

# PowerScale Introduction for NetApp Admins

## PowerScale nomenclature from a NetApp background

January 2022

H15757.2

## White Paper

### Abstract

This document provides an overview of PowerScale scale-out NAS including the NetApp Clustered Data ONTAP background architecture and features.

Dell Technologies

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2017–2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA January 2022 H15757.2.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Contents

- Executive summary ..... 4
- PowerScale overview..... 5
- OneFS features..... 6
- References..... 12

## Executive summary

### Overview

For enterprises to take advantage of the advanced storage technologies of Dell PowerScale and begin transitioning from their platform, they must determine how the new architecture fits into their existing infrastructure. This document provides an overview of PowerScale architecture, features, and nomenclature for enterprises migrating from traditional NAS platforms, such as NetApp Clustered Data ONTAP (cDOT).

### Revisions

Date	Description
May 2017	Initial release
June 2020	Updated with PowerScale branding and converted to new template
January 2022	Updated document content and template

### We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

**Author:** Aqib Kazi

---

**Note:** For links to other documentation for this topic, see the [PowerScale Info Hub](#).

---

# PowerScale overview

## Introduction

The PowerScale OneFS operating system is based on a distributed architecture and is built as a clustered system. Each node provides compute, memory, networking, and storage. The notion of controllers, HA, active/standby, or disk shelves are not present in a pure scale-out architecture. When a node is added to a cluster, the cluster performance and capacity increase collectively.

With this scale-out distributed architecture, system management is far simpler than traditional NAS platforms, due to the single namespace, single volume, single file system, and single pane of management. Also, the data protection is software-based rather than RAID-based, eliminating the associated complexities including configuration, maintenance, and additional storage utilization. Administrators do not have to be concerned with RAID groups or load distribution.

NetApp Clustered Data ONTAP (cDOT) has evolved into a clustered system with controllers. cDOT introduced FlexGroups composed of aggregates and FlexVols across nodes. However, each NetApp filer has a file system beneath it. For example, if a NetApp cluster contains 24 nodes, it has 24 underlying file systems.

PowerScale is a single volume, which makes cluster management simple. As the cluster grows in capacity, the single volume automatically grows. Administrators are no longer required to manually migrate data between volumes. OneFS repopulates and balances data between all nodes when a new node is added, making the node part of the global namespace. All the nodes in a PowerScale cluster are equal in the hierarchy. Drives share data intra-node and inter-node.

Also, PowerScale is easy to deploy, operate, and manage. Most enterprises require only one full-time employee to manage a PowerScale cluster.

## Architecture

PowerScale offers node types based on performance and storage requirements. To form a cluster, a minimum of three nodes are required. The nodes are connected to the back-end through private switches in an active/active state, and are responsible for distributing data. The front-end or client-facing connections support multiple subnets, which are split between pools of IP addresses. Also, PowerScale SmartConnect provides load-balancing between the nodes for client connections.




Figure 1. PowerScale architecture

## File system

OneFS is based on FreeBSD. OneFS has a root file path, `/ifs`, which is globally accessible. Access permissions and hidden shares may be applied as required.

## Cluster expansion

Expanding storage on a NetApp system requires adding disk shelves to an HA pair or adding HA pairs to the cluster.

Adding a new node to a PowerScale cluster takes sixty seconds. After you cable the node, select **Add to Existing Cluster** on the LCD panel. OneFS automatically balances data evenly across all nodes. The OneFS single volume capacity grows automatically without any other actions.

# OneFS features

## Overview

The single volume and single namespace of PowerScale OneFS also lead to a unique feature set. Since the entire NAS is a single file system, the concept of FlexVols, shares, qtrees, or FlexGroups do not apply. Each NetApp volume has specific properties associated with limited storage space. Adding storage space tends to be an onerous process depending on the current architecture. As soon as a node is added to a PowerScale cluster, the cluster rebalances automatically, leading to minimal administrator management.

The fact that OneFS is a single volume means that many features are not volume-dependent, but span the entire cluster. SnapshotIQ, NDMP backups, and SmartQuotas do not have limits based on volumes, but they are cluster- or directory-specific.

With a single volume NAS, features like SmartConnect, SmartPools, and CloudPools utilize the entire cluster by differentiating node performance. You can use SmartConnect to configure access zones that are mapped to specific node performance. SmartPools can tier cold data to nodes with deep archive storage, and CloudPools can store frozen data in the cloud. Regardless of where the data is residing, it is presented as a single namespace to the end user.

Table 1. PowerScale and NetApp feature nomenclature

Feature	PowerScale	NetApp	Details
Data protection	Software-based Reed-Solomon Error Correction Up to N+4 Protection Hybrid Protection Options available - Cover Node and/or drive failure	NetApp RAID-TEC NetApp RAID-DP	OneFS provides data protection against more simultaneous hardware failures and is software-based, providing a higher storage utilization. NetApp only supports a maximum of triple parity.
Snapshots	SnapshotIQ	NetApp Snapshot	<b>OneFS supports:</b> 1,024 snapshots/directory 20,000 snapshots/cluster <b>NetApp Snapshot:</b> 255 snapshots/volume
Replication	SyncIQ	SnapMirror MetroCluster	SyncIQ supports clusters in compliance mode. SnapMirror supports SnapLock volumes.
Backup	NDMP	SnapVault NDMP	SnapVault limits are volume- and snapshot-dependent. 255 maximum per volume. OneFS supports 64 NDMP connections per node.
Load balancing or Multi-tenancy	SmartConnect Access Zones	MultiStore – vFiler Storage Virtual Machines	OneFS SmartConnect Access Zones are performance-based and also support multi-tenancy. NetApp does not support load balancing since all connections are through the controller, but multi-tenancy is supported through vFiler or SVM.
Compliance	SmartLock	SnapLock	
Compression	Compression	Compression	
Deduplication	SmartDedupe	Inter and Intra Volume dedupe	

Feature	PowerScale	NetApp	Details
Data reduction	In-Line Data Reduction	In-Line Data Compaction	
Data tiering	SmartPools	FabricPools	
Quotas	SmartQuotas	Quotas are user, tree, volume, or group-based.	OneFS provides configurable options to monitor and enforce storage limits at the user, group, cluster, directory, or subdirectory level.
Monitoring	InsightIQ CloudIQ DataIQ	OnCommand Suite ActiveIQ	
Cloud tiering	CloudPools	FabricPools	

**NetApp FlexVols, shares, and Qtrees**

NetApp requires administrators to create space manually and explicitly define aggregates and flexible volumes. The concept of FlexVols, Shares, and Qtrees are nonexistent in OneFS, since the file system is a single volume and namespace, spanning the entire cluster.

SMB shares and NFS exports are created through the web or command-line interface in OneFS. Both methods allow the user to create either within seconds with security options. SmartQuotas can manage storage limits with accounting, warning messages, or hard limits of enforcements. You can apply limits by directory, user, or group.

**Data replication**

PowerScale SyncIQ provides data replication. Similar to NetApp SnapMirror, this replication could be for a disaster recovery site or a form of backup. SyncIQ provides options for ensuring all data is retained during failover and failback from the disaster recovery cluster. SyncIQ is fully configurable with options for execution times and bandwidth management. Only changed data blocks are transferred with SyncIQ, allowing for minimal resource consumption. SyncIQ can be directory-based or for entire cluster replication, rather than volume-based with NetApp SnapMirror. A SyncIQ target cluster may be configured as a target for several source clusters.

**Load balancing**

Since OneFS is a distributed architecture across a collection of nodes, client connectivity to these nodes requires load-balancing. PowerScale SmartConnect provides options for balancing the client connections to the nodes within a cluster. Balancing options are round-robin or based on the current load. Also, SmartConnect Zones can be configured to have clients connect based on group and performance needs. For example, the engineering group may require high-performance nodes. A zone can be configured, forcing connections to those nodes.

**Compliance**

PowerScale SmartLock offers similar features to NetApp SnapLock, including WORM files and meeting strict SEC requirements.



**Role-based access control**

PowerScale supports role-based access control (RBAC), allowing you to configure administrative tasks without a root or administrator account. A role is a collection of OneFS privileges that are limited to an area of administration. You may provision custom roles for security, auditing, storage, or backup tasks with RBACs. Privileges are assigned to roles. As a user logs in to the cluster through the platform API, the OneFS command-line interface, or the OneFS web administration interface, they are granted privileges based on their role membership.

**Storage utilization and data protection**

Storage utilization is the amount of storage available after the NAS system overhead is deducted. The overhead consists of the space required for data protection and the operating system. NetApp data protection is RAID-TEC or RAID-DP, depending on the configuration. RAID-TEC and RAID-DP require three disks, or two disks, respectively, for parity of each additional RAID group.

PowerScale uses software-based data protection rather than RAID. The software-based data protection stripes data across nodes in stripe units. Based on the Reed-Solomon algorithm, some of the stripe units are Forward Error Correction (FEC) or parity units. The FEC units provide a variable to reformulate the data, if there is a drive or node failure. Data protection is customizable to be for node loss or a hybrid protection of node and drive failure.

With software-based data protection, the protection scheme is not per-cluster but has more granularity that allows you to make data protection specific to a file or directory. This capability does not require you to create more storage volumes or manually migrate data. Instead, OneFS runs a job in the background, moving data as configured.

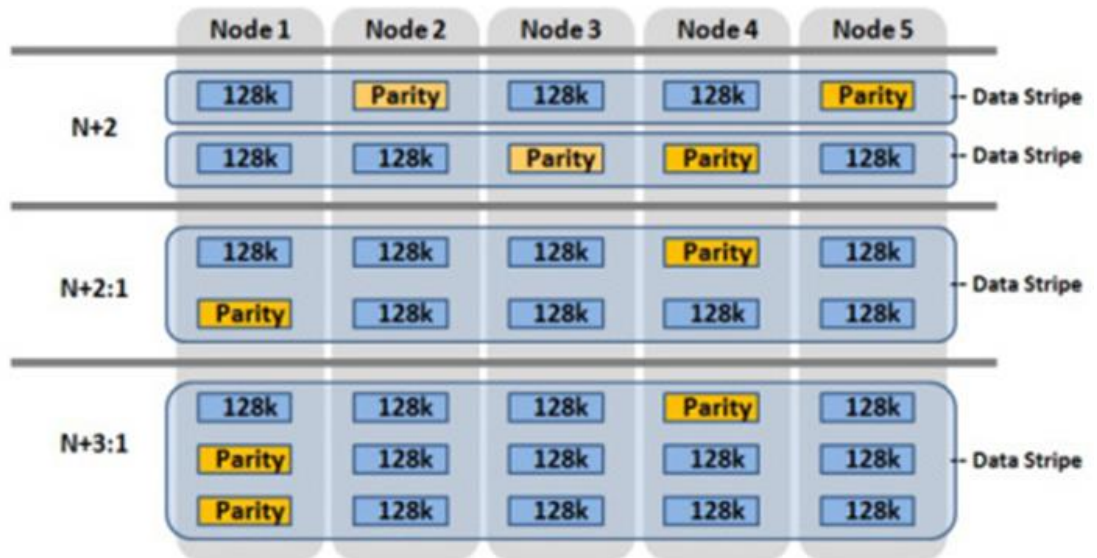


Figure 2. PowerScale data protection

**Snapshots**

SnapshotIQ provides options for taking on-demand or scheduled snapshots of directories, rather than taking volume-based snapshots with NetApp Snapshot. These snapshots have a small overhead with little to no impact on performance and storage. The limits are cluster-based rather than volume-based. PowerScale supports up to 1,024 snapshots per-

directory and up to 20,000 per cluster. An option is also provided for restoring using Windows Volume Manager, allowing users to perform restores without requiring administrator assistance, minimizing help desk calls.

### Backups

OneFS supports backups through NDMP to several third-party vendors offering NDMP support. You can also add an accelerator node to the cluster to process backups faster.

### Data tiering

PowerScale has integrated features to tier data based on its age or file type. NetApp does not have integrated functionality to tier data. Instead, you must use a third-party tool. Implementation and management of an additional tool requires administrator resources.

You can configure PowerScale SmartPools to move data to a set of nodes automatically. For example, if a file has not been accessed in the last 90 days, you can migrate it to a node with deeper storage. You can also define the value of storage based on performance.

PowerScale CloudPools migrates data to a cloud provider, with only a stub remaining on the PowerScale cluster, based on similar policies. Typically, CloudPools migrates frozen data, or data that is not accessed in the last year.




Figure 3. PowerScale SmartPools and CloudPools

### Authentication and access control

NetApp and PowerScale both support several methods for user authentication and access control. PowerScale OneFS supports UNIX and Windows permissions for data-level access control. OneFS is designed for a mixed environment that allows the configuration of both Windows access control lists (ACLs) and standard UNIX permissions on the cluster file system. Also, OneFS provides user and identity mapping, permission mapping, and merging between Windows and UNIX environments.

OneFS supports local and remote authentication providers. Anonymous access is supported for protocols that allow it. Concurrent use of multiple authentication provider types, including Active Directory, LDAP, and NIS, is supported. For example, OneFS is often configured to authenticate Windows clients with Active Directory and to authenticate UNIX clients with LDAP.

## Multi-tenancy

NetApp Clustered Data ONTAP (cDOT) supports multi-tenancy with Storage Virtual Machines (SVMs), formerly vServers. SVMs isolate storage and network resources across a cluster of controller HA pairs. vFilers and SVMs require managing protocols, shares, and volumes to provision successfully. Volumes cannot be non-disruptively moved between SVMs. PowerScale OneFS provides multi-tenancy through SmartConnect and Access Zones. Management is simple as the file system is one volume and access is provided by hostname and directory, rather than by volume.

SmartConnect zones allow administrators to provision DNS hostnames specific to IP pools, subnets, and network interfaces. If only a single authentication provider is required, all SmartConnect zones map to a default Access Zone. However, if directory access and authentication providers vary, multiple Access Zones are provisioned, mapping to a directory, authentication provider, and SmartConnect Zone. As a result, authenticated users of an Access Zone only have visibility into their respective directory. Conversely, an administrator with complete file system access can migrate data non-disruptively between directories.

## Monitoring

PowerScale InsightIQ, CloudIQ, and DataIQ provide performance monitoring and reporting tools, similar to the NetApp OnCommand Suite and ActiveIQ.

## References

### Related documentation

The following documentation provides other information related to this document. Access to the Dell Technologies documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [PowerScale Info Hub](#)
- [PowerScale OneFS Info Hubs](#)
- [PowerScale Network Design Considerations](#)
- [Superna Eyeglass](#)
- [High Availability and Data Protection with Dell EMC PowerScale Scale-Out NAS](#)
- [PowerScale OneFS CloudPools Administration Guide](#)
- [PowerScale CloudPools and ECS Solution Guide](#)
- [Dell EMC PowerScale OneFS Best Practices](#)