

TECHNICAL OVERVIEW OF THE NEW AND IMPROVED FEATURES OF ONEFS 8.2 OPERATING SYSTEM

Abstract

This introductory white paper provides a technical overview of the new and improved features of the Isilon OneFS 8.2 operating system. OneFS 8.2 delivers increased node counts per cluster, improved cloud integration, increased security, accelerated performance and simplified management at scale.

April 2019

TABLE OF CONTENTS

| | |
|---|---|
| Introduction | 3 |
| Increased Cluster Capacity and Performance..... | 3 |
| CloudPools Enhancements..... | 3 |
| Networking – OneFS Networking Enhancements | 4 |
| Security – Multi-Factor Authentication (MFA)..... | 4 |
| Security – Improved Audit Granularity | 4 |
| Security - Role Based Access Control Flexibility..... | 4 |
| Security – HDFS Transparent Data Encryption..... | 4 |
| Security – SyncIQ Replication Encryption..... | 5 |
| Quota Enhancements | 5 |
| Quota - Default Directory Quotas..... | 5 |
| Quota - Logical Quotas | 5 |
| Quota - Scalability | 5 |
| Quota - Miscellaneous Features | 5 |
| Backup – Commvault Backup Optimization | 5 |
| Backup - New Hardware | 6 |
| Backup - Performance and Scalability | 6 |
| Backup - CloudPools Integration..... | 6 |
| OneFS Upgrade and Node Firmware Updates | 6 |
| SmartLock Enhancements | 6 |
| Improved Workload Performance Monitoring..... | 6 |
| SyncIQ Bandwidth Throttling..... | 7 |
| File Pool Policy | 7 |
| Improved Performance for EDA Workloads | 7 |
| Small File Efficiency for Archive Workloads | 7 |
| Summary..... | 7 |
| TAKE THE NEXT STEP | 7 |

Introduction

Isilon OneFS 8.2 is a major update to the OneFS operating system and it introduces several new and improved features to simplify storage management, enhance performance and increase security. This latest version of OneFS will enable you to modernize IT, drive new workloads and maximizes the capital of data within your data lake. Two significant features include enhancements to the Isilon CloudPools tiering software and up to a 70% increase in cluster size and performance.

Increased Cluster Capacity and Performance

For over a decade, Isilon nodes have been using InfiniBand switches for the backend intra-cluster node-to-node communication. OneFS 8.1 introduced the use of Ethernet switches for the backend communication while OneFS 8.1.1 added a choice of Dell Ethernet switches like the Z9100 and Z9264 to simplify configurations and provide a Dell-on-Dell solution. A “flat” backend configuration for a network switch or chassis is supported with OneFS. A single Isilon cluster running OneFS 8.2 can now scale up to 252 nodes – up from 144 nodes thus delivering up to 70% increase in capacity and performance.

CloudPools Enhancements

CloudPools delivers cost-optimized data placement for cold or frozen data to a choice of public, private or hybrid cloud providers. The re-architected version of CloudPools in OneFS 8.2 delivers improved quality and performance while adding several new features.

- CloudPools now provides more intelligent snapshot integration by eliminating the copy-on-write after tiering a file, resulting in space saving on the Isilon cluster.
- Isilon SnapshotIQ is now aware of a tiered file and does not place the original file into the snapshot area.
- CloudPools has improved how sparse files are stored in the cloud by not storing the sparse region, which in turn reduces space requirements in the cloud.
- CloudPools is increasingly aware of SmartLock WORM capabilities and prevents a file from being tiered that belongs to a Compliance mode SmartLock domain. Files in Enterprise mode SmartLock domain can still be tiered to the cloud.
- CloudPools has better integration with anti-virus products. Tiered files can now be checked and processed for viruses by anti-virus products, resulting in reduced risk and increased security. By default, anti-virus software ignores CloudPools tiered files.
- CloudPools has increased the chunk size used when it tiers files to the cloud. The chunk size has changed from 1 MB to 2 MB which enables CloudPools to deliver improved recall performance - recalls of tiered files are faster because data is recalled in bigger chunks which means fewer roundtrips to the cloud.
- CloudPools has been optimized to deliver improved I/O access patterns and caching of data and metadata for tiered files.
- Platform Support: CloudPools delivers increased security by now providing Federal Commercial Cloud Services (C2S) support. Federal workloads mandate the use of C2S capabilities such as those available in the private portion of Amazon AWS cloud. C2S support provides a more secure authentication mechanism using client certificates, which provides a secure way to use cloud technology. CloudPools also includes a secure way to store certificates along with a way to validate them using a CAP server that is maintained by the customer. CloudPools also supports AWS V4 authentication which is available across all AWS regions. CloudPools negotiates starting at the highest signature level (V4) and works its way down – first try V4 and if that fails, then switch to V2. CloudPools now delivers new support for Alibaba Aliyun as a cloud target and it is set up, configured and managed similar to Google Cloud or Amazon AWS cloud targets.
- CloudPools simplifies management with improved reporting capabilities. You can now view network statistics associated with different cloud providers or even at the policy level. These statistics like bytes written/read from the cloud are available to InsightIQ and can be segmented based on the Cloud Account. CloudPools reporting like the number of bytes read/written to cloud and the number of read/write requests to cloud can be reported on a per-policy basis, which provides increased granularity of reporting.

Networking – OneFS Networking Enhancements

With support for increased node counts, OneFS 8.2 now includes several features to improve the performance and reliability of data access over the network. An internal OneFS network component called Flexnet delivers faster network configuration at scale, and a rewrite of the SmartConnect code takes advantage of the changes in Flexnet. Flexnet provides OneFS an understanding of network configuration and the network state (eg node information and pool members). Flexnet now stores these independently, increasing stability and responsiveness of the cluster while reducing contention especially as nodes increase in count. This allows Flexnet consumers like SmartConnect to only load specific items of interest which improves performance.

SmartConnect has been redesigned and re-architected to now support 252 nodes. It delivers better handling of DNS functionality by delivering it as a separate component. The SSIP – the service that handles queries from DNS - is much more robust by allowing nodes within a group to choose which node has control of SSIP. Using multiple-SSIPs provides SSIP failover capabilities which minimizes failure points and increases availability. At least 2 SSIPs are now recommended with a maximum of 6 SSIPs per subnet that can be configured with SmartConnect Advanced. Nodes may request and own multiple SSIPs in a subnet to deliver high availability.

Security – Multi-Factor Authentication (MFA)

OneFS 8.2 includes a number of security enhancements to reduce risk and provide tighter access control.

- The SSH protocol can be configured using the CLI to store public keys in LDAP rather than stored in the user's home directory. This simplifies setup and configuration and reduces the chance of misconfiguration.
- The SSH protocol now includes native support for the external Duo (Cisco) service for unified access and security to improve trust between the users and the storage resources accessed. The Duo service that provides MFA access can be used in conjunction with password and/or keys to provide additional security. The Duo service delivers maximum flexibility by including support for the Duo App, SMS, voice and USB keys. As a fallback, specific users and groups in an exclusion list may be allowed to bypass MFA, if specified on the Duo server. It is also possible to generate one-time access to users to accommodate events like a forgotten phone or fallback mode associated with the unavailability of the Duo service. Carefully consider the advantages and disadvantages of providing MFA support in your environment prior to deployment.

Security – Improved Audit Granularity

As clusters grow, the number of audit events increases, resulting in events that are collected but do not need processing by external audit applications. OneFS audit now allows more control in selecting what protocol activity is offloaded for audit. It eliminates the collection of unneeded audit events for which no third-parties have registered, thus improving performance and reliability.

Security - Role Based Access Control Flexibility

Role Based Access Control (RBAC) enables well defined roles and privileges within the cluster, fine-tuning the scope of administrators while simultaneously increasing security and reducing risk. RBAC has been extended to support multi-tenancy, providing fine grained control at an access zone level rather than at the system level. Delegating access to departmental administrators now enables them to have the flexibility and control to manage their own resources with zone level roles and privileges. Apart from built-in roles and privileges, authentication providers can also be configured at a zone level

Security – HDFS Transparent Data Encryption

As big data deployments become more mainstream, Hadoop workflow security becomes more important. HDFS Transparent Data Encryption (TDE) delivers data in-flight and data-at-rest encryption for Hadoop datasets for the Apache, Hortonworks and Cloudera Hadoop stacks. HDFS TDE provides enterprise and government customer the necessary features for regulatory and compliance requirements. Encryption is performed on the client side and can use an external key management server. Data written to and read from these HDFS encryption zones can only be accessed over the HDFS protocol.

Security – SyncIQ Replication Encryption

End to end encryption of replicated data between Isilon clusters increases security and reduces risk. Data in flight encryption of SyncIQ traffic makes data transfer between OneFS clusters more secure, enabling regulatory compliance. X.509 Transport Layer Security (TLS) certificates used by the nodes in the replicating clusters are managed via the SyncIQ certificate store. SyncIQ encryption settings can be set on a global basis or on a per-policy basis.

Quota Enhancements

New or improved SmartQuota features include default directory quotas, quota templates, logical quotas, percentage-based quota limits, multi-email recipients, reporting and alerting. These capabilities simplify quota management and offer administrators improved usability, flexibility, scalability and performance.

Quota - Default Directory Quotas

Prior to OneFS 8.2, default quotas as a management simplification were available only for users and groups. OneFS 8.2 introduces a new type of quota called default directory quota. A default directory quota set on a folder causes its immediate children folders to inherit that directory quota. The inherited quotas on the child folders are setup using linked quotas which inherit the configuration from the parent folder. Any change in the default directory quota of the parent automatically propagates to the children. Home directories and project shares are workflows that will significantly benefit from the use of templates to setup default directory quotas, eliminating the need to manually setup quotas on each directory.

Quota - Logical Quotas

Prior to OneFS 8.2, quota accounting metrics are aligned to account for the physical blocks required to store files on the cluster. OneFS 8.2 introduces a new accounting feature called application logical quotas. Application logical quotas allow users to view and report on quotas and free space as seen by an application in terms of how much capacity is available to store data regardless of any data reduction or tiering technology that is in place. The logical quotas report on total data stored across different tiers including data tiered to the cloud via CloudPools.

Quota - Scalability

Prior to OneFS 8.2, 20,000 quotas were supported per cluster. OneFS 8.2 now supports 500,000 quotas per cluster. The notification daemon has also been enhanced to support up to 20X more emails per seconds than earlier versions.

Quota - Miscellaneous Features

SmartQuotas has been enhanced to send quota notifications to multiple email recipients. Smartquotas can also be used to set advisory and soft quotas a percentage of the hard quota, thus providing improved convenience, flexibility and usability.

Backup – Commvault Backup Optimization

Isilon is compatible with several leading backup software vendors and is a very popular backup target. A new feature of OneFS 8.2 now supports sparse files for Commvault backups which delivers improved storage efficiency, eliminates the need to over-provision storage and hence lowers costs. When files are deleted from the Commvault backup catalog, sparse file reclamation reduces the storage consumption on Isilon. Commvault v11 SP11 and later along with OneFS 8.2 are required to take advantage of this storage optimization capability. The feature is configurable and enabled on a per-share basis on Isilon and is setup and managed from the Commvault application.

Backup - New Hardware

As a replacement to the A100 Backup Accelerator, OneFS 8.2 includes new support for a multi-function front-end network adapter. The new hybrid 10GbE/Fibre Channel NIC has a FPGA module which also provides FC connectivity and supports the NDMP protocol. The NIC has 4 ports – two 10GbE and two 8Gb FC – and can be installed on any “Gen 6” nodes. The card supports two-way NDMP backup to tape drives and can support front-end network traffic from clients at the same time. The cards must be installed in pairs—either 2 or 4 cards per chassis. Nodes within a cluster will route backup requests to nodes that have the hybrid NIC. A major advantage of the hybrid NIC over the dedicated A100 appliance is that the hybrid NIC doesn’t consume any additional rack space.

Backup - Performance and Scalability

As data sets stored on Isilon continue to grow in size and backup windows continue to shrink, optimizing the backup workflow is key for customers. Enhancements in the NDMP protocol layer improve performance and enable faster backups. The NDMP redirector automatically distributes NDMP local backup or restore operations to nodes with less load. It checks the CPU consumption on the nodes along with the number of NDMP sessions already running on each node to select the best node for backup use. The NDMP redirector supports 2-way NDMP backups only. Once the backup starts, an NDMP throttler manages the resources consumed by the backup and restore operations on a node and ensures it doesn’t overwhelm the node, allowing the node to also service other activity.

Backup - CloudPools Integration

NDMP backup also seamlessly integrates with CloudPools to backup and restore smartlinks. Based on a user specified selection, NDMP can back up full files or smartlinks. NDMP will correctly restore the file or the smartlink based on the target version of OneFS and the options selected during the restore.

OneFS Upgrade and Node Firmware Updates

With increased node counts in a cluster comes the possibility of longer OneFS upgrades and node firmware updates. With OneFS 8.2, upgrades can be paused and resumed allowing upgrades to accommodate and span maintenance windows. When an upgrade is paused, the node currently being upgraded will continue until completion but no new node upgrade will start until the process is resumed. It is also possible to install a patch as a post-reboot activity of an upgrade prior to the upgrade being committed. Rolling upgrades to OneFS 8.2 are supported for OneFS 8.0 through OneFS 8.1.2. Note that there is no upgrade path from OneFS 8.1.3 to OneFS 8.2.

Prior to OneFS 8.2, node firmware updates had to be installed per node, a time-consuming operation especially in large clusters. It is now possible to apply node firmware updates on multiple nodes simultaneously by providing a list of nodes to be updated. An upgrade helper tool is used to select an appropriate combination of nodes to be updated as well as a list of nodes that should not be updated at the same time. For example, nodes in a node-pair cannot be simultaneously updated.

SmartLock Enhancements

OneFS 8.2 introduces several optimizations to improve the processing of WORM files especially in a Compliance mode domain. It is possible for conflicts to exist between versions of a file in a Compliance mode domain especially if they are associated with a SyncIQ failover and failback situation. A new SmartLock compliance store maintenance job now helps in cleaning up such orphan files when they expire, thus reclaiming space. By default, the garbage collection job runs monthly.

Improved Workload Performance Monitoring

As clusters increase in size and the number of competing workloads places demands on system resources, administrators asked for visibility to support better cluster resource sharing. OneFS 8.2 introduces a new feature called partitioned performance. Working over the SMB protocol, partitioned performance management allows administrators to define workloads (any combination of directories, shares, users, clients and access zones) and view performance statistics (protocols / disk ops, read / write bandwidth, CPU, IOPs etc.) to identify rogue workloads that consume excessive system resources. Statistics are updated every 30 seconds. Monitoring workloads can also be used with chargeback and resource accounting.

SyncIQ Bandwidth Throttling

Prior to OneFS 8.2, SyncIQ only allowed bandwidth settings at a global level with reservations split evenly among all running policies. With OneFS 8.2, in addition to the global settings, administrators can make bandwidth reservations on a per-policy basis, thus providing fine grained bandwidth control. SyncIQ will try to meet these reservation requests based on what is already running and on the existing bandwidth rules and schedules. If a policy doesn't have a specified reservation, it is allocated from the reserve specified in the global configuration. If there is not enough bandwidth available for a policy to get the requested amount, the system will evenly split the bandwidth across all running policies until they reach the requested reservation.

File Pool Policy

SmartPools delivers optimized policy-based data placement to the appropriate storage tier based on business rules. Scanning a large file system takes time and system resources. New jobs called File Policy and Index Update have been introduced to reduce the CPU and disk impact of applying file pool policies by improving job scan performance. Using snapshot change lists, the File Policy job enforces file pool policies significantly faster and reduces system impact.

Improved Performance for EDA Workloads

OneFS 8.2 includes several performance improvements specifically targeted to EDA workloads based on data create, access and delete patterns. In some EDA applications, large scale file deletes and creates cause impact to the workload until such tasks complete. OneFS 8.2 provides significant improvements in performance with two workflows – delete with snapshots and small file delete. Improved snapshot metadata caching reduces latency by reducing the total cluster snapshot IO. Small file deletes for EDA workloads benefited significantly from asynchronous file removal.

Small File Efficiency for Archive Workloads

OneFS 8.1.1 introduced small file efficiency feature for compressing healthcare PACS workloads. Containerization is managed by the SmartPools job. A storage efficiency estimation tool can provide guidance as to the benefit of containerizing an archive workload. OneFS 8.2 now includes a defragmenter for that compression feature. The defragmenter runs on a schedule and addresses any fragmentation, thus maintaining overall compression efficiency.

Summary

OneFS 8.2 is a major upgrade to the OneFS operating system and includes 70% larger clusters, enhanced CloudPools support, improvements in scalability, performance, security along with simplified management.

TAKE THE NEXT STEP

Contact your Dell EMC sales representative or authorized reseller to learn more about how Isilon scale-out NAS storage solutions can benefit your organization.

[Shop Dell EMC Isilon](#) to compare features and get more information.



[Learn more](#) about Dell EMC Isilon solutions



[Contact](#) a Dell EMC Expert



[View more](#) resources



Join the conversation
with #DellEMCStorage