



ONTAP concepts

ONTAP 9

NetApp
September 18, 2024

Table of Contents

- ONTAP concepts 1
 - ONTAP platforms 1
 - ONTAP user interfaces 1
 - Cluster storage 2
 - High-availability pairs 3
 - AutoSupport and Active IQ Digital Advisor 4
 - Network architecture 5
 - Client protocols 8
 - Disks and aggregates 10
 - Volumes, qtrees, files, and LUNs 16
 - Storage virtualization 17
 - Path failover 21
 - Load balancing 23
 - Replication 26
 - Storage efficiency 33
 - Security 43
 - ONTAP and VMware vSphere 48
 - Application aware data management 50
 - FabricPool 51

ONTAP concepts

ONTAP platforms

ONTAP data management software offers unified storage for applications that read and write block or file data. Options in storage configurations range from high-speed flash to lower-priced spinning media to cloud-based object storage.

ONTAP implementations run on the following:

- **NetApp-engineered systems:** [FAS hybrid-flash systems](#), [All-Flash FAS \(AFF\) A-Series and C-Series](#), and [All-Flash SAN Array \(ASA\) platforms](#)
- **Commodity hardware:** [ONTAP Select](#)
- **Private, public, or hybrid clouds:** [Cloud Volumes ONTAP](#)
- **Specialized implementations**, including [FlexPod Datacenter](#), which offers best-in-class converged infrastructure

Together these implementations form the basic framework of the *NetApp data fabric*, with a common software-defined approach to data management and fast, efficient replication across platforms.

ONTAP user interfaces

ONTAP data management software offers multiple interfaces you can use to manage your ONTAP clusters. These interface options provide different levels of access and functionality and give you the flexibility to manage your ONTAP clusters as appropriate based on your environment.

You can use any of these interfaces to administer your ONTAP clusters and perform data management operations

ONTAP System Manager

ONTAP System Manager is a web-based user interface that provides a simplified and intuitive way to manage your cluster. You can administer common operations such as storage configuration, data protection, and network setup and management. System Manager also provides risk and cluster performance monitoring and insights to help you react to cluster issues and get ahead of issues before they occur. [Learn more](#).

The ONTAP System Manager interface was redesigned in ONTAP 9.7 to provide a more streamlined and intuitive user experience. The procedures in this documentation site describe the redesigned System Manager interface, which continues to offer more options with each ONTAP release.



The older System Manager interface is referred to in ONTAP documentation as System Manager Classic. The last ONTAP release in which the System Manager Classic interface is available is ONTAP 9.7.

BlueXP

Beginning with ONTAP 9.12.1, you can use the BlueXP web-based interface to manage your hybrid multicloud infrastructure from a single control plane while retaining the familiar System Manager dashboard. BlueXP enables you to create and administer cloud storage (for example, Cloud Volumes ONTAP), use NetApp's data services (for example, Cloud Backup), and control many on-premise and edge storage devices. Adding on-premises ONTAP systems to BlueXP enables you to manage all your storage and data assets from a single interface. [Learn more](#).

ONTAP command line interface

The [ONTAP command line interface \(CLI\)](#) is a text-based interface that allows you to interact with a cluster, node, SVM, and more using [commands](#). CLI commands are available based on [role type](#). You can access the ONTAP CLI through SSH or a console connection to a node in the cluster.

ONTAP REST API

Beginning with ONTAP 9.6, you can access a RESTful API that allows you to programmatically manage and automate cluster operations. Use the API to perform various ONTAP administrative tasks, such as creating and managing volumes, snapshots, and aggregates, as well as monitoring cluster performance. You can access the ONTAP REST API directly using a utility such as curl or with any programming language that supports a REST client, such as Python, PowerShell, and Java. [Learn more](#).



ONTAPI is a proprietary ONTAP API that precedes the ONTAP REST API. The ONTAPI interface will be disabled in future versions of ONTAP. If you are using ONTAPI, you should plan your [migration to the ONTAP REST API](#).

NetApp toolkits and frameworks

NetApp provides client toolkits for specific development languages and environments that abstract the ONTAP REST API and make it easier to create automation code. [Learn more](#).

In addition to these toolkits, you can create and deploy automation code using frameworks. [Learn more](#).

Cluster storage

The current iteration of ONTAP was originally developed for NetApp's scale out *cluster* storage architecture. This is the architecture you typically find in datacenter implementations of ONTAP. Because this implementation exercises most of ONTAP's capabilities, it's a good place to start in understanding the concepts that inform ONTAP technology.


Datacenter architectures usually deploy dedicated FAS or AFF controllers running ONTAP data management software. Each controller, its storage, its network connectivity, and the instance of ONTAP running on the controller is called a *node*.

Nodes are paired for high availability (HA). Together these pairs (up to 12 nodes for SAN, up to 24 nodes for NAS) comprise the cluster. Nodes communicate with each other over a private, dedicated cluster interconnect.

Depending on the controller model, node storage consists of flash disks, capacity drives, or both. Network ports on the controller provide access to data. Physical storage and network connectivity resources are virtualized, visible to cluster administrators only, not to NAS clients or SAN hosts.

Nodes in an HA pair must use the same storage array model. Otherwise you can use any supported combination of controllers. You can scale out for capacity by adding nodes with like storage array models, or for performance by adding nodes with higher-end storage arrays.

Of course you can scale up in all the traditional ways as well, upgrading disks or controllers as needed. ONTAP's virtualized storage infrastructure makes it easy to move data nondisruptively, meaning that you can scale vertically or horizontally without downtime.



You can scale out for capacity by adding nodes with like controller models, or for performance by adding nodes with higher-end storage arrays, all while clients and hosts continue to access data.

High-availability pairs


Cluster nodes are configured in *high-availability (HA) pairs* for fault tolerance and nondisruptive operations. If a node fails or if you need to bring a node down for routine maintenance, its partner can *take over* its storage and continue to serve data from it. The partner *gives back* storage when the node is brought back on line.

HA pairs always consist of like controller models. The controllers typically reside in the same chassis with redundant power supplies.

The HA pairs are fault tolerant nodes that can communicate with each other in different ways to allow each node to continually check whether its partner is functioning and to mirror log data for the other's nonvolatile memory. When a write request is made to a node, it is logged in NVRAM on both nodes before a response is sent back to the client or host. On failover, the surviving partner commits the failed node's uncommitted write requests to disk, ensuring data consistency.

Connections to the other controller's storage media allow each node to access the other's storage in the event of a takeover. Network path failover mechanisms ensure that clients and hosts continue to communicate with the surviving node.

To assure availability, you should keep performance capacity utilization on either node at 50% to accommodate the additional workload in the failover case. For the same reason, you may want to configure no more than 50% of the maximum number of NAS virtual network interfaces for a node.



On failover, the surviving partner commits the failed node's uncommitted write requests to disk, ensuring data consistency.

Takeover and giveback in virtualized ONTAP implementations

Storage is not shared between nodes in virtualized “shared-nothing” ONTAP implementations like Cloud Volumes ONTAP for AWS or ONTAP Select. When a node goes down, its partner continues to serve data from a synchronously mirrored copy of the node's data. It does not take over the node's storage, only its data serving function.

AutoSupport and Active IQ Digital Advisor

ONTAP offers artificial intelligence-driven system monitoring and reporting through a web portal and through a mobile app. The AutoSupport component of ONTAP sends telemetry that is analyzed by Active IQ Digital Advisor.

Active IQ enables you to optimize your data infrastructure across your global hybrid cloud by delivering actionable predictive analytics and proactive support through a cloud-based portal and mobile app. Data-driven insights and recommendations from Active IQ are available to all NetApp customers with an active SupportEdge contract (features vary by product and support tier).

Here are some things you can do with Active IQ:

- Plan upgrades. Active IQ identifies issues in your environment that can be resolved by upgrading to a

newer version of ONTAP and the Upgrade Advisor component helps you plan for a successful upgrade.

- View system wellness. Your Active IQ dashboard reports any issues with wellness and helps you correct those issues. Monitor system capacity to make sure you never run out of storage space.
- Manage performance. Active IQ shows system performance over a longer period than you can see in System Manager. Identify configuration and system issues that are impacting your performance.
- Maximize efficiency. View storage efficiency metrics and identify ways to store more data in less space.
- View inventory and configuration. Active IQ displays complete inventory and software and hardware configuration information. See when service contracts are expiring to ensure you remain covered.

Related information

[NetApp Documentation: Active IQ Digital Advisor](#)

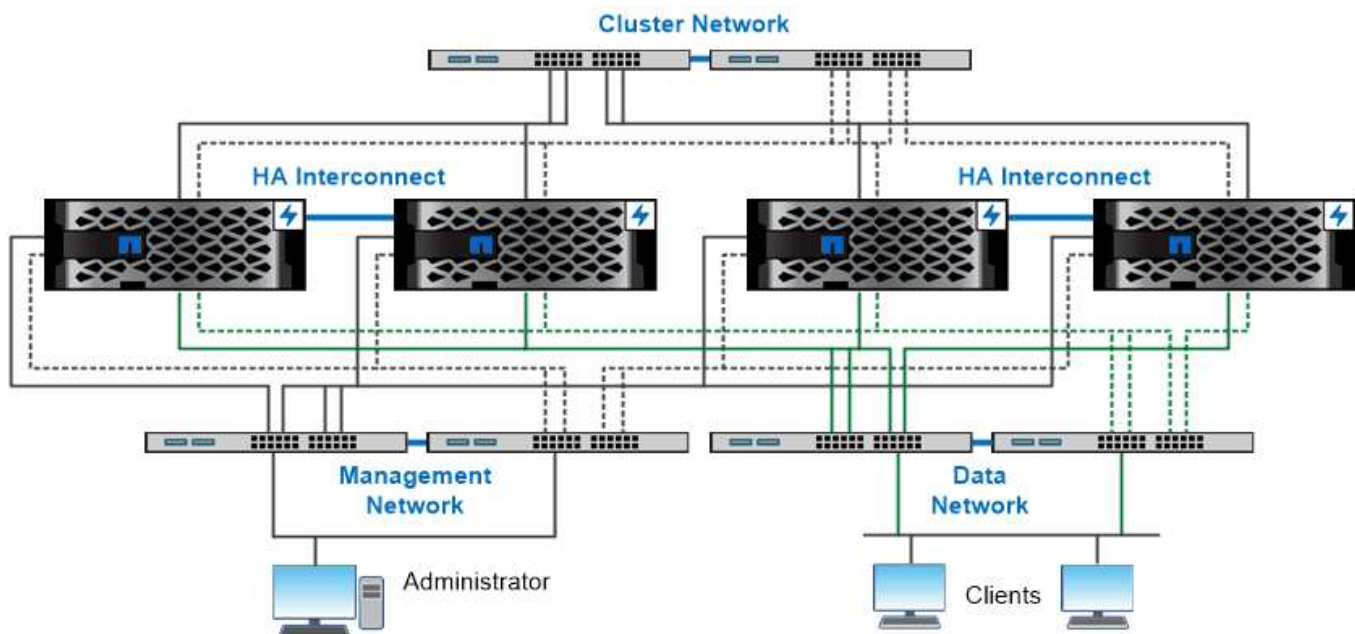
[Launch Active IQ](#)

[SupportEdge Services](#)

Network architecture

Network architecture overview

The network architecture for an ONTAP datacenter implementation typically consists of a cluster interconnect, a management network for cluster administration, and a data network. NICs (network interface cards) provide physical ports for Ethernet connections. HBAs (host bus adapters) provide physical ports for FC connections.



The network architecture for an ONTAP datacenter implementation typically consists of a cluster interconnect, a management network for cluster administration, and a data network.

Logical ports

In addition to the physical ports provided on each node, you can use *logical ports* to manage network traffic. Logical ports are interface groups or VLANs.

Interface groups


Interface groups combine multiple physical ports into a single logical “trunk port”. You might want to create an interface group consisting of ports from NICs in different PCI slots to ensure against a slot failure bringing down business-critical traffic.

An interface group can be single-mode, multimode, or dynamic multimode. Each mode offers differing levels of fault tolerance. You can use either type of multimode interface group to load-balance network traffic.

VLANs

VLANs separate traffic from a network port (which could be an interface group) into logical segments defined on a switch port basis, rather than on physical boundaries. The *end-stations* belonging to a VLAN are related by function or application.

You might group end-stations by department, such as Engineering and Marketing, or by project, such as release1 and release2. Because physical proximity of the end-stations is irrelevant in a VLAN, the end-stations can be geographically remote.



You can use VLANs to segregate traffic by department.

Support for industry-standard network technologies

ONTAP supports all major industry-standard network technologies. Key technologies include IPspaces, DNS load balancing, and SNMP traps.

Broadcast domains, failover groups, and subnets are described in [NAS path failover](#).

IPspaces

You can use an *IPspace* to create a distinct IP address space for each virtual data server in a cluster. Doing so enables clients in administratively separate network domains to access cluster data while using overlapping IP addresses from the same IP address subnet range.

A service provider, for example, could configure different IPspaces for tenants using the same IP addresses to access a cluster.

DNS load balancing

You can use *DNS load balancing* to distribute user network traffic across available ports. A DNS server dynamically selects a network interface for traffic based on the number of clients that are mounted on the interface.

SNMP traps

You can use *SNMP traps* to check periodically for operational thresholds or failures. SNMP traps capture system monitoring information sent asynchronously from an SNMP agent to an SNMP manager.

FIPS compliance

ONTAP is compliant with the Federal Information Processing Standards (FIPS) 140-2 for all SSL connections. You can turn on and off SSL FIPS mode, set SSL protocols globally, and turn off any weak ciphers such as RC4.

RDMA overview

ONTAP’s Remote Direct Memory Access (RDMA) offerings support latency-sensitive and high-bandwidth workloads. RDMA allows data to be copied directly between storage system memory and host system memory, circumventing CPU interruptions and overhead.

NFS over RDMA

Beginning with ONTAP 9.10.1, you can configure [NFS over RDMA](#) to enable the use of NVIDIA GPUDirect Storage for GPU-accelerated workloads on hosts with supported NVIDIA GPUs.

RDMA cluster interconnect

RDMA cluster interconnect reduces latency, decreases failover times, and accelerates communication between nodes in a cluster.

Beginning with ONTAP 9.10.1, cluster interconnect RDMA is supported for certain hardware systems when used with X1151A cluster NICs. Beginning with ONTAP 9.13.1, X91153A NICs also support cluster interconnect RDMA. Consult the table to learn what systems are supported in different ONTAP releases.

Systems	Supported ONTAP versions
<ul style="list-style-type: none">• AFF A400• ASA A400	ONTAP 9.10.1 and later
<ul style="list-style-type: none">• AFF A900• ASA A900• FAS9500	ONTAP 9.13.1 and later

Given the appropriate storage system set up, no additional configuration is needed to use RDMA interconnect.

Client protocols

ONTAP supports all major industry-standard client protocols: NFS, SMB, FC, FCoE, iSCSI, NVMe, and S3.

NFS

NFS is the traditional file access protocol for UNIX and LINUX systems. Clients can access files in ONTAP volumes using the following protocols.

- NFSv3
- NFSv4
- NFSv4.2
- NFSv4.1
- pNFS

You can control file access using UNIX-style permissions, NTFS-style permissions, or a mix of both.

Clients can access the same files using both NFS and SMB protocols.

SMB

SMB is the traditional file access protocol for Windows systems. Clients can access files in ONTAP volumes using the SMB 2.0, SMB 2.1, SMB 3.0, and SMB 3.1.1 protocols. Just like with NFS, a mix of permission styles are supported.

SMB 1.0 is available but disabled by default in ONTAP 9.3 and later releases.

FC

Fibre Channel is the original networked block protocol. Instead of files, a block protocol presents an entire virtual disk to a client. The traditional FC protocol uses a dedicated FC network with specialized FC switches, and requires the client computer to have FC network interfaces.

A LUN represents the virtual disk, and one or more LUNs are stored in an ONTAP volume. The same LUN can be accessed through the FC, FCoE, and iSCSI protocols, but multiple clients can access the same LUN only if they are part of a cluster that prevents write collisions.

FCoE

FCoE is basically the same protocol as FC, but uses a datacenter-grade Ethernet network in place of the traditional FC transport. The client still requires an FCoE-specific network interface.

iSCSI

iSCSI is a block protocol that can run on standard Ethernet networks. Most client operating systems offer a software initiator that runs over a standard Ethernet port. iSCSI is a good choice when you need a block protocol for a particular application, but do not have dedicated FC networking available.

NVMe/FC and NVMe/TCP

The newest block protocol, NVMe, is specifically designed to work with flash-based storage. It offers scalable sessions, a significant reduction in latency, and an increase in parallelism, making it well suited to low-latency and high-throughput applications such as in-memory databases and analytics.

Unlike FC and iSCSI, NVMe does not use LUNs. Instead it uses namespaces, which are stored in an ONTAP volume. NVMe namespaces can be accessed only through the NVMe protocol.

S3

Beginning with ONTAP 9.8, you can enable an ONTAP Simple Storage Service (S3) server in an ONTAP cluster, allowing you to serve data in object storage using S3 buckets.

ONTAP supports two on-premises use case scenarios for serving S3 object storage:

- FabricPool tier to a bucket on local cluster (tier to a local bucket) or remote cluster (cloud tier).
- S3 client app access to a bucket on the local cluster or a remote cluster.



ONTAP S3 is appropriate if you want S3 capabilities on existing clusters without additional hardware and management. For deployments larger than 300TB, NetApp StorageGRID software continues to be the NetApp flagship solution for object storage. Learn about [StorageGRID](#).

Disks and aggregates


Disks and local tiers (aggregates) overview

You can manage ONTAP physical storage using System Manager and the CLI. You can create, expand, and manage local tiers (aggregates), work with Flash Pool local tiers (aggregates), manage disks, and manage RAID policies.

What local tiers (aggregates) are

Local tiers (also called *aggregates*) are containers for the disks managed by a node. You can use local tiers to isolate workloads with different performance demands, to tier data with different access patterns, or to segregate data for regulatory purposes.

- For business-critical applications that need the lowest possible latency and the highest possible performance, you might create a local tier consisting entirely of SSDs.
- To tier data with different access patterns, you can create a *hybrid local tier*, deploying flash as high-performance cache for a working data set, while using lower-cost HDDs or object storage for less frequently accessed data.
 - A *Flash Pool* consists of both SSDs and HDDs.
 - A *FabricPool* consists of an all-SSD local tier with an attached object store.
- If you need to segregate archived data from active data for regulatory purposes, you can use a local tier consisting of capacity HDDs, or a combination of performance and capacity HDDs.



Datacenter



Cloud

You can use a FabricPool to tier data with different access patterns, deploying SSDs for frequently accessed “hot” data and object storage for rarely accessed “cold” data.

Working with local tiers (aggregates)

You can perform the following tasks:

- [Manage local tiers \(aggregates\)](#)
- [Manage disks](#)
- [Manage RAID configurations](#)
- [Manage Flash Pool tiers](#)

You perform these tasks if the following are true:

- You do not want to use an automated scripting tool.
- You want to use best practices, not explore every available option.
- You have a MetroCluster configuration and you are following the procedures in the [MetroCluster](#) documentation for initial configuration and guidelines for local tiers (aggregates) and disk management.

Related information

- [Manage FabricPool cloud tiers](#)

Local tiers (aggregates) and RAID groups

Modern RAID technologies protect against disk failure by rebuilding a failed disk’s data on a spare disk. The system compares index information on a “parity disk” with data on the remaining healthy disks to reconstruct the missing data, all without downtime or a significant performance cost.

A local tier (aggregate) consists of one or more *RAID groups*. The *RAID type* of the local tier determines the number of parity disks in the RAID group and the number of simultaneous disk failures that the RAID configuration protects against.

The default RAID type, RAID-DP (RAID-double parity), requires two parity disks per RAID group and protects against data loss in the event of two disks failing at the same time. For RAID-DP, the recommended RAID group size is between 12 and 20 HDDs and between 20 and 28 SSDs.

You can spread out the overhead cost of parity disks by creating RAID groups at the higher end of the sizing recommendation. This is especially the case for SSDs, which are much more reliable than capacity drives. For local tiers that use HDDs, you should balance the need to maximize disk storage against countervailing factors like the longer rebuild time required for larger RAID groups.

Mirrored and unmirrored local tiers (aggregates)


ONTAP has an optional feature called *SyncMirror* which you can use to synchronously mirror local tier (aggregate) data in copies, or *plexes*, stored in different RAID groups. Plexes ensure against data loss if more disks fail than the RAID type protects against, or if there is a loss of connectivity to RAID group disks.

When you create a local tier with System Manager or using the CLI, you can specify that the local tier is mirrored or unmirrored.

How unmirrored local tiers (aggregates) work

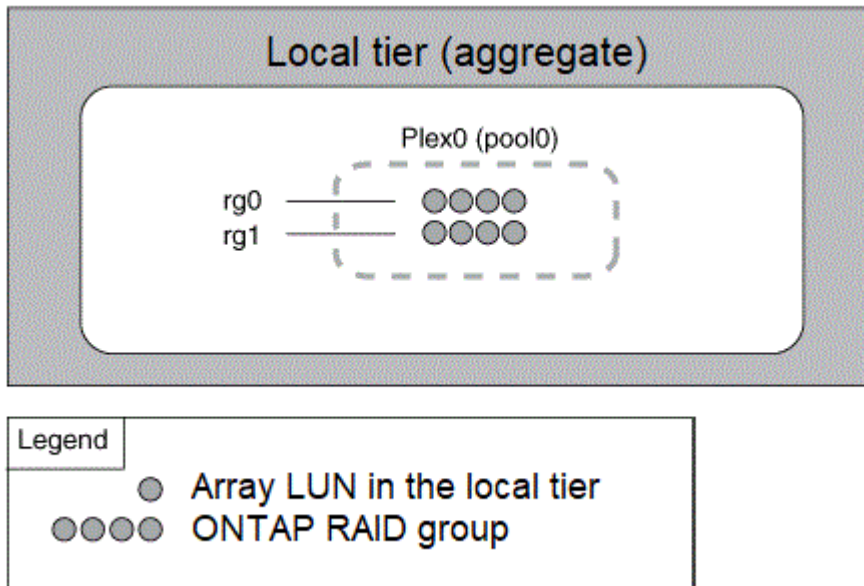
If you do not specify that the local tiers are mirrored, then they are created as unmirrored local tiers (aggregates). Unmirrored local tiers have only one *plex* (a copy of their data), which contains all of the RAID groups belonging to that local tier.

The following diagram shows an unmirrored local tier composed of disks, with its one plex. The local tier has four RAID groups: rg0, rg1, rg2, and rg3. Each RAID group has six data disks, one parity disk, and one dparity (double parity) disk. All disks used by the local tier come from the same pool, “pool0”.



The following diagram shows an unmirrored local tier with array LUNs, with its one plex. It has two RAID

groups, rg0 and rg1. All array LUNs used by the local tier come from the same pool, “pool0”.



How mirrored local tiers (aggregates) work


Mirrored aggregates have two *plexes* (copies of their data), which use the SyncMirror functionality to duplicate the data to provide redundancy.

When you create a local tier, you can specify that it is a mirrored local tier. Also, you can add a second plex to an existing unmirrored local tier to make it a mirrored tier. Using SyncMirror functionality, ONTAP copies the data in the original plex (plex0) to the new plex (plex1). The plexes are physically separated (each plex has its own RAID groups and its own pool), and the plexes are updated simultaneously.

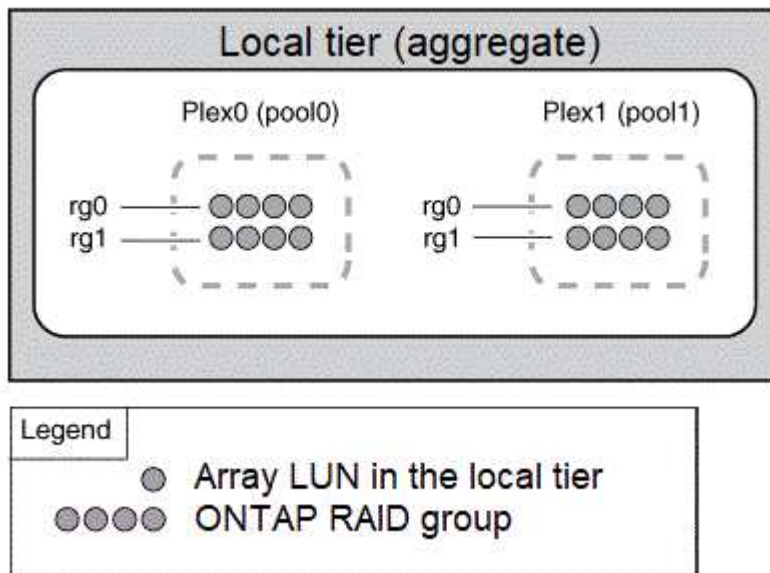
This configuration provides added protection against data loss if more disks fail than the RAID level of the aggregate protects against or if there is a loss of connectivity, because the unaffected plex continues to serve data while you fix the cause of the failure. After the plex that had a problem is fixed, the two plexes resynchronize and reestablish the mirror relationship.

The disks and array LUNs on the system are divided into two pools: “pool0” and “pool1”. Plex0 gets its storage from pool0 and plex1 gets its storage from pool1.

The following diagram shows a local tier composed of disks with the SyncMirror functionality enabled and implemented. A second plex has been created for the local tier, “plex1”. The data in plex1 is a copy of the data in plex0, and the RAID groups are also identical. The 32 spare disks are allocated to pool0 or pool1 using 16 disks for each pool.



The following diagram shows a local tier composed of array LUNs with the SyncMirror functionality enabled and implemented. A second plex has been created for the local tier, “plex1”. Plex1 is a copy of plex0, and the RAID groups are also identical.



It's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space may be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices may have a negative impact on performance.


Root-data partitioning

Every node must have a root aggregate for storage system configuration files. The root aggregate has the RAID type of the data aggregate.

System Manager does not support root-data or root-data-data partitioning.


A root aggregate of type RAID-DP typically consists of one data disk and two parity disks. That's a significant "parity tax" to pay for storage system files, when the system is already reserving two disks as parity disks for each RAID group in the aggregate.

Root-data partitioning reduces the parity tax by apportioning the root aggregate across disk partitions, reserving one small partition on each disk as the root partition and one large partition for data.



Root-data partitioning creates one small partition on each disk as the root partition and one large partition on each disk for data.

As the illustration suggests, the more disks used to store the root aggregate, the smaller the root partition. That's also the case for a form of root-data partitioning called *root-data-data partitioning*, which creates one small partition as the root partition and two larger, equally sized partitions for data.



Root-data-data partitioning creates one small partition as the root partition and two larger, equally sized partitions for data.

Both types of root-data partitioning are part of the ONTAP *Advanced Drive Partitioning (ADP)* feature. Both are configured at the factory: root-data partitioning for entry-level FAS2xxx, FAS9000, FAS8200, FAS80xx and AFF systems, root-data-data partitioning for AFF systems only.

Learn more about [Advanced Drive Partitioning](#).

Drives partitioned and used for the root aggregate

The drives that are partitioned for use in the root aggregate depend on the system configuration.

Knowing how many drives are used for the root aggregate helps you to determine how much of the drives' capacity is reserved for the root partition, and how much is available for use in a data aggregate.

The root-data partitioning capability is supported for entry-level platforms, All Flash FAS platforms, and FAS platforms with only SSDs attached.

For entry-level platforms, only the internal drives are partitioned.

For All Flash FAS platforms and FAS platforms with only SSDs attached, all drives that are attached to the controller when the system is initialized are partitioned, up to a limit of 24 per node. Drives that are added after system configuration are not partitioned.

Volumes, qtrees, files, and LUNs

ONTAP serves data to clients and hosts from logical containers called *FlexVol volumes*. Because these volumes are only loosely coupled with their containing aggregate, they offer greater flexibility in managing data than traditional volumes.

You can assign multiple FlexVol volumes to an aggregate, each dedicated to a different application or service. You can expand and contract a FlexVol volume, move a FlexVol volume, and make efficient copies of a FlexVol volume. You can use *qtrees* to partition a FlexVol volume into more manageable units, and *quotas* to limit volume resource usage.

Volumes contain file systems in a NAS environment and LUNs in a SAN environment. A LUN (logical unit

number) is an identifier for a device called a *logical unit* addressed by a SAN protocol.

LUNs are the basic unit of storage in a SAN configuration. The Windows host sees LUNs on your storage system as virtual disks. You can nondisruptively move LUNs to different volumes as needed.

In addition to data volumes, there are a few special volumes you need to know about:

- A *node root volume* (typically “vol0”) contains node configuration information and logs.
- An *SVM root volume* serves as the entry point to the namespace provided by the SVM and contains namespace directory information.
- *System volumes* contain special metadata such as service audit logs.

You cannot use these volumes to store data.



Volumes contain files in a NAS environment and LUNs in a SAN environment.

FlexGroup volumes

In some enterprises a single namespace may require petabytes of storage, far exceeding even a FlexVol volume's 100TB capacity.

A *FlexGroup volume* supports up to 400 billion files with 200 constituent member volumes that work collaboratively to dynamically balance load and space allocation evenly across all members.

There is no required maintenance or management overhead with a FlexGroup volume. You simply create the FlexGroup volume and share it with your NAS clients. ONTAP does the rest.

Storage virtualization

Storage virtualization overview

You use *storage virtual machines (SVMs)* to serve data to clients and hosts. Like a virtual machine running on a hypervisor, an SVM is a logical entity that abstracts physical resources. Data accessed through the SVM is not bound to a location in storage. Network access to the SVM is not bound to a physical port.



SVMs were formerly called "vservers." The ONTAP command line interface still uses the term "vserver".


An SVM serves data to clients and hosts from one or more volumes, through one or more network *logical interfaces (LIFs)*. Volumes can be assigned to any data aggregate in the cluster. LIFs can be hosted by any physical or logical port. Both volumes and LIFs can be moved without disrupting data service, whether you are performing hardware upgrades, adding nodes, balancing performance, or optimizing capacity across aggregates.

The same SVM can have a LIF for NAS traffic and a LIF for SAN traffic. Clients and hosts need only the address of the LIF (IP address for NFS, SMB, or iSCSI; WWPN for FC) to access the SVM. LIFs keep their addresses as they move. Ports can host multiple LIFs. Each SVM has its own security, administration, and namespace.

In addition to data SVMs, ONTAP deploys special SVMs for administration:

- An *admin SVM* is created when the cluster is set up.
- A *node SVM* is created when a node joins a new or existing cluster.
- A *system SVM* is automatically created for cluster-level communications in an IPspace.

You cannot use these SVMs to serve data. There are also special LIFs for traffic within and between clusters, and for cluster and node management.



Data accessed through an SVM is not bound to a physical storage location. You can move a volume without disrupting data service.

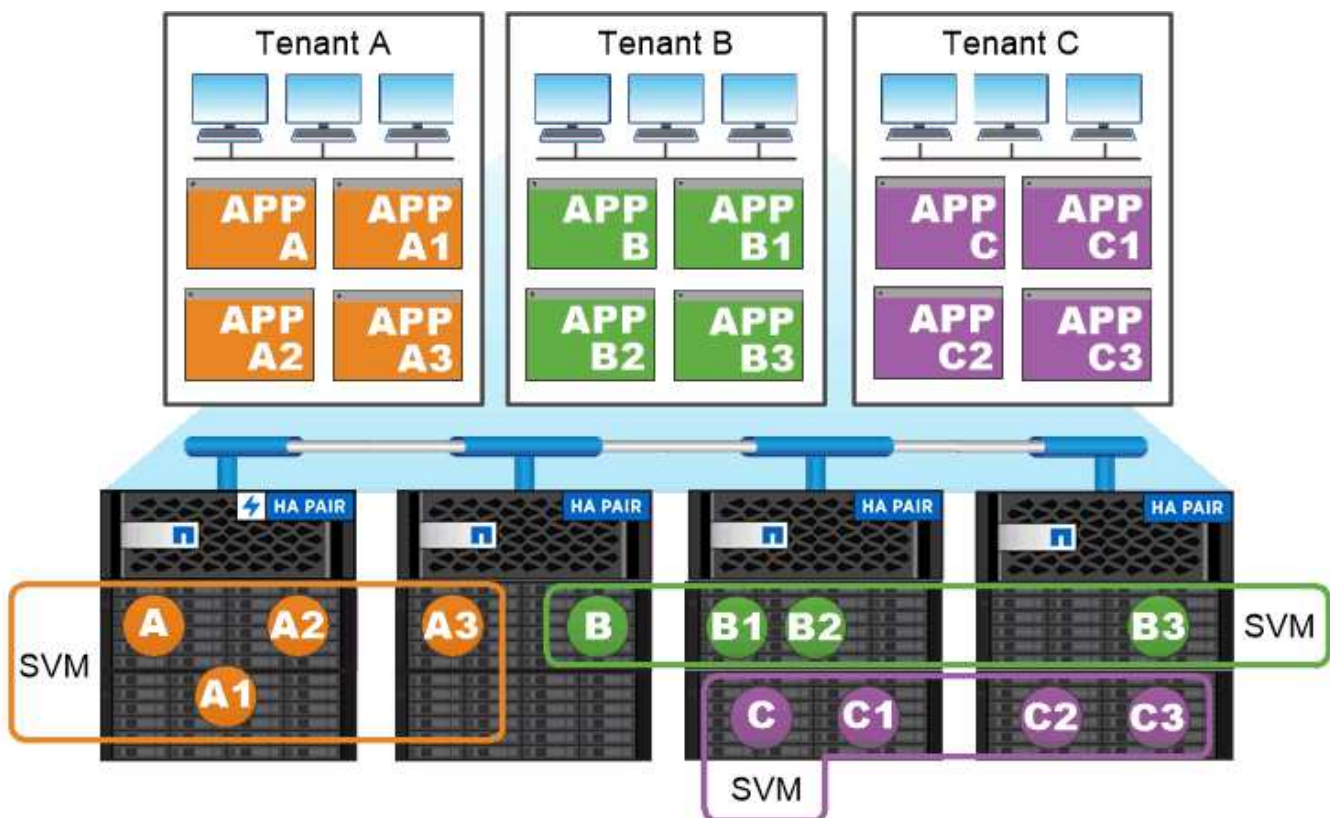
Why ONTAP is like middleware

The logical objects ONTAP uses for storage management tasks serve the familiar goals of a well-designed middleware package: shielding the administrator from low-level implementation details and insulating the configuration from changes in physical characteristics like nodes and ports. The basic idea is that the administrator should be able to move volumes and LIFs easily, reconfiguring a few fields rather than the entire storage infrastructure.

SVM use cases

Service providers use SVMs in secure multitenancy arrangements to isolate each tenant's data, to provide each tenant with its own authentication and administration, and to simplify chargeback. You can assign multiple LIFs to the same SVM to satisfy different customer needs, and you can use QoS to protect against tenant workloads “bullying” the workloads of other tenants.

Administrators use SVMs for similar purposes in the enterprise. You might want to segregate data from different departments, or keep storage volumes accessed by hosts in one SVM and user share volumes in another. Some administrators put iSCSI/FC LUNs and NFS datastores in one SVM and SMB shares in another.



Service providers use SVMs in multitenant environments to isolate tenant data and simplify chargeback.

Cluster and SVM administration

A *cluster administrator* accesses the admin SVM for the cluster. The admin SVM and a cluster administrator with the reserved name `admin` are automatically created when the cluster is set up.

A cluster administrator with the default `admin` role can administer the entire cluster and its resources. The cluster administrator can create additional cluster administrators with different roles as needed.

An *SVM administrator* accesses a data SVM. The cluster administrator creates data SVMs and SVM administrators as needed.

SVM administrators are assigned the `vsadmin` role by default. The cluster administrator can assign different roles to SVM administrators as needed.

Role-Based Access Control (RBAC)

The *role* assigned to an administrator determines the commands to which the administrator has access. You assign the role when you create the account for the administrator. You can assign a different role or define custom roles as needed.

Namespaces and junction points

A NAS *namespace* is a logical grouping of volumes joined together at *junction points* to create a single file system hierarchy. A client with sufficient permissions can access files in the namespace without specifying the location of the files in storage. Junctioned volumes can reside anywhere in the cluster.


Rather than mounting every volume containing a file of interest, NAS clients mount an NFS *export* or access an SMB *share*. The export or share represents the entire namespace or an intermediate location within the namespace. The client accesses only the volumes mounted below its access point.

You can add volumes to the namespace as needed. You can create junction points directly below a parent volume junction or on a directory within a volume. A path to a volume junction for a volume named “vol3” might be `/vol1/vol2/vol3`, or `/vol1/dir2/vol3`, or even `/dir1/dir2/vol3`. The path is called the *junction path*.

Every SVM has a unique namespace. The SVM root volume is the entry point to the namespace hierarchy.



To ensure that data remains available in the event of a node outage or failover, you should create a *load-sharing mirror* copy for the SVM root volume.



A namespace is a logical grouping of volumes joined together at junction points to create a single file system hierarchy.

Example

The following example creates a volume named “home4” located on SVM vs1 that has a junction path /eng/home:

```
cluster1::> volume create -vserver vs1 -volume home4 -aggregate aggr1
-size 1g -junction-path /eng/home
[Job 1642] Job succeeded: Successful
```

Path failover

Path failover overview


There are important differences in how ONTAP manages path failover in NAS and SAN topologies. A NAS LIF automatically migrates to a different network port after a link failure. A SAN LIF does not migrate (unless you move it manually after the failure). Instead, multipathing technology on the host diverts traffic to a different LIF—on the same SVM, but accessing a different network port.

NAS path failover

A NAS LIF automatically migrates to a surviving network port after a link failure on its current port. The port to which the LIF migrates must be a member of the *failover group* for the LIF. The *failover group policy* narrows the failover targets for a data LIF to ports on the node that owns the data and its HA partner.

For administrative convenience, ONTAP creates a failover group for each *broadcast domain* in the network architecture. Broadcast domains group ports that belong to the same layer 2 network. If you are using VLANs, for example, to segregate traffic by department (Engineering, Marketing, Finance, and so on), each VLAN defines a separate broadcast domain. The failover group associated with the broadcast domain is automatically updated each time you add or remove a broadcast domain port.

It is almost always a good idea to use a broadcast domain to define a failover group to ensure that the failover group remains current. Occasionally, however, you may want to define a failover group that is not associated with a broadcast domain. For example, you may want LIFs to fail over only to ports in a subset of the ports defined in the broadcast domain.



A NAS LIF automatically migrates to a surviving network port after a link failure on its current port.

Subnets

A *subnet* reserves a block of IP addresses in a broadcast domain. These addresses belong to the same layer 3 network and are allocated to ports in the broadcast domain when you create a LIF. It is usually easier and less error-prone to specify a subnet name when you define a LIF address than it is to specify an IP address and network mask.

SAN path failover


A SAN host uses ALUA (Asymmetric Logical Unit Access) and MPIO (multipath I/O) to reroute traffic to a surviving LIF after a link failure. Predefined paths determine the

possible routes to the LUN served by the SVM.

In a SAN environment, hosts are regarded as *initiators* of requests to LUN *targets*. MPIO enables multiple paths from initiators to targets. ALUA identifies the most direct paths, called *optimized paths*.

You typically configure multiple optimized paths to LIFs on the LUN's owning node, and multiple non-optimized paths to LIFs on its HA partner. If one port fails on the owning node, the host routes traffic to the surviving ports. If all the ports fail, the host routes traffic over the non-optimized paths.

ONTAP Selective LUN Map (SLM) limits the number of paths from the host to a LUN by default. A newly created LUN is accessible only through paths to the node that owns the LUN or its HA partner. You can also limit access to a LUN by configuring LIFs in a *port set* for the initiator.



A SAN host uses multipathing technology to reroute traffic to a surviving LIF after a link failure.

Moving volumes in SAN environments

By default, ONTAP *Selective LUN Map (SLM)* limits the number of paths to a LUN from a SAN host. A newly created LUN is accessible only through paths to the node that owns the LUN or its HA partner, the *reporting nodes* for the LUN.

This means that when you move a volume to a node on another HA pair, you need to add reporting nodes for the destination HA pair to the LUN mapping. You can then specify the new paths in your MPIO setup. After the volume move is complete, you can delete the reporting nodes for the source HA pair from the mapping.

Load balancing

Performance of workloads begins to be affected by latency when the amount of work on a node exceeds the available resources. You can manage an overloaded node by increasing the available resources (upgrading disks or CPU), or by reducing load (moving volumes or LUNs to different nodes as needed).

You can also use ONTAP *storage quality of service (QoS)* to guarantee that performance of critical workloads is not degraded by competing workloads:

- You can set a QoS throughput *ceiling* on a competing workload to limit its impact on system resources (QoS Max).
- You can set a QoS throughput *floor* for a critical workload, ensuring that it meets minimum throughput targets regardless of demand by competing workloads (QoS Min).
- You can set a QoS ceiling and floor for the same workload.


Throughput ceilings

A throughput ceiling limits throughput for a workload to a maximum number of IOPS or MB/s. In the figure below, the throughput ceiling for workload 2 ensures that it does not “bully” workloads 1 and 3.

A *policy group* defines the throughput ceiling for one or more workloads. A workload represents the I/O operations for a *storage object*: a volume, file, or LUN, or all the volumes, files, or LUNs in an SVM. You can specify the ceiling when you create the policy group, or you can wait until after you monitor workloads to specify it.



Throughput to workloads might exceed the specified ceiling by up to 10 percent, especially if a workload experiences rapid changes in throughput. The ceiling might be exceeded by up to 50% to handle bursts.



The throughput ceiling for workload 2 ensures that it does not “bully” workloads 1 and 3.

Throughput floors

A throughput floor guarantees that throughput for a workload does not fall below a minimum number of IOPS. In the figure below, the throughput floors for workload 1 and workload 3 ensure that they meet minimum throughput targets, regardless of demand by workload 2.

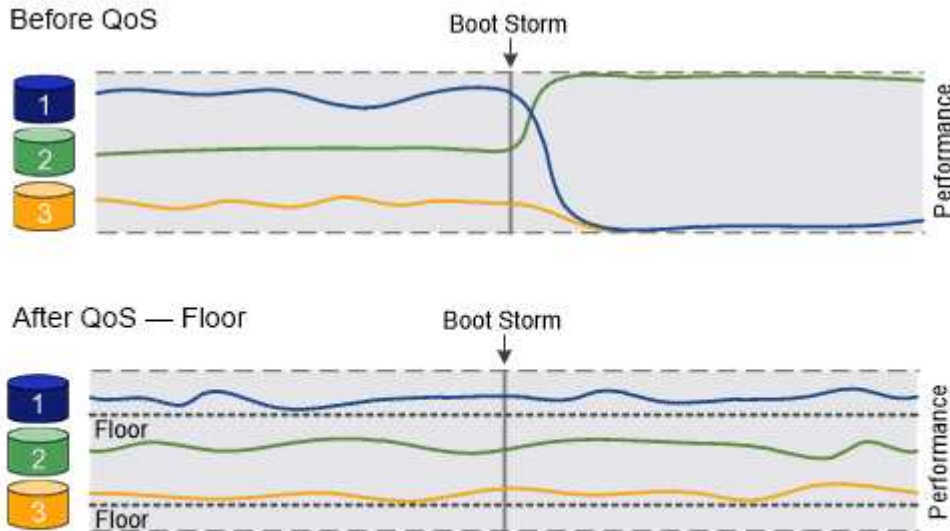


As the examples suggest, a throughput ceiling throttles throughput directly. A throughput floor throttles throughput indirectly, by giving priority to the workloads for which the floor has been set.

A workload represents the I/O operations for a volume, LUN, or, beginning with ONTAP 9.3, file. A policy group that defines a throughput floor cannot be applied to an SVM. You can specify the floor when you create the policy group, or you can wait until after you monitor workloads to specify it.



Throughput to a workload might fall below the specified floor if there is insufficient performance capacity (headroom) on the node or aggregate, or during critical operations like `volume move trigger-cutover`. Even when sufficient capacity is available and critical operations are not taking place, throughput to a workload might fall below the specified floor by up to 5 percent.



The throughput floors for workload 1 and workload 3 ensure that they meet minimum throughput targets, regardless of demand by workload 2.

Adaptive QoS

Ordinarily, the value of the policy group you assign to a storage object is fixed. You need to change the value manually when the size of the storage object changes. An increase in the amount of space used on a volume, for example, usually requires a corresponding increase in the throughput ceiling specified for the volume.

Adaptive QoS automatically scales the policy group value to workload size, maintaining the ratio of IOPS to TBs|GBs as the size of the workload changes. That's a significant advantage when you are managing hundreds or thousands of workloads in a large deployment.

You typically use adaptive QoS to adjust throughput ceilings, but you can also use it to manage throughput floors (when workload size increases). Workload size is expressed as either the allocated space for the storage object or the space used by the storage object.



Used space is available for throughput floors in ONTAP 9.5 and later. It is not supported for throughput floors in ONTAP 9.4 and earlier.

+ Beginning in ONTAP 9.13.1, you can use adaptive QoS to set throughput floors and ceilings at the SVM level.

- An *allocated space* policy maintains the IOPS/TB|GB ratio according to the nominal size of the storage object. If the ratio is 100 IOPS/GB, a 150 GB volume will have a throughput ceiling of 15,000 IOPS for as long as the volume remains that size. If the volume is resized to 300 GB, adaptive QoS adjusts the

throughput ceiling to 30,000 IOPS.

- A *used space* policy (the default) maintains the IOPS/TB|GB ratio according to the amount of actual data stored before storage efficiencies. If the ratio is 100 IOPS/GB, a 150 GB volume that has 100 GB of data stored would have a throughput ceiling of 10,000 IOPS. As the amount of used space changes, adaptive QoS adjusts the throughput ceiling according to the ratio.

Replication

Snapshot copies

Traditionally, ONTAP replication technologies served the need for disaster recovery (DR) and data archiving. With the advent of cloud services, ONTAP replication has been adapted to data transfer between endpoints in the NetApp data fabric. The foundation for all these uses is ONTAP Snapshot technology.


A *Snapshot copy* is a read-only, point-in-time image of a volume. After a Snapshot copy is created, the active file system and the Snapshot copy point to the same disk blocks; therefore, the Snapshot copy doesn't use extra disk space. Over time, the image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last Snapshot copy was made.

Snapshot copies owe their efficiency to ONTAP's core storage virtualization technology, its *Write Anywhere File Layout (WAFL)*. Like a database, WAFL uses metadata to point to actual data blocks on disk. But, unlike a database, WAFL does not overwrite existing blocks. It writes updated data to a new block and changes the metadata.

Snapshot copies are efficient because, rather than copy data blocks, ONTAP references metadata when creating a Snapshot copy. Doing so eliminates both the "seek time" that other systems incur in locating the blocks to copy and the cost of making the copy itself.

You can use a Snapshot copy to recover individual files or LUNs, or to restore the entire contents of a volume. ONTAP compares pointer information in the Snapshot copy with data on disk to reconstruct the missing or damaged object, without downtime or a significant performance cost.

A *Snapshot policy* defines how the system creates Snapshot copies of volumes. The policy specifies when to create the Snapshot copies, how many copies to retain, how to name them, and how to label them for replication. For example, a system might create one Snapshot copy every day at 12:10 a.m., retain the two most recent copies, name them "daily" (appended with a timestamp), and label them "daily" for replication.



A Snapshot copy records only changes to the active file system since the last Snapshot copy.

SnapMirror disaster recovery and data transfer

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. As its name implies, SnapMirror creates a replica, or *mirror*, of your working data in secondary storage from which you can continue to serve data in the event of a catastrophe at the primary site.

Data is mirrored at the volume level. The relationship between the source volume in primary storage and the destination volume in secondary storage is called a *data protection relationship*. The clusters in which the volumes reside and the SVMs that serve data from the volumes must be *peered*. A peer relationship enables clusters and SVMs to exchange data securely.



You can also create a data protection relationship between SVMs. In this type of relationship, all or part of the SVM's configuration, from NFS exports and SMB shares to RBAC, is replicated, as well as the data in the volumes the SVM owns.


Beginning with ONTAP 9.10.1, you can create data protection relationships between S3 buckets using SnapMirror S3. Destination buckets can be on local or remote ONTAP systems, or on non-ONTAP systems such as StorageGRID and AWS.

The first time you invoke SnapMirror, it performs a *baseline transfer* from the source volume to the destination volume. The baseline transfer typically involves the following steps:

- Make a Snapshot copy of the source volume.
- Transfer the Snapshot copy and all the data blocks it references to the destination volume.
- Transfer the remaining, less recent Snapshot copies on the source volume to the destination volume for use in case the “active” mirror is corrupted.

Once a baseline transfer is complete, SnapMirror transfers only new Snapshot copies to the mirror. Updates are asynchronous, following the schedule you configure. Retention mirrors the Snapshot policy on the source. You can activate the destination volume with minimal disruption in case of a disaster at the primary site, and reactivate the source volume when service is restored.

Because SnapMirror transfers only Snapshot copies after the baseline is created, replication is fast and nondisruptive. As the failover use case implies, the controllers on the secondary system should be equivalent or nearly equivalent to the controllers on the primary system to serve data efficiently from mirrored storage.



A SnapMirror data protection relationship mirrors the Snapshot copies available on the source volume.

Using SnapMirror for data transfer

You can also use SnapMirror to replicate data between endpoints in the NetApp data fabric. You can choose between one-time replication or recurring replication when you create the SnapMirror policy.

SnapMirror cloud backups to object storage

SnapMirror cloud is a backup and recovery technology designed for ONTAP users who want to transition their data protection workflows to the cloud. Organizations moving away from legacy backup-to-tape architectures can use object storage as an alternative repository for long-term data retention and archiving. SnapMirror cloud provides ONTAP-to-object storage replication as part of an incremental forever backup strategy.


SnapMirror cloud was introduced in ONTAP 9.8 as an extension to the family of SnapMirror replication technologies. While SnapMirror is frequently used for ONTAP-to-ONTAP backups, SnapMirror cloud uses the same replication engine to transfer Snapshot copies for ONTAP to S3-compliant object storage backups.

Targeted for backup use cases, SnapMirror cloud supports both long-term retention and archives workflows. As with SnapMirror, the initial SnapMirror cloud backup performs a baseline transfer of a volume. For subsequent backups, SnapMirror cloud generates a snapshot copy of the source volume and transfers the snapshot copy with only the changed data blocks to an object storage target.

SnapMirror cloud relationships can be configured between ONTAP systems and select on-premises and public cloud object storage targets - including Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage. Additional on-premises object storage targets include StorageGRID and ONTAP S3.

SnapMirror cloud replication is a licensed ONTAP feature and requires an approved application to orchestrate data protection workflows. Several orchestration options are available for managing SnapMirror cloud backups:

- Multiple 3rd party backup partners who offer support for SnapMirror cloud replication. Participating vendors are available on the [NetApp blog](#).
- BlueXP backup and recovery for a NetApp-native solution for ONTAP environments
- APIs for developing custom software for data protection workflows or leveraging automation tools



SnapVault archiving

The SnapMirror license is used to support both SnapVault relationships for backup, and SnapMirror relationships for disaster recovery. Beginning with ONTAP 9.3, SnapVault licenses are deprecated, and SnapMirror licenses can be used to configure vault, mirror, and mirror-and-vault relationships. SnapMirror replication is used for ONTAP-to-ONTAP replication of Snapshot copies, supporting both backup and disaster recovery use cases.

SnapVault is archiving technology, designed for disk-to-disk Snapshot copy replication for standards compliance and other governance-related purposes. In contrast to a SnapMirror relationship, in which the destination usually contains only the Snapshot copies currently in the source volume, a SnapVault destination typically retains point-in-time Snapshot copies created over a much longer period.


You might want to keep monthly Snapshot copies of your data over a 20-year span, for example, to comply with government accounting regulations for your business. Since there is no requirement to serve data from vault storage, you can use slower, less expensive disks on the destination system.

As with SnapMirror, SnapVault performs a baseline transfer the first time you invoke it. It makes a Snapshot copy of the source volume, then transfers the copy and the data blocks it references to the destination volume. Unlike SnapMirror, SnapVault does not include older Snapshot copies in the baseline.

Updates are asynchronous, following the schedule you configure. The rules you define in the policy for the relationship identify which new Snapshot copies to include in updates and how many copies to retain. The labels defined in the policy (“monthly,” for example) must match one or more labels defined in the Snapshot policy on the source. Otherwise, replication fails.



SnapMirror and SnapVault share the same command infrastructure. You specify which method you want to use when you create a policy. Both methods require peered clusters and peered SVMs.



A SnapVault data protection relationship typically retains point-in-time Snapshot copies created over a longer period than the Snapshot copies on the source volume.

Cloud backup and support for traditional backups

In addition to SnapMirror and SnapVault data protection relationships, which were disk-to-disk only for ONTAP 9.7 and earlier, there are now several backup solutions that offer a less expensive alternative for long-term data retention.

Numerous third-party data protection applications offer traditional backup for ONTAP-managed data. Veeam, Veritas, and Commvault, among others, all offer integrated backup for ONTAP systems.

Beginning with ONTAP 9.8, SnapMirror cloud provides asynchronous replication of Snapshot copies from ONTAP instances to object storage endpoints. SnapMirror cloud replication requires a licensed application for orchestration and management of data protection workflows. SnapMirror cloud relationships are supported from ONTAP systems to select on-premises and public cloud object storage targets — including AWS S3, Google Cloud Storage Platform, or Microsoft Azure Blob Storage — which provides enhanced efficiency with vendor backup software. Contact your NetApp representative for a list of supported certified applications and object storage vendors.

If you are interested in cloud-native data protection, BlueXP can be used to configure SnapMirror or SnapVault relationships between on-premises volumes and Cloud Volumes ONTAP instances in the public cloud.

BlueXP also provides backups of Cloud Volumes ONTAP instances using a Software as a Service (SaaS) model. Users can back up their Cloud Volumes ONTAP instances to S3 and S3-compliant public cloud object storage using Cloud Backup found on NetApp Cloud Central.

[Cloud Volumes ONTAP and BlueXP documentation resources](#)

[NetApp Cloud Central](#)


MetroCluster continuous availability

MetroCluster configurations protect data by implementing two physically separate, mirrored clusters. Each cluster synchronously replicates the data and SVM configuration of the other. In the event of a disaster at one site, an administrator can activate the mirrored SVM and begin serving data from the surviving site.

- *Fabric-attached MetroCluster* configurations support metropolitan-wide clusters.
- *Stretch MetroCluster* configurations support campus-wide clusters.

Clusters must be peered in either case.

MetroCluster uses an ONTAP feature called *SyncMirror* to synchronously mirror aggregate data for each cluster in copies, or *plexes*, in the other cluster's storage. If a switchover occurs, the remote plex on the surviving cluster comes online and the secondary SVM begins serving data.



When a MetroCluster switchover occurs, the remote plex on the surviving cluster comes online and the secondary SVM begins serving data.

Using SyncMirror in non-MetroCluster implementations

You can optionally use SyncMirror in a non-MetroCluster implementation to protect against data loss if more disks fail than the RAID type protects against, or if there is a loss of connectivity to RAID group disks. The feature is available for HA pairs only.

Aggregate data is mirrored in plexes stored on different disk shelves. If one of the shelves becomes unavailable, the unaffected plex continues to serve data while you fix the cause of the failure.

Keep in mind that an aggregate mirrored using SyncMirror requires twice as much storage as an unmirrored aggregate. Each plex requires as many disks as the plex it mirrors. You would need 2,880 GB of disk space, for example, to mirror a 1,440 GB aggregate, 1,440 GB for each plex.

With SyncMirror, it's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space may be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices may have a negative impact on SyncMirror resynchronization performance, which indirectly impacts operational workflows such as NDU for non-shared cloud deployments and switchback for MetroCluster deployments.



SyncMirror is also available for FlexArray Virtualization implementations.

Storage efficiency

ONTAP storage efficiency overview

Storage efficiency is the measure of how effectively a storage system utilizes available space by optimizing storage resources, minimizing wasted space, and reducing written data's physical footprint. Higher storage efficiency enables you to store the maximum amount of data within the smallest possible space at the lowest possible cost. For example, utilizing storage efficient technologies that detect and eliminate duplicate data blocks and data blocks filled with zeros decreases the overall amount of physical storage you need and reduces your overall cost.

ONTAP offers a wide range of storage efficiency technologies that reduce the amount of physical hardware or cloud storage consumed by your data and that also yield significant improvements to system performance, including faster reads of data, faster copies of datasets, and faster VM provisioning.

ONTAP storage efficiency technologies include:

- **Thin provisioning**

[Thin provisioning](#) enables you to allocate storage in a volume or LUN as it is needed instead of reserving it in advance. This reduces the amount of physical storage you need by allowing you to over-allocate your volumes or LUNs based on potential usage without reserving space that isn't currently being used.

- **Deduplication**

[Deduplication](#) reduces the amount of physical storage required for a volume in three distinct ways.

- **Zero block deduplication**

Zero block deduplication detects and eliminates data blocks filled with all zeros and only updates metadata. 100% of the space typically used by zero blocks is then saved. Zero block deduplication is enabled by default on all deduplicated volumes.

- **Inline deduplication**

Inline deduplication detects duplicate data blocks and replaces them with references to a unique shared block before data is written to disk. Inline deduplication speeds up VM provisioning by 20% to 30%. Depending upon your version of ONTAP and your platform, inline deduplication is available at the volume or aggregate level. It is enabled by default on AFF and ASA systems. You need to manually enable inline deduplication on FAS systems.

- **Background deduplication**

Background deduplication also detects duplicate data blocks and replaces them with references to a unique shared block, but further enhances storage efficiency by doing so after data is written to the disk. You can set up background deduplication to run when certain criteria are met on your storage system. For example, you might enable background deduplication to occur when your volume reaches 10% utilization. You can also manually trigger background deduplication or set it to run on a specific schedule. It is enabled by default on AFF and ASA systems. You need to manually enable background deduplication on FAS systems.

Deduplication is supported within volumes and across volumes within an aggregate. Reads of deduplicated data typically incur no performance charge.

- **Compression**

Compression reduces the amount of physical storage required for a volume by combining data blocks in compression groups, each of which is stored as a single block. When a read or an overwrite request is received, only a small group of blocks is read, not the entire file. This process optimizes read and overwrite performance and enables greater scalability in the size of the files being compressed.

Compression can be run inline or postprocess. Inline compression provides immediate space savings by compressing data in memory before it is written to disk. Postprocess compression first writes the blocks to disk as uncompressed and then at a scheduled time compresses the data. It is enabled by default on AFA systems. You need to manually enable compression on all other systems.

- **Compaction**

Compaction reduces the amount of physical storage required for a volume by taking data chunks that are stored in 4 KB blocks, but that are less than 4 KB in size, and combining them into a single block. Compaction takes place while data is still in memory so unnecessary space is never consumed on disks. It is enabled by default on AFF and ASA systems. You need to manually enable compaction on FAS systems.

- **FlexClone volumes, files and LUNs**

FlexClone technology leverages Snapshot metadata to create writable, point-in-time copies of a volume, file or LUN. Copies share data blocks with their parents, consuming no storage except what is required for metadata until changes are written to a copy or its parent. When a change is written, only the delta is stored.

Where traditional dataset copies can take minutes or even hours to create, FlexClone technology lets you copy even the largest datasets almost instantaneously.

- **Temperature-sensitive storage efficiency**

ONTAP provides [temperature-sensitive storage efficiency](#) benefits by assessing how often your volume's data is accessed and mapping that frequency to the degree of compression applied to that data. For cold data that is accessed infrequently, larger data blocks are compressed. For hot data that is accessed frequently and is overwritten more often, smaller data blocks are compressed, making the process more efficient.

Temperature-sensitive storage efficiency (TSSE), introduced in ONTAP 9.8, is enabled automatically on newly created thinly provisioned AFF volumes. It is not enabled on [AFF A70, AFF A90, and AFF A1K platforms](#) that are introduced in ONTAP 9.15.1, which use a hardware offload processor.

- **CPU or dedicated offload processor storage efficiency**

Beginning with ONTAP 9.15.1, ONTAP provides [CPU or dedicated offload processor storage efficiency](#) and data compaction on AFF A70, AFF A90, and AFF A1K platforms. Storage efficiency is enabled automatically and requires no configuration.

You can realize the benefit of these technologies in your day-to-day operations with minimal effort. For example, suppose you need to supply 5,000 users with storage for home directories, and you estimate that the maximum space needed by any user is 1 GB. You could reserve a 5 TB aggregate in advance to meet the total potential storage need. However, you also know that home directory capacity requirements vary greatly across your organization. Instead of reserving 5 TB of total space for your organization, you can create a 2 TB aggregate. Then you can use thin provisioning to nominally assign 1 GB of storage to each user but allocate the storage only as needed. You can actively monitor the aggregate over time and increase the actual physical size as necessary.

In another example, suppose you are using a Virtual Desktop Infrastructure (VDI) with a large amount of duplicate data among your virtual desktops. Deduplication reduces your storage usage by automatically eliminating duplicate blocks of information across the VDI, replacing them with a pointer to the original block. Other ONTAP storage efficiency technologies, such as compression, can also run in the background without your intervention.

ONTAP disk partitioning technology delivers greater storage efficiency as well. RAID DP technology protects against double disk failure without sacrificing performance or adding disk-mirroring overhead. Advanced SSD partitioning with ONTAP 9 increases usable capacity by almost 20%.

NetApp provides the same storage efficiency features available with on-premises ONTAP in the cloud. When you migrate data from on-premises ONTAP to the cloud, the existing storage efficiency is preserved. For example, suppose you have an SQL database containing business-critical data that you want to move from an on-premises system to the cloud. You can use data replication in BlueXP to migrate your data and, as part of the migration process, you can enable your latest on-premises policy for Snapshot copies in the cloud.

Thin provisioning

ONTAP offers a wide range of storage efficiency technologies in addition to Snapshot copies. Key technologies include thin provisioning, deduplication, compression, and FlexClone volumes, files, and LUNs. Like Snapshot copies, all are built on ONTAP's Write Anywhere File Layout (WAFL).

A *thin-provisioned* volume or LUN is one for which storage is not reserved in advance. Instead, storage is allocated dynamically, as it is needed. Free space is released back to the storage system when data in the volume or LUN is deleted.

Suppose that your organization needs to supply 5,000 users with storage for home directories. You estimate that the largest home directories will consume 1 GB of space.

In this situation, you could purchase 5 TB of physical storage. For each volume that stores a home directory, you would reserve enough space to satisfy the needs of the largest consumers.

As a practical matter, however, you also know that home directory capacity requirements vary greatly across your community. For every large user of storage, there are ten who consume little or no space.


Thin provisioning allows you to satisfy the needs of the large storage consumers without having to purchase storage you might never use. Since storage space is not allocated until it is consumed, you can “overcommit” an aggregate of 2 TB by nominally assigning a size of 1 GB to each of the 5,000 volumes the aggregate contains.

As long as you are correct that there is a 10:1 ratio of light to heavy users, and as long as you take an active role in monitoring free space on the aggregate, you can be confident that volume writes won’t fail due to lack of space.

Deduplication

Deduplication reduces the amount of physical storage required for a volume (or all the volumes in an AFF aggregate) by discarding duplicate blocks and replacing them with references to a single shared block. Reads of deduplicated data typically incur no performance charge. Writes incur a negligible charge except on overloaded nodes.

As data is written during normal use, WAFL uses a batch process to create a catalog of *block signatures*. After deduplication starts, ONTAP compares the signatures in the catalog to identify duplicate blocks. If a match exists, a byte-by-byte comparison is done to verify that the candidate blocks have not changed since the catalog was created. Only if all the bytes match is the duplicate block discarded and its disk space reclaimed.



Deduplication reduces the amount of physical storage required for a volume by discarding duplicate data blocks.

Compression

Compression reduces the amount of physical storage required for a volume by combining data blocks in *compression groups*, each of which is stored as a single block. Reads of compressed data are faster than in traditional compression methods because ONTAP

decompresses only the compression groups that contain the requested data, not an entire file or LUN.

You can perform inline or postprocess compression, separately or in combination:

- *Inline compression* compresses data in memory before it is written to disk, significantly reducing the amount of write I/O to a volume, but potentially degrading write performance. Performance-intensive operations are deferred until the next postprocess compression operation, if any.
- *Postprocess compression* compresses data after it is written to disk, on the same schedule as deduplication.

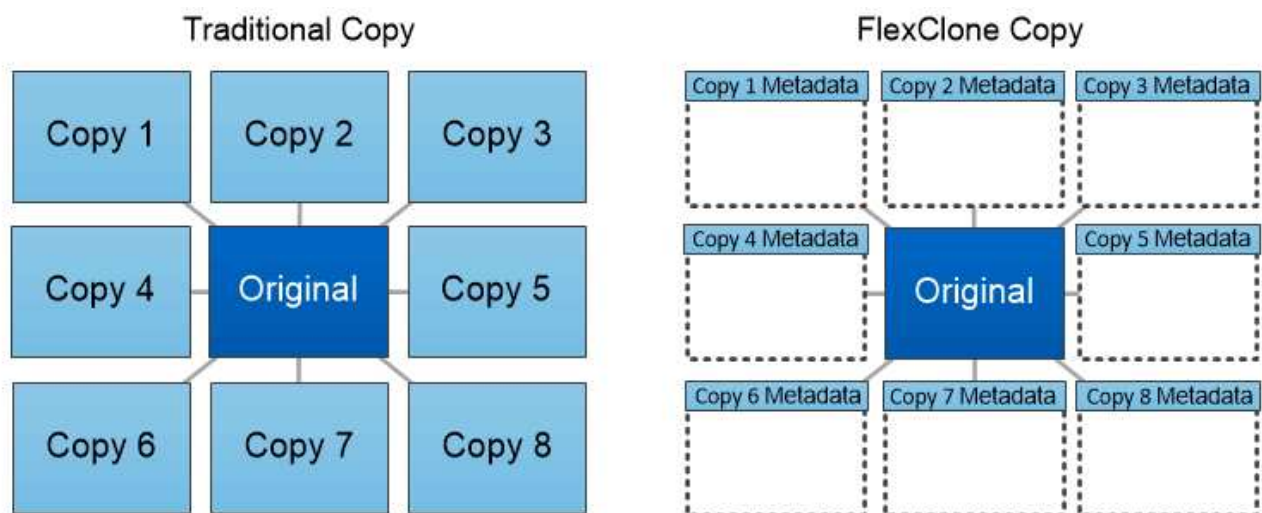
Inline data compaction Small files or I/O padded with zeros are stored in a 4 KB block whether or not they require 4 KB of physical storage. *Inline data compaction* combines data chunks that would ordinarily consume multiple 4 KB blocks into a single 4 KB block on disk. Compaction takes place while data is still in memory, so it is best suited to faster controllers.

FlexClone volumes, files, and LUNs

FlexClone technology references Snapshot metadata to create writable, point-in-time copies of a volume. Copies share data blocks with their parents, consuming no storage except what is required for metadata until changes are written to the copy. FlexClone files and FlexClone LUNs use identical technology, except that a backing Snapshot copy is not required.

Where traditional copies can take minutes or even hours to create, FlexClone software lets you copy even the largest datasets almost instantaneously. That makes it ideal for situations in which you need multiple copies of identical datasets (a virtual desktop deployment, for example) or temporary copies of a dataset (testing an application against a production dataset).

You can clone an existing FlexClone volume, clone a volume containing LUN clones, or clone mirror and vault data. You can *split* a FlexClone volume from its parent, in which case the copy is allocated its own storage.



FlexClone copies share data blocks with their parents, consuming no storage except what is required for metadata.

Capacity measurements in System Manager

System capacity can be measured as either physical space or logical space. Beginning with ONTAP 9.7, System Manager provides measurements of both physical and logical capacity.

The differences between the two measurements are explained in the following descriptions:

- **Physical capacity:** Physical space refers to the physical blocks of storage used in the volume or local tier. The value for physical used capacity is typically smaller than the value for logical used capacity due to the reduction of data from storage efficiency features (such as deduplication and compression).
- **Logical capacity:** Logical space refers to the usable space (the logical blocks) in a volume or local tier. Logical space refers to how theoretical space can be used, without accounting for results of deduplication or compression. The value for logical space used is derived from the amount of physical space used plus the savings from storage efficiency features (such as deduplication and compression) that have been configured. This measurement often appears larger than the physical used capacity because it includes Snapshot copies, clones, and other components, and it does not reflect the data compression and other reductions in the physical space. Thus, the total logical capacity could be higher than the provisioned space.



In System Manager, capacity representations do not account for root storage tier (aggregate) capacities.

Measurements of used capacity

Measurements of used capacity are displayed differently depending on the version of System Manager you are using, as explained in the following table:

Version of System Manager	Term used for capacity	Type of capacity referred to
9.9.1 and later	Logical Used	Logical space used (if storage efficiency settings have been enabled)
9.7 and 9.8	Used	Logical space used (if storage efficiency settings have been enabled)
9.5 and 9.6 (Classic view)	Used	Physical space used

Capacity measurement terms

The following terms are used when describing capacity:

- **Allocated capacity:** The amount of space that has been allocated for volumes in a storage VM.
- **Available:** The amount of physical space available to store data or to provision volumes in a storage VM or on a local tier.
- **Capacity across volumes:** The sum of the used storage and available storage of all the volumes on a storage VM.


- **Client data:** The amount of space used by client data (either physical or logical).
 - Beginning with ONTAP 9.13.1, the capacity used by client data is referred to as **Logical Used**, and the capacity used by Snapshot copies is displayed separately.
 - In ONTAP 9.12.1 and earlier, the capacity used by client data added to the capacity used by Snapshot copies is referred to as **Logical Used**.
- **Committed:** The amount of committed capacity for a local tier.
- **Data reduction:**
 - Beginning with ONTAP 9.13.1, data reduction ratios are displayed as follows:
 - The data reduction value displayed on the **Capacity** panel is the ratio of logical used space compared to physical used space without considering the significant reductions achieved when using storage efficiency features, such as Snapshot copies.
 - When you display the details panel, you see both the ratio that was displayed on the overview panel and the overall ratio of all logical used space compared to physical used space. Referred to as **With Snapshot copies**, this value includes the benefits derived from using Snapshot copies and other storage efficiency features.
 - In ONTAP 9.12.1 and earlier, data reduction ratios are displayed as follows:
 - The data reduction value displayed on the **Capacity** panel is the overall ratio of all logical used space compared to physical used space, and it includes the benefits derived from using Snapshot copies and other storage efficiency features.
 - When you display the details panel, you see both the **Overall** ratio that was displayed on the overview panel and ratio of logical used space used only by client data compared to physical used space used only by client data, referred to as **Without Snapshot copies and clones**.
- **Logical used:**
 - Beginning with ONTAP 9.13.1, the capacity used by client data is referred to as **Logical Used**, and the capacity used by Snapshot copies is displayed separately.
 - In ONTAP 9.12.1 and earlier, the capacity used by client data added to capacity used by Snapshot copies is referred to as **Logical Used**.
- **Logical used %:** The percentage of the current logical used capacity compared to the provisioned size, excluding Snapshot reserves. This value can be greater than 100%, because it includes efficiency savings in the volume.
- **Maximum capacity:** The maximum amount of space allocated for volumes on a storage VM.
- **Physical used:** The amount of capacity used in the physical blocks of a volume or local tier.
- **Physical used %:** The percentage of capacity used in the physical blocks of a volume compared to the provisioned size.
- **Provisioned capacity:** A file system (volume) that has been allocated from a Cloud Volumes ONTAP system and is ready to store user or application data.
- **Reserved:** The amount of space reserved for already provisioned volumes in a local tier.
- **Used:** The amount of space that contains data.
- **Used and reserved:** The sum of physical used and reserved space.

Capacity of a storage VM

The maximum capacity of a storage VM is determined by the total allocated space for volumes plus the remaining unallocated space.

- The allocated space for volumes is the sum of the used capacity and the sum of available capacity of FlexVol volumes, FlexGroup volumes, and FlexCache volumes.
- The capacity of volumes is included in the sums, even when they are restricted, offline, or in the recovery queue after deletion.
- If volumes are configured with auto-grow, the maximum autosize value of the volume is used in the sums. Without auto-grow, the actual capacity of the volume is used in the sums.

The following chart explains how the measurement of the capacity across volumes relates to the maximum capacity limit.



Beginning with ONTAP 9.13.1, cluster administrators can [enable a maximum capacity limit for a storage VM](#). However, storage limits cannot be set for a storage VM that contains volumes that are for data protection, in a SnapMirror relationship, or in a MetroCluster configuration. Also, quotas cannot be configured to exceed the maximum capacity of a storage VM.

After the maximum capacity limit is set, it cannot be changed to a size that is less than the currently allocated capacity.

When a storage VM reaches its maximum capacity limit, certain operations cannot be performed. System Manager provides suggestions for next steps in [Insights](#).

Capacity measurement units

System Manager calculates storage capacity based on binary units of 1024 (2^{10}) bytes.

- Beginning with ONTAP 9.10.1, storage capacity units are displayed in System Manager as KiB, MiB, GiB, TiB, and PiB.

- In ONTAP 9.10.0 and earlier, these units are displayed in System Manager as KB, MB, GB, TB, and PB.



The units used in System Manager for throughput continue to be KB/s, MB/s, GB/s, TB/s, and PB/s for all releases of ONTAP.

Capacity unit displayed in System Manager for ONTAP 9.10.0 and earlier	Capacity unit displayed in System Manager for ONTAP 9.10.1 and later	Calculation	Value in bytes
KB	KiB	1024	1024 bytes
MB	MiB	1024 * 1024	1,048,576 bytes
GB	GiB	1024 * 1024 * 1024	1,073,741,824 bytes
TB	TiB	1024 * 1024 * 1024 * 1024	1,099,511,627,776 bytes
PB	PiB	1024 * 1024 * 1024 * 1024 * 1024	1,125,899,906,842,624 bytes

Related information

[Monitor capacity in System Manager](#)

[Logical space reporting and enforcement for volumes](#)

Temperature-sensitive storage efficiency overview

ONTAP provides temperature-sensitive storage efficiency benefits by assessing how often your volume's data is accessed and mapping that frequency to the degree of compression applied to that data. For cold data that is accessed infrequently, larger data blocks are compressed, and for hot data, which is accessed frequently and is overwritten more often, smaller data blocks are compressed, making the process more efficient.

Temperature-sensitive storage efficiency (TSSE) is introduced in ONTAP 9.8 and is enabled automatically on newly created thinly provisioned AFF volumes. You can enable temperature-sensitive storage efficiency on existing AFF volumes and on thinly provisioned non-AFF DP volumes.

Introduction of "default" and "efficient" modes

Beginning with ONTAP 9.10.1, *default* and *efficient* volume-level storage efficiency modes are introduced for AFF systems only. The two modes provide a choice between file compression (default), which is the default mode when creating new AFF volumes, or temperature-sensitive storage efficiency (efficient), which enables temperature-sensitive storage efficiency. With ONTAP 9.10.1, [temperature-sensitive storage efficiency must be explicitly set](#) to enable auto adaptive compression. However, other storage efficiency features like data-compaction, auto dedupe schedule, inline deduplication, cross volume inline deduplication, and cross volume background deduplication are enabled by default on AFF platforms for both default and efficient modes.

Both storage efficiency modes (default and efficient) are supported on FabricPool-enabled aggregates and with all tiering policy types.

Temperature-sensitive storage efficiency enabled on C-Series platforms

Temperature-sensitive storage efficiency is enabled by default on AFF C-Series platforms and when migrating volumes from a non-TSSE platform to a TSSE-enabled C-Series platform using volume move or SnapMirror with the following releases installed on the destination:

- ONTAP 9.12.1P4 and later
- ONTAP 9.13.1 and later

For more information, see [Storage efficiency behavior with volume move and SnapMirror operations](#).

For existing volumes, temperature-sensitive storage efficiency is not enabled automatically, however, you can [modify the storage efficiency mode](#) manually to change to efficient mode.



Once you change the storage efficiency mode to efficient you cannot change it back.

Improved storage efficiency with sequential packing of contiguous physical blocks

Beginning with ONTAP 9.13.1, temperature-sensitive storage efficiency adds sequential packing of contiguous physical blocks to further improve storage efficiency. Volumes that have temperature-sensitive storage efficiency enabled automatically have sequential packing enabled when you upgrade systems to ONTAP 9.13.1. After sequential packing is enabled, you must [manually repack existing data](#).

Upgrade considerations

When upgrading to ONTAP 9.10.1 and later, existing volumes are assigned a storage efficiency mode based on the type of compression currently enabled on the volumes. During an upgrade, volumes with compression enabled are assigned the default mode, and volumes with temperature-sensitive storage efficiency enabled are assigned the efficient mode. If compression is not enabled, storage efficiency mode remains blank.

CPU or dedicated offload processor storage efficiency

Beginning with ONTAP 9.15.1, ONTAP provides storage efficiency and data compaction on AFF A70, AFF A90, and AFF A1K platforms. Depending on the platform, compression is performed using either the main CPU or with a dedicated offload processor. Storage efficiency is enabled automatically and requires no configuration. Storage efficiency is enabled by default on newly created thin-provisioned volumes and is applied to existing data, including volumes moved from other platforms to AFF A70, AFF A90, or AFF A1K platforms.

Temperature-sensitive storage efficiency is not applied on AFF A70, AFF A90, and AFF A1K platforms. Compression is not based on hot or cold data on these platforms, so compression begins without waiting for data to become cold.

Storage efficiency on AFF A70, AFF A90, and AFF A1K platforms uses sequential packing of contiguous physical blocks to further improve storage efficiency for compressed data.

Data that's migrated to AFF A70, AFF A90, or AFF A1K platforms using volume move or SnapMirror technology is converted automatically to 32k inline compression. Storage efficiency for data migrated from an AFF A70, A90, or A1K platform to an earlier platform is transformed automatically based on the destination platform's efficiency capabilities.

For information about upgrading a controller to an AFF A70, AFF A90, or AFF A1K, see the [ONTAP Hardware](#)

Security

Client authentication and authorization

ONTAP uses standard methods to secure client and administrator access to storage and to protect against viruses. Advanced technologies are available for encryption of data at rest and for WORM storage.

ONTAP authenticates a client machine and user by verifying their identities with a trusted source. ONTAP authorizes a user to access a file or directory by comparing the user's credentials with the permissions configured on the file or directory.

Authentication

You can create local or remote user accounts:

- A local account is one in which the account information resides on the storage system.
- A remote account is one in which account information is stored on an Active Directory domain controller, an LDAP server, or a NIS server.

ONTAP uses local or external name services to look up host name, user, group, netgroup, and name mapping information. ONTAP supports the following name services:

- Local users
- DNS
- External NIS domains
- External LDAP domains

A *name service switch table* specifies the sources to search for network information and the order in which to search them (providing the equivalent functionality of the `/etc/nsswitch.conf` file on UNIX systems). When a NAS client connects to the SVM, ONTAP checks the specified name services to obtain the required information.

Kerberos support Kerberos is a network authentication protocol that provides “strong authentication” by encrypting user passwords in client-server implementations. ONTAP supports Kerberos 5 authentication with integrity checking (krb5i) and Kerberos 5 authentication with privacy checking (krb5p).

Authorization

ONTAP evaluates three levels of security to determine whether an entity is authorized to perform a requested action on files and directories residing on an SVM. Access is determined by the effective permissions after evaluation of the security levels:

- Export (NFS) and share (SMB) security

Export and share security applies to client access to a given NFS export or SMB share. Users with administrative privileges can manage export and share-level security from SMB and NFS clients.

- Storage-Level Access Guard file and directory security

Storage-Level Access Guard security applies to SMB and NFS client access to SVM volumes. Only NTFS access permissions are supported. For ONTAP to perform security checks on UNIX users for access to data on volumes for which Storage-Level Access Guard has been applied, the UNIX user must map to a Windows user on the SVM that owns the volume.

- NTFS, UNIX, and NFSv4 native file-level security

Native file-level security exists on the file or directory that represents the storage object. You can set file-level security from a client. File permissions are effective regardless of whether SMB or NFS is used to access the data.

Authentication with SAML

ONTAP supports Security Assertion Markup Language (SAML) for authentication of remote users. Several popular identity providers (IdPs) are supported. For more information on supported IdPs and instructions for enabling SAML authentication, refer to [Configure SAML authentication](#).

OAuth 2.0 with ONTAP REST API clients

Support for the Open Authorization (OAuth 2.0) framework is available beginning with ONTAP 9.14. You can only use OAuth 2.0 to make authorization and control access decisions when the client uses the REST API to access ONTAP. However, you can configure and enable the feature with any of the ONTAP administrative interfaces, including the CLI, System Manager, and REST API.

The standard OAuth 2.0 capabilities are supported along with several popular authorization servers. You can further enhance ONTAP security by using sender-constrained access tokens based on Mutual TLS. And there is a wide variety of authorization options available, including self-contained scopes as well as integration with the ONTAP REST roles and local user definitions. See [Overview of the ONTAP OAuth 2.0 implementation](#) for more information.

Administrator authentication and RBAC

Administrators use local or remote login accounts to authenticate themselves to the cluster and SVM. Role-Based Access Control (RBAC) determines the commands to which an administrator has access.

Authentication

You can create local or remote cluster and SVM administrator accounts:

- A local account is one in which the account information, public key, or security certificate resides on the storage system.
- A remote account is one in which account information is stored on an Active Directory domain controller, an LDAP server, or a NIS server.

Except for DNS, ONTAP uses the same name services to authenticate administrator accounts as it uses to authenticate clients.

RBAC

The *role* assigned to an administrator determines the commands to which the administrator has access. You assign the role when you create the account for the administrator. You can assign a different role or define custom roles as needed.

Virus scanning

You can use integrated antivirus functionality on the storage system to protect data from being compromised by viruses or other malicious code. ONTAP virus scanning, called *Vscan*, combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when.

Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. The *ONTAP Antivirus Connector*, provided by NetApp and installed on the external server, handles communications between the storage system and the antivirus software.

- You can use *on-access scanning* to check for viruses when clients open, read, rename, or close files over SMB. File operation is suspended until the external server reports the scan status of the file. If the file has already been scanned, ONTAP allows the file operation. Otherwise, it requests a scan from the server.

On-access scanning is not supported for NFS.


- You can use *on-demand scanning* to check files for viruses immediately or on a schedule. You might want to run scans only in off-peak hours, for example. The external server updates the scan status of the checked files, so that file-access latency for those files (assuming they have not been modified) is typically reduced when they are next accessed over SMB.

You can use on-demand scanning for any path in the SVM namespace, even for volumes that are exported only through NFS.

You typically enable both scanning modes on an SVM. In either mode, the antivirus software takes remedial action on infected files based on your settings in the software.

Virus scanning in disaster recovery and MetroCluster configurations

For disaster recovery and MetroCluster configurations, you must set up separate Vscan servers for the local and partner clusters.



The storage system offloads virus scanning operations to external servers hosting antivirus software from third-party vendors.

Encryption

ONTAP offers both software- and hardware-based encryption technologies to ensure that data at rest cannot be read if the storage medium is repurposed, returned, misplaced, or stolen.

ONTAP is compliant with the Federal Information Processing Standards (FIPS) 140-2 for all SSL connections. You can use the following encryption solutions:

- Hardware solutions:

- NetApp Storage Encryption (NSE)

NSE is a hardware solution that uses self-encrypting drives (SEDs).

- NVMe SEDs

ONTAP provides full disk encryption for NVMe SEDs that do not have FIPS 140-2 certification.

- Software solutions:

- NetApp Aggregate Encryption (NAE)

NAE is a software solution that enables encryption of any data volume on any drive type where it is enabled with unique keys for each aggregate.

- NetApp Volume Encryption (NVE)

NVE is a software solution that enables encryption of any data volume on any drive type where it is

enabled with a unique key for each volume.

Use both software (NAE or NVE) and hardware (NSE or NVMe SED) encryption solutions to achieve double encryption at rest. Storage efficiency is not affected by NAE or NVE encryption.

NetApp Storage Encryption

NetApp Storage Encryption (NSE) supports SEDs that encrypt data as it is written. The data cannot be read without an encryption key stored on the disk. The encryption key, in turn, is accessible only to an authenticated node.

On an I/O request, a node authenticates itself to an SED using an authentication key retrieved from an external key management server or Onboard Key Manager:

- The external key management server is a third-party system in your storage environment that serves authentication keys to nodes using the Key Management Interoperability Protocol (KMIP).
- The Onboard Key Manager is a built-in tool that serves authentication keys to nodes from the same storage system as your data.

NSE supports self-encrypting HDDs and SSDs. You can use NetApp Volume Encryption with NSE to double encrypt data on NSE drives.



If you are using NSE on a system with a Flash Cache module, you should also enable NVE or NAE. NSE does not encrypt data that resides on the Flash Cache module.

NVMe self-encrypting drives

NVMe SEDs do not have FIPS 140-2 certification, however, these disks use AES 256-bit transparent disk encryption to protect data at rest.

Data encryption operations, such as generating an authentication key, are performed internally. The authentication key is generated the first time the disk is accessed by the storage system. After that, the disks protect data at rest by requiring storage system authentication each time data operations are requested.

NetApp Aggregate Encryption

NetApp Aggregate Encryption (NAE) is a software-based technology for encrypting all data on an aggregate. A benefit of NAE is that volumes are included in aggregate level deduplication, whereas NVE volumes are excluded.

With NAE enabled, the volumes within the aggregate can be encrypted with aggregate keys.

Beginning with ONTAP 9.7, newly created aggregates and volumes are encrypted by default when you have the [NVE license](#) and onboard or external key management.

NetApp Volume Encryption

NetApp Volume Encryption (NVE) is a software-based technology for encrypting data at rest one volume at a time. An encryption key accessible only to the storage system ensures that volume data cannot be read if the underlying device is separated from the system.

Both data, including Snapshot copies, and metadata are encrypted. Access to the data is given by a unique XTS-AES-256 key, one per volume. A built-in Onboard Key Manager secures the keys on the same system with your data.

You can use NVE on any type of aggregate (HDD, SSD, hybrid, array LUN), with any RAID type, and in any supported ONTAP implementation, including ONTAP Select. You can also use NVE with NetApp Storage Encryption (NSE) to double encrypt data on NSE drives.

When to use KMIP servers Although it is less expensive and typically more convenient to use the Onboard Key Manager, you should set up KMIP servers if any of the following are true:

- Your encryption key management solution must comply with Federal Information Processing Standards (FIPS) 140-2 or the OASIS KMIP standard.
- You need a multi-cluster solution. KMIP servers support multiple clusters with centralized management of encryption keys.

KMIP servers support multiple clusters with centralized management of encryption keys.

- Your business requires the added security of storing authentication keys on a system or in a location different from the data.

KMIP servers stores authentication keys separately from your data.

Related information


[FAQ - NetApp Volume Encryption and NetApp Aggregate Encryption](#)

WORM storage

SnapLock is a high-performance compliance solution for organizations that use *write once, read many (WORM)* storage to retain critical files in unmodified form for regulatory and governance purposes.

A single license entitles you to use SnapLock in strict *Compliance mode*, to satisfy external mandates like SEC Rule 17a-4, and a looser *Enterprise mode*, to meet internally mandated regulations for the protection of digital assets. SnapLock uses a tamper-proof *ComplianceClock* to determine when the retention period for a WORM file has elapsed.

You can use *SnapLock for SnapVault* to WORM-protect Snapshot copies on secondary storage. You can use SnapMirror to replicate WORM files to another geographic location for disaster recovery and other purposes.



SnapLock uses a tamper-proof ComplianceClock to determine when the retention period for a WORM file has elapsed.

ONTAP and VMware vSphere

You can integrate ONTAP and related NetApp products with VMware vSphere. There are several options available depending on your technology environment and business

needs.

Selected concepts and terminology

As you begin to explore using ONTAP and related NetApp products in a VMware environment, it's helpful to first be familiar with some of the key terminology and concepts.

Logical unit number

A LUN is a number used to identify a *logical unit* within a Storage Area Network (SAN). These addressable devices are typically logical disks accessed through the Small Computer System Interface (SCSI) protocol or one of its encapsulated derivatives.

VMware vSphere virtual volume

A virtual volume (vVol) provides a volume-level abstraction for storage used by a virtual machine. It includes several benefits and provides an alternative to using a traditional LUN.

Persistent reservations

Persistent reservations are supported with SCSI-3 and an improvement over the earlier SCSI-2 reservations. They allow multiple client initiators to communicate with a single target while locking out other nodes. The reservations can persist even if the bus is reset for error recovery.



Beginning with ONTAP 9.15.1, you can create a persistent reservation for a virtual volume using SCSI-3. This feature is only supported using ONTAP Tools for VMware vSphere 9 with a Windows Server Failover Cluster (WSFC).

Windows Server Failover Clustering

Microsoft WSFC is a feature of the Windows Server operating system providing fault tolerance and high availability. A set of server nodes (physical or virtual) are joined together as a cluster to provide resiliency in case of a failure. WSFC is commonly used to deploy infrastructure services include database, file, and namespaces servers.

VMware vSphere Storage APIs - Storage Awareness

VASA is a set of APIs providing integration of the storage arrays with vCenter for management and administration. The architecture is based on several components including the *VASA Provider* which handles communication between VMware vSphere and the storage systems. With ONTAP, the provider is implemented as part of ONTAP tools for VMware vSphere.

VMware vSphere Storage APIs - Array Integration

VAAI is a set of APIs that enable communication between VMware vSphere ESXi hosts and the storage devices. The API includes a set of primitive operations used by the hosts to offload storage operations to the array. VAAI can provide significant performance improvements for storage-intensive tasks.

NetApp SnapCenter

SnapCenter is a centralized and scalable platform providing data protection for applications, databases, host file systems, and virtual machines using ONTAP storage systems. It leverages the native ONTAP technologies including Snapshot, SnapRestore, FlexClone, SnapMirror, and SnapVault.

NetApp plug-ins and related technologies

NetApp provides robust support for integrating ONTAP and related products with VMware vSphere technologies.

ONTAP tools for VMware vSphere

ONTAP tools for VMware vSphere is a set of tools for integrating ONTAP and vSphere. It implements the provider functionality of the VASA API framework. ONTAP tools also includes the vCenter plug-in, a storage replication adapter (SRA) for VMware Site Recovery Manager, and a REST API server you can use to build automation applications.

NFS Plug-In for VMware VAAI

The NetApp NFS Plug-In for VMware VAAI provides access to the VAAI features. The plug-in can be installed on ESXi hosts and allows the hosts to leverage VAAI with the NFS datastores on ONTAP. It provides several operations including cloning, space reservations, and snapshot offloading.

VMware Site Recovery Manager

VMware Site Recovery Manager (SRM) provides a disaster recovery capability. SRM integrates with ONTAP tools for VMware vSphere to access and leverage the ONTAP data management functionality.

vSphere Metro Storage Cluster

vSphere Metro Storage Cluster (vMSC) is a technology that enables and supports vSphere in a *stretched cluster* deployment. vMSC solutions are supported with NetApp MetroCluster and SnapMirror active sync (formerly SMBC). These solutions provide enhanced business continuity in the case of domain failure. The resiliency model is based on your specific configuration choices.

SnapCenter Plug-in for VMware vSphere

The SnapCenter Plug-in for VMware vSphere (SCV) is a Linux-based virtual appliance you can deploy together with the SnapCenter server or as a standalone application. In both cases, SCV provides backup and restore operations for VMs, datastores, and VMDKs. The operations are fast, space-efficient, crash-consistent, and VM-consistent.

Get more information

There are several additional resources available to help you prepare to deploy ONTAP in a VMware vSphere environment.

- [ONTAP tools for VMware vSphere documentation](#)
- [Enterprise applications: VMware vSphere with ONTAP](#)
- [NetApp KB: What are SCSI Reservations and SCSI Persistent Reservations?](#)
- [SnapCenter Plug-in for VMware vSphere documentation](#)

Application aware data management

Application aware data management enables you to describe the application that you want to deploy over ONTAP in terms of the application, rather than in storage terms. The application can be configured and ready to serve data quickly with minimal inputs by using System Manager and REST APIs.

The application aware data management feature provides a way to set up, manage, and monitor storage at the level of individual applications. This feature incorporates relevant ONTAP best practices to optimally provision

applications, with balanced placement of storage objects based on desired performance service levels and available system resources.

The application aware data management feature includes a set of application templates, with each template consisting of a set of parameters that collectively describe the configuration of an application. These parameters, which are often preset with default values, define the characteristics that an application administrator could specify for provisioning storage on an ONTAP system, such as database sizes, service levels, protocol access elements such as LIFs as well as local protection criteria and remote protection criteria. Based on the specified parameters, ONTAP configures storage entities such as LUNs and volumes with appropriate sizes and service levels for the application.

You can perform the following tasks for your applications:

- Create applications by using the application templates
- Manage the storage associated with the applications
- Modify or delete the applications
- View applications
- Manage Snapshot copies of the applications
- Create [consistency groups](#) to provide data protection capabilities by selecting multiple LUNs in the same or in different volumes

FabricPool

Many NetApp customers have significant amounts of stored data that is rarely accessed. We call that *cold* data. Customers also have data that is frequently accessed, which we call *hot* data. Ideally, you want to keep your hot data on your fastest storage for best performance. Cold data can move to slower storage as long as it is immediately available if needed. But how do you know which parts of your data are hot and which are cold?

FabricPool is an ONTAP feature that automatically moves data between a high-performance local tier (aggregate) and a cloud tier based on access patterns. Tiering frees up expensive local storage for hot data while keeping cold data readily available from low-cost object storage in the cloud. FabricPool constantly monitors data access and moves data between tiers for best performance and maximum savings.

Using FabricPool to tier cold data to the cloud is one of the easiest ways to gain cloud efficiency and create a hybrid cloud configuration. FabricPool works at the storage block level, so it works with both file and LUN data.

But FabricPool is not just for tiering on-premises data to the cloud. Many customers use FabricPool in Cloud Volumes ONTAP to tier cold data from more-expensive cloud storage to lower-cost object storage within the cloud provider. Beginning with ONTAP 9.8, you can capture analytics on FabricPool-enabled volumes with [File System Analytics](#) or [temperature-sensitive storage efficiency](#).

The applications using the data are not aware that data is tiered, so no changes to your applications are needed. Tiering is fully automatic, so there is no ongoing administration needed.

You can store cold data in object storage from one of the major cloud providers. Or choose NetApp StorageGRID to keep your cold data in your own private cloud for highest performance and complete control over your data.

Related information

[FabricPool System Manager doc](#)

[BlueXP tiering](#)

[FabricPool playlist on NetApp TechComm TV](#)

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.