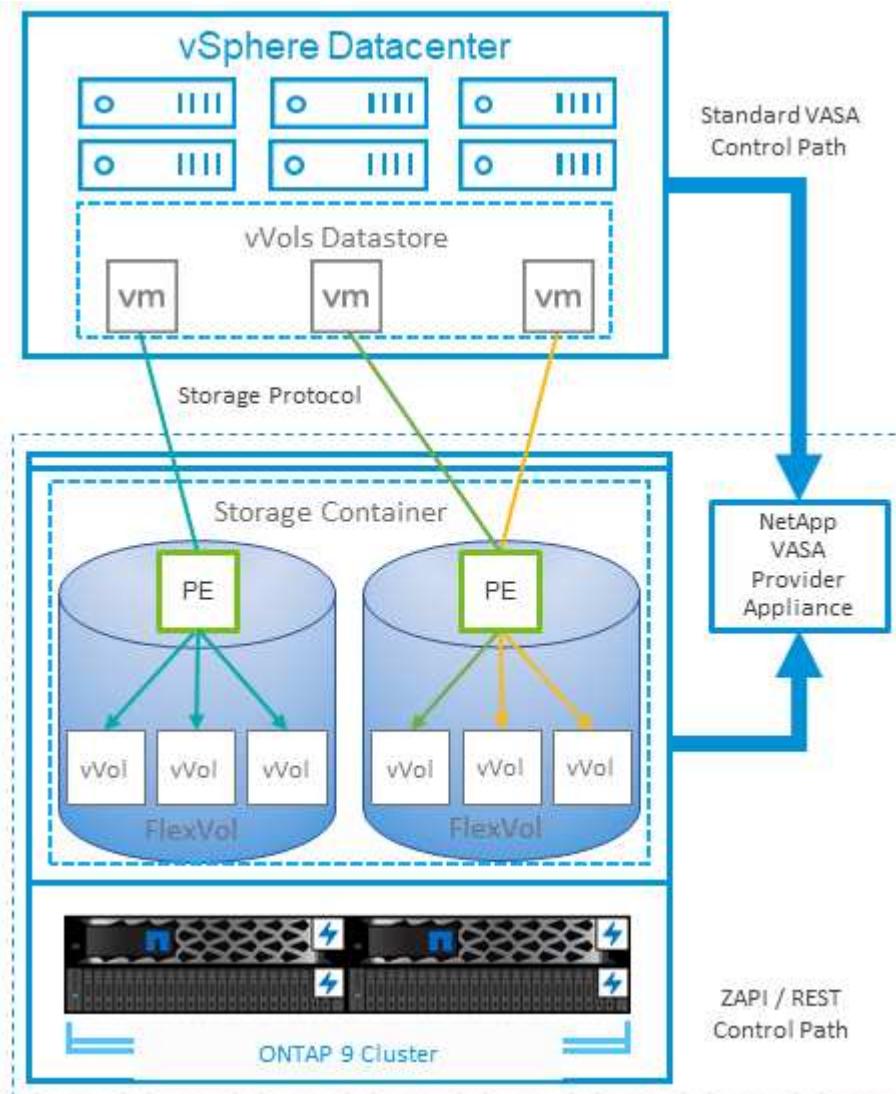
After the storage capability profile is defined, it can be used to provision VMs using the storage policy that identifies its requirements. The mapping between the VM storage policy and the datastore storage capability profile allows vCenter to display a list of compatible datastores for selection. This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

- **First-party offerings.** Amazon FSx for NetApp ONTAP, Google Cloud NetApp Volumes, and Azure NetApp Files provide high-performance, multi-protocol managed storage services in the leading public cloud environments. They can be used directly by VMware Cloud on AWS (VMC on AWS), Azure VMware Solution (AVS), and Google Cloud VMware Engine (GCVE) as datastores or storage for guest operating systems (GOS) and compute instances.
- **Cloud Services.** Use NetApp Backup and Recovery or SnapMirror Cloud to protect data from on-premises systems using public cloud storage. NetApp Copy and Sync helps migrate and keep your data synchronized across NAS, and object stores. NetApp Disaster Recovery provides a cost-effective and efficient solution for leveraging NetApp technologies as the foundation for a robust and capable disaster recovery solution for DR to cloud, DR to on-prem, and on-prem to on-prem.
- **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.

For a deep dive into more security topics, refer to the following resources.

- [ONTAP Select documentation](#)
- [Backup and Recovery documentation](#)
- [Disaster Recovery documentation](#)
- [Amazon FSx for NetApp ONTAP](#)
- [VMware Cloud on AWS](#)
- [What is Azure NetApp Files?](#)
- [Azure VMware Solution](#)
- [Google Cloud VMware Engine](#)
- [What is Google Cloud NetApp Volumes?](#)

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 make it easy to protect any data with at-rest encryption. NetApp Storage Encryption (NSE) uses self-encrypting drives (SEDs) 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.

For a deep dive into more security topics, refer to the following resources.

- [Security technical reports](#)
- [Security hardening guides](#)
- [ONTAP security and data encryption product documentation](#)

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.

Virtual Machines

Name	Status	Power State	Protocol	Capacity (Used Allocated)	A IOPS	VM Latency (ms)	Host IOPS	Host Latency (ms)	Network Latency (ms)	Datastore IOPS	Datastore Latency (ms)
vCenter7	ON	NFS		160 GB 712 GB	183	0	243	0	0	831	0.3

POWER: ON

TOPOLOGY VIEW

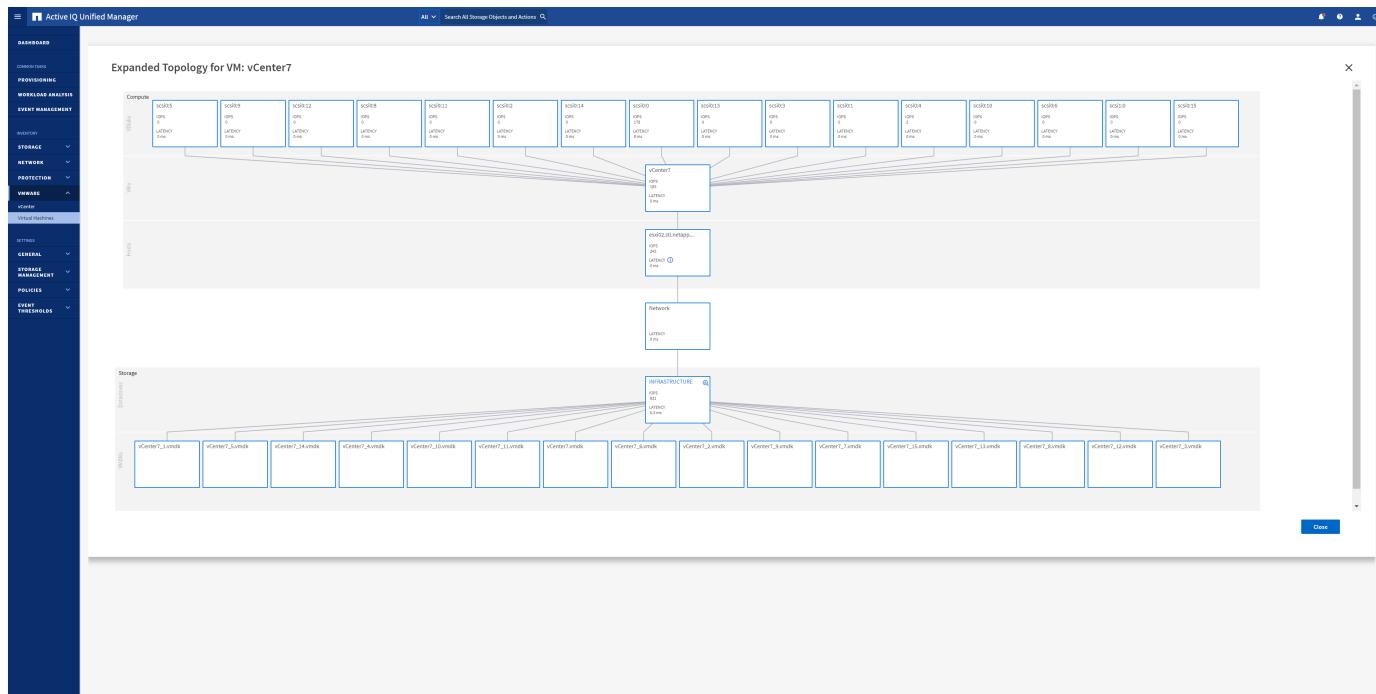
Compute

Storage

Showing all 44 Virtual Machines

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

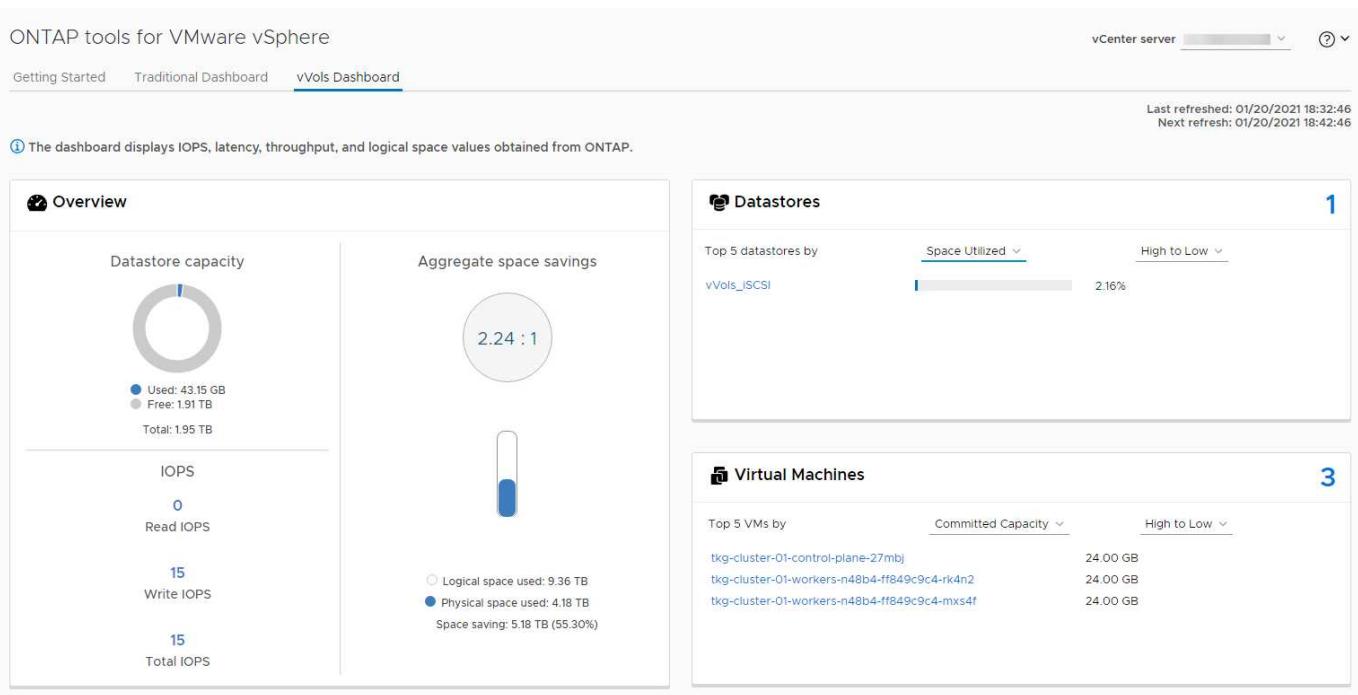
VMware vSphere APIs for Storage Awareness (VASA) make it easy for a storage administrator to configure datastores with well-defined capabilities and enable the VM administrator to use those whenever needed to provision VMs without having to interact with each other.

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.

Before 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.



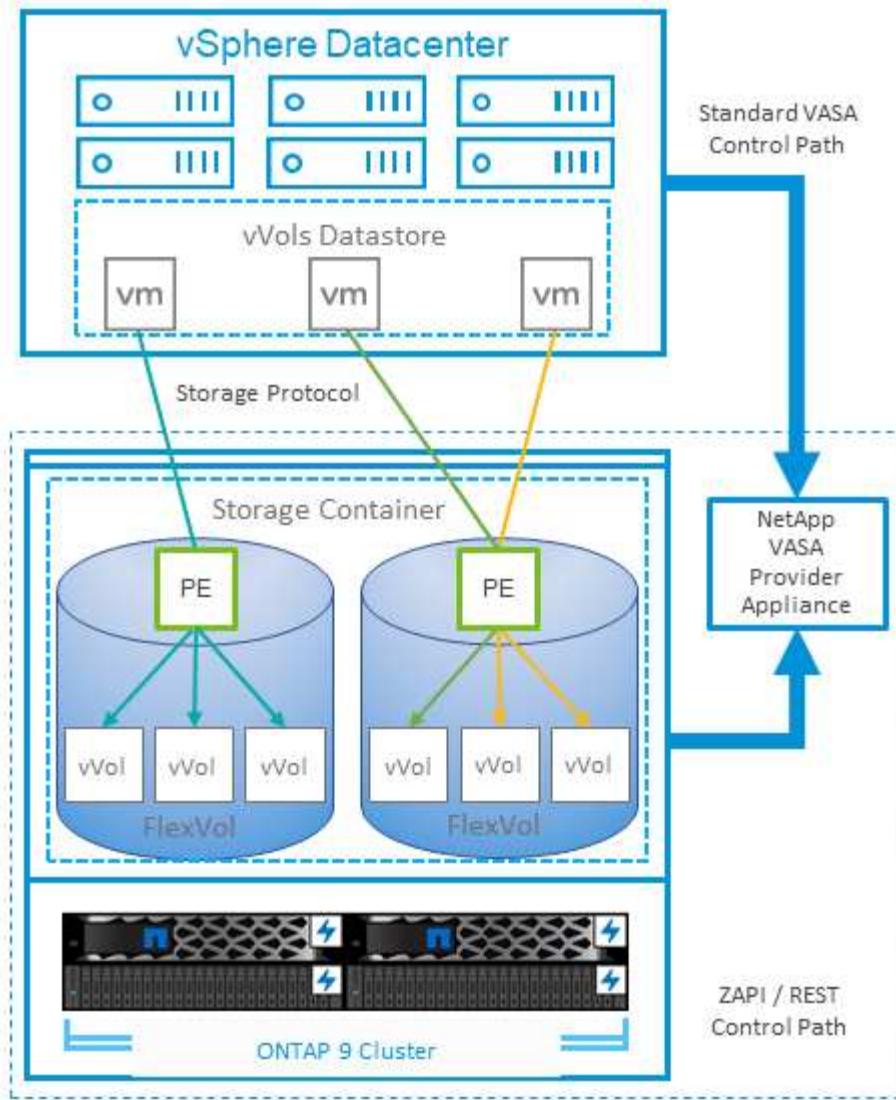
After the storage capability profile is defined, it can be used to provision VMs using the storage policy that identifies its requirements. The mapping between the VM storage policy and the datastore storage capability profile allows vCenter to display a list of compatible datastores for selection. This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



VMware Storage Distributed Resource Scheduler

VMware Storage Distributed Resource Scheduler (SDRS) is a vSphere feature that automatically places VMs in a datastore cluster 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:

- Don't use SDRS unless you have a specific requirement to do so.

- SDRS is not needed when using ONTAP. SDRS is not aware of ONTAP storage efficiency features such as deduplication and compression, so it might make decisions that are not optimal for your environment.
- SDRS is not aware of ONTAP QoS policies, so it might make decisions that are not optimal for performance.
- SDRS is not aware of ONTAP snapshot copies, so it might make decisions that cause snapshots to grow exponentially. For example, moving a VM to another datastore creates new files in the new datastore, which causes the snapshot to grow. This is especially true for VMs with large disks or many snapshots. Then, should the VM be moved back to the original datastore, the snapshot on the original datastore will grow even larger.

If you do use SDRS, consider the following best practices:

- 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 may be reduced depending on how well it deduplicates or compresses on the destination.
 - 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.

More information about SDRS can be found in the VMware documentation at [Storage DRS FAQ](#).

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 ONTAP tools overview page, scroll down to the bottom and click apply recommended Settings in the ESXi Host compliance portlet.

Here are the recommended host settings for all currently supported versions of ONTAP.

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 Space Reclamation for VMFS5 Virtual Machines	No

Host Setting	NetApp Recommended Value	Reboot Required
VMFS3.EnableVMFS6U nmap	Keep default (1) For more information, see VMware vSphere APIs: Array Integration (VAAI)	No
NFS Settings		
newSyncInterval	If you are not using the vSphere CSI for Kubernetes, set per VMware KB 386364	No
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 default (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.MaxConnPerIP	vSphere 7.0 to 8.0, set to 128. This setting is ignored in ESXi releases after 8.0.	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		

Host Setting	NetApp Recommended Value	Reboot Required
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



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. 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:

- When using non-ASA systems 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 latency option is now also available when using the HPP (High Performance Plugin) with NVMe namespaces, and with vSphere 8.0u2 and later, iSCSI and FCP connected LUNs. The new option considers I/O bandwidth and path latency when selecting the optimal path for I/O. NetApp recommends using the latency option in environments with non-equivalent path connectivity, such as cases with more network hops on one path than another, or when using a NetApp ASA system. See [Change Default Parameters for Latency Round Robin](#) 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](#)

Copyright information

Copyright © 2026 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.