

# TECHNICAL OVERVIEW OF NEW AND IMPROVED FEATURES OF EMC ISILON ONEFS 7.1.1

## **ABSTRACT**

This introductory white paper provides a technical overview of the new and improved enterprise grade features introduced in EMC Isilon OneFS 7.1.1. OneFS 7.1.1 introduces key improvements in performance, security, manageability, and data protection

July, 2014

Copyright © 2014 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided "as is." EMC Corporation 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 EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com).

All other trademarks used herein are the property of their respective owners.

Part Number H13173

## **TABLE OF CONTENTS**

<b>EXECUTIVE SUMMARY</b>	<b>4</b>
AUDIENCE	4
<b>INTRODUCTION</b>	<b>4</b>
<b>PERFORMANCE DRIVES NEW AND IMPROVED WORKLOADS</b>	<b>4</b>
SmartFlash (L3 cache):	4
SMB 3.0 multi-channel:	5
NDMP:	5
SyncIQ:	5
<b>SECURITY ENHANCEMENTS IMPROVE PROTECTION</b>	<b>5</b>
Self-encrypting drives:	5
Access zones:	6
Audit:	6
Role based Access Control (RBAC):	6
<b>MANAGEABILITY</b>	<b>6</b>
MMC integration:	6
Drive Support Package (DSP):	6
Non-Disruptive Drive Firmware Update (NDFU):	6
Hadoop:	6
<b>SUMMARY</b>	<b>6</b>

## EXECUTIVE SUMMARY

There is an explosive growth in unstructured data across all enterprises driven by the internet of things. By many industry estimates, over 80% of the new storage capacity installed in organizations around the world will be for unstructured data.

EMC® Isilon® OneFS® 7.1.1 is the industry's leading Scale-out NAS platform that is known for its massive capacity, operational simplicity, extreme performance, and unmatched storage utilization of unstructured data. This white paper provides a technical overview of the features introduced in OneFS 7.1.1. OneFS 7.1.1 allows customers to drive improved workloads to achieve new levels of performance through innovative features like SmartFlash leveraging flash as cache and SMB 3.0 multi-channel support.

This release also introduces the scale-out data lake which allows customers to store, manage, and secure their unstructured data for both traditional, second platform workloads as well as emerging, third platform ones.

## AUDIENCE

This white paper is intended for EMC Customers, Partners, and Employees. It is aimed at Storage and Windows Administrators that want a high-level overview of the technical features that are new or updated in OneFS 7.1.1.

## INTRODUCTION

EMC® Isilon® is a scale-out enterprise data lake that stores, manages, protects, secures, and enables reporting and analytics on unstructured data for both traditional and for emerging workloads. The Isilon core building blocks which drive today and tomorrow's workloads include the OneFS operating system, the scale-out NAS architecture, the scale-out data lake, and enterprise-grade software features. In addition, Isilon leverages capabilities like data protection, data management, performance management, and security features to provide best in class storage services and data services surrounding the scale-out data lake.

Isilon OneFS 7.1.1 (aka "Jaws") drives innovation to a new level with a focus around performance, enterprise features, security, and manageability. As data grows exponentially, customers continue to expand their second platform traditional workloads of file shares, home directories, NAS archives, and HPC. In addition, customers are looking for comprehensive storage solutions that can help them to expand into third platform emerging workloads of cloud, big data analytics, and mobility. By blending industry-leading storage efficiency, security, scalability, and simplicity, OneFS 7.1.1 provides a compelling scale-out data lake storage repository for today and the future.

## Performance drives new and improved workloads

The overarching innovation driving the OneFS 7.1.1 release is an emphasis on scaling performance to run new workloads across different verticals and industries. In addition, there is a distinct need for delivering the performance needs of existing workloads and driving them to the next level. The performance features in this release include support for a new SmartFlash leveraging flash as cache, SMB 3.0 multi-channel, NDMP improvements, and SyncIQ® enhancements.

**SMARTFLASH (L3 CACHE):** Solid-state drives (SSD) can be deployed on the Isilon cluster to increase cache size and speed up file system performance across larger working file sets. Like most operating systems, OneFS uses available random access memory (RAM) as cache to hold frequently accessed file and metadata. This enables OneFS to better optimize data protection and file system performance. The benefits of RAM cache are limited only by capacity. Normally, when RAM cache is full, OneFS removes the oldest cached data and processes new data requests by accessing the storage drives. This cycle is repeated each time RAM cache reaches capacity. To mitigate the cache cycling issue, you can specify that SSDs in your node pools should be used as SmartFlash scale-out flash tier. SmartFlash adds significantly to the operating system's available cache space and is much faster than accessing data on hard disk drives (HDD).

Caching occurs in OneFS at multiple levels and for multiple types of data. The following table describes the types of file system cache available on an Isilon cluster and shows how the SmartFlash tier fits into the overall OneFS cache structure.

Name	Type	Profile	Description
<b>L1 cache</b>	RAM	Volatile	Also known as front-end cache. It holds clean copies of file system meta data and data requested by the front-end network (through NFS, SMB, HTTP, and so on).
<b>L2 cache</b>	RAM	Volatile	Also known as back-end cache. It holds clean copies of file system meta data and data on the node that owns the data
<b>SmartCache</b>	NVRAM	Nonvolatile	Holds any pending changes to front-end files waiting to be written to storage. If there is a power failure, after power is restored, OneFS writes those changes to storage
<b>SmartFlash (L3 cache)</b>	SSD	Nonvolatile	Holds file data and meta data evicted from L2 cache, effectively increasing L2 cache capacity

During normal operation, as L2 cache reaches capacity, OneFS evicts older data to free up L2 cache space. OneFS evaluates data to be evicted and, depending on your workflow, moves the data to the SmartFlash tier. In this way, much more of the most frequently accessed data is held in cache, and overall file system performance is improved. SmartFlash tier accelerates file access only on the nodes where the SSDs reside. The SmartFlash tier extends caching capacity and therefore increases overall file system performance. In addition, in case of a power failure, the data on the SmartFlash tier is retained and still available after power is restored. SmartFlash provides the most benefit in workflows that involve random file access and where blocks are reused.

**SMB 3.0 MULTI-CHANNEL:** The SMB protocol allows Windows® clients to connect to shares on an Isilon cluster. In the previous OneFS releases, Isilon has supported SMB 1.0 and 2.x. With the SMB 3.0 multi-channel support, an appropriately hardware configured Windows 8, Windows Server 2012, or later clients can connect to an Isilon cluster and take advantage of the performance and availability features. No manual SMB configurations are needed on the Windows machine or on Isilon to enable this support. SMB 3.0 multi-channel offers simultaneous SMB client connections to a single Isilon session thus providing increased throughput, connection failure tolerance, and automatic discovery. The Windows client must have two or more network cards, or one or more network cards that support Receive Side Scaling (RSS), or one or more network cards with link aggregation enabled. The SMB connection is limited to the single node and connections aren't shared across nodes. SMB 3.0 multi-channel can significantly increase throughput for the appropriate workload running on the Windows machine.

**NDMP:** NDMP backup and restore performance have been improved with the parallelism of the data transfer operations. From the earlier, predominantly single threaded architecture, parallel restore means that multiple threads are used to restore files, thus attempting to read and write data as fast as the tape can deliver it. Additional enhancements include changes to the NDMP server's internal data structures and code flows that enable improved throughput for transferring small files. Another enhancement added to the backup is the ability to specify a file list of files to be copied. This list could potentially be an output of an application or a project of content to be archived on to Isilon. In addition to the flexibility, using a file list eliminates a costly tree walk to identify and gather the files and hence shortens the time for a backup.

**SYNCIQ:** Replication of data from one Isilon cluster to another using SyncIQ has been enhanced by improving the underlying technology. File splitting allows a large file to be partitioned across multiple threads and replicated in parallel. In earlier versions of OneFS, the SyncIQ implementation works at a file level granularity. It was possible for a single large file at the end of the replication to be processed by just one worker thread even though there were multiple threads that could simultaneously process the file. In this release, file sub-ranges can be divided and split across multiple threads resulting in potentially shorter replication cycles, and higher efficiency. This would likely result in more consistent throughput than previous versions of OneFS.

## Security enhancements improve protection

OneFS provides enhanced security and robust enterprise class features. Earlier versions of OneFS supported auditing, role based access control (RBAC), data at rest encryption via Self-Encrypting Drives (SEDs), SmartLock®, and access zones, which form a strong foundation for a security platform while providing the needed flexibility to work at the speed of business. This release of OneFS introduces new SEDs and provides incrementally improved features for access zones, audit, and RBAC security realms.

**SELF-ENCRYPTING DRIVES:** This release of OneFS expands the availability of SEDs to provide Data at-rest Encryption (DARE) capabilities across the entire node family. In addition to the 3TB and 4TB SEDs, OneFS 7.1.1 introduces a 900GB SAS SED HDD for the S-Series and an 800GB SED SSD for all supported nodes.

**ACCESS ZONES:** Access zones provide the mechanism to partition the cluster into logical units enabling secure segregation of resources. Zones are a simplified version of multi-tenancy in that it provides a way to carve the resource stack into virtualized data stores. In this OneFS release, access zones have been restructured to enforce best practices by requiring a root or a base directory to be designated for each access zone. A base directory designates a top level directory inside the access zone, defines the file system view of the access zone, and is used to define file system separation. SMB shares are now tied to a single access zone and zones can no longer be used to share data. All this provides better separation of resources while maintaining enhanced levels of security. In addition to SMB support, access zones now support a HDFS namespace per access zone which means you can now run multiple separate HDFS namespaces in the same cluster. Lastly, OneFS also prevents access to non-system zones through NFS, SSH, and the web administrative interface, thus enhancing security.

**AUDIT:** OneFS 7.1 introduced the concept of forwarding SMB and configuration audit events to a Common Event Enabler (CEE) server for processing. OneFS 7.1.1 has improved the audit capabilities by enabling the forwarding of audit system configuration information to syslog for storage and analysis. Syslog can be used to filter and display needed information.

**ROLE BASED ACCESS CONTROL (RBAC):** RBAC has been enhanced to add a greater degree of granularity to the roles and the actions these roles can perform. A new web UI controls the creation, delegation, setting, and changing the roles and providing simplicity of operations. Several new roles have been added like `ISI_PRIV_IFS_BACKUP` and `ISI_PRIV_IFS_RESTORE` that allow a user to backup files without having explicit permissions to these files. Essentially, this feature allows you to have the role of Backup Administrator assigned to a designated set of users. This feature is only enabled over SMB protocol. A new utility has been added that enables an administrator to determine the effective permission on a file or directory, hence providing enhanced troubleshooting capability. Additionally, the Platform API has been enhanced to streamline the handling of roles.

## Manageability

OneFS 7.1.1 extends the enterprise manageability features to make it simpler to manage large amounts of storage. The use of Microsoft Management Console (MMC) to manage shares, a new drive support package, and an advanced disk firmware management system has been added to this release.

**MMC INTEGRATION:** MMC integration allows for Windows administrators to use familiar tools to flexibly manage the creation of shares on Isilon. A suitably privileged administrator can remotely administer a share through the MMC shared folders snap-in, which allows you to connect to an access zone and directly manage all the shares for that zone. The typical operations that can be performed include creating and deleting shared folders, and configuring access permissions to a shared folder. You can also view a list of active SMB connections and close any open SMB session. The Isilon cluster must be joined to an Active Directory domain from which the MMC console must be invoked.

**DRIVE SUPPORT PACKAGE (DSP):** A new DSP has been added to OneFS to simplify and streamline the management of existing and new drives across different vendors. By removing any hardcoding of drive information and managing that information via the DSP, the administrator can maintain the drive firmware much easier. No longer is a patch or reboot required to swap a replacement drive from a different vendor, or the ability to add a new drive. This ultimately reduces the administrator's reluctance to keep the firmware up to date and also reduces the likelihood of failures due to known drive issues. These DSP's are created and managed by Isilon.

**NON-DISRUPTIVE DRIVE FIRMWARE UPDATE (NDFU):** this new feature eliminates the need for a reboot when the drive firmware is updated. It automatically determines and applies the update for the drive's firmware without taking a node down. Alerts can be configured to indicate the need for a drive firmware update. NDFU leverages the DSP described above. NDFU is a big step towards lights-out operation and maximum uptime.

**HADOOP:** OneFS 7.1.1 includes some enhancements to HDFS™ especially in the support of WebHDFS, providing greater simplicity. WebHDFS is a REST-based API that allows clients to access Hadoop® without having to talk the native HDFS protocol. Apache™ Hue is a common utility to access WebHDFS. HDFS is also integrated with access zone enabling multiple namespaces in a single deployment. Per zone authentication is possible with Active Directory or LDAP. HDFS is now integrated with the Platform API, which provides integration points with third-party management and configuration tools.

## Summary

This release of OneFS 7.1.1 drives increased performance to enable new workloads and run existing workloads faster. Innovation in the use of SmartFlash leveraging flash as cache enables simplicity and efficiency in the face of storage growth.