



Concepts

Element Software

NetApp
April 17, 2024

This PDF was generated from https://docs.netapp.com/us-en/element-software/concepts/concept_intro_product_overview.html on April 17, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Concepts 1
 - Find more information 1
 - Product overview 1
 - SolidFire architecture overview 2
- Nodes 7
- Clusters 8
- Security 10
- Accounts and permissions 12
- Storage 14
- Data protection 16
- Performance and quality of service 20

Concepts

Learn basic concepts related to Element software.

- [Product overview](#)
- [SolidFire architecture overview](#)
- [Nodes](#)
- [Clusters](#)
- [Security](#)
- [Accounts and permissions](#)
- [Volumes](#)
- [Data protection](#)
- [Performance and quality of service](#)

Find more information

- [SolidFire all-flash storage overview](#)
- [SolidFire and Element Software Documentation](#)

Product overview

A SolidFire all-flash storage system is comprised of discrete hardware components (drive and nodes) that are combined into a single pool of storage resources. This unified cluster presents as a single storage system for use by external clients and is managed with NetApp Element software.

Using the Element interface, API, or other management tools, you can monitor SolidFire cluster storage capacity and performance, and manage storage activity across a multi-tenant infrastructure.

SolidFire features

A Solidfire system provides the following features:

- Offers high performance storage for your large scale, private cloud infrastructure
- Provides a flexible scale that lets you meet changing storage needs
- Uses an API-driven storage management Element software interface
- Guarantees performance using Quality of Service policies
- Includes automatic load balancing across all nodes in the cluster
- Rebalances clusters automatically when nodes are added or subtracted

SolidFire deployment

Use storage nodes provided by NetApp and integrated with NetApp Element software.

Find more information


- [SolidFire all-flash storage overview](#)
- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

SolidFire architecture overview

A SolidFire all-flash storage system is comprised of discrete hardware components (drive and nodes) that are combined into a pool of storage resources with NetApp Element software running independently on each node. This single storage system is managed as a single entity by using the Element software UI, API and other management tools.

A SolidFire storage system includes the following hardware components:

- **Cluster:** The hub of the SolidFire storage system that is a collection of nodes.
- **Nodes:** The hardware components grouped into a cluster. There are two types of nodes:
 - Storage nodes, which are servers containing a collection of drives
 - Fibre Channel (FC) nodes, which you use to connect to FC clients
- **Drives:** Used in storage nodes to store data for the cluster. A storage node contains two types of drives:
 - Volume metadata drives store information that defines the volumes and other objects within a cluster.
 - Block drives store data blocks for volumes.



You can manage, monitor, and update the system using the Element web UI and other compatible tools:

- [SolidFire software interfaces](#)
- [SolidFire Active IQ](#)
- [Management node for Element software](#)
- [Management services](#)

Common URLs

These are the common URLs you use with a SolidFire all-flash storage system:

URL	Description
<code>https://[storage cluster MVIP address]</code>	Access the NetApp Element software UI.
<code>https://activeiq.solidfire.com</code>	Monitor data and receive alerts to any performance bottlenecks or potential system issues.
<code>https://[management node IP address]</code>	Access NetApp Hybrid Cloud Control to upgrade your storage installation and update management services.
<code>https://[IP address]:442</code>	From the per-node UI, access network and cluster settings and utilize system tests and utilities. Learn more.

URL	Description
<code>https://[management node IP address]/mnode</code>	Use management services REST API and other functionality from the management node. Learn more.
<code>https://[management node IP address]:9443</code>	Register the vCenter Plug-in package in the vSphere Web Client. Learn more.

Find more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

SolidFire software interfaces

You can manage a SolidFire storage system using different NetApp Element software interfaces and integration utilities.

Options

- [NetApp Element software user interface](#)
- [NetApp Element software API](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp Hybrid Cloud Control](#)
- [Management node UIs](#)
- [Additional integration utilities and tools](#)

NetApp Element software user interface

Enables you to set up Element storage, monitor cluster capacity and performance, and manage storage activity across a multi-tenant infrastructure. Element is the storage operating system at the heart of a SolidFire cluster. Element software runs independently on all nodes in the cluster and enables the nodes of the cluster to combine resources that are presented as a single storage system to external clients. Element software is responsible for all cluster coordination, scale and management of the system as a whole. The software interface is built upon the Element API.

[Manage storage with Element software](#)

NetApp Element software API

Enables you to use a set of objects, methods, and routines to manage Element storage. The Element API is based on the JSON-RPC protocol over HTTPS. You can monitor API operations in the Element UI by enabling the API Log; this enables you to see the methods that are being issued to the system. You can enable both requests and responses to see how the system replies to the methods that are issued.

[Manage storage with the Element API](#)

NetApp Element Plug-in for vCenter Server

Enables you to configure and manage storage clusters running Element software using an alternative interface for the Element UI within VMware vSphere.

NetApp Hybrid Cloud Control

Enables you to upgrade Element storage and management services and manage storage assets using the NetApp Hybrid Cloud Control interface.

[Manage and monitor storage with NetApp Hybrid Cloud Control overview](#)

Management node UIs

The management node contains two UIs: a UI for managing REST-based services and a per-node UI for managing network and cluster settings and operating system tests and utilities. From the REST API UI, you can access a menu of service-related APIs that control service-based system functionality from the management node.

Additional integration utilities and tools

Although you typically manage your storage with NetApp Element, NetApp Element API, and NetApp Element Plug-in for vCenter Server, you can use additional integration utilities and tools to access storage.

Element CLI

[Element CLI](#) enables you to control a SolidFire storage system using a command-line interface without having to use the Element API.

Element PowerShell Tools

[Element PowerShell Tools](#) enable you to use a collection of Microsoft Windows PowerShell functions that use the Element API to manage a SolidFire storage system.

Element SDKs

[Element SDKs](#) enable you to manage your SolidFire cluster using these tools:

- Element Java SDK: Enables programmers to integrate the Element API with the Java programming language.
- Element .NET SDK: Enables programmers to integrate the Element API with the .NET programming platform.
- Element Python SDK: Enables programmers to integrate the Element API with the Python programming language.

SolidFire Postman API testing suite

Enables programmers to use a collection of [Postman](#) functions that test Element API calls.

SolidFire Storage Replication Adapter

[SolidFire Storage Replication Adapter](#) integrates with the VMware Site Recovery Manager (SRM) to enable communication with replicated SolidFire storage clusters and execute supported workflows.

SolidFire vRO

[SolidFire vRO](#) provides a convenient way to use the Element API to administer your SolidFire storage system

with VMware vRealize Orchestrator.

SolidFire VSS Provider

[SolidFire VSS Provider](#) integrates VSS shadow copies with Element snapshots and clones.

Find more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

SolidFire Active IQ

[SolidFire Active IQ](#) is a web-based tool that provides continually updated historical views of cluster-wide data. You can set up alerts for specific events, thresholds, or metrics. SolidFire Active IQ enables you to monitor system performance and capacity, as well as stay informed about cluster health.

You can find the following information about your system in SolidFire Active IQ:

- Number of nodes and status of the nodes: healthy, offline, or fault
- Graphical representation of CPU, memory usage, and node throttling
- Details about the node, such as serial number, slot location in the chassis, model, and version of NetApp Element software running on the storage node
- CPU and storage-related information about the virtual machines

To learn about SolidFire Active IQ, see the [SolidFire Active IQ documentation](#).

For more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp Support Site > Tools for Active IQ](#)

Management node for Element software

The [management node \(mNode\)](#) is a virtual machine that runs in parallel with one or more Element software-based storage clusters. It is used to upgrade and provide system services including monitoring and telemetry, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting.

The management node interacts with a storage cluster to perform management actions, but is not a member of the storage cluster. Management nodes periodically collect information about the cluster through API calls and report this information to Active IQ for remote monitoring (if enabled). Management nodes are also responsible for coordinating software upgrades of the cluster nodes.

As of the Element 11.3 release, the management node functions as a microservice host, allowing for quicker updates of select software services outside of major releases. These microservices or [management services](#) are updated frequently as service bundles.

Management services for SolidFire all-flash storage

As of the Element 11.3 release, **management services** are hosted on the [management node](#), allowing for quicker updates of select software services outside of major releases.

Management services provide central and extended management functionality for SolidFire all-flash storage. These services include [NetApp Hybrid Cloud Control](#), Active IQ system telemetry, logging, and service updates, as well as the QoSSIOC service for the Element Plug-in for vCenter.



Learn more about [management services releases](#).

Nodes

Nodes are hardware or virtual resources that are grouped into a cluster to provide block storage and compute capabilities.

NetApp Element software defines various node roles for a cluster. The types of node roles are the following:

- [Management node](#)
- [Storage node](#)
- [Fibre Channel node](#)

[Nodes states](#) vary depending on cluster association.

Management node

A management node is a virtual machine used to upgrade and provide system services including monitoring and telemetry, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting. [Learn more](#)

Storage node

A SolidFire storage node is a server containing a collection of drives that communicate with each other through the Bond10G network interface. Drives in the node contain block and metadata space for data storage and data management. Each node contains a factory image of NetApp Element software.

Storage nodes have the following characteristics:

- Each node has a unique name. If a node name is not specified by an administrator, it defaults to SF-XXXX, where XXXX is four random characters generated by the system.
- Each node has its own high-performance non-volatile random access memory (NVRAM) write cache to improve overall system performance and reduce write latency.
- Each node is connected to two networks, storage and management, each with two independent links for redundancy and performance. Each node requires an IP address on each network.
- You can create a cluster with new storage nodes, or add storage nodes to an existing cluster to increase storage capacity and performance.
- You can add or remove nodes from the cluster at any time without interrupting service.

Fibre Channel node

SolidFire Fibre Channel nodes provide connectivity to a Fibre Channel switch, which you can connect to Fibre Channel clients. Fibre Channel nodes act as a protocol converter between the Fibre Channel and iSCSI protocols; this enables you to add Fibre Channel connectivity to any new or existing SolidFire cluster.

Fibre Channel nodes have the following characteristics:

- Fibre Channel switches manage the state of the fabric, providing optimized interconnections.
- The traffic between two ports flows through the switches only; it is not transmitted to any other port.
- Failure of a port is isolated and does not affect operation of other ports.
- Multiple pairs of ports can communicate simultaneously in a fabric.

Node states of operation

A node can be in one of several states depending on the level of configuration.

- **Available**

The node has no associated cluster name and is not yet part of a cluster.

- **Pending**

The node is configured and can be added to a designated cluster.

Authentication is not required to access the node.

- **Pending Active**

The system is in the process of installing compatible Element software on the node. When complete, the node will move to the Active state.

- **Active**

The node is participating in a cluster.

Authentication is required to modify the node.

In each of these states, some fields are read only.

Find more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Clusters

A cluster is the hub of a SolidFire storage system and is made up of a collection of nodes. You must have at least four nodes in a cluster for SolidFire storage efficiencies to be realized. A cluster appears on the network as a single logical group and can then be accessed as block storage.

Creating a new cluster initializes a node as communications owner for a cluster and establishes network communications for each node in the cluster. This process is performed only once for each new cluster. You can create a cluster using the Element UI or the API.

You can scale out a cluster by adding additional nodes. When you add a new node, there is no interruption of service and the cluster automatically uses the performance and capacity of the new node.

Administrators and hosts can access the cluster using virtual IP addresses. Any node in the cluster can host the virtual IP addresses. The management virtual IP (MVIP) enables cluster management through a 1GbE connection, while the storage virtual IP (SVIP) enables host access to storage through a 10GbE connection. These virtual IP addresses enable consistent connections regardless of the size or makeup of a SolidFire cluster. If a node hosting a virtual IP address fails, another node in the cluster begins hosting the virtual IP address.



Beginning in Element version 11.0, nodes can be configured with IPv4, IPv6, or both addresses for their management network. This applies to both storage nodes and management nodes, except for management node 11.3 and later which does not support IPv6. When creating a cluster, only a single IPv4 or IPv6 address can be used for the MVIP and the corresponding address type must be configured on all nodes.

More on clusters

- [Authoritative storage clusters](#)
- [Rule of thirds](#)
- [Stranded capacity](#)
- [Storage efficiency](#)
- [Storage cluster quorum](#)

Authoritative storage clusters

The authoritative storage cluster is the storage cluster that NetApp Hybrid Cloud Control uses to authenticate users.

If your management node only has one storage cluster, then it is the authoritative cluster. If your management node has two or more storage clusters, one of those clusters is assigned as the authoritative cluster and only users from that cluster can log into NetApp Hybrid Cloud Control. To find out which cluster is the authoritative cluster, you can use the `GET /mnode/about` API. In the response, the IP address in the `token_url` field is the management virtual IP address (MVIP) of the authoritative storage cluster. If you attempt to log into NetApp Hybrid Cloud Control as a user that is not on the authoritative cluster, the login attempt will fail.

Many NetApp Hybrid Cloud Control features are designed to work with multiple storage clusters, but authentication and authorization have limitations. The limitation around authentication and authorization is that the user from the authoritative cluster can execute actions on other clusters tied to NetApp Hybrid Cloud Control even if they are not a user on the other storage clusters.

Before proceeding with managing multiple storage clusters, you should ensure that users defined on the authoritative clusters are defined on all other storage clusters with the same permissions. You can manage users from the [Element software user interface](#).

See [create and manage storage cluster assets](#) for more information on working with management node storage cluster assets.

Rule of thirds

When you mix storage node types in a NetApp SolidFire storage cluster, no single storage node can contain more than 33% of the total storage cluster capacity.

Stranded capacity

If a newly added node accounts for more than 50 percent of the total cluster capacity, some of the capacity of this node is made unusable ("stranded"), so that it complies with the capacity rule. This remains the case until more storage capacity is added. If a very large node is added that also disobeys the capacity rule, the previously stranded node will no longer be stranded, while the newly added node becomes stranded. Capacity should always be added in pairs to avoid this from happening. When a node becomes stranded, an appropriate cluster fault is thrown.

Storage efficiency

Netapp SolidFire storage clusters make use of deduplication, compression, and thin provisioning to reduce the amount of physical storage needed for storing a volume.

- **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.

- **Deduplication**

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

- **Thin provisioning**

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.

Storage cluster quorum

Element software creates a storage cluster from selected nodes, which maintains a replicated database of the cluster configuration. A minimum of three nodes are required to participate in the cluster ensemble to maintain quorum for cluster resiliency.

Security

When you use your SolidFire all-flash storage system, your data is protected by industry-standard security protocols.

Encryption at Rest (hardware)

All drives in storage nodes are capable of encryption leverage AES 256-bit encryption at the drive level. Each drive has its own encryption key, which is created when the drive is first initialized. When you enable the encryption feature, a cluster-wide password is created, and chunks of the password are then distributed to all nodes in the cluster. No single node stores the entire password. The password is then used to password-

protect all access to the drives. The password is needed to unlock the drive and then not needed unless power is removed from the drive or the drive is locked.

[Enabling the hardware encryption at rest feature](#) does not affect performance or efficiency on the cluster. If an encryption-enabled drive or node is removed from cluster configuration with the Element API or Element UI, encryption at rest will be disabled on the drives. After the drive is removed, the drive can be secure erased by using the `SecureEraseDrives` API method. If a physical drive or node is forcibly removed, the data remains protected by the cluster-wide password and the drive's individual encryption keys.

Encryption at Rest (software)

Another type of encryption-at-rest, software encryption-at-rest enables all data written to SSDs in a storage cluster to be encrypted. [When enabled](#), it encrypts all data written and decrypts all data read automatically in software. Software encryption at rest mirrors the Self-Encrypting Drive (SED) implementation in hardware to provide data security in the absence of SED.



For SolidFire all-flash storage clusters, software encryption at rest must be enabled during cluster creation and cannot be disabled after the cluster has been created.

Both software and hardware-based encryption-at-rest can be used independently or in combination with one another.

External key management

You can configure Element software to use a third-party KMIP-compliant key management service (KMS) to manage storage cluster encryption keys. When you enable this feature, the storage cluster's cluster-wide drive access password encryption key is managed by a KMS that you specify.

Element can use the following key management services:

- Gemalto SafeNet KeySecure
- SafeNet AT KeySecure
- HyTrust KeyControl
- Vormetric Data Security Manager
- IBM Security Key Lifecycle Manager

For more information on configuring external key management, see [the get started with external key management](#) documentation.

Multi-factor authentication

Multi-factor authentication (MFA) enables you to require users to present multiple types of evidence to authenticate with the NetApp Element web UI or storage node UI upon login. You can configure Element to accept only multi-factor authentication for logins integrating with your existing user management system and identity provider. You can configure Element to integrate with an existing SAML 2.0 identity provider which can enforce multiple authentication schemes, such as password and text message, password and email message, or other methods.

You can pair multi-factor authentication with common SAML 2.0 compatible identity providers (IdPs), such as Microsoft Active Directory Federation Services (ADFS) and Shibboleth.

To configure MFA, see [the enable multi-factor authentication](#) documentation.

FIPS 140-2 for HTTPS and data at rest encryption

NetApp SolidFire storage clusters support encryption that complies with the Federal Information Processing Standard (FIPS) 140-2 requirements for cryptographic modules. You can enable FIPS 140-2 compliance on your SolidFire cluster for both HTTPS communications and drive encryption.

When you enable FIPS 140-2 operating mode on your cluster, the cluster activates the NetApp Cryptographic Security Module (NCSM) and leverages FIPS 140-2 Level 1 certified encryption for all communication via HTTPS to the NetApp Element UI and API. You use the `EnableFeature` Element API with the `fips` parameter to enable FIPS 140-2 HTTPS encryption. On storage clusters with FIPS-compatible hardware, you can also enable FIPS drive encryption for data at rest using the `EnableFeature` Element API with the `FipsDrives` parameter.

For more information about preparing a new storage cluster for FIPS 140-2 encryption, see [Create a cluster supporting FIPS drives](#).

For more information about enabling FIPS 140-2 on an existing, prepared cluster, see [the EnableFeature Element API](#).

For more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Accounts and permissions

To administer and provide access to storage resources on your system, you'll need to set up accounts for system resources.

Using Element storage, you can create and manage the following types of accounts:

- [Administrator user accounts for the storage cluster](#)
- [User accounts for storage volume access](#)
- [Authoritative cluster user accounts for NetApp Hybrid Cloud Control](#)

Storage cluster administrator accounts

There are two types of administrator accounts that can exist in a storage cluster running NetApp Element software:

- **Primary cluster administrator account:** This administrator account is created when the cluster is created. This account is the primary administrative account with the highest level of access to the cluster. This account is analogous to a root user in a Linux system. You can change the password for this administrator account.
- **Cluster administrator account:** You can give a cluster administrator account a limited range of administrative access to perform specific tasks within a cluster. The credentials assigned to each cluster administrator account are used to authenticate API and Element UI requests within the storage system.



A local (non-LDAP) cluster administrator account is required to access active nodes in a cluster via the per-node UI. Account credentials are not required to access a node that is not yet part of a cluster.

You can [manage cluster administrator accounts](#) by creating, deleting, and editing cluster administrator accounts, changing the cluster administrator password, and configuring LDAP settings to manage system access for users.

User accounts

User accounts are used to control access to the storage resources on a NetApp Element software-based network. At least one user account is required before a volume can be created.

When you create a volume, it is assigned to an account. If you have created a virtual volume, the account is the storage container.

Here are some additional considerations:

- The account contains the CHAP authentication required to access the volumes assigned to it.
- An account can have up to 2000 volumes assigned to it, but a volume can belong to only one account.
- User accounts can be managed from the NetApp Element Management extension point.

Authoritative cluster user accounts

Authoritative cluster user accounts can authenticate against any storage asset associated with the NetApp Hybrid Cloud Control instance of nodes and clusters. With this account, you can manage volumes, accounts, access groups, and more across all clusters.

Authoritative user accounts are managed from the top right menu User Management option in NetApp Hybrid Cloud Control.

The [authoritative storage cluster](#) is the storage cluster that NetApp Hybrid Cloud Control uses to authenticate users.

All users created on the authoritative storage cluster can log into the NetApp Hybrid Cloud Control. Users created on other storage clusters *cannot* log into Hybrid Cloud Control.

- If your management node only has one storage cluster, then it is the authoritative cluster.
- If your management node has two or more storage clusters, one of those clusters is assigned as the authoritative cluster and only users from that cluster can log into NetApp Hybrid Cloud Control.

While many NetApp Hybrid Cloud Control features work with multiple storage clusters, authentication and authorization have necessary limitations. The limitation around authentication and authorization is that users from the authoritative cluster can execute actions on other clusters tied to NetApp Hybrid Cloud Control even if they are not a user on the other storage clusters. Before proceeding with managing multiple storage clusters, you should ensure that users defined on the authoritative clusters are defined on all other storage clusters with the same permissions. You can manage users from NetApp Hybrid Cloud Control.

Volume accounts

Volume-specific accounts are specific only to the storage cluster on which they were created. These accounts enable you to set permissions on specific volumes across the network, but have no effect outside of those volumes.

Volume accounts are managed within the NetApp Hybrid Cloud Control Volumes table.

Storage

Volumes

The NetApp Element storage system provisions storage using volumes. Volumes are block devices accessed over the network by iSCSI or Fibre Channel clients.

Element storage enables you to create, view, edit, delete, clone, backup or restore volumes for user accounts. You can also manage each volume on a cluster, and add or remove volumes in volume access groups.

Persistent volumes

Persistent volumes allow management node configuration data to be stored on a specified storage cluster, rather than locally with a VM, so that data can be preserved in the event of management node loss or removal. Persistent volumes are an optional yet recommended management node configuration.

An option to enable persistent volumes is included in the installation and upgrade scripts when [deploying a new management node](#). Persistent volumes are volumes on an Element software-based storage cluster that contain management node configuration information for the host management node VM that persists beyond the life of the VM. If the management node is lost, a replacement management node VM can reconnect to and recover configuration data for the lost VM.

Persistent volumes functionality, if enabled during installation or upgrade, automatically creates multiple volumes. These volumes, like any Element software-based volume, can be viewed using the Element software web UI, NetApp Element Plug-in for vCenter Server, or API, depending on your preference and installation. Persistent volumes must be up and running with an iSCSI connection to the management node to maintain current configuration data that can be used for recovery.



Persistent volumes that are associated with management services are created and assigned to a new account during installation or upgrade. If you are using persistent volumes, do not modify or delete the volumes or their associated account

Virtual volumes (vVols)

vSphere Virtual Volumes is a storage paradigm for VMware that moves much of the storage management for vSphere from the storage system to VMware vCenter. With Virtual Volumes (vVols), you can allocate storage according to the requirements of individual virtual machines.

Bindings

The NetApp Element cluster chooses an optimal protocol endpoint, creates a binding that associates the ESXi host and virtual volume with the protocol endpoint, and returns the binding to the ESXi host. After it is bound, the ESXi host can perform I/O operations with the bound virtual volume.

Protocol endpoints

VMware ESXi hosts use logical I/O proxies known as protocol endpoints to communicate with virtual volumes. ESXi hosts bind virtual volumes to protocol endpoints to perform I/O operations. When a virtual machine on the host performs an I/O operation, the associated protocol endpoint directs I/O to the virtual volume with which it is paired.

Protocol endpoints in a NetApp Element cluster function as SCSI administrative logical units. Each protocol endpoint is created automatically by the cluster. For every node in a cluster, a corresponding protocol endpoint is created. For example, a four-node cluster will have four protocol endpoints.

iSCSI is the only supported protocol for NetApp Element software. Fibre Channel protocol is not supported. Protocol endpoints cannot be deleted or modified by a user, are not associated with an account, and cannot be added to a volume access group.

Storage containers

Storage containers are logical constructs that map to NetApp Element accounts and are used for reporting and resource allocation. They pool raw storage capacity or aggregate storage capabilities that the storage system can provide to virtual volumes. A VVol datastore that is created in vSphere is mapped to an individual storage container. A single storage container has all available resources from the NetApp Element cluster by default. If more granular governance for multi-tenancy is required, multiple storage containers can be created.

Storage containers function like traditional accounts and can contain both virtual volumes and traditional volumes. A maximum of four storage containers per cluster is supported. A minimum of one storage container is required to use VVols functionality. You can discover storage containers in vCenter during VVols creation.

VASA provider

To make vSphere aware of the vVol feature on the NetApp Element cluster, the vSphere admin must register the NetApp Element VASA Provider with vCenter. The VASA provider is the out-of-band control path between vSphere and the Element cluster. It is responsible for executing requests on the Element cluster on behalf of vSphere, such as creating VMs, making VMs available to vSphere, and advertising storage capabilities to vSphere.

The VASA provider runs as part of the cluster master in Element software. The cluster master is a highly available service that fails over to any node in the cluster as needed. If the cluster master fails over, the VASA provider moves with it, ensuring high availability for the VASA provider. All provisioning and storage management tasks use the VASA provider, which handles any changes needed on the Element cluster.



For Element 12.5 and earlier, do not register more than one NetApp Element VASA provider to a single vCenter instance. Where a second NetApp Element VASA provider is added, this renders all VVOL datastores inaccessible.



VASA support for up to 10 vCenters is available as an upgrade patch if you have already registered a VASA provider with your vCenter. To install, follow the directions in the VASA39 manifest and download the .tar.gz file from the [NetApp Software Downloads](#) site. The NetApp Element VASA provider uses a NetApp certificate. With this patch, the certificate is used unmodified by vCenter to support multiple vCenters for VASA and VVols use. Do not modify the certificate. Custom SSL certificates are not supported by VASA.

Find more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Volume access groups

By creating and using volume access groups, you can control access to a set of volumes.

When you associate a set of volumes and a set of initiators with a volume access group, the access group grants those initiators access to that set of volumes.

Volume access groups in NetApp SolidFire storage enable iSCSI initiator IQNs or Fibre Channel WWPNs to access a collection of volumes. Each IQN that you add to an access group can access each volume in the group without using CHAP authentication. Each WWPN that you add to an access group enables Fibre Channel network access to the volumes in the access group.

Volume access groups have the following limits:

- A maximum of 128 initiators per volume access group.
- A maximum of 64 access groups per volume.
- An access group can be made up of a maximum of 2000 volumes.
- An IQN or WWPN can belong to only one volume access group.
- For Fibre Channel clusters, a single volume can belong to a maximum of four access groups.

Initiators

Initiators enable external clients access to volumes in a cluster, serving as the entry point for communication between clients and volumes. You can use initiators for CHAP-based rather than account-based access to storage volumes. A single initiator, when added to a volume access group, allows volume access group members to access all storage volumes added to the group without requiring authentication. An initiator can belong to only one access group.

Data protection

Data protection features include remote replication, volume snapshots, volume cloning, Protection Domains, and high availability with double Helix technology.

Element storage data protection includes the following concepts:

- [Remote replication types](#)
- [Volume snapshots for data protection](#)
- [Volume clones](#)
- [Backup and restore process overview for Element storage](#)
- [Protection Domains](#)
- [Custom Protection Domains](#)
- [Double Helix high availability](#)

Remote replication types

Remote replication of data can take the following forms:

- [Synchronous and asynchronous replication between clusters](#)
- [Snapshot-only replication](#)

- [Replication between Element and ONTAP clusters using SnapMirror](#)

For more information, see [TR-4741: NetApp Element Software Remote Replication](#).

Synchronous and asynchronous replication between clusters

For clusters running NetApp Element software, real-time replication enables the quick creation of remote copies of volume data.

You can pair a storage cluster with up to four other storage clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios.

Synchronous replication

Synchronous replication continuously replicates data from the source cluster to the target cluster and is affected by latency, packet loss, jitter, and bandwidth.

Synchronous replication is appropriate for the following situations:

- Replication of several systems over a short distance
- A disaster recovery site that is geographically local to the source
- Time-sensitive applications and the protection of databases
- Business continuity applications that require the secondary site to act as the primary site when the primary site is down

Asynchronous replication

Asynchronous replication continuously replicates data from a source cluster to a target cluster without waiting for the acknowledgments from the target cluster. During asynchronous replication, writes are acknowledged to the client (application) after they are committed on the source cluster.

Asynchronous replication is appropriate for the following situations:

- The disaster recovery site is far from the source and the application does not tolerate latencies induced by the network.
- There are bandwidth limitations on the network connecting the source and target clusters.

Snapshot-only replication

Snapshot-only data protection replicates changed data at specific points of time to a remote cluster. Only those snapshots that are created on the source cluster are replicated. Active writes from the source volume are not.

You can set the frequency of the snapshot replications.

Snapshot replication does not affect asynchronous or synchronous replication.

Replication between Element and ONTAP clusters using SnapMirror

With NetApp SnapMirror technology, you can replicate snapshots that were taken using NetApp Element software to ONTAP for disaster recovery purposes. In a SnapMirror relationship, Element is one endpoint and ONTAP is the other.

SnapMirror is a NetApp Snapshot replication technology that facilitates disaster recovery, designed for failover

from primary storage to secondary storage at a geographically remote site. SnapMirror technology creates a replica, or mirror, of the working data in secondary storage from which you can continue to serve data if an outage occurs 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 are referred to as endpoints in which the volumes reside and the volumes that contain the replicated data must be peered. A peer relationship enables clusters and volumes to exchange data securely.

SnapMirror runs natively on the NetApp ONTAP controllers and is integrated into Element, which runs on NetApp HCI and SolidFire clusters. The logic to control SnapMirror resides in ONTAP software; therefore, all SnapMirror relationships must involve at least one ONTAP system to perform the coordination work. Users manage relationships between Element and ONTAP clusters primarily through the Element UI; however, some management tasks reside in NetApp ONTAP System Manager. Users can also manage SnapMirror through the CLI and API, which are both available in ONTAP and Element.

See [TR-4651: NetApp SolidFire SnapMirror Architecture and Configuration](#) (login required)

You must manually enable SnapMirror functionality at the cluster level by using Element software. SnapMirror functionality is disabled by default, and it is not automatically enabled as part of a new installation or upgrade.

After enabling SnapMirror, you can create SnapMirror relationships from the Data Protection tab in the Element software.

NetApp Element software 10.1 and above supports SnapMirror functionality to copy and restore snapshots with ONTAP systems.

Systems running Element 10.1 and above include code that can communicate directly with SnapMirror on ONTAP systems running 9.3 or higher. The Element API provides methods to enable SnapMirror functionality on clusters, volumes, and snapshots. Additionally, the Element UI includes functionality to manage SnapMirror relationships between Element software and ONTAP systems.

Starting with Element 10.3 and ONTAP 9.4 systems, you can replicate ONTAP originated volumes to Element volumes in specific use cases with limited functionality.

For more information, see ONTAP documentation.

Volume snapshots for data protection

A volume snapshot is a point-in-time copy of a volume that you could later use to restore a volume to that specific time.

While snapshots are similar to volume clones, snapshots are simply replicas of volume metadata, so you cannot mount or write to them. Creating a volume snapshot also takes only a small amount of system resources and space, which makes snapshot creation faster than cloning.

You can replicate snapshots to a remote cluster and use them as a backup copy of the volume. This enables you to roll back a volume to a specific point in time by using the replicated snapshot; you can also create a clone of a volume from a replicated snapshot.

You can back up snapshots from a Element cluster to an external object store, or to another Element cluster. When you back up a snapshot to an external object store, you must have a connection to the object store that allows read/write operations.

You can take a snapshot of an individual volume or multiple for data protection.

Volume clones

A clone of a single volume or multiple volumes is point-in-time copy of the data. When you clone a volume, the system creates a snapshot of the volume and then creates a copy of the data referenced by the snapshot.

This is an asynchronous process, and the amount of time the process requires depends on the size of the volume you are cloning and the current cluster load.

The cluster supports up to two running clone requests per volume at a time and up to eight active volume clone operations at a time. Requests beyond these limits are queued for later processing.

Backup and restore process overview for Element storage

You can back up and restore volumes to other SolidFire storage, as well as to secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

You can back up a volume to the following:

- A SolidFire storage cluster
- An Amazon S3 object store
- An OpenStack Swift object store

When you restore volumes from OpenStack Swift or Amazon S3, you need manifest information from the original backup process. If you are restoring a volume that was backed up on a SolidFire storage system, no manifest information is required.

Protection Domains

A Protection Domain is a node or a set of nodes grouped together such that any part or even all of it might fail, while maintaining data availability. Protection Domains enable a storage cluster to heal automatically from the loss of a chassis (chassis affinity) or an entire domain (group of chassis).

You can manually enable Protection Domain monitoring using the NetApp Element Configuration extension point in the NetApp Element Plug-in for vCenter Server. You can select a Protection Domain threshold based on node or chassis domains. You can also enable Protection Domain monitoring using the Element API or web UI.

A Protection Domain layout assigns each node to a specific Protection Domain.

Two different Protection Domain layouts, called Protection Domain levels, are supported.

- At the node level, each node is in its own Protection Domain.
- At the chassis level, only nodes that share a chassis are in the same Protection Domain.
 - The chassis level layout is automatically determined from the hardware when the node is added to the cluster.
 - In a cluster where each node is in a separate chassis, these two levels are functionally identical.

When creating a new cluster, if you are using storage nodes that reside in a shared chassis, you might want to consider designing for chassis-level failure protection using the Protection Domains feature.

Custom Protection Domains

You can define a custom Protection Domain layout that matches your specific chassis and node layout, and where each node is associated with one and only one custom Protection Domain. By default, each node is assigned to the same default custom Protection Domain.

If no custom Protection Domains are assigned:

- Cluster operation is unaffected.
- Custom level is neither tolerant nor resilient.

When you configure custom Protection Domains for a cluster, there are three possible levels of protection, which you can see from the Element web UI dashboard:

- Not protected: The storage cluster is not protected from the failure of one of its custom Protection Domains. To fix this, add additional storage capacity to the cluster or reconfigure the cluster's custom Protection Domains to protect the cluster from possible data loss.
- Fault tolerant: The storage cluster has enough free capacity to prevent data loss after the failure of one of its custom Protection Domains.
- Fault resilient: The storage cluster has enough free capacity to self-heal after the failure of one of its custom Protection Domains. After the healing process has completed, the cluster will be protected from data loss if additional domains were to fail.

If more than one custom Protection Domain is assigned, each subsystem will assign duplicates to separate custom Protection Domains. If this is not possible, it reverts to assigning duplicates to separate nodes. Each subsystem (for example, bins, slices, protocol endpoint providers, and ensemble) does this independently.

You can use the Element UI to [configure custom Protection Domains](#), or you can use the following API methods:

- [GetProtectionDomainLayout](#) - shows which chassis and which custom Protection Domain each node is in.
- [SetProtectionDomainLayout](#) - enables a custom Protection Domain to be assigned to each node.

Double Helix high availability

Double Helix data protection is a replication method that spreads at least two redundant copies of data across all drives within a system. The “RAID-less” approach enables a system to absorb multiple, concurrent failures across all levels of the storage system and repair quickly.

Performance and quality of service

A SolidFire storage cluster has the ability to provide Quality of Service (QoS) parameters on a per-volume basis. You can guarantee cluster performance measured in inputs and outputs per second (IOPS) using three configurable parameters that define QoS: Min IOPS, Max IOPS, and Burst IOPS.



SolidFire Active IQ has a QoS recommendations page that provides advice on optimal configuration and set up of QoS settings.

Quality of Service parameters

IOPS parameters are defined in the following ways:

- **Minimum IOPS** - The minimum number of sustained inputs and outputs per second (IOPS) that the storage cluster provides to a volume. The Min IOPS configured for a volume is the guaranteed level of performance for a volume. Performance does not drop below this level.
- **Maximum IOPS** - The maximum number of sustained IOPS that the storage cluster provides to a volume. When cluster IOPS levels are critically high, this level of IOPS performance is not exceeded.
- **Burst IOPS** - The maximum number of IOPS allowed in a short burst scenario. If a volume has been running below the Max IOPS, burst credits are accumulated. When performance levels become very high and are pushed to maximum levels, short bursts of IOPS are allowed on the volume.

Element software uses Burst IOPS when a cluster is running in a state of low cluster IOPS utilization.

A single volume can accrue Burst IOPS and use the credits to burst above their Max IOPS up to their Burst IOPS level for a set "burst period." A volume can burst for up to 60 seconds if the cluster has the capacity to accommodate the burst. A volume accrues one second of burst credit (up to a maximum of 60 seconds) for every second that the volume runs below its Max IOPS limit.

Burst IOPS are limited in two ways:

- A volume can burst above its Max IOPS for a number of seconds equal to the number of burst credits that the volume has accrued.
 - When a volume bursts above its Max IOPS setting, it is limited by its Burst IOPS setting. Therefore, the burst IOPS never exceeds the burst IOPS setting for the volume.
- **Effective Max Bandwidth** - The maximum bandwidth is calculated by multiplying the number of IOPS (based on the QoS curve) by the IO size.

Example: QoS parameter settings of 100 Min IOPS, 1000 Max IOPS, and 1500 Burst IOPS have the following effects on quality of performance:

- Workloads are able to reach and sustain a maximum of 1000 IOPS until the condition of workload contention for IOPS becomes apparent on the cluster. IOPS are then reduced incrementally until IOPS on all volumes are within the designated QoS ranges and contention for performance is relieved.
- Performance on all volumes is pushed toward the Min IOPS of 100. Levels do not drop below the Min IOPS setting but could remain higher than 100 IOPS when workload contention is relieved.
- Performance is never greater than 1000 IOPS, or less than 100 IOPS for a sustained period. Performance of 1500 IOPS (Burst IOPS) is allowed, but only for those volumes that have accrued burst credits by running below Max IOPS and only allowed for a short periods of time. Burst levels are never sustained.

QoS value limits

Here are the possible minimum and maximum values for QoS.

Parameters	Min value	Default	4 4KB	5 8KB	6 16KB	262KB
Min IOPS	50	50	15,000	9,375*	5556*	385*
Max IOPS	100	15,000	200,000**	125,000	74,074	5128
Burst IOPS	100	15,000	200,000**	125,000	74.074	5128

*These estimations are approximate.


**Max IOPS and Burst IOPS can be set as high as 200,000; however, this setting is allowed only to effectively uncap the performance of a volume. Real-world maximum performance of a volume is limited by cluster usage and per-node performance.

QoS performance

The QoS performance curve shows the relationship between block size and the percentage of IOPS.

Block size and bandwidth have a direct impact on the number of IOPS that an application can obtain. Element software takes into account the block sizes it receives by normalizing block sizes to 4k. Based on workload, the system might increase block sizes. As block sizes increase, the system increases bandwidth to a level necessary to process the larger block sizes. As bandwidth increases the number of IOPS the system is able to attain decreases.

The QoS performance curve shows the relationship between increasing block sizes and the decreasing percentage of IOPS:



As an example, if block sizes are 4k, and bandwidth is 4000 KBps, the IOPS are 1000. If block sizes increase to 8k, bandwidth increases to 5000 KBps, and IOPS decrease to 625. By taking block size into account, the system ensures that lower priority workloads that use higher block sizes, such as backups and hypervisor activities, do not take too much of the performance needed by higher priority traffic using smaller block sizes.

QoS policies

A QoS policy enables you to create and save a standardized quality of service setting that can be applied to many volumes.

QoS policies are best for service environments, for example, with database, application, or infrastructure servers that rarely reboot and need constant equal access to storage. Individual volume QoS is best for light use VMs, such as virtual desktops or specialized kiosk-type VMs, that may be rebooted, powered on, or

powered off daily or several times a day.

QoS and QoS policies should not be used together. If you are using QoS policies, do not use custom QoS on a volume. Custom QoS will override and adjust QoS policy values for volume QoS settings.



The selected cluster must be Element 10.0 or later to use QoS policies; otherwise, QoS policy functions are not available.

Find more information

- [SolidFire and Element Software Documentation](#)

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.