



Datastores and protocols

Enterprise applications

NetApp
April 25, 2024

Table of Contents

- Datstores and protocols 1
 - vSphere datastore and protocol features overview 1
 - SAN (FC, FCoE, NVMe/FC, iSCSI), RDM 7
 - NFS 9
 - FlexGroup volumes 11

Datstores and protocols

vSphere datastore and protocol features overview

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

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

FCP, FCoE, NVMe/FC, NVMe/TCP, and iSCSI are block protocols that use the vSphere Virtual Machine File System (VMFS) to store VMs inside ONTAP LUNs or NVMe namespaces that are contained in an ONTAP FlexVol volume. Note that, starting from vSphere 7.0, VMware no longer supports software FCoE in production environments. NFS is a file protocol that places VMs into datstores (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 datstores, but it does generally applies to vSphere 6.x and later releases using supported ONTAP releases. You can also consult [VMware configuration maximums](#) for specific vSphere releases to confirm specific limits.

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

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
Optimal queue depth per LUN or file system	64-256	64-256	Autonegotiated	Refer to NFS.MaxQueueDepth in Recommended ESXi host and other ONTAP settings .

The following table lists supported VMware storage-related functionalities.

Capacity/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
vMotion	Yes	Yes	Yes	Yes
Storage vMotion	Yes	Yes	Yes	Yes
VMware HA	Yes	Yes	Yes	Yes
Storage Distributed Resource Scheduler (SDRS)	Yes	Yes	Yes	Yes
VMware vStorage APIs for Data Protection (VADP)–enabled backup software	Yes	Yes	Yes	Yes
Microsoft Cluster Service (MSCS) or failover clustering within a VM	Yes	Yes*	Yes*	Not supported
Fault Tolerance	Yes	Yes	Yes	Yes
Site Recovery Manager	Yes	Yes	No**	V3 only**
Thin-provisioned VMs (virtual disks)	Yes	Yes	Yes	Yes This setting is the default for all VMs on NFS when not using VAAI.
VMware native multipathing	Yes	Yes	Yes, using the new High Performance Plugin (HPP)	NFS v4.1 session trunking requires ONTAP 9.14.1 and later

The following table lists supported ONTAP storage management features.

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
Data deduplication	Savings in the array	Savings in the array	Savings in the array	Savings in the datastore
Thin provisioning	Datastore or RDM	Datastore or RDM	Datastore	Datastore

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

The following table lists supported backup features.

Capability/Feature	FC/FCoE	iSCSI	NVMe-oF	NFS
ONTAP Snapshots	Yes	Yes	Yes	Yes
SRM supported by replicated backups	Yes	Yes	No**	V3 only**
Volume SnapMirror	Yes	Yes	Yes	Yes
VMDK image access	VADP-enabled backup software	VADP-enabled backup software	VADP-enabled backup software	VADP-enabled backup software, vSphere Client, and vSphere Web Client datastore browser
VMDK file-level access	VADP-enabled backup software, Windows only	VADP-enabled backup software, Windows only	VADP-enabled backup software, Windows only	VADP-enabled backup software and third-party applications
NDMP granularity	Datastore	Datastore	Datastore	Datastore or VM

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

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

Selecting a storage protocol

Systems running ONTAP software support all major storage protocols, so customers can choose what is best

for their environment, depending on existing and planned networking infrastructure and staff skills. NetApp testing has generally shown little difference between protocols running at similar line speeds, so it is best to focus on your network infrastructure and staff capabilities over raw protocol performance.

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

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

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 VSC 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 correlating 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) 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. That being said, deploying datastores as multiple qtrees in a single volume might be useful for highly automated environments that can benefit from datastore-level quotas or VM file clones.
- 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 100TB). 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 use 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 that are accessed by FC, iSCSI, or FCoE. VMFS allows traditional LUNs 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 16TB on most systems, and 128TB on All-SAN-Array systems. Therefore, a maximum size VMFS 5 datastore on most ONTAP systems can be created by using four 16TB LUNs. While there can be a performance benefit for high-I/O workloads with multiple LUNs (with high-end FAS or AFF systems), this benefit is offset by added management complexity to create, manage, and protect the datastore LUNs and increased availability risk. NetApp generally recommends using a single, large LUN for each datastore and only span if there is a special need to go beyond a 16TB datastore. As with NFS, consider using multiple datastores (volumes) to maximize performance on a single ONTAP controller.
- Older guest operating systems (OSs) needed alignment with the storage system for best performance and storage efficiency. However, modern vendor-supported OSs from Microsoft and Linux distributors such as Red Hat no longer require adjustments to align the file system partition with the blocks of the underlying storage system in a virtual environment. If you are using an old OS that might require alignment, search the NetApp Support Knowledgebase for articles using “VM alignment” or request a copy of TR-3747 from a NetApp sales or partner contact.
- Avoid the use of defragmentation utilities within the guest OS, as this offers no performance benefit and affects storage efficiency and snapshot space usage. Also consider turning off search indexing in the guest OS for virtual desktops.
- ONTAP has led the industry with innovative storage efficiency features, allowing you to get the most out of your usable disk space. AFF systems take this efficiency further with default inline deduplication and compression. Data is deduplicated across all volumes in an aggregate, so you no longer need to group similar operating systems and similar applications within a single datastore to maximize savings.
- In some cases, you might not even need a datastore. For the best performance and manageability, avoid using a datastore for high-I/O applications such as databases and some applications. Instead, consider guest-owned file systems such as NFS or iSCSI file systems managed by the guest or with RDMS. For specific application guidance, see NetApp technical reports for your application. For example, [Oracle Databases on ONTAP](#) has a section about virtualization with helpful details.
- First Class Disks (or Improved Virtual Disks) allow for vCenter-managed disks independent of a VM with vSphere 6.5 and later. While primarily managed by API, they can be useful with vVols, especially when managed by OpenStack or Kubernetes tools. They are supported by ONTAP as well as ONTAP tools for VMware vSphere.

Datastore and VM migration

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

- Use Storage vMotion to move the bulk of your virtual machines to ONTAP. Not only is this approach nondisruptive to running VMs, it also allows ONTAP storage efficiency features such as inline deduplication and compression to process the data as it migrates. Consider using vCenter capabilities to select multiple VMs from the inventory list and then schedule the migration (use Ctrl key while clicking Actions) at an appropriate time.
- While you could carefully plan a migration to appropriate destination datastores, it is often simpler to

migrate in bulk and then organize later as needed. You might want to use this approach to guide your migration to different datastores if you have specific data protection needs, such as different Snapshot schedules.

- Most VMs and their storage may be migrated while running (hot), but migrating attached (not in datastore) storage such as ISOs, LUNs, or NFS volumes from another storage system might require cold migration.
- Virtual machines that need more careful migration include databases and applications that use attached storage. In general, consider the use of the application's tools to manage migration. For Oracle, consider using Oracle tools such as RMAN or ASM to migrate the database files. See [TR-4534](#) for more information. Likewise, for SQL Server, consider using either SQL Server Management Studio or NetApp tools such as SnapManager for SQL Server or SnapCenter.

ONTAP tools for VMware vSphere

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

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

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

In general, NetApp recommends using the ONTAP tools for VMware vSphere interface within vCenter to provision traditional and vVols datastores to make sure best practices are followed.

General Networking

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

- Separate storage network traffic from other networks. A separate network can be achieved by using a dedicated VLAN or separate switches for storage. If the storage network shares physical paths such as uplinks, you might need QoS or additional uplink ports to make sure of sufficient bandwidth. Don't connect hosts directly to storage; use switches to have redundant paths and allow VMware HA to work without intervention. See [Direct connect networking](#) for additional information.
- Jumbo frames can be used if desired and supported by your network, especially when using iSCSI. If they are used, make sure they are configured identically on all network devices, VLANs, and so on in the path between storage and the ESXi host. Otherwise, you might see performance or connection problems. The MTU must also be set identically on the ESXi virtual switch, the VMkernel port, and also on the physical ports or interface groups of each ONTAP node.
- NetApp only recommends disabling network flow control on the cluster network ports within an ONTAP cluster. NetApp makes no other recommendations for best practices for the remaining network ports used for data traffic. You should enable or disable as necessary. See [TR-4182](#) for more background on flow control.
- When ESXi and ONTAP storage arrays are connected to Ethernet storage networks, NetApp recommends

configuring the Ethernet ports to which these systems connect as Rapid Spanning Tree Protocol (RSTP) edge ports or by using the Cisco PortFast feature. NetApp recommends enabling the Spanning-Tree PortFast trunk feature in environments that use the Cisco PortFast feature and that have 802.1Q VLAN trunking enabled to either the ESXi server or the ONTAP storage arrays.

- NetApp recommends the following best practices for link aggregation:
 - Use switches that support link aggregation of ports on two separate switch chassis using a multi-chassis link aggregation group approach such as Cisco's Virtual PortChannel (vPC).
 - Disable LACP for switch ports connected to ESXi unless you are using dvSwitches 5.1 or later with LACP configured.
 - Use LACP to create link aggregates for ONTAP storage systems with dynamic multimode interface groups with port or IP hash. Refer to [Network Management](#) for further guidance.
 - Use an IP hash teaming policy on ESXi when using static link aggregation (e.g., EtherChannel) and standard vSwitches, or LACP-based link aggregation with vSphere Distributed Switches. If link aggregation is not used, then use "Route based on the originating virtual port ID" instead.

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

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

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

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

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

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

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

- With VMFS datastores

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

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

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

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

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

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

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


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

Other best practices to consider:

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

teaming when multiple virtual switches are present. You can also use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput. The following figure provides an example of multipath connectivity. In ONTAP, configure either a single-mode interface group for failover with two or more links that are connected to two or more switches, or use LACP or other link-aggregation technology with multimode interface groups to provide HA and the benefits of link aggregation.

- If the Challenge-Handshake Authentication Protocol (CHAP) is used in ESXi for target authentication, it must also be configured in ONTAP using the CLI (`vserver iscsi security create`) or with System Manager (edit Initiator Security under Storage > SVMs > SVM Settings > Protocols > iSCSI).
- Use ONTAP tools for VMware vSphere to create and manage LUNs and igroups. The plug-in automatically determines the WWPNs of servers and creates appropriate igroups. It also configures LUNs according to best practices and maps them to the correct igroups.
- Use RDMs with care because they can be more difficult to manage, and they also use paths, which are limited as described earlier. ONTAP LUNs support both [physical and virtual compatibility mode](#) RDMs.
- For more on using NVMe/FC with vSphere 7.0, see this [ONTAP NVMe/FC Host Configuration guide](#) and [TR-4684](#). The following figure depicts multipath connectivity from a vSphere host to an ONTAP LUN.



NFS

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

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


- Use a single logical interface (LIF) for each SVM on each node in the ONTAP cluster. Past recommendations of a LIF per datastore are no longer necessary. While direct access (LIF and datastore on same node) is best, don't worry about indirect access because the performance effect is generally minimal (microseconds).
- VMware has supported NFSv3 since VMware Infrastructure 3. vSphere 6.0 added support for NFSv4.1, which enables some advanced capabilities such as Kerberos security. Where NFSv3 uses client-side locking, NFSv4.1 uses server-side locking. Although an ONTAP volume can be exported through both protocols, ESXi can only mount through one protocol. This single protocol mount does not preclude other ESXi hosts from mounting the same datastore through a different version. Make sure to specify the protocol version to use when mounting so that all hosts use the same version and, therefore, the same

locking style. Do not mix NFS versions across hosts. If possible, use host profiles to check compliancy.

- Because there is no automatic datastore conversion between NFSv3 and NFSv4.1, create a new NFSv4.1 datastore and use Storage vMotion to migrate VMs to the new datastore.
- Please refer to the NFS v4.1 Interoperability table notes in the [NetApp Interoperability Matrix tool](#) for specific ESXi patch levels required for support.
- VMware supports nconnect with NFSv3 beginning in vSphere 8.0U2. More information on nconnect can be found at the [NFSv3 nConnect feature with NetApp and VMware](#)
- NFS export policies are used to control access by vSphere hosts. You can use one policy with multiple volumes (datastores). With NFSv3, ESXi uses the sys (UNIX) security style and requires the root mount option to execute VMs. In ONTAP, this option is referred to as superuser, and when the superuser option is used, it is not necessary to specify the anonymous user ID. Note that export policy rules with different values for `-anon` and `-allow-suid` can cause SVM discovery problems with the ONTAP tools. Here's a sample policy:
 - Access Protocol: nfs (which includes both nfs3 and nfs4)
 - Client Match Spec: 192.168.42.21
 - RO Access Rule: sys
 - RW Access Rule: sys
 - Anonymous UID
 - Superuser: sys
- If the NetApp NFS Plug-In for VMware VAAI is used, the protocol should be set as `nfs` instead of `nfs3` when the export policy rule is created or modified. The VAAI copy offload feature requires the NFSv4 protocol to function, even if the data protocol is NFSv3. Specifying the protocol as `nfs` includes both the NFSv3 and NFSv4 versions.
- NFS datastore volumes are junctioned from the root volume of the SVM; therefore, ESXi must also have access to the root volume to navigate and mount datastore volumes. The export policy for the root volume, and for any other volumes in which the datastore volume's junction is nested, must include a rule or rules for the ESXi servers granting them read-only access. Here's a sample policy for the root volume, also using the VAAI plug-in:
 - Access Protocol: nfs (which includes both nfs3 and nfs4)
 - Client Match Spec: 192.168.42.21
 - RO Access Rule: sys
 - RW Access Rule: never (best security for root volume)
 - Anonymous UID
 - Superuser: sys (also required for root volume with VAAI)
- Use ONTAP tools for VMware vSphere (the most important best practice):
 - Use ONTAP tools for VMware vSphere to provision datastores because it simplifies management of export policies automatically.
 - When creating datastores for VMware clusters with the plug-in, select the cluster rather than a single ESX server. This choice triggers it to automatically mount the datastore to all hosts in the cluster.
 - Use the plug-in mount function to apply existing datastores to new servers.
 - When not using ONTAP tools for VMware vSphere, use a single export policy for all servers or for each cluster of servers where additional access control is needed.
- Although ONTAP offers a flexible volume namespace structure to arrange volumes in a tree using

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

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



The following table lists NFS versions and supported features.

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

FlexGroup volumes

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

ONTAP 9.8, along with the ONTAP tools for VMware vSphere 9.8 and SnapCenter plugin for VMware 4.4

releases added support for FlexGroup volume-backed datastores in vSphere. FlexGroup volumes simplify the creation of large datastores and automatically create the necessary distributed constituent volumes across the ONTAP cluster to get the maximum performance from an ONTAP system.

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

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

Copy offload

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

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

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

Consider the following scenario:

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

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

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

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

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

QoS settings

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

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

in the datastore.

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

Metrics

ONTAP 9.8 also added new file-based performance metrics (IOPS, throughput, and latency) for FlexGroup files, and these metrics can be viewed in the ONTAP tools for VMware vSphere dashboard and VM reports. The ONTAP tools for VMware vSphere plug-in also allows you to set Quality of Service (QoS) rules using a combination of maximum and/or minimum IOPS. These can be set across all VMs in a datastore or individually for specific VMs.

Best practices

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

Copyright information

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

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

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

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

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

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

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

Trademark information

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