

Dell PowerScale OneFS : Powerful Software for Modern Storage

NAS simple, parallel fast, and fully featured for scalable storage management

ESSENTIALS

- Scales up to 252 nodes in a cluster and up to PBs of raw capacity in a single namespace
- Automated policy-driven tiered storage to optimize resources
- Seamless public cloud integration with a software-defined offer
- Resilient data protection for a highly available environment
- Comprehensive security and compliance options
- Optimized storage consumption with flexible quotas
- Seamless load balancing of client connections for maximum availability
- Storage efficiency, deduplication and compression to reduce cost

The power of the OneFS operating system

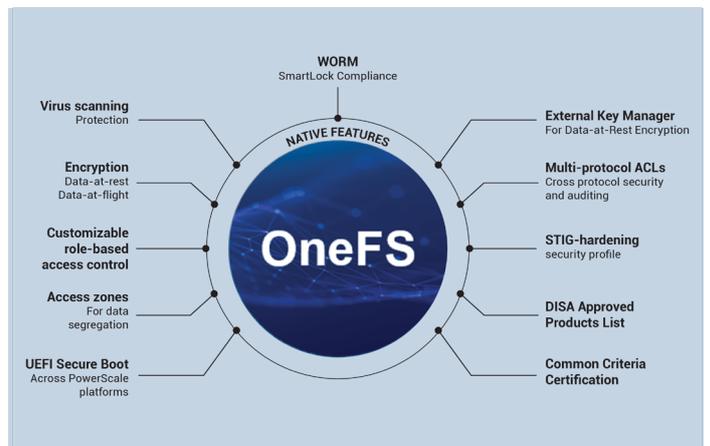
PowerScale OneFS is the foundation of the world's most flexible¹, efficient², and secure³ scale-out storage platform. Designed to simplify data management, OneFS allows you to seamlessly store, protect, and analyze your data while supporting a wide range of workloads, including AI and HPC. By consolidating data into a single system with NAS-like ease, OneFS eliminates silos with a unified volume, filesystem, and namespace, delivering parallel performance and effortless scalability. With support for protocols like NFS, SMB, S3, and HDFS, OneFS enables simultaneous file and object applications within the same namespace, making it a versatile solution for modern data needs. Its software-defined architecture and intuitive management allow you to expand capacity and performance—from terabytes to exabytes—in minutes, without disruption. Whether deployed as an appliance, as-a-service, or a software-defined solution for AWS and Microsoft Azure, OneFS provides the flexibility to scale your data infrastructure to match your ambitions, no matter where your data resides—at the edge, core, or cloud.

OneFS software features

OneFS offers powerful software modules designed to simplify storage and data management at scale. Its advanced storage management features—such as auto-balancing, snapshots, data protection, backup, replication, and disaster recovery—streamline and automate the management of OneFS-powered clusters, reducing complexity for administrators. In addition, OneFS provides robust data management capabilities, including quotas and deduplication, which help administrators and data owners maximize the value of their data investments. Features like remote monitoring further enhance operational efficiency by lowering costs and minimizing risks, ensuring a seamless and secure storage experience.

PowerScale OneFS built-in security features

The PowerScale OneFS operating system comes with a comprehensive set of built-in security features including customizable role-based access control, access zones for data segregation, data at rest and data in flight encryption, virus scanning tool, WORM SmartLock compliance, multi-protocol ACLs for cross protocol security and auditing, external key manager for data at rest encryption, STIG-hardened security profile across PowerScale platforms, and certifications including Common Criteria and DISA Approved Product List. OneFS security was further reinforced with features like multi-factor authentication with support for CAC/PIV, software-based firewall functionality, SSO for administrative WebUI, FIPS-compliant data in flight in the latest software release. Additional cyber-resiliency tools are available for data isolation, automated threat detection and rapid recovery in PowerScale Cyber Protection Solution.

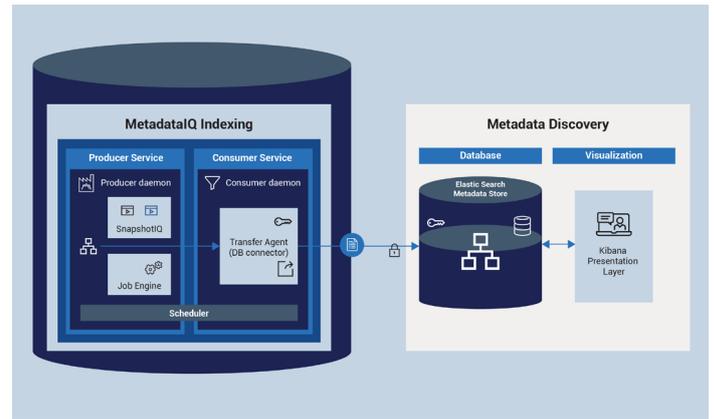


1. Based on internal analysis of publicly available information sources, August 2021.
2. Based on Dell analysis comparing efficiency-related features: data reduction, storage capacity, data protection, hardware, space, lifecycle management efficiency, and ENERGY STAR certified configurations, June 2023.
3. Based on Dell analysis comparing cyber-security software capabilities offered for Dell PowerScale vs. competitive products, September 2022.

MetadatalQ

The OneFS MetadatalQ framework is designed to provide advanced metadata capabilities, allowing customers to index and discover the data they need for their workflows and analytics needs. This metadata catalog may be used for queries, data visualization and data lifecycle management. As customers add analytics workflows, the ability to simply and efficiently query data, wherever it may reside, is vital for the time-to-results they require.

Internally, MetadatalQ is predicated on the OneFS Job Engine's ChangeListCreate job, which tracks the delta (changelist) between two snapshots. MetadatalQ parses entries in each changelist in batches, updating the metadata index residing off-cluster in an ElasticSearch database. This database can store the metadata from multiple PowerScale clusters, providing a global catalog of an organization's unstructured data repositories.

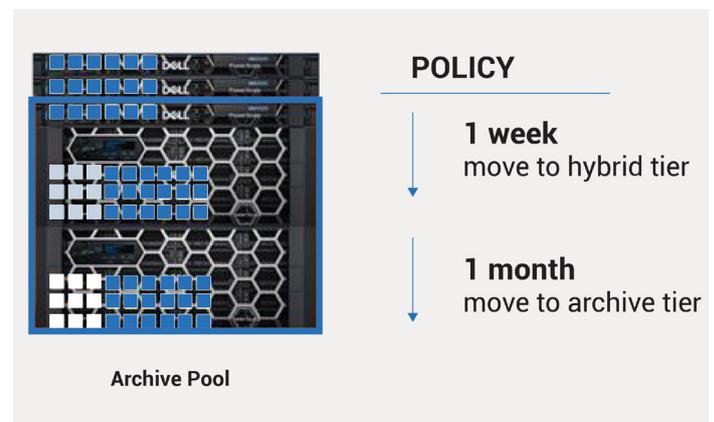


PowerScale SmartPools

OneFS comprises a single file system single namespace that spans all the nodes of a cluster. SmartPools allow multiple storage tiers to exist within a single file system to aggregate and consolidate applications within a single storage pool. This gives you workflow isolation, higher utilization, and independent scalability – from a single point of management.

SmartPools allows you to define the value of the data within your workflows based on policies and automatically aligns the data to the appropriate price/performance tier over time. Utilizing file-level granularity and control with automatic policies, you can tune performance and data layout, storage tier alignment and protection settings – all without impact to your end-users. SmartPools allows for unprecedented flexibility, granularity, and ease of management. SmartPools aligns the business value of data with optimal

storage performance and cost by optimized data placement including flash, hybrid, archive, and cloud storage. Policy triggers that determine optimized data placement include criteria like file age, size, type, owner, location, or date fields. By default, the SmartPools jobs run nightly to apply new policies to the selected data and seamlessly tier the files to the appropriate location based on your requirements. Capacity and performance-based data placement provide even further granularity while improving storage efficiency.

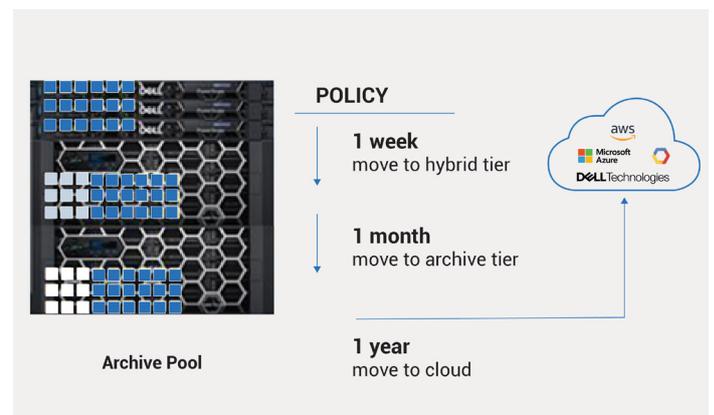


PowerScale CloudPools

CloudPools software provides policy-based automated tiering that lets you seamlessly integrate with the cloud as an additional storage tier for the cluster. This allows you to address rapid data growth, reduce storage costs and optimize data center storage resources by using the cloud for frozen data. In this way, your more valuable storage may be used for more active data and applications, while frozen data may be retained at minimum cost for compliance, historical or other business reasons.

With CloudPools, you have a flexible choice of tiering data to public, private or hybrid cloud options. You can select from Amazon Web Services (AWS) S3, Google Cloud Platform (GCP), Alibaba Aliyun, Federal C2S clouds, Microsoft Azure or ECS.

With CloudPools and SmartPools, you may define policies that identify the data to be tiered, the criteria for tiering and the choice of the public or private cloud target. The policies are dynamic, flexible, and scalable, which provides you with granular control of the data placement. A policy can be based on any combination of file metadata attributes such as timestamps, file name or type or file size.



When a file is tiered, it is replaced by a SmartLink that contains the maps to the contents in the cloud. Users access the data the same way as before without changes to policies and procedures – you do not have to do anything different to access the data. If you access a tiered file, only the relevant blocks are retrieved without the need to retrieve the entire file from the cloud. When you modify a tiered file, only the relevant portions of the file are rewritten to the cloud, thereby optimizing the bandwidth. CloudPools allows you to encrypt or compress data that is transmitted.

PowerScale SmartQuotas

SmartQuotas allow you to control and limit storage usage by assigning quotas at the cluster, directory, subdirectory, user, and group levels. SmartQuotas span across the entire cluster thus enabling you to easily administer storage from a single interface. With its thin provisioning capability, SmartQuotas allows you to present more storage capacity to applications and users than is physically installed. In this way, you can limit their actual physical storage resources to what is only needed today and automatically add storage resources on demand to meet changing business requirements in the future. Storage capacity can be automatically increased with minimal administrative overhead, so that you can purchase less storage capacity up front, defer capacity upgrades to match actual business usage and save on power and cooling costs associated with keeping unused disk capacity spinning.

When actual capacity begins to approach its designated threshold, nodes can be added to the cluster quickly and easily – typically in minutes. The result is unprecedented grow-as-you-go flexibility and value if you are looking to minimize costs while keeping pace with data growth.

SmartQuotas can be used to establish hard, soft and advisory storage capacity limits that can be set across your organization for specific users and groups, as well as across the various directory structures. Setting up advisory and soft quotas as a percentage of the hard quota provides improved convenience, flexibility, and usability. You can also use SmartQuotas to configure alerts and send email notifications to end users, letting them know that quota limits are approaching, enforcing hard stops on writes, or providing a grace period of several days before maintaining thresholds.

PowerScale SmartConnect

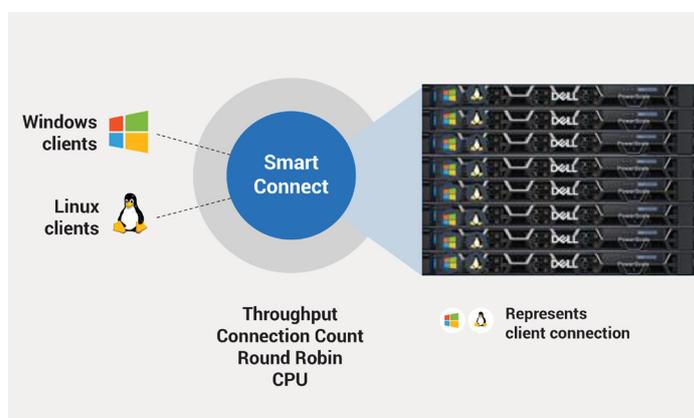
SmartConnect delivers intelligent, automatic client connection load balancing and failover capabilities to optimize storage performance and data availability. Through a single host name, SmartConnect enables client connection load balancing and dynamic NFS protocol failover and fallback of client connections across nodes to provide optimal utilization of the cluster. Without the need to install client-side drivers, you can easily manage several clients even in the face of system failures. SmartConnect balances client connections across nodes based on policies that help ensure optimal usage of your cluster resources. By leveraging your existing network infrastructure, SmartConnect provides a layer of intelligence that allows all client and user resources to point to a single host name, enabling easy management of a large and growing numbers of clients. Based on user configurable policies, SmartConnect applies intelligent algorithms (e.g., CPU utilization, aggregate throughput, connection count, or round robin) and distributes clients across the cluster to optimize client performance and end-user experience.

SmartConnect uses a virtual IP failover scheme that does not require any client-side drivers. The cluster shares a “pool” of virtual IPs that is distributed across all nodes of the cluster. The cluster distributes an IP address across NFS (Linux and UNIX) clients based on the policy. If a node within a cluster is brought down for any reason including a failure, the cluster automatically distributes the IPs of that node to the remaining nodes and the clients can keep using the same IP of the failed node. The virtual IPs on the client connection will seamlessly failover to another node in the cluster. This ensures that when a node failure occurs, all in-flight reads and writes are handed off to another node in the cluster to finish its operation without any user or application interruption.

	DIRECTORY	USER	GROUP
Advisory quota	✓	✓	✓
Soft quota	✓	✓	✓
Hard quota	✓	✓	✓

EXAMPLE

	Apps	Employees	Executives
	4T each	10GB each	25GB each

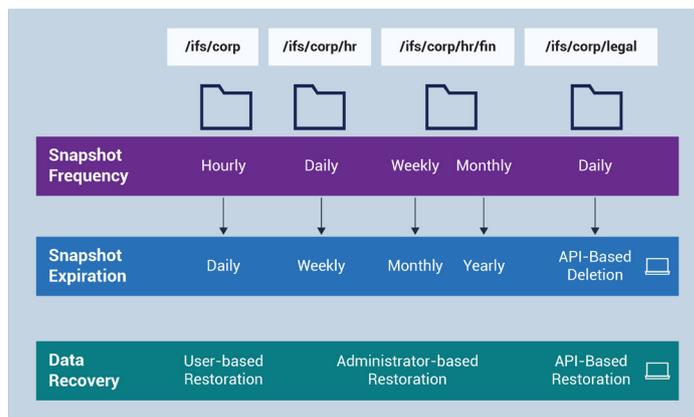


PowerScale SnapshotIQ

SnapshotIQ software provides simple point-in-time data protection and recovery by making frequent, user-recoverable backups of files. SnapshotIQ backups data automatically and as frequently as required to meet your RPO regardless of the size of the file system or directory.

SnapshotIQ also offers an extremely fast snapshot capability – typically less than one second to create. When needed, near-immediate data restoration is available to easily meet your RTO. Snapshots can be taken at a granular level, and you can take up to 1024 snapshots per directory. Because SnapshotIQ is globally coherent and spans across all nodes regardless of the cluster size, you can administer snapshots from a single point of view. With SnapshotIQ, you no longer need to worry about managing snapshot capacity and performance. With the flexibility of your cluster, storage capacity and performance can be added on-the-fly, quickly and transparently, without having to replicate or delete snapshots. Since snapshots are an integral part of the OneFS file system there is no need to pre-allocate dedicated snapshot reserve space. Once your baseline snapshot has been established, only changes to blocks that make up a file are reflected in updates to the current version of snapshots.

The automated SnapRevert functionality of SnapshotIQ also makes restoration to a specific recovery point extremely easy. One of the largest IT costs associated with backup and restore is the sheer number of help desk calls from end-users who accidentally delete a file or directory. To reduce these costs, SnapshotIQ can be used to empower end-users by enabling them to easily find and restore their own accidentally deleted files or folders – without any IT intervention.

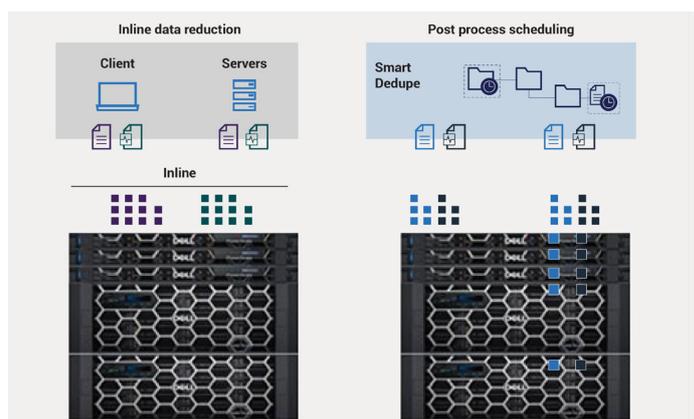


PowerScale SmartDedupe

SmartDedupe maximizes the storage efficiency of a cluster by decreasing the amount of physical storage by scanning the on-disk data for identical blocks and then eliminating the duplicates. This approach is commonly referred to as post-process or asynchronous deduplication. After duplicate blocks are discovered, SmartDedupe moves a single copy of those blocks to a special set of files known as shadow stores. During this process, duplicate blocks are removed from the actual files and replaced with pointers to the shadow stores.

With post-process deduplication, new data is first stored on the storage device and then a subsequent process analyzes the data looking for commonality. This means that the initial file-write or modify performance is not impacted since no additional computation is required in the write path. The process of sampling, fingerprinting, and matching the data is used to create an index that helps with the matching of duplicate blocks.

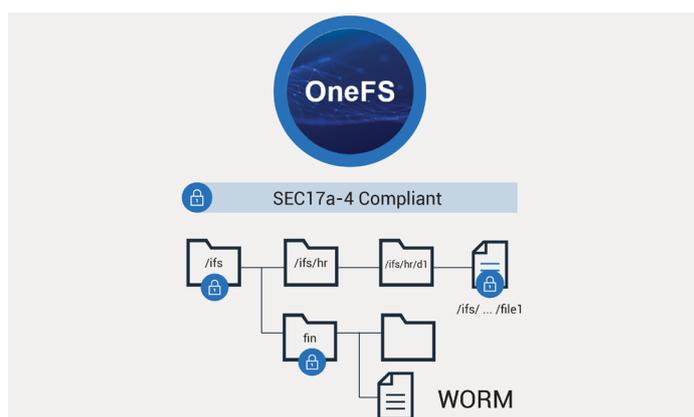
SmartDedupe can be configured all the way from a volume to a directory level granularity. You can schedule when and how frequently the SmartDedupe job runs. You can monitor and report on the status and progress of the SmartDedupe job. An assessment job can be run in estimation mode to predict the potential space savings of the dedupe process. The efficiency is dependent on the type of data scanned and the potential compressibility of that data. Various PowerScale and Isilon nodes provide the capability of inline compression and deduplication.



PowerScale SmartLock

Protecting critical data from alteration is a key business imperative for most organizations. SmartLock helps you protect your critical data against accidental, premature or malicious alteration or deletion. Because SmartLock is a software-based approach to Write Once Read Many (WORM), you can store SmartLock-protected data alongside other data types in your cluster with no effect on performance or availability and without the added cost of purchasing and maintaining specialty WORM-capable hardware.

SmartLock operates in either one of two modes – in an Enterprise mode or in a Compliance mode. You must choose the desired mode of operation during the initial cluster configuration.



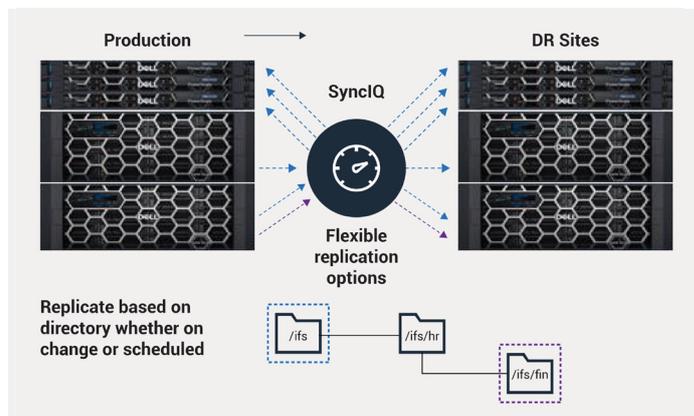
In Compliance mode, login by the root user is disabled, providing the extra level of protection to meet regulatory requirements. With Compliance mode, SmartLock can help you to meet regulatory compliance requirements to provide absolute retention and protection of data – including the most stringent SEC 17a-4 requirements. Data protected with SmartLock cannot be altered by anyone. In Enterprise mode, this data can be deleted by an authorized administrator. Retention times set under SmartLock are absolute, elapsed time and thereby preclude the impact of potential time zones changes, leap years or other time and calendar-related events which might occur during the retention period.

With SmartLock, you can protect your data at the directory-level and thereby eliminate the wasted space and complexity of managing WORM protections across multiple devices or volumes. You can set customized retention times for specific files. SmartLock is tightly integrated with OneFS and provides efficient storage for your WORM data.

PowerScale SyncIQ

SyncIQ offers powerful, flexible, and easy to manage the secure replication of data for disaster recovery, business continuance, disk-to-disk backup and remote disk archive. SyncIQ delivers replication performance because every node can send and receive data. Replication gets faster the larger your data store grows since SyncIQ can take advantage of any available network bandwidth. By default, replication between nodes is encrypted.

A simple and intuitive web-based UI allows you to easily organize SyncIQ replication job rates and priorities to match business continuance priorities. You can configure a directory, file system or even specific files for replication based on the business criticality. You can also create remote archive copies of data that needs to be retained so you can reclaim valuable capacity in your production system. SyncIQ can support end-to-end encryption of data along with seamless integration with Dell PowerScale Cybersecurity Solution powered by Superna.



Performance of SyncIQ include incremental transfer where only changed blocks are replicated, snapshot integration, bandwidth metering, monitoring, throttling and flexible scheduling. For availability, SyncIQ jobs can be configured for alerting and logging along with failure and recovery.

PowerScale SmartSync

PowerScale SmartSync is a built-in data mobility and replication feature in OneFS that unifies file-to-file replication between PowerScale clusters with file-to-object replication to ECS and public clouds (AWS, Microsoft Azure, or other S3 targets), managed through centralized policies and workflows. It supports one-time copies or ongoing incremental transfers, moving only changed data to reduce bandwidth and improve efficiency. SmartSync can operate in either a push or pull mode, offers policy-based retention and archiving, and integrates with S3 Intelligent Tiering to automatically place replicated data in the most cost-effective S3 storage tier based on access patterns. By providing flexible, efficient file-to-file and file-to-object replication, SmartSync enables organizations to protect, archive, and move unstructured data across on-premises and cloud environments.

InsightIQ

PowerScale InsightIQ provides advanced performance monitoring and reporting to help maximize PowerScale cluster performance and efficiency. It collects and analyzes detailed metrics such as throughput, latency, I/O patterns, and capacity trends, then presents them through dashboards, reports, and alerts that help teams troubleshoot issues, identify workflow bottlenecks, optimize application performance, correlates workflow and network events, and plan for future growth.

InsightIQ provides historical analysis, customizable reporting, and scalability to monitor large multi-node environments –monitoring multiple clusters—with practical guidance of up to 20 clusters or 504 nodes per instance and ongoing improvements toward monitoring 1,000+ nodes.

These capabilities accelerate troubleshooting, right-size high-performance tiers, and improve capacity planning to reduce risk and TCO, ensuring consistent performance and operational efficiency across PowerScale deployments.

SmartQoS

Dell PowerScale delivers Quality of Service (QoS) through both network-level tagging and workload-level controls in OneFS. At the network layer, QoS is implemented via DSCP tagging in the host-based firewall, which classifies outgoing front-end traffic into logical classes (for example, transactional file/object I/O, network management, bulk data) and assigns each class a DSCP value. This enables downstream switches and routers to prioritize latency-sensitive traffic such as NFS/SMB/S3 data access over less critical flows like replication, improving client experience on shared links. QoS is configured cluster-wide through CLI, WebUI, or REST API and

is preserved across upgrades. At the workload level, SmartQoS allows administrators to set protocol operations-per-second limits on specific workloads—identified by protocol, path, share/export, user, client IP, or access zone—across NFS, SMB, and S3. This prevents “noisy neighbor” effects, helps enforce performance SLAs, and protects business-critical applications by throttling lower-priority jobs, with configuration and monitoring available via CLI, WebUI, and API.

Take the next step

Contact your Dell sales representative or authorized reseller to learn more about PowerScale OneFS.



[Learn more](#) about Dell PowerScale OneFS



[Contact](#) a Dell Technologies Expert



[View more](#) PowerScale Cyber Protection Solution resources



[Join the conversation](#) with #HashTag

Copyright © Dell Inc.. All Rights Reserved. Dell Technologies, Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.