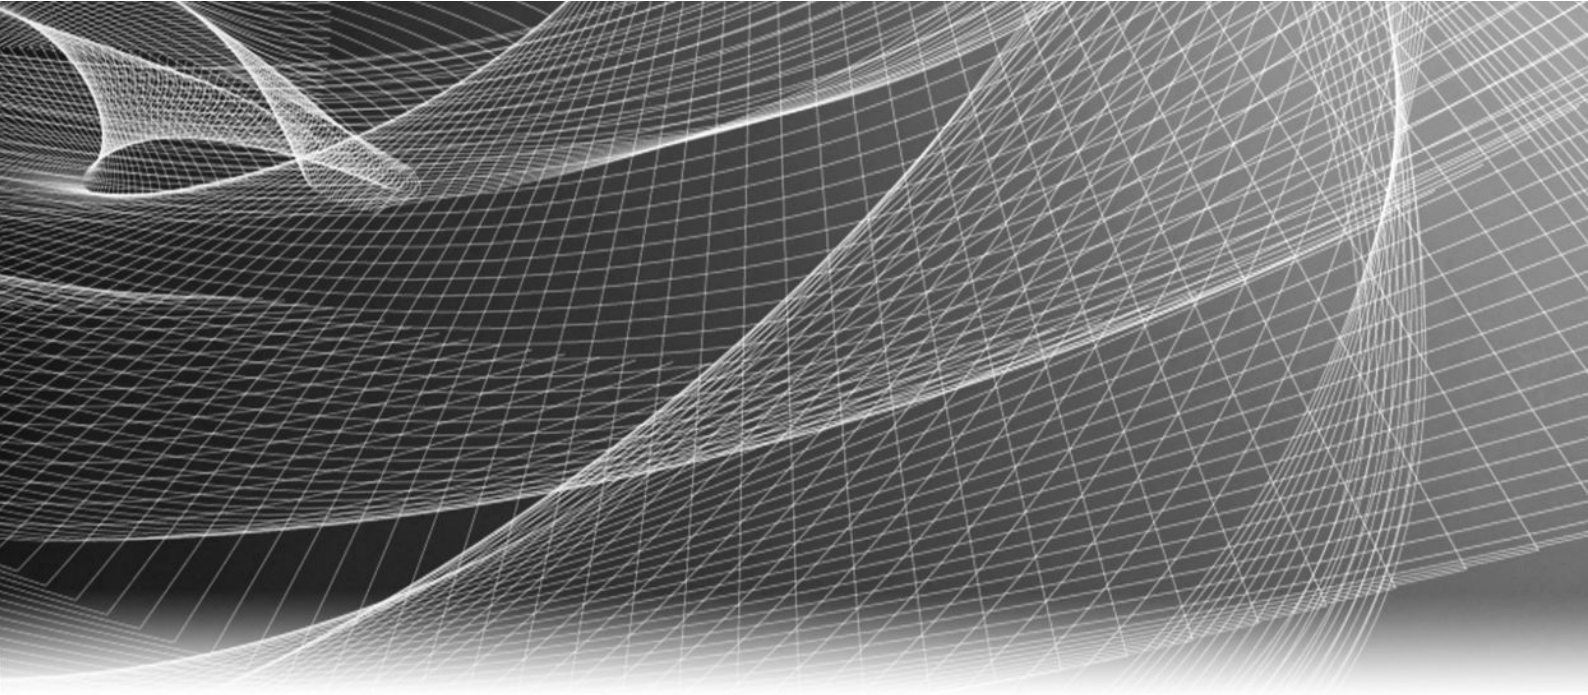


Isilon OneFS

Version 7.2.1.0 - 7.2.1.2

Release Notes

EMC²

Copyright © 2016 EMC Isilon. All rights reserved. Published in the USA.

Published February 11, 2016

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.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (https://support.emc.com).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Chapter 1 OneFS Release Notes 7

OneFS 7.2.1 Release notes ... 8

Chapter 2 Upgrading OneFS 9

Target Code ... 10
Supported upgrade paths ... 10

Chapter 3 New features, software support, logging, and controls 13

New and changed in OneFS 7.2.1 - Highlights ... 14
File system ... 14
Hardware ... 14
Networking ... 15
OneFS API ... 15
Security Profiles ... 15
New and changed in OneFS 7.2.1.2 ... 15
Authentication ... 15
Events, alerts, and cluster monitoring ... 16
File system ... 16
Networking ... 17
NFS ... 17
OneFS web administration interface ... 17
SmartLock ... 18
SMB ... 18
New and changed in OneFS 7.2.1.1 ... 18
Authentication ... 18
Backup, recovery, and snapshots ... 19
Cluster configuration ... 19
Events, alerts, and cluster monitoring ... 19
Hardware ... 20
HDFS ... 21
Networking ... 21
Security ... 22
SMB ... 23
Upgrade and installation ... 23
New and changed in OneFS 7.2.1.0 ... 25
Authentication ... 25
Networking ... 25
Security ... 26

Chapter 4 New hardware and firmware support 27

New hardware and firmware support in OneFS 7.2.1.2 ... 28
New hardware and firmware support in OneFS 7.2.1.1 ... 28
New hardware and firmware support in OneFS 7.2.1.0 ... 28

Chapter 5 Resolved issues 29

Resolved in OneFS 7.2.1.2 ... 30

Authentication 30
Backup, recovery, and snapshots 31
Cluster configuration 33
Events, alerts, and cluster monitoring 34
File system 35
File transfer 36
Hardware 36
HDFS 37
Job engine 38
Migration 40
Networking 40
NFS 41
OneFS web administration interface 42
Security 43
SMB 44
Upgrade 45
Resolved in OneFS 7.2.1.1 45
Antivirus 45
Authentication 46
Backup, recovery, and snapshots 48
Cluster configuration 54
Diagnostic tools 54
Events, alerts, and cluster monitoring 55
File system 58
File transfer 63
Hardware 63
HDFS 66
Job engine 68
Migration 68
Networking 69
NFS 71
OneFS API 74
OneFS web administration interface 75
Security 75
SmartLock 75
SmartQuotas 76
SMB 77
Upgrade 81
Resolved in OneFS 7.2.1.0 82
Events, alerts, and cluster monitoring 82
Hardware 82
Networking 82
NFS 83
OneFS web administration interface 84
Upgrade and installation 84

Chapter 6 OneFS patches included in this release 85

Patches included in OneFS 7.2.1.2 86
Patches included in OneFS 7.2.1.1 86

Chapter 7 Known issues 89

Antivirus 90
Authentication 90
Backup, recovery, and snapshots 92

Cluster configuration 95
Command-line interface 96
Diagnostic tools 96
Events, alerts, and cluster monitoring 97
File system 98
File transfer 101
Hardware 101
HDFS 102
iSCSI 103
Job engine 103
Migration 105
Networking 106
NFS 107
OneFS API 108
OneFS web administration interface 109
Security 109
Security profiles 109
SmartLock 111
SmartQuotas 111
SMB 111
Upgrade and installation 112
Virtual plug-ins 114

Chapter 8 OneFS Release Resources 115

OneFS information and documentation 116
Functional areas in the OneFS release notes 117
Where to go for support 120
Provide feedback about this document 121

CHAPTER 1

OneFS Release Notes

The OneFS release notes contain information about new features, changes in functionality, issues that are resolved, support for new hardware and firmware, and known issues and limitations in the Isilon OneFS 7.2.1 operating system.

- OneFS 7.2.1 Release notes 8

OneFS 7.2.1 Release notes

The OneFS 7.2.1 release notes contain descriptions of all of the enhancements, functionality changes, new features, support for hardware, support for firmware, and resolved issues that are included in the release.

- OneFS 7.2.1.2 released: February 11, 2016 **(Target Code)**
- OneFS 7.2.1.1 released: November 23, 2015
- OneFS 7.2.1.0 released: July 27, 2015

The new features, functionality changes, resolved issues, and known issues listed in the release notes are categorized by functional area. For a list of the functional areas used to categorize the release notes and a brief description of what each functional area typically contains, see the *Functional areas in the OneFS release notes* section in the *OneFS release resources* section at the end of this document.

For a list of available OneFS releases and information about target code releases and general availability (GA) releases, see Current Isilon Software Releases on the EMC Online Support site.

CHAPTER 2

Upgrading OneFS

OneFS upgrades comprise a full operating system upgrade and require that the Isilon cluster be rebooted. To help ensure that the version of OneFS to which you upgrade contains all of the resolved issues included in the version you are upgrading from, upgrades are supported only from designated previous releases of OneFS.

Before upgrading OneFS, review the *Supported upgrade paths* section of this document to verify that the cluster can be upgraded from your current version of OneFS directly to this release.

See the OneFS Upgrade Planning and Process Guide on the EMC Online Support site for detailed upgrade instructions and additional upgrade information.

To download the installer for this maintenance release, see the OneFS Downloads page on the EMC Online Support site.

- Target Code 10
- Supported upgrade paths 10

Target Code

OneFS 7.2.1.2 is the current 7.2.1.x target code version. A OneFS release is designated as Target Code after it satisfies specific criteria, which includes production time in the field, deployments across all supported node platforms, and additional quality metrics. For information about upgrading to OneFS Target Code, see Upgrading to OneFS Target Code on the Isilon EMC Community Network (ECN) pages.

Supported upgrade paths

To ensure that the version of OneFS you are upgrading to contains all of the bug fixes included in the version of OneFS you are upgrading from, upgrades are only supported from specified versions of OneFS.

Upgrade resources

For more information about the differences between simultaneous and rolling upgrades and other important details about the OneFS upgrade process—including recommended pre-upgrade tasks—see the OneFS Upgrades - Isilon Info Hub on the EMC.com web site.

Upgrades to OneFS 7.2.1.2 (Target Code)

Simultaneous and rolling upgrades to OneFS 7.2.1.2 are supported from the following OneFS versions:

- OneFS 7.2.1.0 through 7.2.1.1
- OneFS 7.2.0.2 through 7.2.0.5
- OneFS 7.1.1.4 through 7.1.1.8

If you plan to upgrade to OneFS 7.2.1.2 from one of the following versions of OneFS, you must first perform a rolling upgrade or a simultaneous upgrade to OneFS 7.2.0.2 through 7.2.0.5:

- OneFS 7.2.0.0 through OneFS 7.2.0.1

If you plan to upgrade to OneFS 7.2.1.2 from one of the following versions of OneFS, you must first perform a rolling upgrade or a simultaneous upgrade to either OneFS 7.2.0.2 through 7.2.0.5 or OneFS 7.1.1.4 through 7.1.1.8:

- OneFS 7.1.1.0 through OneFS 7.1.1.3
- OneFS 7.1.0.0 through OneFS 7.1.0.6

If you plan to upgrade to OneFS 7.2.1.2 from one of the following versions of OneFS, you must first perform a simultaneous upgrade to either OneFS 7.2.0.2 through 7.2.0.5 or OneFS 7.1.1.4 through 7.1.1.8:

- OneFS 7.0.2.0 through OneFS 7.0.2.13
- OneFS 7.0.1.0 through OneFS 7.0.1.10
- OneFS 7.0.0.0 (7.0)

CHAPTER 3

New features, software support, logging, and controls

This section contains descriptions of new features, new software support, new protocol and protocol version support, additional logging, and new controls such as command-line options and sysctl parameters.

New features enable you to perform tasks or implement configurations that were previously unavailable.

These new features include:

- New software and protocol support
- Updated software and protocol version support
- New logging
- New controls such as command options, sysctl parameters, and OneFS web administration controls

Functionality changes include modifications and enhancements to OneFS that enable you to perform preexisting tasks in new ways, or that improve underlying OneFS functionality or performance. These changes also include removing support for deprecated protocols and software.

The functionality changes documented in the release notes include:

- Changes to the formatting or syntax of a pre-existing command
- Changes to underlying code to improve performance
- Updates to integrated OneFS components such as OpenSSL and GNU bash
- Changes to enable functionality in the OneFS web administration interface that was previously available only from the command-line interface
- Changes to remove support for old and deprecated protocols

- New and changed in OneFS 7.2.1 - Highlights 14
- New and changed in OneFS 7.2.1.2 15
- New and changed in OneFS 7.2.1.1 18
- New and changed in OneFS 7.2.1.0 25

New and changed in OneFS 7.2.1 - Highlights

File system

Global namespace acceleration and L3 cache interaction

In OneFS, solid-state drives (SSDs) can be used either for SSD strategies with Global namespace acceleration (GNA) mirroring or for L3 cache. In OneFS 7.2.1, any SSDs that are used for L3 cache are not counted against the GNA requirements, which are as follows:
GNA requires that at least 20% of the nodes in the cluster contain at least one SSD, and that at least 1.5% of the total cluster storage is SSD-based. If the ratio of SSDs to non-SSDs on the cluster falls below the 1.5% threshold, GNA becomes inactive even if enabled. GNA is reactivated when the ratio is corrected. Because L3 cache offers acceleration using a caching mechanism rather than mirroring, and is either simply on or off, you might want to consider transitioning to L3 cache exclusively for your SSDs.

Hardware

Node class compatibility

Node class compatibility enables you to deploy a newer-generation node to a cluster containing nodes from an earlier generation. For example, if your cluster already has a node pool made up of Isilon X200 nodes, you can define a compatibility that enables you to add an X210 node to the same node pool. Compatibilities for S200/S210, X200/X210, X400/X410, and NL400/NL410 nodes are supported. Compatible node generations must have identical HDD and SSD layouts and must have a compatible memory configuration as shown in the following table:

S200/S210 Compatibility		X200/X210 Compatibility		X400/X410 Compatibility		NL400/NL410 Compatibility	
S200 RAM	S210 RAM	X200 RAM	X210 RAM	X400 RAM	X410 RAM	NL400 RAM	NL410 RAM
24 GB	32 GB	6 GB	Not available	24 GB	32 GB	12 GB	Not available
48 GB	64 GB	12 GB		48 GB	64 GB	24 GB	24 GB
96 GB	128 GB	24 GB	24 GB	96 GB	128 GB	48 GB	48 GB
	256 GB	48 GB	48 GB	192 GB	256 GB	–	–

SSD compatibility

SSD compatibility enables nodes with different capacity SSDs to be added to the same node pool. For example, if your cluster has a node pool made up of Isilon S200 nodes with 100 GB SSDs, and you purchase new S200 nodes with 200 GB SSDs, you can create an SSD compatibility so that the new S200 nodes can be provisioned into the existing S200 node pool. For nodes of different generations—for example, S200 and S210 nodes—you might have to create both a node class compatibility and an SSD compatibility to enable the S210 nodes to be provisioned into an S200 node pool.

Support for new platforms

OneFS 7.2.1 supports two new platforms, the X210 and the NL410. These nodes support node class compatibility; X210 nodes can be added to existing X200 node pools and NL410 nodes can be added to existing NL400 node pools as described in the table above.

Networking

Enhanced IPv6 support

Dual-stack support: Dual-stack support enables you to configure the cluster to accept client connections through IPv4 and IPv6 concurrently.

IPv6 environments: Beginning in OneFS 7.2.1, IPv6-only environments are supported.

Note

The following protocols and tools are not supported over IPv6: HDFS, InsightIQ 3.1 and earlier, MMC, NIS, and Swift

OneFS API

Enhanced Swift support

Isilon Swift is now supported over both the HTTP and secure HTTP (HTTPS) protocols.

Note

Isilon Swift supports the HTTP protocol with IPv6. Isilon Swift also supports the HTTP protocol with IPv4 if STIG hardening is disabled. Isilon Swift does not support the HTTPS protocol with IPv6.

Security Profiles

Security Technical Implementation Guide (STIG)

This version of OneFS introduces a license-based security hardening feature that enables you to apply and revert a hardening policy on the EMC Isilon cluster to meet the Defense Information Systems Agency (DISA) STIG configuration standards.

New and changed in OneFS 7.2.1.2

Authentication

New and changed in OneFS 7.2.1.2	ID
The `/usr/bin/isi_hwtools/isi_hdfw_update` command was added to the `sudoers` file. The `sudoers` file defines the commands that a user with sudo privileges is permitted to run. This addition allows the compadmin user to update drive firmware on clusters running in compliance mode by running the `/usr/bin/isi_hwtools/isi_hdfw_update` command.	163612

Events, alerts, and cluster monitoring

New and changed in OneFS 7.2.1.2	ID

Because the compadmin user does not have privileges to run the `rm` command, the compadmin user was unable to clear the CELOG database in OneFS 7.2.1.1 and earlier. Beginning in OneFS 7.2.1.2 a script was added that enables you to clear the CELOG database. This script can be run on cluster in compliance mode or non-compliance mode.
To run the script to clear the database on a cluster running in compliance mode, log in using as the compadmin user and run the following command:
`sudo /usr/bin/isi_celog/clear_aged_events.py`
To run the script to clear the database on a cluster running in non-compliance mode, log in as the root user and run the following command:
`/usr/bin/isi_celog/clear_aged_events.py`
By default, the script will clear events that have been retained for more than 90 days. To clear the database of events that have been retained for more or fewer than 90 days, specify the desired number of days when you run the script. For example, if you want to clear events that have been retained for more than 45 days, run the following command:
`/usr/bin/isi_celog/clear_aged_events.py 45`
You can also specify the `--dryrun` option when you run the script. The `--dryrun` option does not clear any events from the database. This option provides information about which events will be cleared when the script is run without the `--dryrun` option. When the option is included, information that describes how many events will be cleared from the database and the specific dates that will be used to determine which events to clear will be returned.
	160308

A new OneFS system event, 400160001, is generated when a configured Common Event Enabler (CEE) server cannot be reached. If the CEE server becomes available, the event is cleared. If a CEE server is configured incorrectly and therefore can never be reached—for example, if there is a typographical error in the CEE server URI—the event must be manually cleared.
`ID: 1.73`
`Type: 400160001`
`Severity: warn`
`Value: 0.0`
`Message: Audit CEE server unavailable: '<CEE_server_address>': Bad CEE response status: 18`
	157378

File system

New and changed in OneFS 7.2.1.2	ID
Logging to the `ec_logstore_recover_file` log file was improved as follows: • The default logging level for the endurant cache logstore recovery process changed from `NOTICE` to `INFO`. • The number of files that are successfully recovered following the recovery process is recorded in the log file.	158951

New and changed in OneFS 7.2.1.2	ID

In the following example of the log file output, 10 files were recovered:
`Logstore 1:0007:0013: 10 files to recover data for (error 0)`
`10/10 files were successfully recovered from logstore`	

Networking

New and changed in OneFS 7.2.1.2	ID
A new kernel driver for the QLogic (formerly Broadcom) NetXtreme Ethernet (BXE) 10 GigE NIC, 1.78.79.ISILON.01, was integrated into OneFS 7.2.1.2.	164251

Logging to the `ec_logstore_recover_file` log file was improved as follows:
• The default logging level for the endurant cache logstore recovery process changed from `NOTICE` to `INFO`.
• The number of files that are successfully recovered following the recovery process is recorded in the log file.
In the following example of the log file output, 10 files were recovered:
`Logstore 1:0007:0013: 10 files to recover data for (error 0)`
`10/10 files were successfully recovered from logstore`
	158951

NFS

New and changed in OneFS 7.2.1.2	ID
Mounting an external NFS file system onto a node that is part of an Isilon cluster is not supported. Beginning in OneFS 7.2.1.2, if an external NFS file system is mounted on a node in the cluster, the following message appears on the console after you run the `mount` command, where ‹PATH› is the path to the external file system: `WARNING: <PATH> is being mounted. Use of Isilon nodes as NFS clients is unsupported.`	162439

OneFS web administration interface

New and changed in OneFS 7.2.1.2	ID
The time zone database that OneFS relies on—when you configure the time zone for the cluster—was updated to Time Zone Data v.2015g. This database is made available by the Internet Assigned Numbers Authority (IANA). For more information about the changes to this version of the time zone database, see the IANA website.	162441

SmartLock

New and changed in OneFS 7.2.1.2	ID
The `/usr/bin/isi_hwtools/isi_ecclist_config` command was added to the `sudoers` file, which is a file that defines which commands a user with sudo privileges is permitted to run. These additional commands enable EMC Isilon Technical Support staff to troubleshoot clusters that are in compliance mode.	162716

SMB

New and changed in OneFS 7.2.1.2	ID
New control: A OneFS registry key was added to the `gconfig` file, MaxSMB2DialectVersion. This key provides the ability to control the highest SMB2 dialect of the SMB2 protocol. The valid options and the effect of setting each option are as follows: • `2.02` • This option enables the SMB 2.02 dialect and disables the SMB 2.10 and 3.0 dialects. • `2.10` • This option enables the SMB 2.02 and 2.10 dialects and disables the SMB 3.0 dialect. • `3.00` • This option enables all three SMB dialects. **Note** The default setting is option `3.00`. If an invalid option is selected, a message similar to the following appears in the `lwiod.log` file: `INFO: Invalid SMB2 protocol max dialect configuration. MaxSmb2DialectVersion is not in allowed list (2.02, 2.10, 3.00). Default MaxSmb2DialectVersion value is in use.` **Note** This option was added to enable support to troubleshoot issues that might be related to differences between the supported SMB2 dialects.	158080

New and changed in OneFS 7.2.1.1

Authentication

New and changed in OneFS 7.2.1.1	ID
A user that attempts to connect to the cluster over SSH, through the OneFS API, or through a serial cable, can no longer be authenticated on clusters running in	156328

New and changed in OneFS 7.2.1.1	ID
compliance mode if any of the following identifiers are assigned to the user as either the user's primary ID or as a supplemental ID: • UID: 0 • SID: S-1-22-1-0 For more information, see ESA-2015-148 on the EMC Online Support site.	
The message that is logged in the `/var/log/lsassd.log` file when a trusted Active Directory domain is offline now includes the name of the domain that cannot be reached. In the example below, *‹domain_name›* is the name of the domain that is offline: `[lsass] Domain '<domain_name>' is offline`	155805

Backup, recovery, and snapshots

New and changed in OneFS 7.2.1.1	ID
If you run the `stat` command to view information about a file, the Snapshot ID of the file is now included in the output. This information appears in the st_snapid field.	154833
Reduces lock contention by changing the lock type used by the SyncIQ coordinator when reading the `siqpolicies.gc` file coordinator from an exclusive lock to a shared lock.	151757

Cluster configuration

New and changed in OneFS 7.2.1.1	ID
More detailed logging was added to help diagnose issues that occur when SmartPools are upgraded during a OneFS upgrade and issues that occur after running the `smartpools-upgrade` command. **Note** The more detailed information appears in the `/var/log/messages` file.	151958
The output from the `sysctl efs.bam.disk_pool_db` command now shows the `equivalence_id` for pool groups. The ID helps Isilon Technical Support identify internal datastructure values when troubleshooting issues related to storage pools.	151748

Events, alerts, and cluster monitoring

New and changed in OneFS 7.2.1.1	ID
Usernames were added to the LOGON and LOGOFF events that the syslog forwarder sends to the `/var/log/audit_protocol.log` file. In the following example	153427

New and changed in OneFS 7.2.1.1	ID

entries, <SID> is the user's security ID, <ADNAME> is the name of the Active Directory (AD) domain to which the user belongs, and <username> is the name of the user that logged onto or logged off of the cluster.
 	
`<SID>|1|10.7.145.202|LOGON|SUCCESS|<ADNAME>\<username>`
`<SID>|1|10.7.145.202|LOGOFF|SUCCESS|<ADNAME>\<username>`
 	
Note
If the AD configuration changes between the time that the LOGON and LOGOFF events are captured and the time that the syslog forwarder sends the events to the `audit_protocol.log` file, the recorded user and domain names might be invalid or unavailable.	

Hardware

New and changed in OneFS 7.2.1.1	ID
Wear life thresholds were added for the system area on the following Sunset Cove Plus SSD drive models: • l HGST HUSMM1620ASS200 • l HGST HUSMM1640ASS200 • l HGST HUSMM1680ASS200 • l HGST HUSMM1680ASS205 • l HGST HUSMM1616ASS200 The addition of these thresholds enables OneFS to generate alerts and log events if the wear life of the system area on these SSD drive models reaches 88 percent (warn), 89 percent (critical), or 90 percent (smartfail).	154585

New control
Options were added to the `isi_dsp_install` command to enable you to display the version number of the most recently installed drive support package (DSP) or to display a list of previously installed DSPs.
To display the version number of the most recently installed DSP, run the following command:
 	
`isi_dsp_install --latest`
 	
Output similar to the following is displayed:
 	
`2015-06-22 15:02:21 || Drive_Support_v1.7.tar`
 	
To display a list of previously installed DSPs, run the following command:
 	
`isi_dsp_install --history`
	154223

New and changed in OneFS 7.2.1.1	ID

Output similar to the following is displayed:
 	
`2015-06-22 15:00:20 || Drive_Support_v1.5.tar`
`2015-06-22 15:01:36 || Drive_Support_v1.6.tar`
`2015-06-22 15:02:21 || Drive_Support_v1.7.tar`	

The error that appears if you run the `isi_dmilog` command on a platform that does not support the command was changed.
Previous error:
 	
`dmilog functions not supported on this platform`
 	
New error:
 	
`dmilog functions not supported on this platform - please consult 'isi_hwmon -h'`
 	
For more information about the `isi_hwmon` command, see article 199270 on the EMC Online Support site.
	154317

To improve the response time on A100, S210, X410, and HD400 nodes, redundant requests for a node's sensor data were removed from the `isi_hw_status` command.	152141
The node format command was improved so that the progress of the node format operation is reported in percentage complete. Prior to this change, dots were displayed on the console until the operation was complete.	151200

HDFS

New and changed in OneFS 7.2.1.1	ID
Support for IBM BigInsights for Apache Hadoop was added.	161774
Support for Ambari 2.1.1 was added.	159647
Support for Ambari 2.0.2 was added.	158607
1.7.0_IBM HDFS was added to the list of supported Ambari servers.	155817
Support for Ambari 2.0.1 and 2.1.0 was added.	154629
Support for Ambari 2.0.0 was added.	154622
The datanode port that HDFS listens on was changed from 1021 to 585 to avoid conflicts with other processes that might have been listening on the same port.	152980
Support for the HDFS truncate remote procedure call was added.	152153

Networking

New and changed in OneFS 7.2.1.1	ID
The default network flow control setting for Isilon nodes that contain Intel network interface cards (ixgbe NICs) was changed. The default flow control setting is now `1`. The new flow control setting configures the ixgbe NIC to receive pause frames but	155825

New and changed in OneFS 7.2.1.1	ID
does not send pause frames. This configuration is consistent with Isilon nodes that contain Chelsio NICs. **Note** Ethernet flow control in a full-duplex physical link provides a mechanism that enables an interface or switch to request a short pause in frame transmission from a sender by issuing a media access control (MAC) control message and PAUSE specification as described in the 802.3x full-duplex supplement standard.	
Support for PTR record lookup for SmartConnect zone member addresses was added.	152265
New control: The following parameters were added to the `isi networks` command: • `--disable-dns-tcp-support` • `--enable-dns-tcp-support` The `--disable-dns-tcp-support` parameter can be used to disable TCP support for SmartConnect; the `---enable-dns-tcp-support` parameter can be used to enable TCP support for SmartConnect. By default, TCP support is enabled and SmartConnect works as expected. If TCP support is disabled, SmartConnect doesn't listen for TCP connections on the DNS port (53), and clients that attempt a DNS query over TCP receive a connection refused error.	152185

The error messages that are logged if the `flx_config.xml` file cannot be read or loaded were updated to facilitate diagnosis of the issue. Beginning in OneFS 7.2.1.1, if the `flx_config.xml` file cannot be read or loaded—for example, if the file cannot be read because a node's network interface card is not accessible—lines similar to the following might appear in the `/var/log/messages` file or the `/var/log/isi_flexnet_d.log` file:
`isi_smartconnect[15482]: Error processing subnet in flexnet config: 7`
`isi_smartconnect[15482]: parameter member iface-class of member nonexistant`
`isi_smartconnect[15482]: /ifs/.ifsvar/modules/flexnet/flx_config.xml is corrupt (configuration errno 7: [/ifs/.ifsvar/modules/flexnet/flx_config.xml] parameter 'member iface-class' of 'member' nonexistant) isi_smartconnect[15482]: Corrupt config found on /ifs`
`isi_smartconnect[15482]: Unable to load FlexNet configurations.`
	151789

Security

New and changed in OneFS 7.2.1.1	ID
On clusters running in compliance mode, you can no longer run the su command to assume the privileges of a user with root-level (UID=0) access to the cluster. If you attempt to run the su command to assume the privileges of a user with root-level privileges, the following message appears on the console: `su: UID 0 denied by compliance mode`	158564

New and changed in OneFS 7.2.1.1	ID
Note Due to this change in behavior, beginning in OneFS 7.2.1.1, clusters running in compliance mode cannot be reimaged by running the `sudo isi_reimage` command.	
The version of Apache installed on the cluster was updated to 2.2.31. For more information, see ESA-2015-114 on the EMC Online Support site.	157834
The version of OpenSSL installed on the cluster was updated to 1.0.1p. For more information, see ESA-2015-155 on the EMC Online Support site.	157834
The network time protocol (NTP) service was updated to version 4.2.8P1. For more information, see ESA-2015-154 on the EMC Online Support site.	154837

SMB

New and changed in OneFS 7.2.1.1	ID
Support was added for the SMB2_CREATE_QUERY_ON_DISK_ID (QFid) SMB create request value.	154609

Logging performed by the srvsvc process was improved as follows:
• The default logging level for the srvsvc process was changed from WARNING to INFO.
• The user name and domain name for the user performing an action is logged in the `/var/log/srvsvcd.log` file, in addition to the SID.
• The action of modifying or deleting an SMB share via the Microsoft Management Console (MMC) snap-in is logged in the `/var/log/srvsvcd.log` file, including the user name.
An example of the new logging output appears below, where *‹USER SID info›* is the name and SID of the user and *‹SMB_SHARE_NAME›* is the name of the share:
`Log level changed to INFO`
`DOMAIN_NAME\USER_NAME <USER SID info> set info on share <SMB_SHARE_NAME>`
`DOMAIN_NAME\USER_NAME <USER SID info> deleted share <SMB_SHARE_NAME>`
	151756, 151754

Upgrade and installation

New and changed in OneFS 7.2.1.1	ID
Beginning in OneFS 7.1.0, the file in which cluster configuration information is stored was changed from a plain text file (`gconfig`) to a database file (`isi_gconfg.db`). In conjunction with this change, the maximum allowed size of the configuration information for an SMB share was limited to 8192 bytes (8 KB). Beginning in OneFS 7.2.1.1, the OneFS pre-upgrade check checks the size of the configuration information for an SMB share prior to upgrading the cluster and the cluster is prevented from being upgraded if the configuration size is greater than 8 KB.	158073

New and changed in OneFS 7.2.1.1	ID

The pre-upgrade check can be run alone, or as part of the upgrade process. In either case, if the configuration size of an SMB share exceeds the maximum size allowed , a message similar to the following appears on the console during the pre-upgrade check:
 	
`Error: The 'share_name' share has too many access permissions and it cannot be upgraded.`
`The suggested resolution for this issue is:`
`1. Remove those users from the share permissions.`
`2. Add those users to a group.`
`3. Add that group to the share permissions.`
`4. Retry the upgrade.`
 	
If the pre-upgrade check detects that the configuration size of an SMB share exceeds the maximum size allowed when it is running as part of the default upgrade process, the pre-upgrade check portion of the upgrade completes, however the OneFS upgrade is not started, and a message similar to the following appears on the console, and in the SMB upgrade log file located in the `/ifs/.ifsvar/tmp` directory:
 	
`Error: The 'share_name' share has too many access permissions and it cannot be upgraded.`
`The suggested resolution for this issue is:`
`1. Remove those users from the share permissions.`
`2. Add those users to a group.`
`3. Add that group to the share permissions.`
`4. Retry the upgrade.`
 	
Under these conditions, the upgrade process cannot be completed until the SMB share configuration information is reduced in size. In most cases, this can be accomplished by following the resolution suggested during the pre-upgrade check. If you encounter this limitation and cannot reduce the size of the SMB configuration information by following these steps, contact EMC Isilon Technical support for assistance.
 	
Note
 	
Prior to the addition of this check, if the configuration size of an SMB share on a cluster that was being upgraded to OneFS 7.1.0 or later exceeded the maximum size allowed, some of the share information might not have been preserved during the upgrade process, and an error similar to the following might have appeared in the `/var/log/isi_gconfig_d.log` file:
 	
`Update error: value for key 'share_name' has size (12324) greater than max allowed value size (8192)`	

In earlier versions of OneFS, you could assign the same IP address range or overlapping IP address ranges to the int-b and failover networks. However, this configuration can cause connectivity issues between nodes on the internal network. Beginning in 7.2.1.1, the OneFS pre-upgrade check verifies whether or not internal network subnets are configured with overlapping IP address ranges before upgrading the cluster. If this configuration is detected, the pre-upgrade check completes, but the OneFS upgrade is not started, and the cluster cannot be upgraded until the int-b and failover networks are configured correctly.	156918

New and changed in OneFS 7.2.1.1	ID

In addition, messages similar to the following appear on the console or on the **Cluster Management** page in the OneFS web administration interface:
 	
`check_internal_network check failed:`
`======================`
`check_internal_network`
`======================`
`ERROR: int-a overlaps with int-b`	

OneFS 7.2.1.x does not support file pool policy names that begin with a number. If you upgraded from OneFS 6.5.5.x, which supported file pool policies with names that began with a number, and if any of your preexisting file pool policies began with a number, following the upgrade the Smart Pool jobs failed, and file pool policies could not be created or modified.
In OneFS 7.2.1.1, a pre-upgrade check was added to halt an upgrade if the cluster configuration being upgraded contains file pool policies that begin with a number. If unsupported names are detected, the following message appears, where `<affected_policy>` is the policy or policies with a name that begins with a number:
 	
`Upgrade would invalidate all file pool policies with names that start with a number.`
`As of OneFS 7.1.1, filepool policy names can no longer have a numeric first character.`
`The following policies must be renamed before upgrade:`
 	
`<affected_policy>`
 	
These messages also appear in the `/var/log/update_handler_<date>_<time>.txt` file, where *‹date›_‹time›* is the date and time the upgrade process was started.
	151961

New and changed in OneFS 7.2.1.0

Authentication

New and changed in OneFS 7.2.1.0	ID
Beginning in OneFS 7.2.1.0, the lsass process retrieves the privilege attribute certificate (PAC) information from the Service-for-User (S4U) Kerberos extensions. This change improves performance when authenticating users connecting to the cluster over SMB and NFS. This change decreases the amount of time it takes to create access tokens for Microsoft Active Directory (AD) users.	136060

Networking

New and changed in OneFS 7.2.1.0	ID
New control: The following sysctl parameter was added to preserve or reset the value of differentiated services code point (DSCP) bits in the IP header:	148525

New and changed in OneFS 7.2.1.0	ID

`net.inet6.ip6.preserve_dscp`
To preserve DSCP bits, run the following command:
`isi_sysctl_cluster net.inet6.ip6.preserve_dscp=1`
To clear DSCP bits, run the following command:
`isi_sysctl_cluster net.inet6.ip6.preserve_dscp=0`	

New control: Beginning in OneFS 7.2.1, you can use the `isi_auto_linklocal` script to enable or disable link-local IP addresses from the command-line interface. OneFS 7.2.1 does not require link-local IP addresses with IPv6. However, in IPv6 environments where compliance with RFC 4291 is required, link-local addresses are mandatory. Link-local IP addressees are not enabled in OneFS by default. You must run the `isi_auto_linklocal` script to enable (or disable) link local IP addresses. For more information and the procedure for enabling link-local IP addresses, see article 204801 on the EMC Online Support site.	139318, 139278

Security

New and changed in OneFS 7.2.1.0	ID
The version of OpenSSL that is installed on the cluster has been updated to OpenSSL 1.0.1j.	140609
SSL v3 is no longer supported for HTTPS connections to the cluster. As a result, HTTP clients cannot connect to the OneFS web administration interface through a connection that relies on SSL v3. Workaround: Enable TLS 1.x for HTTP connections to the web administration interface. For more information, see ESA-2015-015 on the EMC Online Support site.	137904

CHAPTER 4

New hardware and firmware support

Newly supported hardware models and firmware revisions are described in this chapter.

- New hardware and firmware support in OneFS 7.2.1.2 28
- New hardware and firmware support in OneFS 7.2.1.1 28
- New hardware and firmware support in OneFS 7.2.1.0 28

New hardware and firmware support in OneFS 7.2.1.2

Hardware	Model Number	Drive Type	Compatible Nodes	Firmware
Support was added for 32GB Smart Modular PS3108 (XP+) boot flash drives. (158655)	SHSLM6B032G HM11EMC	Boot flash	NL400, S200, X200, X400, 108NL	S8FM08.2

New hardware and firmware support in OneFS 7.2.1.1

Hardware	Model Number	Drive Type	Compatible Nodes	Firmware
Support was added for the Ultrastar 800 GB drives with D252 firmware. (152252)	HGST HUSMM1680A SS205	SED SSD	HD400, NL400, X200, X400, X410	D252
Fixes upgrade path from MKAOA580 firmware, which is installed on 3 TB drives. (152252)	HGST HUA723030AL A640	HDD	X200, X400, NL400, IQ 108NL, IQ 108000X	MKAOA580

New hardware and firmware support in OneFS 7.2.1.0

Support for new platforms

OneFS 7.2.1 supports two new platforms: the X210 and the NL410. These nodes support node class compatibility; X210 nodes can be added to existing X200 node pools, and NL410 nodes can be added to existing NL400 node pools as described in the table above.

CHAPTER 5

Resolved issues

OneFS 7.2.1 contains all of the fixes that were included in OneFS 7.2.0.2 and OneFS 7.1.1.4. The following additional issues are resolved in OneFS 7.2.1.

- Resolved in OneFS 7.2.1.2 30
- Resolved in OneFS 7.2.1.1 45
- Resolved in OneFS 7.2.1.0 82

Resolved in OneFS 7.2.1.2

Authentication

Authentication issues resolved in OneFS 7.2.1.2	ID
When a user connected to the cluster, authentication information about the user was not retrieved from the cache. As a result, OneFS always queried the external authentication providers for authentication information. Querying the authentication providers increased network traffic. In addition, authentication of the user might have taken longer than expected.	163735, 163607
If you configured the user or group query settings for an LDAP authentication provider, so that user and group queries were restricted to a distinguished name that differed from the base distinguished name (base_dn), the configuration setting was not honored. As a result, user and group LDAP queries returned the full list of users or groups found in the base_dn.	162437
If the selective authentication setting was enabled for a Windows trusted domain, and if a user was assigned to a group in that domain to which the ISI_PRIV_LOGIN_SSH role-based access privilege was assigned, and if the user attempted to log in through an SSH connection, the user was denied access to the cluster. This issue occurred because the selective authentication setting prevented OneFS from resolving the user's group membership.	161272
If a cluster could not contact a writable Windows 2008 R2 or later domain controller (DC), attempts to fail over to a read-only domain controller were unsuccessful, and a `KRB_AP_ERR_BAD_INTEGRITY` error was returned to the cluster. This issue occurred because the name-type that was sent in the Kerberos Ticket-Granting Ticket request was not KRB5-NT-SRV-INST, which is the required name-type for compatibility in this environment. As a result, clients that were connected to the cluster might have experienced intermittent authentication issues.	160417

If the `sudoers` file contained role-based access control (RBAC) role names with spaces or hyphens, the `sudoers` file was invalid. As a result, `sudo` commands could not be run on the cluster. If this issue occurred on a cluster running in compliance mode, commands that required the `sudo` command—that is, commands that require root-level access—could not be run. This included `sudo` commands executed by the compadmin user. If any user attempted to run a `sudo` command, messages similar to the following appeared on the console:
`sudo: >>> /usr/local/etc/sudoers:syntax error near line 162 <<<`
`sudo: >>> /usr/local/etc/sudoers:syntax error near line 165 <<<`
`sudo: parse error in /usr/local/etc/sudoers near line 162`
`sudo: no valid sudoers sources found, quitting`
`sudo: unable to initialize policy plugin`
	159518

In order to generate a binary password file, a file that is required by a file provider, you must run the `pwd_mkdb` command. Beginning in OneFS 7.2.1.2, the `/usr/sbin/pwd_mkdb` command is included in the `sudoers` file. The addition of the command to the `sudoers` file allows the command to be run by users with sudo privileges. Prior to this change, it was not possible for a privileged user to create a local file provider on a cluster running in compliance mode.	158115

Backup, recovery, and snapshots

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.2	ID
After upgrading to OneFS 7.2.1.1, the NDMP backup services consumed 100 percent of a cluster's CPU resources when no backup or restore jobs were running. This issue occurred because a variable for an NDMP backup socket was incorrectly initialized.	163867
If you modified an existing SyncIQ policy, and if the policy was already configured with a snapshot expiration time, the date might have been shortened when you saved the changes to the policy. By design, when you configure a snapshot expiration time for a SyncIQ policy, SyncIQ internally records the snapshot expiration time in seconds and then converts the seconds to days, weeks, months, or years. For example, if you enter a value of 84 days, after you save the policy, the following text is displayed on the **View SyncIQ Policy Details** page: `Yes, snapshots expire after 2 months 3 weeks 3 days.` However, in OneFS 7.2.1.1 and earlier, if you modified an existing SyncIQ policy in the OneFS web administration interface, and if the policy was configured with a snapshot expiration time, the policy's snapshot expiration time might have been incorrectly rounded to the largest unit of measurement in the displayed text. For example, if you created a new SyncIQ policy with a snapshot expiration value of 84 days, the expiration time was displayed as `2 months 3 weeks 3 days`. If you modified the policy in the Web administration interface—for example, if you changed the value assigned to the **Worker Threads Per Node** field—when you saved your changes to the policy, the expiration time was shortened to 2 months.	162442

If the `--skip_bb_hash` option of a SyncIQ policy was set to `no` (the default setting) and if a SyncIQ file split work item was split between pworkers, it was possible for the pworker that was handling the file split work item to attempt to transfer data that had already been transferred to the target cluster. If this occurred, the isi_migr_pworker process repeatedly restarted and the SyncIQ policy failed. In addition, the following lines appeared in the `/var/log/messages` file:
`isi_migrate[45328]: isi_migr_pworker: *** FAILED ASSERTION cur_len != 0 @ /usr/src/isilon/bin/isi_migrate/pworker/handle_dir.c:463:`
`/boot/kernel.amd64/kernel: [kern_sig.c:3376](pid 45328="isi_migr_pworker")(tid=100957) Stack trace:`
`/boot/kernel.amd64/kernel: Stack:`
`--------------------------------------------------`
`/boot/kernel.amd64/kernel:`
`/lib/libc.so.7:__sys_kill+0xc`
`/boot/kernel.amd64/kernel:`
`/usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/boot/kernel.amd64/kernel:`
`/usr/bin/isi_migr_pworker:migr_continue_file+0x1507`
`/boot/kernel.amd64/kernel:`
`/usr/bin/isi_migr_pworker:migr_continue_generic_file+0x9a`
`/boot/kernel.amd64/kernel:`
`/usr/bin/isi_migr_pworker:migr_continue_work+0x70`
`/boot/kernel.amd64/kernel:`
`/usr/lib/libisi_migrate_private.so.2:migr_process+0xf1`
`/boot/kernel.amd64/kernel:`
`/usr/bin/isi_migr_pworker:main+0x606`
`/boot/kernel.amd64/kernel:`
`/usr/bin/isi_migr_pworker:_start+0x8c`
`/boot/kernel.amd64/kernel:`
`--------------------------------------------------`
	161221

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.2	ID
`/boot/kernel.amd64/kernel: pid 45328 (isi_migr_pworker), uid 0: exited on signal 6 (core dumped)`	
If the SyncIQ target password file `/ifs/.ifsvar/modules/tsm/passwd` existed on a cluster, then SyncIQ failover or failback jobs, domain mark jobs, or snapshot revert jobs failed. If this occurred, a message similar to the following appeared in the `/var/log/messages` file: `Authentication with target failed.`	161214
During differential SyncIQ jobs, the SyncIQ bandwidth process might have failed when new differential sync hash workers were started. If this issue occurred, a message similar to the following appeared in the `/var/log/messages` file: `isi_migr_bandwidth: *** FAILED ASSERTION type < MAX_MSG @/build/mnt/src/isilon/lib/isi_migrate/migr/msg.c`	160316

If you changed a SyncIQ policy **Action** setting from **Copy** to **Synchronize**, and if you also enabled the **Record when a synchronization deletes files or directories** option, the isi_migr_sworker process might have stopped unexpectedly if files were removed from the source before the last incremental sync. If this issue occurred, the policy stopped running and the following error message appeared on the source cluster in the `/var/log/isi_migrate.log` and `/var/log/messages` files:
 	
`Error Message: (policy name: <policy name> target: <target name>) SyncIQ policy failed. A work item has been restarted too many times. This is usually caused by a network failure or a persistent worker crash.`
 	
Additionally, the following output appears in the `/var/log/messages` file on the target cluster:
 	
`isi_migrate[49583]: isi_migr_sworker: *** FAILED ASSERTION index != NULL @ listfuncs.c:22:`
`/boot/kernel.amd64/kernel: [kern_sig.c:3349](pid 49583=""isi_migr_sworker"")(tid=100175) Stack trace:`
`/boot/kernel.amd64/kernel: Stack:`
`--------------------------------------------------`
`/boot/kernel.amd64/kernel: /lib/libc.so.7:__sys_kill+0xc`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:reset_list_head+0x81`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:listadd+0x471`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:delete+0x49f`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:process_entry+0x1f9`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:process_list+0x9a7`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:list2_callback+0xb4`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:generic_msg_unpack+0x8b9`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:migr_process+0x2fc`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:main+0xa53`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:_start+0x8c`
`/boot/kernel.amd64/kernel:`
`--------------------------------------------------`
	160313

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.2	ID
`/boot/kernel.amd64/kernel: pid 49583 (isi_migr_sworker), uid 0: exited on signal 6 (core dumped)`	

If you tried to resume a failed resync-prep SyncIQ job, a `FAILED ASSERTION` error occurred, the resync-prep job failed again, and messages similar to the following appeared in the `/var/log/isi_migrate.log` file on the source cluster, where `<policyname>` was the name of the affected policy:
`(policy name: <policyname>) SyncIQ is unable to connect to a resource on the target cluster. Use ping to determine if target is reachable. Verify that SyncIQ is licensed and enabled on the target cluster, and that the SyncIQ daemons are running. Unable to finalize resync prep on target.`
In addition, a message similar to the following appeared in the `/var/log/isi_migrate.log` file on the target cluster:
`*** FAILED ASSERTION found @ /usr/src/isilon/lib/isi_migrate/config/siq_source.c:797:`
Note
Although a FAILED ASSERTION error occurred, nodes did not restart as a result of this issue.
	158082

Because SyncIQ source snapshots had a default expiration date of one year after creation of the snapshot, if an incremental SyncIQ job started more than one year after the prior SyncIQ job started, the incremental job failed. Starting in OneFS 7.2.1.2, SyncIQ source directory snapshots never expire.	155195

Cluster configuration

Cluster configuration issues resolved in OneFS 7.2.1.2	ID
The isi_rpc_d process might have been affected by a memory leak of network configuration data. If this memory leak occurred, the Remote Procedure Call service could have restarted unexpectedly and messages similar to the following appeared in the `/var/log/messages` file: `net.rpc … Cannot allocate memory`	161716

If the isi_cpool_d process could not read the configuration settings that were stored on a node, the process stopped unexpectedly. If this issue occurred, one of the following lines appeared in the `/var/log/messages` file:
`_ZN23cpool_d_config_settings18set_retry_settingsE15task_type_valueRK22cpool_d_retry_settings`
`_ZNK23cpool_d_config_settings18get_retry_settingsE15task_type_value`
`_ZN23cpool_d_config_settings20set_thread_intervalsE15task_type_valueRK24cpool_d_thread_intervals`
`_ZNK23cpool_d_config_settings17get_csi_intervalsEv`
`_ZNK23cpool_d_config_settings17get_ooi_intervalsEv`
`_ZNK23cpool_d_config_settings17get_wbi_intervalsEv`
	159521

Cluster configuration issues resolved in OneFS 7.2.1.2	ID

`_ZNK23cpool_d_config_settings20get_thread_intervalsE15task_type_value`
 	
Beginning in OneFS 7.2.1.2, if the isi_cpool_d process cannot read the configuration settings, after 10 tries the process restarts, and the following message appears in the `/var/log/isi_cpool_d.log` file:
 	
`failed to load configuration settings after maximum number of attempts`	

If the file-matching criteria for a file pool policy included a path name that ended with a forward slash (/), and if the policy was run the operation returned no matches. As a result, the specified SmartPools actions were not applied to the files.	154607

Events, alerts, and cluster monitoring

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.2	ID
If you had a licensed OneFS module—for example, InsightIQ—and the date of the license expiration for that module was in 30 days or less, a message similar to the following appeared on the console after you ran the `isi events` or the `isi status` command, where *‹license›* was the name of the product license, and *‹exp_date›* was the date of expiration: `All Software license <license> will expire on <exp_date>`	159522
If the `/var/log/isi_audit_cee.log` file was rotated, error messages were not written to the new log file. As a result, new error messages related to auditing were not logged.	158072
If you renamed a directory on a cluster from an NFS client, the OneFS protocol auditing system recorded that a file was renamed. As a result, the renamed event could not be audited correctly.	157979
A name string, **SMB upgrade failure**, was added to the **Windows Networking events** list on the **Events › Event Notification Rules** page in the OneFS web administration interface. This list becomes visible after you try to add a notification rule. Before this fix, the item was present, but the label was blank, which created confusion.	157638
If protocol auditing was enabled and configured to send messages to a Common Event Enabler (CEE) server, and if the CEE server could not be reached or was in an error state, the following message was logged in the `/var/log/isi_audit_cee.log` file every ten seconds: `ERROR: Heartbeat error to <CEE server>` Beginning in 7.2.1.2, under these conditions, the error is logged once every fifteen minutes. This change reduces the affect that logging this message has on the wear life of the boot disks on the affected nodes.	157377

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.2	ID

SNMP traps were not sent for event ID 100010015, `There is at least one smartpool at or over capacity`. If the event threshold was reached and you had configured a notification rule that sent events through SNMP, the system failed to send that notification and the following messages were logged to the `/var/log/isi_celog_notification.log` file:
`INFO: Can't find varbind 'name' in event 100010015`
`WARNING: Invalid varbind(s) found for event 100010015. SNMP Trap will not be sent.`
	157376

File system

File system issues resolved in OneFS 7.2.1.2	ID
On nodes with 6 GB of RAM installed, the sysctl parameter that supplied OneFS with information about the amount of RAM installed in the node reported that the node had 8 GB of RAM installed instead of 6 GB. As a result, it was possible for OneFS to allocate too much memory to the lwio process. If this issue occurred, the overall system performance was impacted.	162434

In certain OneFS configurations, if a memory allocation call that pertained to the `/usr/share/zoneinfo/` directory failed, the lsass process might have leaked file descriptors. If the number of open file descriptors exceeded the normal maximum of 16381, the lsass process stopped responding and the cluster stopped accepting new connections. If this issue occurred, messages similar to the following appeared on the console after you ran the `procstat -f` command:
`22472 lsass 95 v r r------- 1 0 - /usr/share/zoneinfo/UTC`
`22472 lsass 96 v r r------- 1 0 - /usr/share/zoneinfo/posixrules`
`22472 lsass 97 v r r------- 1 0 - /usr/share/zoneinfo/America/Chicago`
	158067

The **Local Storage Usage** graph on the **File System Management › Storage Pools › Summary** page might have displayed the tick marks on the y-axis with values other than the expected values of 0%, 25%, 50%, 75%, and 100%. For example, 41.66666666666667% might have appeared instead of 25%.	157924
If a file name included extended, non-standard ASCII characters, the OneFS antivirus client could not process the file, even if the cluster's file system character encoding was set to UTF-8. If the OneFS antivirus client encountered an infected file whose name contained extended characters, the OneFS web administration interface might have displayed only a partial list of detected threats on the **Antivirus › Detected Threats** tab or the web interface might not have displayed any information on the **Detected Threats** tab. In addition, if you ran the `isi avscan report threat` command under these conditions, a message similar to the following appeared: `Could not decode to UTF-8 column 'FileInfected' with text 'test? file f?r encode.doc'`	157388
If you tried to shut down a node by using the `shutdown` command, and an external script failed during the shutdown process, the shutdown process might	153473

File system issues resolved in OneFS 7.2.1.2	ID
have stopped responding. Logging was enhanced so that the external scripts that the `shutdown` command is trying to run are included in the `/var/log/messages` log file. This enhancement helps troubleshoot the shutdown process of a node. Before the fix, only the failed external scripts, similar to the following, were logged: `Trying to send shutdown event`	

File transfer

File transfer issues resolved in OneFS 7.2.1.2	ID

You can now access the cluster through the Apache service over HTTPS.
To enable HTTPS access, run the following command:
`isi_gconfig -t http-config https_enabled=true`
To disable HTTPS access, run the following command:
`isi_gconfig -t http-config https_enabled=false`
Note
It might take up to ten seconds for the configuration change to take effect, and data transfers to and from the cluster that are in progress over HTTP might be interrupted.
	149709

Hardware

Hardware issues resolved in OneFS 7.2.1.2	ID
In rare cases, when the node firmware on a node in a OneFS cluster was updated, the update for the Chassis Management Controller (CMC) Sensor Data Record (SDR) failed. If the CMCSDR update failed, the following new event (event type 900010013) appeared on the **Events › Summary** page in the OneFS web administration interface where `{device}` was a reference to the CMCSDR device that could not be updated: `Firmware update failed for device: {device}.`	160753
If you tried to run the `isi devices -a smartfail -d` command to smartfail a drive that had already been smartfailed, physically removed, and replaced, but still appeared in the output, the operation failed and a message similar to the following appeared on the console: `!! Error: the smartfail action is invalid for a missing drive.`	159412

Hardware issues resolved in OneFS 7.2.1.2	ID

If the device driver for the front panel of a node did not load, the mcp process repeatedly restarted the isi_lcd_d process. If this issue occurred, the following lines appeared on the console and in the `/var/log/isi_lcd_d.log` file:
`ERROR:lcd.daemon:exiting in _start: Could not detect LCD device`
`INFO:lcd.daemon:LCDDaemon.None (pid 2565) exiting via SystemExit`
`WARNING:lcd.library:LCD Daemon not responding (LCD hardware may be missing)`
Note
This issue did not affect nodes from which the front panel was removed.
	156670

If you tried to update the firmware of a drive on a node that was set to read-only mode, by using the `isi dev -f -a fwupdate` command, and L3 cache was enabled on the node, the node restarted unexpectedly.	155489

HDFS

HDFS issues resolved in OneFS 7.2.1.2	ID
Beginning in OneFS 7.2.1.2, OneFS checks only the first two digits of the Ambari version number when verifying compatibility with a version of Apache Ambari. This check enables OneFS to maintain compatibility with minor releases of Ambari.	163360
When Kerberos authentication was enabled, a race condition might have occurred in the isi_hdfs_d process if two HDFS clients simultaneously tried to establish a Simple Authentication Security Layer (SASL) connection to the cluster. If this issue occurred, the isi_hdfs_d process failed and an `hdfs.core.gz` file was created in the `/var/crash` directory.	162584
Beginning in Ambari 2.0, when Ambari is upgraded, the Ambari Server sends an install_package request to each host's Ambari agent. OneFS was not designed to respond to this request. During an Ambari upgrade, when the request was sent to the cluster, the Ambari upgrade stopped progressing until the Ambari database was modified to allow the upgrade to continue. OneFS now responds correctly to install_package requests, and Ambari 2.0 and later upgrades progress as expected without manual intervention.	162583
Under some circumstances, the HDFS process on a node might have run out of system memory while DataNodes were transferring data to the HDFS server. If this issue occurred, the HDFS process on the affected node restarted, and the HDFS clients connected to an affected node were disconnected.	161650

A change in the implementation of the `–appendtofile` parameter of the `hdfs dfs` command in Apache Hadoop 2.7 might have caused the command to fail when run on a OneFS cluster. If this issue occurred, messages similar to the following appeared on the console:
`-appendToFile: Fatal internal error`
`java.lang.NullPointerException`
`at org.apache.hadoop.hdfs.DFSOutputStream.isLazyPersist (DFSOutputStream.java:1709)`
`at org.apache.hadoop.hdfs.DFSOutputStream.getChecksum4Compute (DFSOutputStream.java:1550)`
`at org.apache.hadoop.hdfs.DFSOutputStream.<init>`
	161579

HDFS issues resolved in OneFS 7.2.1.2	ID
(DFSOutputStream.java:1560) at org.apache.hadoop.hdfs.DFSOutputStream.<init> (DFSOutputStream.java:1667) at org.apache.hadoop.hdfs.DFSOutputStream.newStreamForAppend (DFSOutputStream.java:1694) at org.apache.hadoop.hdfs.DFSClient.callAppend (DFSClient.java: 1824) at org.apache.hadoop.hdfs.DFSClient.append (DFSClient.java: 1885) at org.apache.hadoop.hdfs.DFSClient.append (DFSClient.java: 1855) at org.apache.hadoop.hdfs.DistributedFileSystem$4.doCall (DistributedFileSystem.java:340) at org.apache.hadoop.hdfs.DistributedFileSystem$4.doCall (DistributedFileSystem.java:336) at org.apache.hadoop.fs.FileSystemLinkResolver.resolve (FileSystemLinkResolver.java:81) at org.apache.hadoop.hdfs.DistributedFileSystem.append (DistributedFileSystem.java:336) at org.apache.hadoop.hdfs.DistributedFileSystem.append (DistributedFileSystem.java:318) at org.apache.hadoop.fs.FileSystem.append (FileSystem.java: 1164) at org.apache.hadoop.fs.shell.CopyCommands$AppendToFile. processArguments(CopyCommands.java:341) at org.apache.hadoop.fs.shell.Command.processRawArguments (Command.java:201) at org.apache.hadoop.fs.shell.Command.run (Command.java: 165) at org.apache.hadoop.fs.FsShell.run(FsShell.java: 287) at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java: 70) at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java: 84) at org.apache.hadoop.fs.FsShell.main(FsShell.java:340)	
A memory leak in the data path between the DataNodes and the cluster might have caused the HDFS process to restart unexpectedly. If this issue occurred, the HDFS clients connected to an affected node were disconnected.	158083

Job engine

Job engine issues resolved in OneFS 7.2.1.2	ID
Under some circumstances, if a file the QuotaScan job was processing returned an error, the QuotaScan job did not process the file and the job failed. This issue might have occurred if files were deleted while they were being scanned by the QuotaScan job. If this issue occurred, the QuotaScan job was prevented from retrieving quota statistics, and configured SmartQuotas could not be correctly governed. In addition, a message similar to the following appeared in the QuotaScan job report: `Error in LIN quota scan 0: open 0: Invalid argument`	162133
Under some circumstances, if a file that the QuotaScan job was processing returned an error, the QuotaScan job did not process the file and the job failed. This issue might have occurred if files were deleted while they were being scanned by the QuotaScan job. If this issue occurred, the QuotaScan job was prevented from retrieving the quota statistics, and the configured SmartQuotas could not be correctly governed.	161868

Job engine issues resolved in OneFS 7.2.1.2	ID
In addition, a message similar to the following appeared in the QuotaScan job report: `Error in LIN quota scan 0: open 0: Invalid argument`	

In rare cases, when you ran a dedupe job on a cluster, the job did not complete. If this issue occurred, messages similar to the following appeared in the `/var/log/isi_job_d.log` file:
 	
`[MINOR] dedupe readInNextUniqueEntry/3503 readInNextUniqueEntry candidate list has fewer entries than checkpoint value`
`[MAJOR] dedupe in BEDRSUpdateTaskSetupParameterSanityCheck, dedupe entries processed can't be greater than candidates processed`
	160044

Due to a file conflict with another running job, after you updated the version of OneFS from 7.0.x to version 7.2.1.0 or 7.2.1.1, the isi_job_d process failed while running one of the following job types:
• AVScan
• ChangelistCreate
• Dedupe
• DedupeAssessment
• FSAnalyze
• PermissionRepair
• QuotaScan
• SmartPoolsTree
• TreeDelete
• Undedupe
If this issue occurred, subsequent jobs of any type could not be run, and messages similar to the following appeared in the `/var/log/messages` log file:
 	
`Stack: --------------------------------------------------`
`/usr/lib/libisi_util.so.1:isi_error_decode+0x12e`
`/usr/lib/libisi_util.so.1:isi_error_decode+0x12`
`/usr/lib/libisi_job.so.1:jd_error_decode+0xe9`
`/usr/lib/libisi_job.so.1:jd_error_vec_de`
`/usr/bin/isi_job_d:job_result_decode+0x105`
`/usr/bin/isi_job_d:lin_job_result_decode+0x18`
`/usr/bin/isi_job_d:fsa_job_restore+0x179`
`/usr/bin/isi_job_d:jcp_re_phase+0x18b`
`/usr/bin/isi_job_d:jcp_coord_re_job+0x309`
`/usr/bin/isi_job_d:job_from_status+0x158`
`/usr/bin/isi_job_d:jd_runq_read_queue+0x6c`
`/usr/bin/isi_job_d:coord_main+0x2db`
`/usr/bin/isi_job_d:main+0xb54`
`/usr/bin/isi_job_d:_start+0x14f`
`----------`
	157492

Migration

Migration issues resolved in OneFS 7.2.1.2	ID
After an isi_vol_copy migration completed, file names that began with the letters LARA, were changed. For example, the file name `LARATE.txt` was changed to `LARAt.ETxt`.	158068

Networking

Networking issues resolved in OneFS 7.2.1.2	ID
If all of the following conditions were met, it was possible for data to become inconsistent while it was being transferred to an Isilon node: • l The node contained a 10 GigE network interface card (NIC) that was using the • QLogic NetXtreme Ethernet (BXE) driver. • l VLAN tagging was enabled for the subnet that was assigned to the NIC. • l Large Receive Offload (LRO) was enabled on the NIC. (By default, LRO is enabled on a BXE NIC.) This issue occurred because, under these conditions, the version of the BXE driver that was installed on the node incorrectly interpreted the size of a TCP frame. For more information, see article 208773 on the EMC Online Support site.	164251, 162042
HD400, NL410, S210, X210 and X410 nodes that were configured to communicate through a 10 GigE network interface card that was using the QLogic NetXtreme Ethernet (BXE) driver which was introduced in OneFS 7.1.1 could have encountered an issue where the link reported `no carrier` in the output of the `ifconfig` command. Toggling the interface did not resolve the issue and the node had to be rebooted to reestablish the link.	164251, 154454
If the cluster contained HD400, NL410, S210, X210 or X410 nodes that had QLogic NetExtreme 10 GigE NIC cards and any external network subnets connected to the cluster were set to 9000 MTU, an error similar to the following appeared in the `/var/log/messages` file and the affected nodes rebooted: `ERROR: mbuf alloc fail for fp[01] rx chain (55)`	164251, 153414
HD400, NL410, S210, X210 and X410 nodes that were configured to communicate through a 10 GigE network interface card that was using the QLogic NetXtreme Ethernet (BXE) and were also configured to use the aggregate interface with LACP experienced connectivity issues with those links after rebooting the node or re-configuring the MTU on those interfaces.	164251, 153414
If an incorrect netmask was assigned to an IP address, excessive logging to the `/var/log/isi_flexnet_d.log` file occurred. Because the `/var` partition is stored on a node's boot disks, in rare and extreme cases, excessive logging to the `isi_flexnet_d.log` file might have decreased the wear life of the boot disks on the affected node. **Note** This message is now logged only if the logging level is set to `DEBUG` or higher.	159795

Networking issues resolved in OneFS 7.2.1.2	ID

In rare cases, if network settings were changed—for example, after IP addresses were rebalanced—a node might have restarted unexpectedly. If this occurred, clients connected to the affected node were disconnected and lines similar to the following appeared in the `/var/log/messages` file:
 	
`Stack:`
`--------------------------------------------------`
`kernel:trap_fatal+0x9f`
`kernel:trap_pfault+0x386`
`kernel:trap+0x303`
`kernel:sysctl_dumpentry+0xcb`
`kernel:rn_walktree+0x7b`
`kernel:sysctl_rtsock+0xe7`
`kernel:sysctl_root+0x132`
`kernel:userland_sysctl+0x18f`
`kernel:__sysctl+0xa9`
`kernel:isi_syscall+0x39`
`kernel:syscall+0x28b`
`--------------------------------------------------`
	150207

NFS

NFS issues resolved in OneFS 7.2.1.2	ID
If you performed a rolling upgrade from OneFS 7.1.x or earlier or to OneFS 7.2.1 or later, the export IDs assigned to NFS exports on the cluster were updated. As a result, operations that required the original export ID failed after the cluster was upgraded. For example, on a VMware ESX host configured with an NFS datastore located on an Isilon cluster, the following operations failed: • VMware vMotion migrations failed with a message similar to the following: `Failed to initialize migration at source…Bad parameter` • Attempting to save a virtual machine snapshot failed with a message similar to the following: `An error occurred while taking a snapshot` • Attempting to power on a virtual machine failed with a message similar to the following: `Failed to create swap file '<path_to_file>': Not found` Beginning in OneFS 7.2.1.2, OneFS preserves the original NFS export ID and responds with that ID if it receives a request that references it.	163932
If a OneFS cluster was configured with both Network Information System (NIS) and Active Directory (AD) providers, and if the NIS provider was configured to process authentication requests before the AD provider, when OneFS queried the system for netgroup updates, the NIS provider was skipped. If this issue occurred, netgroups were not automatically refreshed and, as a result, some NFS clients might have been unable to mount NFS exports until the netgroups were manually updated. In addition, if you ran the `isi_ilog -a isi_netgroup_d --level trace` command to configure the isi_netgroup_d process to be logged at the TRACE log level, a message similar to the following appeared on the console and in the `/var/log/isi_netgroup_d.log` file: `nis unhealthy, deferring expiration`	163615

NFS issues resolved in OneFS 7.2.1.2	ID

If an NFSv3 or NFSv4 client attempted to move a subdirectory from one directory to another within a parent directory to which a directory SmartQuota was applied, the file could not be moved and messages similar to the following appeared on the console:
`cannot move `directory_name1' to a subdirectory of itself, `directory_name2'`
OR
`cannot move `directory_name1' to `directory_name2': Input/output error`
This issue occurred even if the `efs.quota.dir_rename_errno` sysctl parameter was set to `18`.
Note
For more information about setting the `efs.quota.dir_rename_errno` sysctl to a value of `18`, see article 90185 on the EMC Online Support site.
For more information about configuring sysctl parameters in OneFS, see article 89232 on the EMC Online Support site.
	159037

A user or member of a group that was connected to the cluster over NFS must have been assigned the read permission in order to run executables on the cluster. This issue did not affect the owner of the executable.	151399

If you ran the `showmount -e` command to view a list of the configured NFS exports, if the output of the command required 2 MB of memory or more to display the full list, the NFS server did not return the information and a message similar to the following appeared on the console:
`RPC: Timed out`
`showmount: can't do exports rpc`
Note
The NFS server continued to function normally.
	149150

OneFS web administration interface

OneFS web administration interface issues resolved in OneFS 7.2.1.2	ID
In the OneFS web administration interface, on the **File System › Storage Pools › File Pool Polices** page if you tried to create a file pool policy by using IF conditions with the "is newer than" or "is older than" parameter, the **Create Policy** button in the **Create a File Pool Policy** dialog box was unavailable.	160163

Security

Security issues resolved in OneFS 7.2.1.2	ID
Versions of OpenSSH earlier than 7.0 might have allowed local users to gain privileges on non-OpenBSD platforms. This issue could have resulted in unauthorized modification or disclosure of information or a denial of service. For more information, see ESA-2015-007 on the EMC Online Support site. **Note** Although the vulnerability is fixed, OneFS is still running a version of OpenSSH that is prior to 7.0. Therefore, signature-based security scanners that check only the reported version of OpenSSH may still incorrectly report that systems containing this fix are vulnerable. These reports should be considered false positives and can safely be ignored.	162431
If all the following conditions existed for a file in the cluster, the Security Technology Implementation Guides (STIG) hardening profile could not be applied to or reverted from the cluster: • The file did not have the correct permissions applied to it. For example, the file was accessible to everyone when it should have been restricted to specific users. • The file was stored on only some nodes in the cluster. • You tried to resolve the permissions issue for the file when you applied or reverted the STIG hardening profile. If this issue occurred, an error similar to the following appeared on the console, where `<node>` was the name of a node where the file was not stored, and `<file>` was the name of the file that did not have the correct permissions: `<node>: chmod: <file>: No such file or directory` To apply or revert the STIG hardening profile, you had to resolve the permissions issue manually on the nodes where the file was stored.	161769

If you configured two external interfaces on a node, the Security Technology Implementation Guides (STIG) hardening profile could not be applied to the node, and the hardening profile could not be applied to the cluster.
If this issue occurred, messages similar to the following appeared on the console, where `<node-x>` was the name of the node to which the hardening profile could not be applied:
`Hardening profile failed on node <node-x>`
In addition, messages similar to the following appeared in the `/var/log/hardening_engine.log` file:
`ERROR isi_WA00555_A22_04: Perform Action Failed- FileContentChange.py -1219`
`ERROR - isi_WA00555_A22_04:perform_action failed - CustomOperation.py - 1732`
	161693

You could not apply the Security Technology Implementation Guides (STIG) hardening profile to a cluster that contained a node that did not have an external interface configured.	161692

Security issues resolved in OneFS 7.2.1.2	ID

If this issue occurred, messages similar to the following appeared on the console, where <node> was the node without an external interface configured:
 	
`More than one occurrence of text that should not be present was found:`
` <node> /etc/mcp/templates/sshd_config`
` At least one issue is catastrophic!!!!!`
` Please resolve this manually before trying to apply hardening.`
` print "Catastrophic issues are marked by failure type UNRECOV`
 	
These messages also appeared in a report (.rpt) file, that was located in the /ifs/.ifsvar/CHE/reports directory.	

The network time protocol (NTP) service was updated to version 4.2.8p3.	160121
Versions of OpenSSH earlier than 6.9p1-2 allowed the number of successive login attempts to exceed the configured limit when keyboard-interactive challenge-response authentication was used. This issue allowed a remote attacker to perform a brute-force attack or cause a denial of service. **Note** This fix does not upgrade the version of OpenSSH installed in OneFS 7.2.1.2. Therefore, security scanners that check only the reported version of OpenSSH will incorrectly report that systems containing this fix are vulnerable. These reports should be considered false positives and can safely be ignored. For more information, see ESA-2015-175 on the EMC Online Support site.	158882

SMB

SMB issues resolved in OneFS 7.2.1.2	ID
If SMB packet signing was enabled, and OneFS received an SMB3 packet with an incorrect signature, OneFS reset the affected TCP connection instead of sending a STATUS_ACCESS_DENIED error message. Because the connection was reset, if this issue occurred, the affected SMB3 client might have been disconnected from the cluster.	161788
If you used Microsoft Management Console (MMC) to configure an SMB share on the cluster from a Windows client and the file path to the share was invalid--for example, if the file path did not exist on the cluster--the share was not created but no error was returned to the Windows client. Beginning in OneFS 7.2.1.2, if you attempt to create an SMB share with an invalid file path through MMC, the following error appears on the client: `The device or directory does not exist.`	160301
Due to a race condition that could have occurred when multiple SMB1 sessions were being opened on the same connection, the lwio process sometimes unexpectedly restarted. If the process restarted, SMB clients connected to the affected node were disconnected from the cluster.	160295
If a symbolic link was migrated from a Microsoft Windows client to the cluster, if the tool that was migrating the data attempted to update the attributes of a symbolic link that had already been migrated, the attributes could not be updated, and the	157629

SMB issues resolved in OneFS 7.2.1.2	ID

migration of the symbolic link failed. For example, if you attempted to migrate data to an Isilon cluster using the EMCopy tool, and if the data contained a symbolic link, the symbolic link was initially migrated but EMCopy could not apply attributes to the symbolic link, and an error similar to the following appeared on the EMCopy client:
 	
`ERROR (5) : \path_to_target\symbolic_link -> Unable to set access time`
 	
In addition, if the EMCopy tool attempt to retry the failed operation, the retry failed and an error similar to the following appeared on the EMCopy client:
 	
`ERROR (4392) : \path_to_target\symbolic_link -> Unable to open, Failed after 1 retries.`	

In some cases, if the security mode of the SMB file sharing service was unchanged from the default configuration in OneFS 6.5.5.x, and if another SMB share setting—for example, `Change Notify`—was also changed from the default setting in OneFS 6.5.5, then, during an upgrade to OneFS 7.2.1.x, the `Impersonate Guest` security setting was changed from `Always` to `Never`. If this issue occurred, following the upgrade, SMB clients might not have been able to access shares on the cluster until the `Impersonate Guest` value for the share was manually set to `Always`.	154826

Upgrade

SMB issues resolved in OneFS 7.2.1.2	ID
During the upgrade process, an MCP error might have prevented the last node in a cluster from upgrading and the `mlist.xml` file might have been missing or damaged. If the `mlist.xml` file was missing or damaged, lines similar to the following appeared in the `/var/log/messages` file: `Error while loading /etc/mcp/mlist.xml: [/etc/mcp/mlist.xml] unknown error`	154824

Resolved in OneFS 7.2.1.1

Antivirus

Antivirus issues resolved in OneFS 7.2.1.1	ID
The OneFS antivirus client could not connect to some ICAP servers if the ICAP URL that you configured on the cluster was not in the following format: `icap://<hostname>:<port>/avscan`	153398

Antivirus issues resolved in OneFS 7.2.1.1	ID
The OneFS web administration interface did not list any files in the **Detected Threats** section of the **Antivirus › Reports** page if any ASCII special characters—for example, an ampersand (&)—were in the path name of any infected file.	153118
If you configured an antivirus scan of a directory in the OneFS web administration interface or from the command-line interface, the forward slash (/) at the end of the designated path was removed from the search string. As a result, the antivirus scanner might have scanned more directories than expected. For example, if the file system included both an `/ifs/data` directory, and an `/ifs/data2` directory, and if you configured the antivirus scanner to scan the `/ifs/ data/` directory, because the forward slash (/) was not included in the path, the antivirus scanner would have scanned both the `/ifs/data` directory and the `/ifs/data2` directory.	151948

Authentication

Authentication issues resolved in OneFS 7.2.1.1	ID
If there were files on the cluster that contained orphaned security identifiers (SIDs), this created authentication issues during the SID2Name lookup operation, since the SIDs were for users and groups that no longer existed. If this occurred, data on the cluster was unavailable and excessive amounts of CPU was consumed on the domain controller.	158867
Although the alternate Security Identities Attribute was configured for an LDAP provider, after upgrading the cluster to OneFS 7.2.1.0, the lsass process failed to look up and map an alternate security identity for users connecting over NFS. As a result, authentication failed and the affected users could not connect to the cluster.	158243
A local user who did not have root privileges could not change their password by running the UNIX `passwd` command. As a result, if an affected user's password expired, they were unable to log on to the cluster until the password was reset through another method.	155581
If the lsassd process was not able to resolve user and group IDs, a message was logged to the `/var/log/messages` file. In rare and extreme cases, excessive logging could decrease the wear life of the boot disks on the affected node. If this occurred, lines similar to the following appeared in the `/var/log/messages` file: `Failed to map token token={UID:10116, GID:100, GROUPS={GID:100, GID:20042}, zone id=-1 }: Failed to lookup uid 10116: LW_ERROR_NO_SUCH_USER`	154610
If an SMB client sent a request to apply an invalid SID to a file or directory on the cluster, the cluster returned a `STATUS_IO_TIMEOUT` response. Depending on the application that was used to send the request, a message similar to the following might have appeared on the client: `The specified network name is no longer available`	154258
Due to a file descriptor (FD) leak that occurred when SMB clients listed files and directories within an SMB share, it was possible for OneFS to eventually run out of available file descriptors. If this occurred, an ACCESS_DENIED or	153772

Authentication issues resolved in OneFS 7.2.1.1	ID

STATUS_TOO_MANY_OPENED_FILES response was sent to SMB clients that attempted to establish a new connection to the cluster or SMB clients that were connected to the cluster that attempted to view or open files. As a result, new SMB connections could not be established, and SMB clients that were connected to the cluster could not view, list, or open files. If this issue occurred, messages similar to the following appeared on the **Dashboard** › **Event summary** page of the OneFS web administration interface, and in the command-line interface when you ran the `isi events list -w | grep -i descriptor` command:
 	
`System is running out of file descriptors`
 	
In addition, messages similar to the following appeared in the `/var/log/lwiod.log` file:
 	
`Could not create socket: Too many open files`
`Failed to accept connection due to too many open files`	

If the cluster was not joined to a Microsoft Active Directory (AD) domain, and you attempted to change the access control list (ACL) of a file on the cluster from a Windows client, the operation failed, and a message similar to the following appeared on the client: `The program cannot open required dialog box because it cannot determine whether the computer named "10.0.1.1: is joined to a domain. Close this message, and try again.` Under these conditions, ACLs could only be modified through the OneFS command-line interface.	153407
If an LDAP server was configured to handle Virtual List View (VLV) search instead of paged search, and if LDAP users were listed, a memory leak occurred when returning more than one page of information. If users were listed a sufficiently large number of times, the lsass process could have run out of memory and then restarted unexpectedly. If this occurred, SMB users could not be authenticated for the several seconds it took for the lsass process to restart.	151935

In environments that relied on Kerberos authentication, if a machine password was changed while there were many active SMB connections to the cluster, a race condition could have been encountered. If this occurred, the lwio process restarted unexpectedly, and lines similar to the following appeared in the `/var/log/messages` file:
 	
`Stack: ------------------------------------------------`
`/usr/lib/libkrb5.so.3:krb5_copy_principal+0x33`
`/usr/lib/kt_isi_pstore.so:krb5_pktd_get_next+0xe6`
`/usr/lib/libkrb5.so.3:krb5_dyn_get_next+0x5e`
`/usr/lib/libkrb5.so.3:krb5_rd_req_decoded_opt+0x4a4`
`/usr/lib/libkrb5.so.3:krb5_rd_req_decoded+0x1d`
`/usr/lib/libkrb5.so.3:krb5_rd_req+0xc1`
`/usr/lib/libgssapi_krb5.so.2:krb5_gss_accept_sec_context+0x8fd`
`/usr/lib/libgssapi_krb5.so.2:gss_accept_sec_context+0x22c`
`/usr/lib/libgssapi_krb5.so.2:spnego_g`
`/boot/kernel.amd64/kernel: ss_accept_sec_context+0x3d6`
`/usr/lib/libgssapi_krb5.so.2:gss_accept_sec_context+0x22c`
`/usr/likewise/lib/lwio-driver/srv.so:SrvGssContinueNegotiate+0x2c5`
`/usr/likewise/lib/lwio-driver/srv.so:SrvGssNegotiate+0xd3`
`/usr/likewise/lib/lwio-driver/`
`srv.so:SrvProcessSessionSetup_SMB_V2+0x6c6`
`/usr/likewise/lib/lwio-driver/`
	151933

Authentication issues resolved in OneFS 7.2.1.1	ID
srv.so:SrvProtocolExecute_SMB_V2+0x1324 /usr/likewise/lib/lwio-driver/srv.so:SrvProtocolExecuteInternal +0x51b /usr/likewise/lib/lwio-driver/ srv.so:SrvProtocolExecuteWorkItemCallback+0x28 /usr/likewise/lib/liblwbase.so.0:WorkThread+0x1f7 /lib/libthr.so.3:_pthread_getprio+0x15d ----------	

Backup, recovery, and snapshots

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.1	ID
If you queried for the date on which a SyncIQ policy would next be run using the `next_run` OneFS API property, the date and time that was returned was incorrect.	158078
During an initial SyncIQ data replication, Access Control Lists (ACLs) applied to symbolic links, pipes, block devices, and character devices were not replicated from the SyncIQ source cluster to the SyncIQ target cluster. As a result, following an initial synchronization, applications and users were prevented from accessing these file system objects and were also prevented from accessing files and directories on the cluster through symbolic links.	158076
When performing an NDMP restore, OneFS verifies the end of the data stream by detecting two consecutive blocks of zeroes. In rare cases, in OneFS 7.2.1.0, if the second block of zeroes was stored in a different buffer than the first block of zeroes, OneFS did not read the second block of zeroes from the other buffer, and instead read the data that followed the first block of zeroes. If this occurred, the restore operation was immediately stopped, and data that was in the process of being restored might have been incompletely restored. This issue did not occur if the RESTORE_OPTIONS NDMP environment variable was set to `1`, specifying that a single-threaded restore operation be performed. **Note** By default, restore operations are multi-threaded.	157386
While synchronizing data between source and target clusters in compliance mode, if the file flags that were applied to a file on the source cluster differed from the file flags that were assigned to the file on the target cluster, SyncIQ attempted to update the file attributes of Write Once Read Many (WORM) committed files on the target cluster even if the retention date for those files had not yet passed. As a result, the synchronization failed. If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file: `bam_ads_setmode error: 30Local error : syncattr error for <path_to_WORM_file>: chfal`	156955
While a SyncIQ policy was running, if a SyncIQ primary worker (pworker) process on the source cluster sent a list of directories to delete to a secondary worker (sworker) on the target cluster, and then the pworker process unexpectedly stopped, the pworker's work range was transferred to another pworker. The other pworker then sent the list of directories to another sworker. This action resulted in two sworker processes on the target cluster trying to delete the same directory at	155811

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.1	ID

the same time. If this issue occurred, the SyncIQ job stopped, and lines similar to the following appeared in the `/var/log/messages` file:
 	
`/boot/kernel.amd64/kernel: [kern_sig.c:3349](pid 70="isi_migr_sworker")(tid=2) Stack trace:`
`/boot/kernel.amd64/kernel: Stack:----------`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:move_dirents +0x1b6`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:delete_lin +0x279`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:delete_lin_callback+0x143`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:generic_msg_unpack+0x8bc`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:migr_process+0x2f1`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:main+0xa18`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_sworker:_start+0x8c`
`/boot/kernel.amd64/kernel: ----------`
`/boot/kernel.amd64/kernel: pid 70 (isi_migr_sworker), uid 0: exited on signal 10 (core dumped)`
 	
Additionally, the following error might have appeared in the `/var/log/isi_migrate.log`:
 	
`Error : Unable to open lin 0:Invalid argument: Invalid argument from remove_entry_from_parent (utils.c:1516) from remove_single_entry (utils.c:1595) from remove_all_parent_dirents (utils.c:1680) from delete_lin (stf_transfer.c:784)`	

If the `--skip_bb_hash` option of an initial SyncIQ policy was set to `no` (the default setting), and if a SyncIQ file split work item was split between pworkers, the pworker that was handling the file split work item might attempt to transfer data that had already been transferred to the target cluster. If this occurred, the isi_migr_pworker process repeatedly restarted and the SyncIQ policy failed. In addition, the following lines appeared in the `/var/log/messages` file:
 	
`isi_migrate[45328]: isi_migr_pworker: *** FAILED ASSERTION cur_len != 0 @ /usr/src/isilon/bin/isi_migrate/pworker/handle_dir.c:463:`
`/boot/kernel.amd64/kernel: [kern_sig.c:3376](pid 45328="isi_migr_pworker")(tid=100957) Stack trace:`
`/boot/kernel.amd64/kernel: Stack:`
`----------`
`/boot/kernel.amd64/kernel: /lib/libc.so.7:__sys_kill+0xc`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_pworker:migr_continue_file+0x1507`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_pworker:migr_continue_generic_file+0x9a`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_pworker:migr_continue_work+0x70`
`/boot/kernel.amd64/kernel: /usr/lib/libisi_migrate_private.so.2:migr_process+0xf1`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_pworker:main+0x606`
`/boot/kernel.amd64/kernel: /usr/bin/isi_migr_pworker:_start+0x8c`
`/boot/kernel.amd64/kernel: ----------`
	155809

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.1	ID
`/boot/kernel.amd64/kernel: pid 45328 (isi_migr_pworker), uid 0: exited on signal 6 (core dumped)`	

If you attempted to run a SyncIQ job from OneFS 5.5 to 7.x without a valid policy ID, the job stopped without dispatching a proper failure message, and an error similar to the following appeared in the `/var/log/messages` file:
 	
`Stack: ------------------------------------------------`
`/lib/libc.so.7:__sys_kill+0xc`
`/usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/usr/lib/libisi_migrate_private.so.2:get_lmap_name+0x54`
`/usr/bin/isi_migr_sworker:work_init_callback+0xacd`
`/usr/bin/isi_migr_sworker:old_work_init4_callback+0x16f`
`/usr/lib/libisi_migrate_private.so.2:generic_msg_unpack+0x8bc`
`/usr/lib/libisi_migrate_private.so.2:migr_process+0x2f1`
`/usr/bin/isi_migr_sworker:main+0xafa`
`/usr/bin/isi_migr_sworker:_start+0x8c`
`------------------------------------------------`
`/boot/kernel.amd64/kernel: pid 24302 (isi_migr_sworker), uid 0: exited on signal 6 (core dumped)`
 	
Note
 	
Starting in OneFS 7.2.1.1, the following message will appear in the `/var/log/isi_migrate.log` file:
 	
`Source version unsupported. 'sync_id' must contain a valid policy id.`
	154831

If the `force_interface` option was enabled on a SyncIQ policy, the SyncIQ scheduler process, isi_migr_sched, leaked memory. If this occurred, scheduled policies stopped running, and the following message appeared in the `/var/log/isi_migrate.log` file: `Cannot allocate memory`	154329
When performing an NDMP backup, if you set the BACKUP_OPTIONS setting to a value of `7` to run incremental token-based backups, OneFS created dump date entries with all levels, rather than creating dump date entries at only level 10. As a result, backup snapshots were kept in the system.	154312
If you used the snapshot-based incremental backup feature during a backup operation and multiple snapshots were created between backups, the feature might have failed to recognize that data had changed during the backup procedure. As a result, some changed files were not backed up. For more information see article 203815 on the EMC Online Support site.	154270
A secondary worker process incorrectly attempted to remove extended user attributes from a WORM-committed file before updating the file retention date. As a result, incremental SyncIQ jobs failed and error messages similar to the following appeared in the `/var/log/isi_migrate.log` file: `Error : Failed to delete user attribute <ATTR>: Read-only file system`	154263
If you configured SyncIQ policies to run when source files were modified by setting the Run Job option to **Whenever the source is modified**, a memory leak could	154260

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.1	ID
have occurred in the SyncIQ scheduler (isi_migr_sched). If this issue occurred, new SyncIQ jobs were not scheduled, some data was unavailable, and a message similar to the following appeared in the `/var/log/isi_migrate.log` file: `Could not allocate parser read buffer: Cannot allocate memory`	
If you performed an NDMP direct access recovery (DAR) or selective restore on an Isilon cluster, OneFS assigned ownership of the restored directories to the root account. Because clusters in compliance mode do not have a root account, the restored directories were inaccessible on clusters in compliance mode, unless the compadmin user was logged on to the compliance cluster.	154251
Although multiple IPv4 and/or IPv6 addresses were defined, NDMP listened to only one IPv4 and/or one IPv6 IP address. For example: • If a node had multiple IPv4 addresses defined, NDMP listened to only one IPv4 address. • If a node had multiple IPv6 addresses defined, NDMP listened to only one IPv6 address. • If a node had both IPv4 addresses and IPv6 addresses defined, NDMP listened to only one IPv4 address and only one IPv6 address.	154249
During a snapshot-based incremental backup, a WORM file might have been backed up as a regular file. If this occurred, and the files were restored, the files were restored as regular files, and they could have been modified after they were restored.	154247
If the isi_ndmp_d process was stopped, the NDMP process ID file was still locked by one or more NDMP child processes. As a result, the mcp process could not restart the isi_ndmp_d process, and no new NDMP connections could be established. If this occurred, a `Failed to spawn NDMP daemon` message appeared in the `/var/log/isi_ndmp_d.log` file.	154245
During Data Management Application (DMA) polling, if no tape was loaded in a backup drive, OneFS set the drive to unbuffered mode. Accessing the drive with a tape loaded in the drive sets the tape drive to buffered mode. As a result, if the initiator does not set a the tape drive to buffered mode, the backup-to-tape performance by non-Isilon initiators might have been very slow. **Note** This was not an issue if the tape drive was used only by Isilon backup accelerators.	153454
If a SyncIQ policy designated a target directory that was nested within the SyncIQ target directory of a preexisting policy, an error occurred during SyncIQ protection domain creation which caused the SyncIQ policy's protection domain to be incomplete. If this occurred, the following message appeared in the `/var/log/isi_migrate.log` file: `create_domain: failed to ifs_domain_add` In addition, if you ran the `isi domain list-lw` command, the **Type** field for the affected SyncIQ target was marked `Incomplete`.	153445

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.1	ID
If you ran an initial sync against a SyncIQ target directory that had a large number of files that did not exist on the source, then the process that removes extra files on the target might have conflicted with the process that domain marks the new target directory. If this occurred, the SyncIQ job failed, and the job had to be restarted.	153438
If you started a restartable backup with a user snapshot, after the backup was completed and the BRE context was removed, the expiration time of the snapshot was changed. As a result, the snapshot might have been deleted prematurely.	153409
If you viewed the details of a snapshot alias in the OneFS web administration interface, the `Most Recent Snapshot Name` was always `No value`, and the `Most Recent Snapshot ID` was always `0`.	153399
If the snmpd process failed to load the `/etc/mcp/sys/lni.xml` file, `/etc/ifs/local.xml` file, or `/etc/ifs/array.xml` file, a memory leak could have occurred. A memory leak in the snmpd process could have caused SNMP monitoring to be interrupted until the snmpd process was manually stopped and then restarted.	152284
If the `Deny permission to modify files with a DOS read-only attribute over both UNIX (NFS) and Windows File Sharing (SMB) ACL policy` option was enabled on the cluster, SyncIQ jobs failed when SyncIQ attempted to synchronize a file or a folder to which the DOS read-only attribute is applied. If a SyncIQ job failed for this reason, an `Operation not permitted` error message appeared in the `isi_migrate.log` file.	152254

If there was a group change on the source cluster while a SyncIQ job was in the process of starting, the SyncIQ scheduler might have stopped unexpectedly and then automatically restarted. If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file:
 	
`Stack:`
`--------------------------------------------------`
`/lib/libc.so.7:__sys_kill+0xc`
`/usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/usr/lib/libisi_migrate.so.2:siq_job_summary_save_new+0x200`
`/usr/bin/isi_migr_sched:sched_main_node_work+0xf3f`
`/usr/bin/isi_migr_sched:main+0xf13`
`/usr/bin/isi_migr_sched:_start+0x8c`
`--------------------------------------------------`
	152193

The physIfaces object identifier (OID) was incorrectly named in the `ISILONTRAP-MIB.txt` file, available in the **General Settings › SNMP Monitoring** tab of the OneFS web administration interface. As a result, it was not always possible to monitor the cluster through SNMP.	152171
In some cases when a SyncIQ job was in progress, the sworker on the target cluster did not send an acknowledgment to the pworker on the source cluster acknowledging that some tasks were completed. Furthermore, if a SyncIQ job was very large, a source pworker could have accumulated a large number of un-acknowledged tasks and then waited for the target sworker to acknowledge work that was already completed. If this occurred, the SyncIQ job would run indefinitely.	152149
If a directory was renamed to a path that had been excluded from a SyncIQ job, the SyncIQ state information for the directory and its children remained stored. However, the directory and its children tree were removed from the SyncIQ target. Any future changes that were made to the directory or its children were treated as	152128

Backup, recovery, and snapshots issues resolved in OneFS 7.2.1.1	ID

changes to included paths. If this occurred, a SyncIQ target error similar to the following appeared in the `/var/log/isi_migrate.log` file:
 	
`Error : Unable to open Lin <LIN>: No such file or directory`
 	
If all of the directories that had been excluded from the SyncIQ job were removed in an incremental SyncIQ job, that incremental SyncIQ job could have failed while trying to delete an excluded directory. If this occurred, an error similar to the following appeared in the `/var/log/messages` or `/var/log/isi_migrate.log` files:
 	
`FAILED ASSERTION found == true`	

All SyncIQ System B-Trees were protected at 8x mirrored, unnecessarily consuming disk space. **Note** Beginning in OneFS 7.2.1.1, the protection policy for SyncIQ SBTs is set to the system disk pool default, which enhances SyncIQ performance. If you want to change the default protection policy for SyncIQ SBTs, contact EMC Isilon Technical Support.	152126
If SyncIQ encountered an issue when processing an alternate data stream for a directory, an incorrect directory path appeared in the error message that was logged in the `/var/log/isi_migrate.log` file.	152115
If you configured or displayed a SyncIQ performance rule in the OneFS web administration interface, the bandwidth limit was described in kilobytes per second (KB/sec). This output did not match the kilobits per second (kbps) value seen in the command-line interface. The web interface and command-line interface now show the bandwidth limit value measured in kilobits per second.	151967
During a SyncIQ job, if the `rm` command that was executed during the cleanup process of the temporary working directory on the target cluster exited with an error, the SyncIQ policy went into an infinite loop, and data could not be synced to the cluster. If this occurred, a message similar to the following appeared in the `/var/log/isi_migrate.log` file: `Unable to cleanup tmp working directory, error is ...`	151943

If a SyncIQ job was interrupted during the STF_PHASE_CC_DIR_DEL (change compute deletion) phase, the Logical Inodes (LINs) could have been incorrectly removed from the SyncIQ job work list. If this occurred, the SyncIQ job failed, and messages similar to the following appeared in the `isi_migrate.log` file:
 	
`Unable to update metadata (inode changes) information for Lin ...`
`Operation failed while trying to detect all deleted lins in ...`
	150614

Cluster configuration

Cluster configuration issues resolved in OneFS 7.2.1.1	ID
Available space remaining on SSDs that are deployed as L3 cache was incorrectly reported in the OneFS web administration interface.	154621
After running the `isi_ntp_config` command to exclude a node from contacting an external NTP server, the following error appeared on the console: `"'str' object has no attribute 'gettext'".` As a result, only one node could be excluded from contacting an external NTP server.	154323
If the CloudPools Redirector was enabled and the FILE_OPEN_REPARSE_POINT flag was also enabled, then, if an SMB client attempted to open a symbolic link, OneFS returned a `STATUS_STOPPED_ON_SYMLINK` response to the client and the symbolic link was inaccessible.	152262
If a file on the cluster was deleted or modified, and the most recent snapshot of that file was deleted, any changes to SmartPools policies might have silently failed to propagate to some snapshot files.	152259

If you attempted to upgrade a SmartPools database that was not successfully upgraded because there were empty node pools, the upgrade did not complete, and an error similar to the following appeared in the OneFS web administration interface and on the console.
`Storage Pool Settings Changes Failed`
`The edit to the existing storage pool settings did not save due to the following error:`
`Changing settings disallowed until SmartPools DB is fully upgraded`
	151955

If you added a preformatted drive to a node, the drive was not properly repurposed for the pool that it was being added to. If this issue occurred, data was not written to the drive, the drive remained unprovisioned until it was reformatted. In addition, messages similar to the following appeared in the `/var/log/messages` file: `isi_drive_repurpose_d[6008]: STORAGE drive (devid:x, lnum:y, bay:z)is not part of any DiskPool. Skipping this drive.`	151749

Diagnostic tools

Diagnostic tools issues resolved in OneFS 7.2.1.1	ID

Because the following EMC Secure Remote Services (ESRS) log files were not listed in the `newsyslog.conf` file—a configuration file that manages log file rotation—over time the files could have grown in size and could have filled the `/var` partition:
`/var/log/GWExt.log`
`/var/log/GWExtHTTPS.log`
	155816

Diagnostic tools issues resolved in OneFS 7.2.1.1	ID
Note If the `/var` partition on a node in the cluster is 90% full, OneFS logs an event warning that a full `/var` partition can lead to system stability issues. Depending on how the cluster is configured, an alert might also be issued for this event.	
If ESRS was configured on the cluster, the ESRS process automatically selected the first available IP address, rather than an IP address from an IP address pool in the System access zone. Because only the System zone allows a user SSH access for remote management, if the selected IP address was not in the System access zone, EMC Isilon Support could not monitor the cluster remotely.	153456
If you ran the `isi_gather_info command` with the `--ftp-proxy-port` and `--save-only` options or with the `--ftp-proxy` and `--save-only` options, the specified FTP proxy port or FTP proxy host values were not saved. As a result, the desired FTP proxy settings had to be specified each time the `isi_gather_info` command was run.	152144

If you ran the `isi_gather_info` command on a node that was encountering back-end network issues, the operation timed out after 3 minutes, and a message similar to the following appeared on the console:
`isi_rdo: [Errno 60] Operation timed out`
`isi_gather_info: FAILED to make required directories on 1 nodes.`
	152109

Events, alerts, and cluster monitoring

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.1	ID
Sometimes, not enough memory was available to run the CELOG process. If this issue occurred, the process failed to log, cancel, or quiet events.	157496
When the isi_audit_cee process was shutting down, the queue producer library (QPL) sometimes closed the file descriptor related to an auditing task before the task was cancelled. Because the QPL reads a task's file descriptor when a task is cancelled, this incorrect ordering of events might have caused the QPL to attempt to read from the closed file descriptor or to read from another file descriptor. If the QPL read data from another file descriptor, the owner of the file or socket that the other file descriptor belonged to might have received invalid data and packet errors because the data that was incorrectly read by the QPL was missing from the file descriptor.	157495
Because isi_rest_server, a component of the Platform API, did not check for the correct error codes when interacting with the OneFS auditing system's QPL, if configuration auditing was enabled and there was an error in the QPL, the error was not handled correctly. If this issue occurred, it might have prevented system configuration changes from being audited.	156401
If auditing is enabled, the audit filter waits for a response from the QPL before sending audit events to the auditing process (isi_audit_d). In OneFS 7.2.0.0 through 7.2.0.3, if the QPL became disconnected from the auditing process, isi_audit_d, while the auditing process was waiting for a response, the QPL failed to send a response to the auditing process. If this	156399

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.1	ID
occurred, auditing events continued to collect in the auditing process until the queue became full. If the auditing process queue became full, processes related to events that were being audited—for example, processes related to file system protocols and configuration changes—might have stopped working. Depending on which related processes were affected, various cluster operations could have been disrupted by this issue—for example, if configuration auditing was enabled, you might have been prevented from making configuration changes through the OneFS web administration interface.	
Because the OneFS auditing system did not correctly convert a POSIX path with multiple path separators (/) into a Microsoft UNC path, if NFS protocol auditing was enabled, incorrect paths could have been recorded in the audit log and applications that rely on the information in the audit log might have been adversely affected.	155804
If file system protocol auditing was enabled and a client opened a parent directory and then opened a subdirectory or file within the parent directory, the auditing system might have incorrectly appended the subdirectory or file path to the parent directory path. If this occurred, the incorrect path might have caused an error in the auditing process and file system protocol events that were in the process of being logged might not have been captured. If the incorrect path was logged, applications that relied on file paths in the audit log might have been adversely affected.	155803
Under some circumstances, multiple isi_papi_d process threads might have called the same code at the same time. If this occurred, the isi_papi_d process might have unexpectedly restarted.	154325

If memory allocated to the clusterwide event log monitoring process (isi_celog_monitor) became very fragmented, the isi_celog_monitor process stopped performing any work. As a result, no new events were recorded, alerts regarding detected events were not sent, and messages similar to the following were repeatedly logged in the `/var/log/isi_celog_monitor.log` file:
 	
`isi_celog_monitor[5723:MainThread:ceutil:92]ERROR: MemoryError`
`isi_celog_monitor[5723:MainThread:ceutil:89]ERROR: Exception in serve_forever()`
 	
Note
 	
Allocated memory is considered fragmented when it is not stored in contiguous blocks. Memory allocated to the CELOG process is more likely to become fragmented in environments with frequent configuration changes and in which many CELOG events are being generated.
	153776

If you ran the `isi statistics client` or `isi statistics heat` command with the `--csv` option, the following error appeared instead of the statistics data: `unsupported operand type(s) for %: 'NoneType' and 'tuple'`	153566
If file system auditing was enabled and you configured the system to audit events in which a user renamed a file, if the user renamed the file from a Mac client connected to the cluster through a virtual private network (VPN), the complete path to the file was not always captured in the audit log. If this occurred, applications that relied on the file paths in the audit logs might have been adversely affected.	153464

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.1	ID
Beginning in OneFS 7.2.1.1, if a user attempts to rename a file and the complete file path to the renamed file is not captured in the audit log, the file is not renamed and an error appears in the audit log.	
Only the root user was permitted to run the `isi_audit_viewer` command. This limitation prevented other users—including users with `sudo` privileges—from viewing configuration audit logs and protocol audit logs on the cluster.	153440
If you enabled auditing on the cluster, only nodes that had the primary external interface (em0) configured could communicate with the Common Event Enabler (CEE) server, even if a secondary interface, such as em1, was configured and active on the node. As a result, the audit logs from these nodes were not collected on the CEE server.	153433
If phase 2 of an FSAnalyze job took longer than 100 minutes to complete, the job might have stopped progressing, progressed very slowly, or might have failed and then resumed. This issue occurred because, during phase 2, the FSAnalyze job updated an SQLite index, and, while the job was updating this index, it could not handle other job engine requests, preventing the job from progressing. In addition, if, while the SQLite index was being created, the queue of requests waiting to be handled grew to greater than 100 (the maximum allowed), the job was terminated and then resumed from a point before the 100 minutes had elapsed.	152253
The stated storage capacities for `/`, `/var`, and `/var/crash` were reported eight times higher than the actual usage amount in the OneFS statistics system. This sometimes caused incorrect capacity sizes to appear in the web administration interface, SNMP queries, or in Platform API-enabled applications.	152215

The isi_papi_d process did not correctly handle CELOG events that referenced pathnames that contained special characters or multibyte characters. If the isi_papi_d process attempted to handle a file containing special characters or multibyte characters, a specifier parse error appeared in the `isi_papi_d.log` file. For example:
`isi_papi_d[37840]: [0x80a403500]: ERROR Event 5.705 specifier parse error:`
`{"enforcement": "advisory", "domain": "directory /ifs/data/\xe8\xa9\xa6\xe9\xa8\x93&dios", "name": "exceeded", "val": 0.0, "devid": 0,"lnn": 0}`
	152184

SNMP monitoring with Nagios fails when using an Isilon-specific Nagios configuration file. The following error appears in Nagios when querying the cluster: `External command error: Timeout: No Response from <IPaddress>:161` The error is caused by an incorrect community name. To resolve this issue, see article 199833 on the EMC Online Support site.	152155
Protocol event logging in the `/var/log/audit_protocol.log` file always showed a value of 0 bytes written for a write event, and close events did not have bytes written or bytes read fields.	152117
Temporary SQLite files were created in the `/var/tmp` directory more frequently than was necessary. Because writes to the `/var` partition can decrease the wear life of boot disks on an affected node, an index was added to the `/ifs/.ifsvar/db/celog/events.db` SQLite database file to reduce the frequency with which these files are written to the `/var/tmp` directory.	152116

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.1	ID
If the CELOG notification master node went down, delivery of event notifications stopped until the down node returned to service or until the CELOG notification subsystem (isi_celog_notification) was restarted, at which point the subsystem would elect a new notification master with the updated group information.	151963

File system

File system issues resolved in OneFS 7.2.1.1	ID
If all of the following SyncIQ conditions existed, the data on the target backup cluster might have become inconsistent with the data on the source cluster when data was replicated by using Isilon SyncIQ: • The source cluster was running OneFS 7.2.1.0, 7.2.0.0–7.2.0.4, or 7.1.1.0–7.1.1.7. • The SyncIQ policy was set to **whenever the source is modified** . • OneFS SmartCache (also referred to as the write cache coalescer) was enabled (SmartCache is enabled by default during initial installation). • The SyncIQ job was an incremental sync (not an initial sync). • Six of the nine phases of the SyncIQ job completed within 15 seconds. This condition occurred if the data transfer involved changes to only a very small number of files, and • Data was transferred during the sync job. Furthermore, if all of the following NDMP conditions existed, the data in sparse files on the tape backup might have become inconsistent with the source cluster data when data was replicated by using NDMP in a three-way backup: • The source cluster was running OneFS 7.2.1.0, 7.2.0.0–7.2.0.4, or 7.1.1.0–7.1.1.7. • OneFS SmartCache (also referred to as the write cache coalescer) was enabled (SmartCache is enabled by default during initial installation). • You were running a three-way backup (either full or incremental). • The data being backed up contained sparse files. • The files that were being backed up were undergoing continuous changes / updates. For more information, see ETA 208772 on the EMC Online Support site.	162707, 85519
In rare cases, if a cluster that contained deduped files had a mirrored protection policy, the data in those files could have become unusable. If this issue occurred, a message similar to the following appeared on the **Dashboard › Event** summary page of the OneFS web administration interface, and on the console after the `isi events list` command was run: `1.983216 09/09 20:07 C 1 458786 Detected IDI failure on LIN 1:014a:005a::HEAD, lbn 19744 (data) 4,10,15073345536:8192 (type user data)`	161470

File system issues resolved in OneFS 7.2.1.1	ID
In addition, a message similar to the following appeared in both the `/var/log/messages` and `/var/log.idi log` files: `/boot/kernel.amd64/kernel: (pid 8221="isi_job_d"): [| IDI_VERIFY=idi_file_data_error:1383 | COND=Calculate data block checksums on read | MSG=Logical block 19744 (data)/120008 of file failed verification: Expected CRC 4db178db, found block CRC c692a790 | LINSNAP=1:014a:005a::HEAD | BADDR=4,10,150733455368192 |]`	
If a node ran for more than 497 days without being rebooted, an issue that affected the OneFS journal buffer sometimes disrupted the drive sync operation. If this issue occurred, OneFS reported that the journal was full, and as a result, resources that were waiting for a response from the journal entered a deadlock state. Any cluster that contained a node that ran for more than 497 consecutive days with no downtime might have rebooted unexpectedly as a result of this issue. For more information, see ETA 202452 on the EMC Online Support site.	158418
If a node ran for eight months or longer without a reboot and the node's internal clock rolled over, the universal memory allocator (UMA) processed an invalid value, which prevented the UMA from reclaiming any of the memory it had allocated. If this issue occurred, the affected node might have run out of memory, causing the node to unexpectedly reboot.	157490
On a compliance mode cluster, if either the retention period or the `DOS Read Only` flag that was applied to a file on a SyncIQ source cluster was changed after the initial synchronization, subsequent incremental SyncIQ jobs failed, and messages similar to the following appeared in the `/var/log/messages` file, where *<path>* was the path to the file on the target cluster: `Local error : syncattr error for <path>: Readonly file system` This issue occurred because, under these conditions, an unnecessary `chown` command was also sent to the target cluster.	156271

If you installed the drive support package (DSP) 1.5 firmware update on a cluster that contained a node with solid-state drives (SSDs) that were configured for use as L3 cache, the node might have rebooted unexpectedly. If a node rebooted for this reason, messages similar to the following appeared in the `/var/log/messages` file:
 	
`Stack: ------------------------------------------------`
`kernel:sched_switch+0x125`
`kernel:mi_switch+0x12e`
`kernel:sleepq_wait+0x3a`
`kernel:_sleep+0x37a`
`efs.ko:l3_mgmt_drive_state+0x9bd`
`efs.ko:drv_change_drive_state+0x178`
`efs.ko:drv_down_drive_prepare+0x1c2`
`efs.ko:drv_down_drive+0x81`
`efs.ko:drv_unmount_drive+0x176`
`efs.ko:drv_modify_drive_state_down+0x1d4`
`efs.ko:ifs_modify_drive_state+0x35a`
`efs.ko:_sys_ifs_modify`
`cpuid = 28`
`Panic occurred in module efs.ko loaded at 0xffffff87bde5a000:`
	154267

File system issues resolved in OneFS 7.2.1.1	ID

If OneFS was not mounted on a node and you ran the `isi_flush --l3-full` command on that node, the node restarted unexpectedly and messages similar to the following appeared in the `/var/log/messages` file:
 	
`Stack: --------------------------------------------------`
`kernel:trap_fatal+0x9f`
`kernel:trap_pfault+0x386`
`kernel:trap+0x303`
`efs.ko:mgmt_finish_super+0x4e`
`efs.ko:l3_mgmt_nuke+0x70`
`efs.ko:sysctl_l3_nuke+0xcb`
`kernel:sysctl_root+0x132`
`kernel:userland_sysctl+0x18f`
`kernel:__sysctl+0xa9`
`kernel:isi_syscall+0x39`
`kernel:syscall+0x28b`
`----------------------------------------------------`
	154265

If you attempted to smartfail multiple nodes that were holding user locks, the lock was held by LK client entries but not present in lock failover (LKF) entries. As a result of this inconsistency, future lock attempts failed, and a manual release of the lock was required to grant the desired access.	153435
If you configured OneFS to send syslog messages to a remote syslog server, the HOSTNAME of the cluster was not included in the messages. The absence of the HOSTNAME entry made it difficult to distinguish messages sent from multiple clusters to the same syslog server.	153418
If you ran a SmartPools job on a file with an alternate data stream (ADS), the job sometimes failed, and continued to fail, even if the job was manually started. If the SmartPools job failed for this reason, the SmartPools process eventually stopped running scheduled jobs, and this might have caused node pools to become full, degrading cluster performance. If this occurred, the SmartPools job reported an error similar to the following in the job history report: `Node 6: pctl2_set_expattr failed: No such file or directory`	153411

In some environments, where there was a heavy workload on the cluster, a node could run out of reserved kernel threads. This condition could have caused the node to restart unexpectedly. As a result, client connectivity to that node was interrupted.
If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file:
 	
`panic @ time 1422835686.820, thread 0xffffff0248243000: ktp: No reserved threads left`
`cpuid = 6`
`Panic occurred in module efs.ko loaded at 0xffffff87b7c84000:`
`Stack: --------------------------------------------------`
`efs.ko:ktp_assign_reserve+0x29f`
`efs.ko:dfq_reassign_cb+0x9b`
`kernel:_sx_xlock_hard+0x276`
`kernel:_sx_xlock+0x4f`
`efs.ko:lki_unlock_impl+0x306`
`efs.ko:lk_unlock+0xbe`
`efs.ko:bam_put_delete_lock_by_lin+0x36`
`efs.ko:_bam_free_free_store+0x34`
`efs.ko:dfq_service_thread+0x139`
`efs.ko:kt_main+0x83`
`kernel:fork_exit+0x7f`
	153397

File system issues resolved in OneFS 7.2.1.1	ID

If you exceeded the number of recommended snapshots on a cluster, nodes in the cluster might have rebooted unexpectedly. If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file:
 	
`/boot/kernel.amd64/kernel:`
`Stack:--------------------------------------------------`
`/boot/kernel.amd64/kernel:kernel:isi_assert_halt+0x42`
`/boot/kernel.amd64/kernel:efs.ko:pset_resize+0x107`
`/boot/kernel.amd64/kernel:efs.ko:pset_add+0x50`
`/boot/kernel.amd64/kernel:efs.ko:bam_data_lock_get_impl+0x1c8`
`/boot/kernel.amd64/kernel:efs.ko:bam_data_lock_get+0x2b`
`/boot/kernel.amd64/kernel:`
`efs.ko:ifm_read_op_init+0xa8`
`/boot/kernel.amd64/kernel:efs.ko:bam_mark_file_data+0xfd`
`/boot/kernel.amd64/kernel:efs.ko:ifs_mark_file_data+0x373`
`/boot/kernel.amd64/kernel:efs.ko:_sys_ifs_mark_file_data+0x166`
`/boot/kernel.amd64/kernel:kernel:isi_syscall+0x53`
`/boot/kernel.amd64/kernel:kernel:syscall+0x1db`
`/boot/kernel.amd64/kernel:-------------------------`
	152675

If a cluster ran for more than 248.5 consecutive days, an issue that affected the OneFS journal buffer could have disrupted the drive sync operation. If this issue occurred, OneFS reported that the journal was full, and as a result, resources that were waiting for a response from the journal entered a deadlocked state. When the journal was in this state, nodes that were affected rebooted to clear the deadlock condition. In addition, a message similar to the following appeared in the `/var/log/messages` file:
 	
`/boot/kernel.amd64/kernel:efs.ko:rbm_buf_timelock_panic_all_cb`
`+0xd0`
 	
For more information, see 209918 on the EMC Online Support site.
	152324

SyncIQ might have consumed excessive amounts of CPU during the phase when SyncIQ is enumerating snapshot directories. This might have caused SyncIQ policies to take longer to complete.	152260

If SmartCache write caching was enabled and if clients were performing synchronous writes to the cluster, it was possible to encounter a runtime error that caused an affected node to unexpectedly restart. If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file:
 	
`Stack: --------------------------------------------------`
`kernel:cregion_issue_write+0xdcb`
`kernel:_cregion_write+0x1f5`
`kernel:cregion_write+0x24`
`kernel:cregion_flush+0xf6`
`kernel:coalescer_flush_overlapping+0x219`
`kernel:coalescer_flush_local_overlap+0x275`
`kernel:bam_coal_flush_local_overlap+0x2d`
`----------------------------------------------`
	152194

While running an initial SyncIQ job, the target root directory and its contents remained in a read-write state instead of read-only until the SyncIQ job completed. As a result, files could have been deleted or modified in the target cluster.	152190
In rare cases, an SMB client released its lease on a file before OneFS received a request to release the lease. If this occurred, the lwio process restarted	152123

File system issues resolved in OneFS 7.2.1.1	ID

unexpectedly, SMB clients connected to the affected node were disconnected, and lines similar to the following appeared in the `/var/log/messages` file:
 	
`Stack: ------------------------------------------------`
`lib/libc.so.7:thr_kill+0xc`
`/usr/likewise/lib/liblwiocommon.so.0:LwIoAssertionFailed+0x9f`
`/usr/likewise/lib/lwio-driver/`
`onefs.so:OnefsOplockBreakFillBuffer_inlock+0xbf`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsOplockComplete_inlock`
`+0x7e`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsOplockBreakToRH+0x187`
`/usr/lib/libisi_ecs.so.1:oplocks_event_dispatcher+0xf3`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsOplockChannelRead+0x8c`
`/usr/likewise/lib/liblwbase.so.0:EventThread+0x333`
`/usr/likewise/lib/liblwbase.so.0:LwRtlThreadRoutine+0xec`
`/lib/libthr.so.3:_pthread_getprio+0x15d`
`------------------------------------------------`	

Under rare circumstances, the lock subsystem did not drain fast enough, causing an assertion failure. When this issue occurred, the node restarted, and the following stack was logged to the `/var/log/messages` file:
 	
`Stack: ------------------------------------------------`
`kernel:isi_assert_halt+0x2ekernel:lki_lazy_drain+0xf76`
`kernel:_lki_split_drain_locks+0xa8`
`kernel:kt_main+0x15ekernel:fork_exit+0x75`
`------------------------------------------------`
`<3>*** FAILED ASSERTION must_drain ==> !pool->lazy_queue_size || !`
`li->mounted @ /b/mnt/src/sys/ifs/lock/lk_initiator.c:13270:`
`lki_lazy_drain_pool on LK_DOMAIN_DATALOCK took 302454934. lazy`
`queue 1870 -> 11. li->llw_count = 0, iter_count=11087431`
`chk_space_time =  0, chk_space_iters = 0 llw_time = 880073`
`llw_iters = 2503 reject_drain_time = 1550050 reject_drain_iters =`
`1 yield_time = 282713930 yield_iters = 11084926`
`shrink_lazy_queue_count = 11087431`
	152104

It was possible for a race condition between the group change and the deadlock probe—a mechanism that attempts to detect and correct deadlock conditions—to cause a node to restart unexpectedly.	151968

Isilon A100 nodes might have restarted unexpectedly during a group change, resulting in data unavailability. If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file:
 	
`Software Watchdog failed on CPU 1 (82353: kt: rtxn_split [-])`
`Stack: ------------------------------------------------ -`
`kernel:isi_hash_resize+0x31f`
`efs.ko:lki_handle_async_reacquire+0x262`
`efs.ko:lki_group_change_commit+0x727`
`efs.ko:lk_group_change_commit_initiator+0x32`
`efs.ko:rtxn_sync_locks_done+0x12e efs.ko:rtxn_split+0x4e9`
`efs.ko:rtxn_split_courtship_thread+0x388`
`efs.ko:kt_main+0x83kernel:fork_exit+0x77`
`------------------------------------------------`
	151956

If you copy configuration files while the isi_mcp process is running, by design, the MD5 command validates the files. If two configuration files with the same file name were copied almost simultaneously, after the second file copy was started, the MD5 process on the first file might have been stopped. If this occurred, the mcp process entered an infinite loop and the isi_mcp child process stopped responding. This issue could be verified by running the `isi_for_array -s 'ps auwxxxHl` command to check the running time of the isi_mcp process. In the	151952

File system issues resolved in OneFS 7.2.1.1	ID

following example output, the isi_mcp process was running for more than 6400 minutes (106 hours).
 	
`isi_for_array -s 'ps auwxxxHl | grep isi_mcp | grep -vi grep'`
`4284 93.0 0.0 55744  8176  ??  R     2Mar14 6425:30.28 isi_mcp: child (isi_mcp)`	

If L3 was enabled in a cluster environment that included SSDs on which L3 cache was previously disabled, the SSDs were smartfailed but not re-added as L3 devices. As a result, if you ran the `isi_devices` command, it was possible to see that the SSDs were not automatically transitioned from the `[REPLACE]` state back to the `[PREPARING]` state, and false drive replacement alerts were generated.	151942
If an SMB client changed the letter case of the name of a file or directory stored on the cluster, the file or directory's ctime (change time) value was not updated. As a result, the affected file or directory was not backed up during incremental backups.	146866

File transfer

File transfer issues resolved in OneFS 7.2.1.1	ID
If a client was connected to the cluster through vsftpd and ran the `ls` or `dir` commands for directories that contained more than 100,000 files, the vsftpd process reached its memory limit, and a memory allocation error occurred. As a result, the files in the affected directories could not be listed.	152267

Hardware

Hardware issues resolved in OneFS 7.2.1.1	ID
In some cases while a node was booting up, a race condition caused the front panel on the node to stop responding. If this issue occurred, the LCD display on the node showed the Isilon logo indefinitely and the buttons on the front panel were unresponsive.	157255
If you installed drive support package (DSP) 1.6 on a cluster that was running OneFS 7.2.1.0 and that contained Sunset Cove Plus SSDs, the SSDs might have been smartfailed prematurely. After DSP 1.6 was installed, this issue might have caused OneFS to incorrectly interpret a system-area wear-life threshold, and then detect that no system-area wear life remained on the SSDs.	157145
In rare cases, a failing dual in-line memory module (DIMM) caused a burst of correctable error correcting code (ECC) errors. If this burst of errors was extreme—for example, if it occurred tens of thousands of times per hour—the performance of the node and the cluster might have been degraded. If this issue occurred, a message similar to the following appeared tens of thousands of times per hour in the `/var/log/messages` file and on the console: `RDIMM P1-DIMM1A (cpu 0, channel 0, dimm 0) non-fatal (correctable) ECC error` This issue continued until the DIMM was replaced.	156347

Hardware issues resolved in OneFS 7.2.1.1	ID
If the hardware abstraction layer (HAL) could not detect the network interface card (NIC) in an Isilon node, the HAL assigned an empty string to the related nic name attribute in the `lni.xml` file, instead of returning an empty list. As a result, when the flexnet configuration file (`flx_config.xml`) was updated with this information, the related ‹nic-name› element in the `flx_config.xml` file was also empty. The empty element was an invalid entry in the file and it rendered the `flx_config.xml` file unusable by the node. Because an updated `flx_config.xml` file is propagated to all nodes in the cluster, this issue could have caused all nodes in the cluster to have a `flx_config.xml` file with invalid entries. If this occurred, client connections to the cluster might have been disrupted until the unusable `flx_config.xml` file was replaced.	155827
If a node with a LOX NVRAM card was unable to communicate with the NVRAM card because the NVRAM card controller was unexpectedly reset, the cluster became unresponsive to all client requests and data on the cluster was unavailable until the affected node was rebooted. **Note** Beginning in OneFS 7.2.1.1, if this issue is encountered, the affected node will be rebooted automatically to prevent the cluster from becoming unresponsive.	155387

If you ran the `isi firmware status` command on a cluster that contained S210 nodes with common form factor power supply units (PSUs) that had part number 071-000-022-00, and if firmware package version 9.3.1 or later was not installed on the cluster, messages similar to the following appeared on the console:
`CFFPS1_Blastoff CFFPS 09.05 2,7`
`CFFPS1_Blastoff_DC CFFPS <CFFPS1_Blastoff_DC> 2,7`
`CFFPS1_Optimus CFFPS <CFFPS1_Optimus_Acbel> 2,7`
`CFFPS2_Blastoff CFFPS 09.05 2,7`
`CFFPS2_Blastoff_DC CFFPS <CFFPS2_Blastoff_DC> 2,7`
`CFFPS2_Optimus CFFPS <CFFPS2_Optimus_Acbel> 2,7`
This issue occurred because earlier versions of OneFS, and earlier versions of the firmware package did not recognize PSU part number 071-000-022-00.
Note
This issue can be resolved in OneFS 7.2.1.0 by installing firmware package version 9.3.1 or later.
	154597

The following event message did not automatically clear after a boot drive was replaced: `Drive at Internal <drive_location> wear_life threshold exceeded: xx (Threshold: xx). Please schedule drive replacement.`	152321
The LED on the chassis turned solid red for a drive prior to completion of the smartfail process. As a result, the drive might have been replaced prematurely, possibly causing data loss.	152189
If you installed a new drive support package (DSP) on a node that already had a DSP installed and you then attempted to update a drive whose update was included only in the new DSP, the `fwupdate` command did not update the drive	152186

Hardware issues resolved in OneFS 7.2.1.1	ID
unless either the isi_drive_d process or the affected node was restarted. If this issue occurred, and you ran the `isi devices -a fwupdate` command before restarting the isi_drive_d process or the node, the following error appeared on the console: `'fwupdate' action complete, 0 drives updated, 0 updates failed`	

If you attempted to install a node firmware package that did not have support for the Chassis Management Controller (CMC) component, on a node that contained a CMC —for example, an S210, X210, X410, NL410, or HD400 node—the installation failed and an Unhandled Exception error similar to the following appeared on the console:
 	
`FAILED : Unhandled exception in safe.id.cmc ('empty_fw_object' object has no attribute 'update')`
 	
Note
 	
Beginning in OneFS 7.2.1.1, if the preceding conditions exist, the following message appears on the console, where `<partnumber>` is the part number of the CMC:
 	
`FW archive does not have support for PN <partnumber>`
	152183

The isi_sasphymon process could potentially close a valid 0 file descriptor. If this issue occurred, any drive associated with the file descriptor would no longer be monitored by the isi_sasphymon process. This issue would also cause excessive logging in the `/var/log/isi_sasphymon.log` file similar to the following:
 	
`isi_sasphymon[3979]: Can't get SCSI Log Sense page 0x18 from Bay 2- scan 6`
`isi_sasphymon[3979]: cam_get_inquiry: error from cam_send_ccb: 9`
`isi_sasphymon[3979]: scsi_get_info: error from scsi_get_inquiry`
	152150

If you ran the `isi_reformat_node` command on a node containing self-encrypting drives (SEDs), sometimes the SEDs could not be released from ownership, and when the node rebooted, the unreleased SEDs came up in a SED_ERROR state. **Note** Beginning in OneFS 7.2.1.1, if you run the `isi_reformat_node` command on a node containing self-encrypting drives (SEDs) that cannot be released from ownership, the following messages appear on the console where `<affected drives>` is a list of the affected drives: `isi_wipe_disk has failed in isi_reformat_node Failed to wipe the following drives: <affected drives> Opening zsh to allow user to revert these drives using the '/usr/bin/isi_hwtools/isi_sed revert' command. To continue with the reformat, enter 'exit' in the shell.`	152130

Hardware issues resolved in OneFS 7.2.1.1	ID
Note If the reformat process continues without reverting the listed drives, it is likely they will be in a SED_ERROR state on the next node boot.	

After replacing boot flash drives in a node and then running the `gmirror status` command, the correct number of active components appeared in the output but a status of `DEGRADED` was incorrectly returned for some components. In the example below, the keystore and mfg mirrors were affected:
`Name Status Components`
`mirror/root0 COMPLETE ad7p4 ad4p4`
`mirror/keystore DEGRADED ad7p11 ad4p12`
`mirror/var-crash COMPLETE ad7p10`
`mirror/mfg DEGRADED ad7p9 ad4p10`
`mirror/journal-backup COMPLETE ad7p8 ad4p8`
`mirror/var1 COMPLETE ad7p7 ad4p7`
`mirror/var0 COMPLETE ad7p6 ad4p6`
`mirror/root1 COMPLETE ad7p5 ad4p5`
Although the operation of the node was unaffected, the incorrect status sometimes led to unnecessary service calls for hardware exchanges.
	152113

HDFS

HDFS issues resolved in OneFS 7.2.1.1	ID
Because the WebHDFS CREATE operation did not explicitly instruct the system to create parent directories, if OneFS received a WebHDFS request to create a file or directory within a parent directory that did not yet exist, the request failed. Beginning in OneFS 7.2.1.1, OneFS will automatically create parent directories if it receives a WebHDFS create request that requires them. **Note** This issue affected Ambari 2.1.0 only.	159567
Because OneFS treated query strings from WebHDFS clients as case-sensitive, some valid queries or operations might have failed. For example, OneFS expected operations such as `GETFILESTATUS` to be upper case, while Boolean arguments and strings were expected to be lower case. As a result, queries similar to the following might have failed because `GetFileStatus` is entered in mixed case: `http://isilon_ip:8082/webhdfs/v1/?op=GetFileStatus&user.name=root` **Note** This issue affected Ambari 2.1.0 only.	159018
If multiple threads attempted to simultaneously update the stored list of blocked IP addresses, the HDFS service restarted and client sessions were disconnected. The service was automatically restored after a few seconds.	158077
If the maximum number of HDFS client connections to the cluster was reached, all worker threads remained busy during processing. As a result, no further cluster connections could be established, namenode remote procedure calls (RPCs) were	154631, 152326

HDFS issues resolved in OneFS 7.2.1.1	ID
queued for long periods of time, and the HDFS server incorrectly appeared to be unavailable.	

If you tried to change ownership of files or directories through the WebHDFS REST API by setting only the owning user or the owning group of a file or directory (but not both), an exception error similar to the following might have appeared in the command-line interface:
 	
`"RemoteException":`
`{`
`"exception" : "SecurityException",`
`"javaClassName": "java.lang.SecurityException",`
`"message" : "Failed to get id rec: 1:"`
`}`
`}`
 	
Additionally, Ambari 2.1 might have failed to install Hortonworks Data Platform 2.3 through the WebHDFS REST API.
	154628

If Kerberos authentication was used with HDFS, the isi_hdfs_d process could eventually have run out of memory and unexpectedly stop. If this issue occurred, an isi_hdfs_d.core file was created in the `/var/log/crash/` directory, and the following lines appeared in the `/var/log/messages` file:
 	
`isi_hdfs_d: isi_hdfs_d: *** FAILED ASSERTION cv->members @ s11n.c:137: oom`
`[kern_sig.c:3376](pid 27685=""isi_hdfs_d"")(tid=102752) Stack trace:`
`Stack: --------------------------------------------------`
`/lib/libc.so.7:thr_kill+0xc`
`/usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/boot/kernel.amd64/kernel: /usr/bin/isi_hdfs_d:file_status_array_append+0x9b`
`/boot/kernel.amd64/kernel: /usr/bin/isi_hdfs_d:util_make_directory_listing+0x90d`
`/boot/kernel.amd64/kernel: /usr/bin/isi_hdfs_d:_rpc2_getListing_ap_2_0_2+0xbf`
`/boot/kernel.amd64/kernel: /usr/bin/isi_hdfs_d:rpc_ver2_2_execute+0x21c`
`/boot/kernel.amd64/kernel: /usr/bin/isi_hdfs_d:_asyncrpctask+0x3a`
`/boot/kernel.amd64/kernel: /usr/bin/isi_hdfs_d:_workerthr+0x257`
`/boot/kernel.amd64/kernel: /lib/libthr.so.3:_pthread_getprio+0x15d`
`/boot/kernel.amd64/kernel:`
`----------------------------------------------------`
	153778

The isi_hdfs_d proccess no longer unnecessarily logs the following message to the `/var/log/isi_hdfs_d.log` file: `RPC getDatanodeReport raised exception: Could not parse 'GetDatanodeReport'`	152250
When creating and writing a file through WebHDFS, the exclusion of Java class names might have caused unexpected errors similar to the following: `mkdir: The requested file or directory does not exist in the filesystem.`	152134
If a Hadoop client tried to export data in Hive to a directory that already existed, and the client did not have permissions on the directory to make the change, the	152131

HDFS issues resolved in OneFS 7.2.1.1	ID

`mkdir` command failed. If the `mkdir` command failed, an error similar to the following appeared on the client:
`FAILED: Execution Error, return code 2 from org.apache.hadoop.hive.ql.exec.CopyTask`
In addition, the following line appeared in the `/var/log/isi_hdfs_d.log` file on the node:
`pfs_mkdir_p failed in mkdirs with unusual errno: Operation not permitted`	

The Ambari server sent a `check_host` command instead of a `host_check` command. If this issue occurred, the following message was logged to the `/var/log/isi_hdfs_d.log` file: `Ambari: Tried to access an undefined component name, which is most likely unsupported: check_host`	152120

Job engine

Job engine issues resolved in OneFS 7.2.1.1	ID
If you attempted to start a PermissionRepair job from the **ClusterManagement** › **Job Operations** › **Job Types** › **Start Job** dialog, and you set the repair type to **Clone: copy permissions from the chosen path to all files and directories** or **Inherit: recursively apply an ACL**, the **Template File or Directory** field did not appear. As a result, you could not configure a **PermissionRepair** job to perform a Clone-type or Inherit-type repair.	154630
In the OneFS web administration interface, the **Edit Job Type Details** page for jobs that had a schedule set to Every Sunday at 12:00am displayed **Close** and **Edit Job Type** buttons instead of **Cancel** and **Save Changes** buttons.	152182

Migration

Migration issues resolved in OneFS 7.2.1.1	ID
If you restarted a full or incremental `isi_vol_copy` migration three or more times, and if a specific file was in the process of being copied to the target cluster each time the `isi_vol_copy` migration was restarted, the file was not successfully copied to the target cluster. **Note** You might still encounter this issue if you restart an `isi_vol_copy` migration of a single, large file three or more times.	154338
If the isi_vol_copy_vnx tool was used to migrate data from a VNX array to a OneFS cluster, and if the data contained any NULL SIDs, the migration process stopped,	151949

Migration issues resolved in OneFS 7.2.1.1	ID

and a message similar to the following appeared in the `/var/log/messages` file:
 	
`/boot/kernel.amd64/kernel:`
`[bam_acl.c:190](pid 83648="isi_vol_copy_vnx")(tid=101308)`
`ifs_verify_acl:`
`Failed verifying`
`security_ace on lin:1:02df:da06. Ace#3. An ACE cannot have a NULL identity type.`	

If you renamed or deleted a directory on the source cluster prior to performing an incremental migration, and if you then created a hard link file with the original name of the deleted or renamed directory, the incremental migration failed. If this occurred, errors similar to the following appeared in the `/var/log/isi_vol_copy.log` file and also on the console:
 	
`[INFO] [isi_vol_copy stdout]: Error:`
`[INFO] [isi_vol_copy stdout]: Failed to create hardlink ./HL_PREFIX_DIR4->./DIR4: err:Operation not permitted[1]`
`…`
`[INFO] [isi_vol_copy stdout]: *** FAILED ASSERTION !"fixupentrytype()" @ /b/mnt/src/isilon/lib/isi_emctar/updated.c:431:`
	151925

During an incremental migration through the isi_vol_copy utility, if a socket file needed to be extracted or migrated, the migration failed and an error similar to the following appeared on the console: `./f2: cannot create file: Operation not supported`	151924
During a full or incremental migration, if midfile checkpoints were enabled or if the `WINDOW_MAX_SIZE > 0` environmental variable was set, an error similar to the following appeared in the `/var/log/isi_vol_copy.log` file and on the console, and the migration had to be restarted from the beginning: `createleaves() - ./file19: not found on tape first = 11988, curfile.ino = 19619`	151922

Networking

Networking issues resolved in OneFS 7.2.1.1	ID
If Source Based Routing (SBR) was enabled on the cluster, client connections that were handled by SBR were disconnected if the MAC address (ARP entry) for the relevant subnet gateway expired. This issue occurred because nodes in the cluster did not send an ARP request to refresh the MAC address and, as a result, attempted to send network traffic to an incorrect destination MAC address for the gateway. **Note** The default expiration time for an ARP entry is 10 minutes.	155824
If the cluster contained X410, S210, or HD400 nodes that had a 10 GigE network interface card that was using the Broadcom NetXtreme Ethernet (BXE) driver and any external network subnets connected to the cluster were set to 9000 MTU, an	154454

Networking issues resolved in OneFS 7.2.1.1	ID
error similar to the following appeared in the `/var/log/messages` file, and the affected nodes rebooted: `ERROR: mbuf alloc fail for fp[01] rx chain (55)` For more information, see article 200096 on the EMC Online Support site.	

Even if a cluster was not configured to use IPv6, the IPv6 traffic on the network might have caused lock contention. If this issue occurred, nodes with locks in contention restarted unexpectedly and messages similar to the following appeared in the `/var/log/messages` file:
 	
`Panic occurred in module kernel loaded at 0xffffffff80200000:`
`Stack: ----------`
`kernel:trap_fatal+0x9f`
`kernel:trap+0x28b`
`kernel:turnstile_wait+0x10b`
`kernel:_rw_rlock+0x2fe`
`kernel:inp_fmt+0x10e`
`kernel:inp_fmt_all+0x2d`
`kernel:mb_audit_fmt+0x18`
`efs.ko:rbm_mbufs_exhausted+0x1e`
`kernel:isi_swatchdog_panic+0xcb`
`kernel:isi_swatchdog_hardclock+0x232`
`kernel:hardclock_cpu+0x13d`
`kernel:hardclock+0x10`
`kernel:lapic_handle_timer+0x10f`
`kernel:cpu_idle_mwait+0x5a`
`kernel:sched_idletd+0x114`
`kernel:fork_exit+0x7f`
`----------`
	154262

In some cases, the Mellanox InfiniBand driver waited for a hardware status register to be cleared, which caused the driver to enter a read and retry loop. If the retry loop timed out, the driver attempted to print out a significant amount of system data three times. Since printing the system data output was enabled by default, and because there was a significant amount of data to be processed, the driver eventually triggered several software watchdog time outs. After five of these time outs, the software watchdog rebooted the affected node and the following lines appeared in the `/var/log/messages` file:
 	
`Consecutive swatchdog state warnings: 5`
`Opt-in swatchdog state warnings: 5`
`Memory pressure swatchdog warnings: 0`
`Majority of swatchdog warnings by opt-in threads!`
`panic @ time 1394782550.534, thread 0xffffff06ebc8a000: Software watchdog timed out`
`cpuid = 3`
`Panic occurred in module kernel loaded at 0xffffffff80200000:`
`Stack: ----------`
`kernel:isi_swatchdog_panic+0x15`
`kernel:isi_swatchdog_hardclock+0x1ea`
`kernel:hardclock_cpu+0xd9`
`kernel:lapic_handle_timer+0x15c`
`kernel:spinlock_exit+0x32kernel:putcons+0x3e`
`kernel:putchar+0x7akernel:kvprintf+0xa3b`
`kernel:__vprintf+0x5bkernel:printf+0x70`
`kernel:_fmt_flush+0x3d`
`kernel:fmt_append+0x47`
`kernel:fmt_print_num+0x1f7`
`kernel:fmt_vprint+0x302`
`kernel:fmt_print+0x5f`
`mthca.ko:_mthca_mst_dump+0xc7`
`mthca.ko:mthca_print_mst_dump+0x56`
`mthca.ko:check_time+0x1c`
	153424

Networking issues resolved in OneFS 7.2.1.1	ID
``` mthca.ko:mthca_cmd_poll+0x105 mthca.ko:mthca_cmd_box+0x65 mthca.ko:mthca_MAD_IFC+0x1cd mthca.ko:mthca_query_port+0x107 kernel:port_active_handler+0x31 kernel:sysctl_root+0xd6 kernel:userland_sysctl+0x15c kernel:__sysctl+0xa9 kernel:isi_syscall+0x53 kernel:syscall+0x1db --- ``` **Note** Beginning in OneFS 7.2.1.1, the system data is not printed by default, allowing the read and retry loop to complete more quickly, and minimizing the chance of a software watchdog time out events.	
In rare cases, a race condition between the networking service and the SmartConnect service caused the SmartConnect service IP to be assigned to a node before the network addresses were updated in the IP pool. If this issue occurred, connection requests to the cluster failed until the dynamic IP addresses in all network pools were manually rebalanced by running the `isi networks` command with the `--sc-rebalance-all` option.	153396
If you performed an extended link flapping test on a node containing a Chelsio network interface card (NIC), the NIC eventually became unresponsive, and had to be manually disabled and then re-enabled before it resumed normal operations. While the NIC was unresponsive, external clients could not communicate with the node; however, because the node's back-end communication was unaffected, data on the node was still available to clients connected to the cluster through other nodes.	151945

NFS

NFS issues resolved in OneFS 7.2.1.1	ID
If an NFS client attempted to create an empty symbolic link—for example, by running the `ln -s "" <filepath>` command to create a symbolic link without a path to a source file—the NFS server restarted unexpectedly and interrupted the workflow. If this issue occurred, messages similar to the following appeared in the `/var/log/messages` file: ``` pid 24071 (nfs), uid 0: exited on signal 11 (core dumped) [kern_sig.c:3376](pid 24141="nfs")(tid=101258) Stack trace: Stack: --- /lib/libc.so.7:strlen+0x2 /usr/likewise/lib/lwio-driver/nfs.so:NfsProtoNfs3Dispatch+0xddf /usr/likewise/lib/lwio-driver/nfs.so:NfsProtoNfs3CallDispatch+0x3e /usr/likewise/lib/liblwbase.so.0:SparkMain+0xb7 --- ```	161286

NFS issues resolved in OneFS 7.2.1.1	ID
In addition, because the MCP process attempted and failed to restart the NFS server, the following message appeared on the **Cluster Events** page in the OneFS web administration interface and was logged in the `/var/log/messages` file: `Process lwsm of service nfs has failed to restart after multiple attempts`	
If you ran the `rmdir` command to remove a directory from an NFS export by using character encoding other than the default encoding—for example, CP932 or ISO-8859-1 encoding—and if the directory name contained a special character, the directory was not removed and a message similar to the following appeared on the console: `failed to remove `\directory_path': Invalid argument`	159373
If all of the following conditions were met, users connected to an NFS export received `Permission denied` errors when they attempted to access file system objects to which they should have had access: • The `--map-lookup-uid` option was enabled (set to `yes`) for the affected NFS export. • The group owner of the affected file system object was one of the user's supplemental groups rather than the user's primary group. • The cluster-side lookup for the user's supplemental groups failed. This issue occurred because, when the lookup for the user's UID failed, OneFS did not correctly apply supplemental group permissions to the user. As a result, the user was denied access to the file system object.	157387

Under some circumstances, after an SMB2 client attempted to access a file on the cluster through a symbolic link, OneFS returned an ESYMLINKSMB2 error (an internal error that is not seen on the client). If this error was returned, the symbolic link was resolved; however, some kernel memory that was allocated in order to complete the process of resolving the symbolic link was not deallocated after the link was resolved. As a result, over time a node's kernel processes might have run out of memory to allocate. If this occurred, the affected node rebooted unexpectedly, and messages similar to the following appeared in the `/var/log/messages` file on the affected node:
`/boot/kernel.amd64/kernel: Pageout daemon can't find enough freepages.`
`System running low on memory. Check for memory pigs`
	157383

If an NFS operation failed because the NFSv3 client that attempted to perform the operation did not have adequate access permissions and then the same NFSv3 client sent a request for file system information, the NFS server unexpectedly restarted and an error message similar to the following was logged in the `/var/log/nfs.log` file: `[lwio] ASSERTION FAILED: Expression = (0), Message = 'Got access denied on stat-only open!'`	157375
If the `Deny permission to modify files with DOS read-only attribute over Windows File Sharing (SMB)` ACL policy option was	155823

NFS issues resolved in OneFS 7.2.1.1	ID
enabled, files to which the `DOS READ-ONLY` flag was applied might have appeared writeable to NFS clients. As a result, a process on an NFS client might have attempted to write a change to a read-only file. If this occurred, the write to the file might have been rejected by the NFS server without sending an error to the client, or a permissions error might have appeared on the client when the file was closed or when the system attempted to move the file's data to persistent storage.	
Although the correct ACLS were assigned to a file—for example, std_delete or modify—NFSv3 and NFSv4 clients could not delete, edit, or move the file unless the `delete_child` permission was set on the parent directory. For more information, see ETA 204898 on the EMC Online Support site.	155800
If the lsass process was not running when NFS configuration information was refreshed on the cluster, it was possible for empty netgroups to be propagated to some or all of the cluster nodes. If this issue occurred, NFS clients were unable to mount NFS exports.	154452
Because OneFS 7.2.1 and later return 64-bit NFS cookies, some older, 32-bit NFS clients were unable to correctly handle read directory (readdir) and extended read directory (readdirplus) responses from OneFS. In some cases, the affected 32-bit clients became unresponsive, and in other cases, the clients could not view all of the directories in an NFS export. In the latter cases, the client could typically view the current directory (".") and its parent directory (".."). For more information, see ETA 205085 on the EMC Online Support site.	154145
In environments with NFS exports rules that referenced hundreds of unresolvable hostnames, the `isi nfs exports list --verbose` command consumed too many reserved privileged socket connections when it was interacting with the isi_netgroup_d process. As a result, commands that used isi_rdo for intra-node communications (for example, `isi_gather_info` or `isi_for_array`) failed to complete for a few seconds. If this occurred, a message similar to the following appeared on the console: `isi_rdo: [Errno 13] TCPTransport.bind_to_reserveport: Unable to bind to privileged port.`	153458
If an NFS client attempted to send an NLM asynchronous request to lock a file and received an error in response to the request, a socket was opened but was not closed. Over time, it was possible for the maximum number of open sockets to be reached. If this occurred, processes could not open new sockets on the affected node. If this issue occurred, affected nodes might have been slow to respond to file lock requests or lock requests sent to an affected node might have timed out, which could have prevented NFS clients from accessing files or applications on the cluster.	153452
If you attempted to manually shut down the NFS server on one or more nodes while NFSv4 clients were connected to those nodes, the operation could have failed. If this issue occurred, lines similar to the following appeared in the `/var/log/messages` file and the affected nodes had to be rebooted: `FAILED on action list 'stopnfs': action 1-2,4-6/6 SERVICE nfs(pid=xyz) returned exit status 1`	152263

NFS issues resolved in OneFS 7.2.1.1	ID
If you created a hard link that contained a colon (:) from an NFSv3 client, the colon and any characters that followed it were removed from the hard link name. As a result, the hard link on the cluster did not have the correct name. If removing the colon and following characters resulted in changing the hard link name to a file name that was already in use in the destination directory on the cluster, a file name conflict resulted, and a `File exists` error appeared on the NFS client.	152257
If an NFS client attempted to list a file or directory at the root of an NFS export mount point directory that began with two dots—for example, `/mnt/nfs_export/..my_folder`— and the requested file or directory did not exist, OneFS returned the contents of the NFS export instead of a `file not found` error message.	152256
If a client held a read lock on a file and an NFS4 client checked the lock status of the file, the response from the cluster incorrectly reported that the original client was holding a write lock on the file. This issue might have caused the program that the NFS client was using to work improperly.	152255
Because the nfs and onefs_nfs drivers (and the flt_audit_nfs driver, if you enabled protocol auditing) share the same process ID, if one of these drivers failed to start, the MCP process did not always detect the failure and did not always restart the stopped drivers.	152172
On the **NFS Export Details** page, if you added a secondary group for either the Map Root User or the Map Non Root User, the value field did not display until you refreshed the web administration interface page.	152142

OneFS API

OneFS API issues resolved in OneFS 7.2.1.1	ID

If you created an SMB share and then created a single user or group with run-as root permissions to the share, the user or group could not be deleted, and the user or group's run-as-root permission could not be modified. If you attempted to delete the user or group, the command appeared to successfully complete; however, the user or group was not deleted. If you attempted to modify the user or group's permissions, the command appeared to successfully complete; however, the original permissions entry was not removed, and an additional entry, with the modified permissions, was added to the share. In the example below, the domain admins group displays the duplicate entries created when the group's permissions run-as-root was modified:
`Account Account Type Run as Root Permission Type Permission`
`---`
`EXAMPLE\domain admins group True allow full`
`EXAMPLE\domain users group False allow change`
`EXAMPLE\domain admins group False allow full`
	157512

OneFS API issues resolved in OneFS 7.2.1.1	ID
If you created an SMB share and then created a single user or group with run-as-root permissions to the share, the user or group could not be deleted, and the user or group's run-as-root permission could not be modified. If you attempted to delete the user or group, the command appeared to successfully complete; however, the user or group was not deleted. If you attempted to modify the user or group's permissions, the command appeared to successfully complete; however, the original permissions entry was not removed, and an additional entry, with the modified permissions, was added to the share. In the example below, the domain admins group displays the duplicate entries created when the group's permissions run-as-root was modified: `Account Account Type Run as Root Permission Type Permission ---------- ---------- FOO\domain admins group True allow full FOO\domain users group False allow change FOO\domain admins group False allow full`	152197

OneFS web administration interface

OneFS web administration interface issues resolved in OneFS 7.2.1.1	ID
In the OneFS web administration interface, if the path to the shared directory for an SMB share was long enough to exceed the width of the SMB shares page, the shared directory **Edit** link was sometimes not visible. **Note** The **Edit** link was accessible if you used the **Tab** key to move to the link.	155798

Security

Security issues resolved in OneFS 7.2.1.1	ID
If a DHE_EXPORT cipher suite was enabled on a server but not on a client, the Transport Layer Security (TLS) protocol version 1.2 and earlier did not properly convey a DHE_EXPORT choice, which could have allowed man-in-the-middle attackers to conduct cipher-downgrade attacks. This issue is also known as the Logjam issue. For more information, see article 205840 on the EMC Online support site.	157834
If RFC 2307 was enabled on the cluster, and if Windows Services for UNIX (SFU) was not configured on the trusted domain, then users in the trusted domain were given root-level privileges (userid=0). For more information, see article 210195 on the EMC Online Support site.	151939

SmartLock

SmartLock issues resolved in OneFS 7.2.1.1	ID
If the compadmin user on a compliance mode cluster ran the `sudo isi_gather_info` command, the `sudo isi_gather_info` command	152119

SmartLock issues resolved in OneFS 7.2.1.1	ID
successfully gathered all of the expected files on the local node, but was unable to gather all of the expected files on remote nodes. This issue occurred because some files on the cluster can be read only by the root user, and the `sudo` command did not enable the compadmin user to run commands as root on remote nodes.	

SmartQuotas

SmartQuotas issues resolved in OneFS 7.2.1.1	ID
If you configured an email mapping rule to route quota limit notifications from your primary domain to (Active Directory) AD users in another (Map-to) domain, OneFS might not have correctly routed the limit notifications to the AD users in the Map-to domain. If this issue occurred, notifications were routed to an incorrect email address that consisted of the name of the owner of the affected directory and providers in the DNS search list.	156727
If you edited the usage limits of an existing directory quota in the OneFS web administration interface, the **Show Available Space as: Size of hard threshold** and **Size of cluster** options were missing from the **Set a hard limit** section. This issue occurred if you chose the **Size of cluster** option when you created the directory quota with a hard limit.	154332

If a SmartQuota threshold was exceeded and then files were moved or deleted to correct the issue, an alert was sometimes sent after the issue was corrected, even though the threshold was no longer exceeded. If this occurred, a false alert similar to the following was generated, where `/ifs/<path>` was the path of the directory that temporarily exceeded the configured threshold:
`Your root quota under /ifs/<path> has been exceeded.`
`Your quota is 12 TB,and 6.7 TB is in use. You must delete files to bring usage below 12 TB before you can create or modify files. Please clean up and free some disk space.`
	153406

If you changed a quota's soft or hard limit through the web administration interface, the **Enforced** parameter changed from `Yes` to `No`, making the quota accounting-only. As a result, configured quota usage limits were not enforced.	152261
If a quota was created with a hard, soft, or advisory threshold that included a decimal point—for example, `isi quota quotas create --hard-threshold=4.5T`—the operation failed, and a message similar to the following appeared on the console: `Unknown suffix '.5T'; expected one of ['b', 'K', 'M', 'G', 'T', 'P', 'B', 'KB', 'MB','GB', 'TB', 'PB']`	152191
In the web administration interface, after clicking **View details** for a quota on the **Quotas & Usage** page, the **%Used** value under **Usage Limits** did not always correctly match the percentage value displayed under **%Used** in the top summary row for the quota.	152110
If you configured a storage quota on a directory with a pathname that contained a single, multibyte character, and if a quota notification email was sent for that directory, the multibyte character in the pathname that appeared in the quota	151953

SmartQuotas issues resolved in OneFS 7.2.1.1	ID
notification email was replaced with an incorrect character, such as a question mark.	

SMB

SMB issues resolved in OneFS 7.2.1.1	ID
On clusters running OneFS 7.1.1.6, 7.1.1.7, 7.2.0.4 or 7.2.1.0, if an SMB client sent a request to initiate an SMB1 session on the cluster and then resent the request, the second SMB1 request was blocked until the first request was processed. On networks where responses to requests were frequently delayed—for example, if a domain controller was slow to process authentication requests—it was possible for the blocked session requests to become permanently stuck . If this occurred, the lwio process on the affected node stopped responding, no new connections to the node could be established, and existing connections might have been disconnected. In some cases, the lwsm and lwio processes had to be restarted in order to resolve the issue.	163485

A change to the way OneFS 7.2.1.1 handles access to the directories within an SMB share might prevent SMB clients from renaming directories or files within SMB shares. This issue occurs if the object the SMB client is attempting to rename is contained within a subdirectory of a directory to which the client has been granted only the Traverse Folder/Execute File permission. Under these conditions, if an SMB client attempts to rename an object within the SMB share, an error appears on the client and the object is not renamed. In some cases—for example, when creating new files within Microsoft Windows Explorer—the client might also be prevented from creating new files because the files cannot be renamed.
The exact error that appears on the client depends on the operating system from which the client is connecting to the cluster, the SMB version the client is using to access the share, and the interface the client is using to view the contents of the share.
For example, if a Microsoft Windows SMB client connected to the cluster over SMB2, attempts to rename a file on the cluster through Windows Explorer, an error similar to the following appears:
`The file 'file' is too large for the destination file system`
However, if a Windows client connected to the cluster over SMB2 attempts to rename a file through a command prompt window, the following error appears:
`The parameter is incorrect.`
	162368

If a Windows client that was connected to the cluster through SMB copied a file from the cluster, the timestamp metadata applied to the file might have become invalid. This issue occurred because OneFS did not properly interpret the value assigned to a file's timestamp metadata if the value was set to `-1`, which is a valid value. Workflows that rely on timestamp metadata might have been negatively affected by this issue. **Note** The SMB protocol specifies that, when file attributes are set, a value of `-1` indicates that the attribute in the corresponding field must not be changed.	161121

SMB issues resolved in OneFS 7.2.1.1	ID
If OneFS SMB2 symbolic link translation was disabled on the cluster, symbolic links to directories could not be deleted.	160739
If you used the **Previous Versions** tab of the file properties dialog to open a previous version of a Microsoft Office file that was stored on the cluster, and if the file resided in a directory for which a OneFS snapshot had been created between saving versions of the file, the previous version could not be opened and an error message similar to the following appeared on the Windows client, where `<file>` was the name of the file that you attempted to open: `Could not open '\\cluster\<file>'`	157726
If an SMB share on the cluster was configured with the `Impersonate Guest` security setting set to `Always`, and if a large number of SMB sessions to the share were being opened and closed, an extra cred file was opened for each SMB session. However, when the SMB session ended, the extra cred file was not correctly closed and, over time, it was possible for the number of open cred files to reach the maximum number of open files allowed. If this occurred, new SMB sessions to the affected node could not be established, and messages similar to the following appeared in the `/var/log/lwiod.log` file: `Failed to accept connection due to too many open files`	157031
Due to a race condition that could occur when multiple SMB1 sessions were being opened on the same connection, the lwio process sometimes restarted unexpectedly. If the process restarted, SMB clients connected to the affected node were disconnected from the cluster.	155818

If a symbolic link was migrated from a Microsoft Windows client to the cluster, if the tool that was migrating the data attempted to update the attributes of a symbolic link that had already been migrated, the attributes could not be updated, and the migration of the symbolic link failed. For example, if you attempted to migrate data to an Isilon cluster using the EMCopy tool, and if the data contained a symbolic link, the symbolic link was initially migrated but EMCopy could not apply attributes to the symbolic link, and an error similar to the following appeared on the EMCopy client:
`ERROR (5) : \path_to_target\symbolic_link -> Unable to set access time`
In addition, if the EMCopy tool attempt to retry the failed operation, the retry failed and an error similar to the following appeared on the EMCopy client:
`ERROR (4392) : \path_to_target\symbolic_link -> Unable to open, Failed after 1 retries.`
	155815

If you attempted to migrate a directory symbolic link from a Microsoft Windows client to an Isilon cluster, OneFS returned a response to the Windows client indicating that the operation was not supported, and the symbolic link was not migrated. Depending on the application that was being used to migrate the data, error messages might have appeared on the client. For example, if you attempted	155808

SMB issues resolved in OneFS 7.2.1.1	ID
to migrate data to an Isilon cluster using the EMCopy tool, the symbolic links were not migrated, and an error similar to the following appeared on the EMCopy client: `ERROR (50) : \path_to_target -> symbolic_link : symlink creation failure`	
SMB clients were unable to display alternate data stream information for files on the cluster that contained alternate data streams.	154839
If SMB auditing was enabled and you set the `--max-cached-messages` parameter to `0` (zero) to disable message caching, the SMB client session and negotiate requests that were waiting to be audited might have prevented new SMB session and negotiate requests from being processed. If this occurred, SMB clients might have been prevented from establishing new connections to the cluster until the backlog of audit messages was processed. **Note** Beginning in OneFS 7.2.1.1, if you set the `--max-cached-messages` parameter to `0` to disable message caching, and the Common Event Enabler (CEE) server becomes unavailable, some audit messages that have not yet been logged might be discarded. This behavior prevents a backlog of requests from disrupting SMB client requests and connections.	154275
If OneFS translation of SMB2 symbolic links was disabled, and a user attempted to delete a symbolic link to a file that was stored on the cluster, the target file might have been deleted instead of the symbolic link. This issue could have resulted in data loss. For more information, see article 207782 on the EMC Online Support site.	153769
If OneFS SMB2 symbolic link translation was disabled on the cluster, symbolic links to directories might have failed and an error similar to the following might have appeared on the client: `The symbolic link cannot be followed because its type is disabled.` For more information about enabling and disabling OneFS SMB2 symbolic link translation, see article 193808 on the EMC Online Support site.	153769
If you queried the contents of a directory in an SMB share from a Microsoft Windows command prompt, and if you included the search string *.* (asterisk–dot–asterisk) immediately after other search characters in the query—for example, `dir do*.*`—the search results did not include the expected files or directories. This issue occurred because OneFS treated the dot as a character rather than as a wildcard. **Note** Searches with only the *.* string listed the entire contents of the directory, as expected.	152340
If you attempted to override the default Windows ACL settings that were applied to an SMB share, by adding custom ACLs to the `/ifs/.ifsvar/smb/isi-share-default-acl/` template directory, the overrides were not implemented. As a result, actual access permissions on the SMB share did not match expected results.	152266

SMB issues resolved in OneFS 7.2.1.1	ID
If the FILE_OPEN_REPARSE_POINT flag was enabled, and an SMB client opened an alternate data stream (ADS) through a symbolic link, the ADS was inaccessible, and the following error appeared on the console: `STATUS_STOPPED_ON_SYMLINK`	152264

If you ran the EMCopy application to migrate data containing symbolic links to the cluster, the SMB process unexpectedly restarted because of an lwio process assertion failure. When the SMB process restarted, clients were disconnected from the cluster and the following error message appeared in the /var/log/lwiod.log file:
 	
`ASSERTION FAILED: Expression = (pFcb->bIsDirectory == bIsDirectory)`
 	
In addition, the following lines appeared in the /var/log/messages file:
 	
`/lib/libc.so.7:thr_kill+0xc`
`/usr/likewise/lib/liblwiocommon.so.0:LwIoAssertionFailed+0xa4`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsCreateFCB+0x896`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsCreateFileCcb+0x3b0`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsCreateInternal+0x90e`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsCreate+0x28d`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsProcessIrpContext +0x12b`
`/usr/likewise/lib/liblwbase.so.0:CompatWorkItem+0x16`
`/usr/likewise/lib/liblwbase.so.0:WorkThread+0x256`
`/usr/likewise/lib/liblwbase.so.0:LwRtlThreadRoutine+0xee`
`/lib/libthr.so.3:_pthread_getprio+0x15d`
 	
On the client, EMCopy might have displayed the following error message:
 	
`ERROR (50) : \\TARGET\symlink ->  folder:symlink creation failure`
	152192

If you ran the isi statistics client command to view information about some SMB1 and SMB2 read and write operations—for example, the namespace_write operation—the word UNKNOWN appeared in the UserName column, instead of a valid user name. As a result, if you ran scripts to filter read/write operations per user, the scripts did not work correctly.	151962
OneFS did not always send the expected share flags to SMB2 clients. As a result, if the inheritable-path-acl option was set for a share, UNC path hyperlinks in Microsoft Outlook emails failed to open.	151936
If the OneFS file system quota was exceeded, an incorrect STATUS_QUOTA_EXCEEDED error was returned during SMB1 and SMB2 write operations instead of STATUS_DISK_FULL. As a result, the client ignored the error and write requests continued, but were not applied, because they were over quota. If any binary files—for example, Microsoft Outlook .pst files—were affected by this issue, they would become unusable.	151929
During an upgrade to OneFS 7.1.1.x, an upgrade script did not properly interpret an empty string value for the HostAcl parameter in the /ifs/.ifsvar/main_config.gc file. This caused SMB shares to be inaccessible after the upgrade was complete, and as a result, the SMB shares had to be re-created. If this occurred, output similar to the following appeared after running the isi_gconfig	151746

SMB issues resolved in OneFS 7.2.1.1	ID
`registry.Services.lwio.Parameters.Drivers.srv.HostAcl` command: `registry.Services.lwio.Parameters.Drivers.srv.HostAcl (char**) = [ "" ]`	

Upgrade

Upgrade issues resolved in OneFS 7.2.1.1	ID
Although the `isi pkg` command was not intended to be used to install a drive support package (DSP), it was possible to install a DSP by running the `isi pkg` command. If a DSP was installed using the `isi pkg` command, the cluster might have exhibited unexpected behavior until the DSP was removed. Beginning in OneFS 7.2.1.1, if you attempt to install a DSP using the `isi pkg` command, the installation fails and a message similar to the following appears in the `/var/log/isi_pkg` log file: `Package <PACKAGE NAME> must be installed with isi_dsp_install.`	153430

If the **Disable access logging** option was set in the OneFS web administration interface, and then you upgraded your cluster from OneFS 6.5.x to OneFS 7.x, the apache2 service failed to start. As a result, client access to HTTP was denied. In addition, an error similar to the following appeared repeatedly in the `/var/log/isi_mcp` file:
 	
`FAILED on action list 'start': action 1/1`
`SERVICE apache2 (pid=3840) returned exit status 1`
	151928

In some cases, after successfully reimaging a node by using a USB flash drive with a bootable image of OneFS, before the node restarted, the following message appeared on the console: `setting read-only with system-shutdown optionisi_read_only: File operation error` This unnecessary error message could have created confusion and potentially delayed completion of the reimage operation. Beginning in OneFS 7.2.1.1, this message no longer appears on the console and the node restarts as expected.	149586

Resolved in OneFS 7.2.1.0

Events, alerts, and cluster monitoring

Events, alerts, and cluster monitoring issues resolved in OneFS 7.2.1.0	ID

The isi_stats_hist_d process sometimes stopped responding when the cluster was under heavy load. If this occurred, the following lines appeared in the `/var/log/messages` file:
 	
`/lib/libc.so.7:thr_kill+0xc`
`/lib/libc.so.7:__assert+0x35`
`/usr/sbin/isi_stats_hist_d:_ZN15stats_hist_ring4initEitb+0x506`
`/usr/sbin/isi_stats_hist_d:_ZN10ring_cache3getEiiiiii+0x228`
`/usr/sbin/`
`isi_stats_hist_d:_ZN11db_mgr_impl5queryER20stats_timeseries_setP10`
`stats_impltRK11query_timesRK14stats_hist_pol+0x33d`
`/usr/sbin/`
`isi_stats_hist_d:_ZN16database_manager5queryER20stats_timeseries_s`
`etP10stats_impltRK11query_timesRK14stats_hist_pol+0x28`
`/usr/sbin/`
`isi_stats_hist_d:_ZN20ecd_query_timeseries8query_dbEP10stats_implt`
`RK11query_timesRK14stats_hist_pol+0x3d`
`/usr/sbin/`
`isi_stats_hist_d:_ZN20ecd_query_timeseries12proc_commandEl+0x56c`
`/usr/sbin/isi_stats_hist_d:main+0xbcd`
`/usr/sbin/isi_stats_hist_d:_start+0x8c`
	148184

Hardware

Hardware issues resolved in OneFS 7.2.1.0	ID
The `isi firmware status` command did not report the firmware version of the Mellanox IB/NVRAM card. This issue affected the S200, X200, X400, and NL400 series nodes.	150298
If a node that was configured for L3 cache had an SED SSD in an `SED_ERROR` state and if that SED SSD was rejoined to the cluster by running a drive `PSID revert` operation followed by an `isi devices format` operation, the SED SSD would be formatted as a storage drive rather than an L3 cache drive. This issue occurred because the `isi devices add` command, a required step in the process, was not run. Beginning in OneFS 7.2.1, the add operation runs automatically to ensure that SED SSD drives configured for L3 cache are reformatted and readded to the cluster correctly.	133696

Networking

Networking issues resolved in OneFS 7.2.1.0	ID
If all of the following conditions existed, an issue sometimes occurred that caused EMC Isilon X410, S210, and HD400 nodes that contain Broadcom NetXtreme Ethernet (BXE) 10 GigE NIC cards to reboot unexpectedly:	153414

Networking issues resolved in OneFS 7.2.1.0	ID
• The cluster contained EMC Isilon X410, S210, or HD400 nodes that contain BXE 10 GigE NIC cards • An external network subnet connected to the cluster contained a node with BXE 10 GigE NIC card is set to 9000 MTU • The cluster was running OneFS 7.2.0.x If this issue occurred, errors similar to the following might have appeared in the `/var/log/messages` file: `ERROR: mbuf alloc fail for fp[01] rx chain (55)` For more information, see article 200096 on the EMC Online Support site.	
A memory leak in the FlexNet process, isi_flexnet_d, could have caused the process to stop running, and could have damaged the `etc/ifs/flx_config.xml` file. If the file was damaged, the file was removed from all nodes, and all clients lost their connection to the cluster.	148110
If a node included two or more network interface cards (NICs) that were assigned IPv6 addresses and if NIC affinity was enabled for those interfaces, one of the two NICs might have stopped sending and receiving network traffic.	135874, 66591, 57880

NFS

NFS issues resolved in OneFS 7.2.1.0	ID

If NFSv4 clients mounted NFS exports on the cluster through NFS aliases, it was possible to encounter a race condition that caused the NFS service to unexpectedly restart. This issue was more likely to occur when many NFSv4 clients were simultaneously mounting exports through NFS aliases. If this race condition was encountered, the NFS service on the affected node unexpectedly restarted, NFS clients connected to the node might have been disconnected, some NFS clients might have been prevented from mounting an export, and the following lines appeared in the `/var/log/messages` file:
`/lib/libc.so.7:thr_kill+0xc`
`/usr/likewise/lib/lwio-driver/nfs.so:NfsAssertionFailed+0xa4`
`/usr/likewise/lib/lwio-driver/nfs.so:Nfs4OpenOwnerAddOpen+0x112`
`/usr/likewise/lib/lwio-driver/nfs.so:NfsProtoNfs4ProcOpen+0x2567`
`/usr/likewise/lib/lwio-driver/nfs.so:NfsProtoNfs4ProcCompound +0x5fe`
`/usr/likewise/lib/lwio-driver/nfs.so:NfsProtoNfs4Dispatch+0x43a`
`/usr/likewise/lib/lwio-driver/nfs.so:NfsProtoNfs4CallDispatch+0x3e`
`/usr/likewise/lib/liblwbase.so.0:SparkMain+0xb7`
	155806, 152211

Due to a memory leak in the isi_papi_d process, if you created or modified a large number of nfs exports, over time the isi_papi_d process might have run out of memory. If this occurred, then, after running either the `isi nfs exports create` or the `isi nfs exports modify` command, an `Incomplete server response` error appeared on the console and the isi_papi_d process unexpectedly restarted.	143519
If an NFSv3 client connection failed-over to a node on which the NFS server service was in the process of starting or restarting, the failover was sometimes unsuccessful and the client was disconnected from the cluster. This issue occurred	142363

NFS issues resolved in OneFS 7.2.1.0	ID
because there was a delay in service while the portmap service was registering ports.	

OneFS web administration interface

OneFS web administration interface issues resolved in OneFS 7.2.1.0	ID

If you selected **Help** › **Help on This Page** or **Help** › **Online Help** from the **General Settings** page of the web administration interface, the following message appeared:
`Not Found`
`The requested URL /onefs/help/GUID-E395ABA6-B63A-4F40-8281-3574CCF6C8B1.html was not found on this server.`
Note
This issue did not affect the SNMP Monitoring and SupportIQ general settings pages.
	152251

If you attempted to access the OneFS web administration interface through an IPv6 IP address, after entering a valid username and password, the web interface behaved in an inconsistent manner. For example: • After entering valid credentials, a `Login unsuccessful` error appeared • After successfully logging in, buttons and controls did not work or did not behave as expected. • After successfully logging in, you were denied access to view pages and the following message was displayed: `No privileges for this functionality`	128808, 102459

Upgrade and installation

Upgrade and installation issues resolved in OneFS 7.2.1.0	ID
During a OneFS upgrade, there was a window of opportunity during which the `array.xml` file on some nodes in the cluster could have contained out-of-date version information. If a node whose `array.xml` file was out-of-date sent messages to a node whose `array.xml` file was current, the affected node exhibited unexpected behavior, such as random group changes. **Note** Although the `array.xml` file on the affected node contained out-of-date information about the version of OneFS installed on the node, the node was successfully upgraded. The unexpected node behavior was resolved when the `array.xml` file was eventually updated.	146443

CHAPTER 6

OneFS patches included in this release

The following section provides a list of patches that address issues that are now fixed in OneFS. If you previously installed one or more of the listed patches, and you upgrade to a release that includes the fix for the issue the patch addressed, you do not need to reinstall those patches after you upgrade.

After upgrading, see Current Isilon OneFS Patches on the EMC Online Support site to find out if any new patches were released that might apply to the version of OneFS you upgraded to.

- Patches included in OneFS 7.2.1.2 86
- Patches included in OneFS 7.2.1.1 86

Patches included in OneFS 7.2.1.2

Functional area	Description	Patch ID
Authentication	If a cluster was not able to contact a writeable Windows 2008 R2 or later Domain Controller (DC), attempts to failover to a read-only domain controller failed and a KRB_AP_ERR_BAD_INTEGRITY error was returned to the cluster. This issue occurred because the name-type sent in the Kerberos Ticket-Granting Ticket request was not KRB5-NT-SRV-INST, which is the required name-type for compatibility in this environment. As a result, clients connected to the cluster could have experienced intermittent authentication issues.	161462
Backup, recovery, and snapshots	After upgrading to 7.2.1.1, the NDMP backup services consumed 100 percent or more of a cluster's core CPU resources when no backup or restore jobs were running. This issue occurred because an NDMP backup socket variable was incorrectly initialized.	163969
HDFS	This patch addressed multiple issues that affected the HDFS service.	164059
Networking	A new kernel driver for the QLogic (formerly Broadcom) NetXtreme Ethernet (BXE)10 GigE NIC, 1.78.79.ISILON.01, was integrated into OneFS 7.2.1.2.	163030
SMB	This patch addressed an issue that occurred if a cluster was not able to contact a writeable Windows 2008 R2 or later domain controller, attempts to failover to a read-only domain controller were unsuccessful, and a KRB_AP_ERR_BAD_INTEGRITY error was returned to the cluster. This patch also addressed an issue that occurred when LDAP could be queried only for the default base_dn field user list, even when a user-specific or group-specific base_dn field existed on the LDAP server.	164327

Patches included in OneFS 7.2.1.1

Functional area	Description	Patch ID
File system	This patch installs a new kernel driver for the QLogic (formerly Broadcom) NetXtreme Ethernet (BXE) 10 GigE NIC. This driver reverts changes that were designed to correct three issues that affect Isilon HD400, NL410, S210, X210 and X410 nodes; installation of this patch may make the cluster susceptible to these issues.	162062
NFS	If an NFS client attempted to create an empty symbolic link—for example, by running the `ln -s "" <filepath>` command to create a symbolic link without a path to a source	159546

Functional area	Description	Patch ID

	file—the NFS server restarted unexpectedly, interrupting the workflow. If this occurred, messages similar to the following appeared in the /var/log/messages file:
	
	`pid 24071 (nfs), uid 0: exited on signal 11 (core dumped)`
	`[kern_sig.c:3376](pid 24141="nfs")(tid=101258)`
	`Stack trace:`
	`Stack:`
	`----------`
	`/lib/libc.so.7:strlen+0x2`
	`/usr/likewise/lib/lwio-driver/ nfs.so:NfsProtoNfs3Dispatch+0xddf`
	`/usr/likewise/lib/lwio-driver/ nfs.so:NfsProtoNfs3CallDispatch+0x3e`
	`/usr/likewise/lib/liblwbase.so.0:SparkMain+0xb7`
	`----------`
	
	In addition, because the MCP process attempted and failed to restart the NFS server, the following message appeared on the Cluster Events page in the OneFS web administration interface and was logged in the `/var/log/messages` file:
	
	`Process lwsm of service nfs has failed to restart after multiple attempts.`	

NFS	If all of the following conditions were met, users connected to an NFS export received `Permission denied` errors when they attempted to access file system objects to which they should have had access: • The `--map-lookup-uid` option is enabled (set to yes) for the affected NFS export. • The group owner of the affected file system object is one of the user's supplemental groups rather than the user's primary group. • The cluster-side lookup for the user's supplemental groups failed. This issue occurred because, when the lookup for the user's UID failed, OneFS did not correctly apply supplemental group permissions to the user. As a result, the user was denied access to the file system object.	157915
SMB	This patch addressed two issues that affected SMB clients if symbolic link translation was disabled on the cluster: in the first, symbolic links to directories might have failed; in the second, the symbolic links to directories could not be deleted.	163539
SMB	This patch addressed two issues that affected SMB clients: in the first, SMB clients could not display information about alternate data streams for files on the cluster that contained them; in the second, after files were copied from the cluster, the timestamp metadata applied to the files might been invalid.	161478

Functional area	Description	Patch ID
SMB	SMB clients were unable to display alternate data stream information for files on the cluster that contained alternate data streams.	160220

CHAPTER 7

Known issues

This section contains the following topics:

- Antivirus 90
- Authentication 90
- Backup, recovery, and snapshots 92
- Cluster configuration 95
- Command-line interface 96
- Diagnostic tools 96
- Events, alerts, and cluster monitoring 97
- File system 98
- File transfer 101
- Hardware 101
- HDFS 102
- iSCSI 103
- Job engine 103
- Migration 105
- Networking 106
- NFS 107
- OneFS API 108
- OneFS web administration interface 109
- Security 109
- Security profiles 109
- SmartLock 111
- SmartQuotas 111
- SMB 111
- Upgrade and installation 112
- Virtual plug-ins 114

Antivirus

Antivirus known issues	ID
If you enable **Scan file when they are opened** in the OneFS web administration interface under **Antivirus › Settings › On Access Scans**, or through the command-line-interface, while files are being scanned, a large backlog of scan requests might occur. This backlog of scan requests might prevent SMB clients from accessing the cluster for long periods of time. Workarounds: • Run the `scan_on_close` command with the `scan_on_open` command to reduce the scan requests and provide improved virus scanning coverage. • Decrease the `scan_on_open` timeout value: 1. Run the following command to confirm the current timeout in seconds: `sysctl efs.bam.av.scan_on_open_timeout` 2. Run the following command where ‹timeout› is the desired new timeout value: `sysctl efs.bam.av.scan_on_open_timeout=`*‹timout›* Note that the `scan_on_open` scan requests fail if they do not complete within the timeout period, so many requests might fail. • Add ICAP servers to better meet high demand for scan requests. • Do not use SMB opportunistic locks with `scan_on_open` on-demand scanning, as this spreads the issue across the cluster and can add further delays.	160302
In the OneFS web administration interface, on the Antivirus **Policies** page, if you double-click the **Start** link for a policy, multiple instances of the AVScan job start.	54477

Authentication

Authentication known issues	ID

If you modify user or group privileges by using the SAMR (Security Account Manager Remote Procedure Call) protocol, the lsass process might unexpectedly restart. If the process restarts for this reason, the following stack will appear in `/var/log/messages`:
`/boot/kernel.amd64/kernel: [kern_sig.c:3361](pid 50230="lsass") (tid=100874) Stack trace:`
`Stack: ----------`
`/boot/kernel.amd64/kernel: /usr/likewise/lib/liblwadvapi_nothr.so.0:LwHashCaselessStringHash+0xe`
`/boot/kernel.amd64/kernel: /usr/likewise/lib/liblwadvapi_nothr.so.0:LwHashGetValue+0x22`
`/boot/kernel.amd64/kernel: /usr/likewise/lib/liblsaserverapi.so.0:MemCacheFindObjectBySid+0x54`
`/boot/kernel.amd64/kernel: /usr/likewise/lib/liblsaserverapi.so.0:MemCacheFindBySid+0x2e`
`/boot/kernel.amd64/kernel: /usr/likewise/lib/liblsaserverapi.so.0:LsaSrvCacheFindObjects+0x2f7`
`/boot/kernel.amd64/kernel: /usr/likewise/lib/liblsaserverapi.so.0:LsaSrvFindObjects+0x1da`
	165414

Authentication known issues	ID
/boot/kernel.amd64/kernel: /usr/likewise/lib/liblsass_id_mapper_default_parser.so.0+0x27110c:0x80869f10c /boot/kernel.amd64/kernel: /usr/likewise/lib/liblsass_id_mapper_default_parser.so.0+0x286e1f:0x8086b4e1f /boot/kernel.amd64/kernel: /usr/likewise/lib/libdcerpc.so.1:rpc__cn_call_executor+0x2ac /boot/kernel.amd64/kernel: /usr/likewise/lib/libdcerpc.so.1:cthread_call_executor+0x244 /boot/kernel.amd64/kernel: /usr/likewise/lib/libdcerpc.so.1:proxy_start+0x3e /boot/kernel.amd64/kernel: /lib/libthr.so.3:_pthread_getprio+0x15d --- While the lsass process is restarting, authentication and access lookup might fail.	
OneFS did not cache the group membership information that was returned when a user was authenticated. As a result, OneFS repeatedly queried the external authentication providers for this information each time the user was authenticated. In environments with an LDAP provider, this issue might have overloaded the provider and caused it to become unavailable. If the provider became unavailable, users could not be authenticated to the cluster until the issue was resolved.	165150
If HTTP clients make a large number of file permission checks through Platform API (PAPI) commands, or if HTTP clients make heavy use of the `isi auth` command, the isi_papi_d process might unexpectedly restart. If this occurs, user authentication might be unavailable, HTTP error code 400 or 500 might be returned when clients try to access the cluster, and messages similar to the following might appear in the `isi_papi_d.log` file: "GET <file path> HTTP/1.1" 500 2548 "-" "Apache8.0_45)"	164877
If you configure public key SSH authentication on a cluster running OneFS 7.1.1.2 through OneFS 7.1.1.5 or OneFS 7.2.0.1 through OneFS 7.2.0.2, and then you upgrade to a later version of OneFS, the root user can no longer log in to the cluster through SSH without entering their password. For more information and a workaround, see article 203215 on the EMC Online Support site. **Note** This issue does not affect upgrades from 7.1.1.6 and later, 7.2.0.3 and later, 7.2.1.0 and later, or 8.0.0.0 and later.	145598
Incorrect sequence numbers during SMB2 traffic could cause the lwio process to fail, resulting in a temporary loss of SMB service.	134247
The lsass process might fail while running the NtlmGetDomainNameFromResponse function due to an incorrectly formed request during NTLM authentication, resulting in a temporary loss of authentication service.	134239
The lsass process might fail while running the NtlmValidateResponse function due to an incorrectly formed request during NTLM authentication, resulting in a temporary loss of authentication service.	134238
The lsass process might fail while running the LwIoSessionAcquireCall function due to an incorrectly formed request during NTLM authentication, resulting in a temporary loss of authentication service.	134237

Authentication known issues	ID
The lsass process might fail while running the AuthenticateNTLMv2 function due to an incorrectly formed request during NTLM authentication, resulting in a temporary loss of authentication service.	134237
The lsass process might fail while running the NtlmGetCStringFromUnicodeBuffer function due to an incorrectly formed request during NTLM authentication, resulting in a temporary loss of authentication service.	134236

If the lsass process is sequentially restarted across nodes in the cluster, the lwio process might unexpectedly restart, causing all SMB and NFS clients to be disconnected from the cluster. If this occurs, the following lines appear in the `/var/log/messages` file:
 	
`/boot/kernel.amd64/kernel: [kern_sig.c:3364](pid 9719="lwio") (tid=102399) Stack trace:`
`/boot/kernel.amd64/kernel: Stack:`
`----------`
`/boot/kernel.amd64/kernel: /lib/libc.so.7:recvfrom+0xc`
`/boot/kernel.amd64/kernel: 0x807f0f2d5 (lookup_symbol: symtab/ strtab not found:2)`
`/boot/kernel.amd64/kernel: 0x807f0f417 (lookup_symbol: symtab/ strtab not found:2)`
`/boot/kernel.amd64/kernel: 0x807f0a819 (lookup_symbol: symtab/ strtab not found:2)`
`/boot/kernel.amd64/kernel: 0x807f0ab44 (lookup_symbol: symtab/ strtab not found:2)`
`/boot/kernel.amd64/kernel: 0x80bb2e474 (lookup_symbol: symtab/ strtab not found:2)`
`/boot/kernel.amd64/kernel: /lib/libthr.so.3:_pthread_getprio+0x15d`
`/boot/kernel.amd64/kernel:`
`----------`
`/boot/kernel.amd64/kernel: pid 9719 (lwio), uid 0: exited on signal 6 (core dumped)`
	131835

If you run the `isi auth local user modify` command with the `--password-never-expires` parameter on one of the default services accounts, you receive an `Invalid Parameter` error. For example, running the following command attempts to set the password to never expire for the insightiq account: `isi auth local users modify --name=insightiq --password-never-expires` For more information, see article 89092 on the EMC Online Support site.	83444

Backup, recovery, and snapshots

Backup, recovery, and snapshots known issues	ID
If you attempt to run a DomainMark job on a SyncIQ source directory that contains the target directory of another SyncIQ policy, the DomainMark job fails. This issue occurs if any of the following conditions exist: • The source directory is contained within the target directory of another replication policy. • The source directory and the target directory of another replication policy are the same. • The source directory contains the target directory of another replication policy.	162515

Backup, recovery, and snapshots known issues	ID
Note Because a DomainMark job must be run on a SyncIQ source directory in order to failback to that directory, configuring a source directory that contains a target directory prevents you from failing back to the source directory.	

If you configure a snapshot schedule expiration time by running the `isi snapshot schedules` command or through the OneFS web administration interface, the expiration time is not always configured correctly and might not display correctly in either interface. For example, if you configure a snapshot to expire in 88 weeks, the web interface might display a snapshot expiration time of 14953 Hours. In addition, the snapshot schedule expiration time that is displayed in the web interface or at the command-line interface might not accurately reflect the actual configured expiration time. This issue occurs because the expiration time is interpreted differently by the web interface, the command-line interface, and the isi_papi_d process.
Workaround: Use the following commands to configure snapshot schedule expiration times from the command-line interface. Values set by running these commands are not interpreted by the isi_papi_d process and therefore are not affected by this issue:
`isi_classic snapshot schedule create`
`isi_classic snapshot schedule modify`
For more information about the preceding commands, run the following commands:
`isi_classic snapshot schedule create –h`
`isi_classic snapshot schedule modify –h`
	139186

When replication jobs can not be run--for example, when a node is being shut down--the isi_migr_sched process might unexpectedly restart.	135744
Restoring a large amount of data (more than 50 TB) might fail due to a memory allocation error. If this occurs, the following message might appear on the console: `<3.3> scdepot-4(id4) isi_ndmp_d[94807]: isi_ndmp_d: *** FAILED ASSERTION rb->datap @ /build/mnt/src/isilon/lib/isi_ndmpbrm2/fast_restore.c:1088: Failed to allocate memory - Cannot allocate memory`	133591
If the cluster is under heavy load, parallel NDMP restore jobs might fail.	130693
If the `Deny permission to modify files with DOS read-only attribute over both UNIX (NFS) and Windows File Sharing (SMB)` ACL policy option is enabled on the cluster, SyncIQ jobs fail when SyncIQ attempts to synchronize a file or folder to which the DOS read-only attribute is applied. If a SyncIQ job fails for this reason, an `Operation not permitted` error appears in the `isi_migrate.log` file. Workaround: Disable the `Deny permission to modify files with DOS read-only attribute over both UNIX (NFS) and Windows File Sharing (SMB)` ACL policy option before running SyncIQ jobs.	127652
If a Backup Accelerator is running multiple NDMP sessions, the NDMP process might run out of memory and the NDMP process might stop running. If this occurs, false sessions might appear on the NDMP session list.	125897

Backup, recovery, and snapshots known issues	ID
Workaround: Perform the following steps to remove stale sessions while retaining current sessions: 1. Open an SSH connection on any node in the cluster and log in using the root account. 2. Run the following command: `rm -rf /ifs/.ifsvar/modules/ndmp/sessions/*`	
A message similar to the following might appear in the `/var/log/messages` file: `isi_migrate[98488]: coord[cert2-long123-d0b]: Problem reading from socket of (null): Connection reset by peer` Workaround: Ignore the error message. This is a transient error from which OneFS automatically recovers.	124767
Because OneFS must build sparse maps for sparse files, and because OneFS cannot back up data while building a map, backup jobs that include large sparse files can take a very long time to complete. In addition, OneFS might run out of memory while backing up a large sparse file with a large number of sparse regions.	124211
The following combination of NDMP environment variables is not supported when backing up data from a OneFS cluster through EMC Networker: • Direct access restores (DAR) is enabled by setting the `DIRECT`NDMP environment variable to `Y` • History format is set to dir/node by setting the `HISTORY` NDMP environment variable to `D` • A backup file list is specified by configuring the `BACKUP_FILE_LIST` NDMP variable If you attempt to run a backup job after setting the preceding variables as described, messages similar to the following appear on the Networker console where *‹file list path›* is the path to a file defined in the file list assigned to the `BACKUP_FILE_LIST` environment variable: `Cannot find file history for <file list path>, therefore this file will not be available for recovery by file browse selection. You may still be able to recover this file with a saveset recovery.`	113999
The SyncIQ scheduler service applies UTF-8 encoding even if the cluster is set with a different encoding. As a result, DomainMark and SnapRevert jobs, which apply cluster encoding, might fail to run.	99383
If you revert a snapshot that contains a SmartLock directory, the operation might fail and the directory might be only partially reverted.	99211
If SyncIQ and SmartQuotas domains overlap, a SyncIQ job might fail with one of the following errors: • `Operation not permitted` • `unable to delete` • `failed to move`	97492

Backup, recovery, and snapshots known issues	ID
• `unable to rename` For more information, see article 88602 on the EMC Online Support site.	
If you are using the Comvault Simpana data management application (DMA), you cannot browse the backup if the data set has file names with non-ASCII characters. As a result, you cannot select single files to restore. Full restoration of the dataset is unaffected. For more information, see article 88714 on the EMC Online Support site.	96545
If you use SyncIQ to synchronize data and some data is freed on the source cluster because a file on the source decreased in size, the data is not freed on the target cluster when the file is synchronized. As a result, the space consumed on the target cluster might be greater than the space consumed on the source.	94614
SyncIQ allows a maximum of five jobs to run at a time. If a SnapRevert job starts while five SyncIQ jobs are running, the SnapRevert job might appear to hang instead of pausing until the SyncIQ job queue can accept the new job.	93061
After performing a successful NDMP backup that contains a large number of files (in the tens of millions), when you restore that backup using Symantec NetBackup, the operation fails and you receive the following error: `error db_FLISTreceive: database system error 220` For more information, see article 88740 on the EMC Online Support site.	87092

Cluster configuration

Cluster configuration known issues	ID
In the OneFS web administration interface, if file names on the cluster contain characters that are not supported by the cluster character encoding, those file names might not display correctly in the File System Explorer in the OneFS web administration interface. Workaround: Rename the files using characters supported by the character encoding that is applied to the cluster.	18901
The cluster configuration wizard requires a default gateway to set up a cluster. You might not have a default gateway if your network uses a local DNS server. Workaround: Enter 0.0.0.0 for your default gateway.	24621
After a drive is replaced, if a mismatch occurs between the hardware listed in the `drive_purpose.db` file and the actual hardware installed on the cluster, some `isi` commands might fail and some processes might stop responding. For more information, see the article 208700 on the EMC Online Support website.	97545
Although the ISI_PRIV_AUDIT privilege does not allow a user to delete file pool policies, a user who is assigned the IS_PRIV_AUDIT privilege can view the controls to delete file pool policies on the **File System** › **Storage Pools** › **File Pool Policies** page.	134378

Command-line interface

Command-line interface known issues	ID
Because the Ranges column in the output of the `isi networks list pools` command is 22 characters wide, if the IP range is more than 22 characters in length, the entire range is not displayed. Workaround: Run the `isi networks list pools` command with the `--verbose` option to display the entire IP range.	140793

If you run an `isi` command with the `--help` option to get more information about the command, the text that is displayed might provide information about the related `isi_classic` command instead of providing information about the command that you typed. For example, if you run the `isi storagepools` command with the `--help` option, the following information appears:
`'isi_classic smartpools health' options are:`
`--verbose, -v Print settings to be applied.`
`--help, -h Print usage help and exit`
	129637

The `isi version osreldate` command returns a random number rather than the expected OneFS release date.	98457, 98452

Diagnostic tools

Diagnostic tools known issues	ID
Because ESRS cannot receive data from a client with an IPv6 address, you cannot configure an ESRS gateway with an IPv6 address. In addition, in environments that include both IPv4 and IPv6 subnets, the ESRS gateway cannot automatically detect the supported IPv4 subnet. As a result, in both cases, the ESRS gateway must be assigned to an IPv4 subnet and at least one node must be assigned to that subnet. **Note** Nodes with only IPv6 addresses cannot be remotely monitored through ESRS.	133882
If you have never uploaded cluster information to EMC Isilon Technical Support, on the **Cluster Management** › **Diagnostics Info** page, the **Gather Status** bar appears gray or black.	133972
In the OneFS web administration interface, on the **Gather Info** page, although the Gather Info process is still running, the **Gather Status** progress bar indicates that the process is complete.	103906

Events, alerts, and cluster monitoring

Events, alerts, and cluster monitoring known issues	ID
If an NFS request specifies an inode, rather than a file name, and more than one hard link to the specified inode exists, OneFS auditing is unable to determine which hard link was intended by the NFS client. If this happens, OneFS auditing might select the incorrect hard link, which can cause client permissions to be misrepresented in audit logs.	136038
The isi_papi_d process might fail when InsightIQ begins monitoring a cluster that contains 80 or more nodes.	135767
The SNMP process (snmpd) might unexpectedly restart after a drive is smartfailed and then replaced.	129711
When alert traffic is high, running the `isi events quiet all` command might time out. As a result, some events might not be quieted and the following error might be displayed: `Error marking events: Error while getting response from isi_celog_coalescer (timed out)` Workaround: Run the `isi events quiet all` command on the master node.	113689, 112774
When the email address list for an event notification rule is modified from the command line, the existing list of email addresses is overwritten by the new email addresses. For more information, see article 88736 on the EMC Online Support site.	89086
Although SNMP requests can reference multiple object IDs, the OneFS subtree responds only to the first object ID.	81183
If a large number of LUNs are active, the event processor might issue a warning about open file descriptors held by the iSCSI daemon. You can safely ignore this warning.	79341
On the **Cluster Overview** page of the OneFS web administration interface, clicking the ID of a node that requires attention, as indicated by a yellow status icon, does not provide details about the status. Workaround: In the list of events, sort the nodes by **Scope** or **Severity** and then click **View details**. Alternatively, run the `isi events list --nodes <id>` command to view the events. For more information, see article 16497 on the EMC Online Support site.	77470
Under some conditions, events might be updated with incorrect information. For example, the event type might not match the event message. If this issue occurs, lines similar to the following are logged in the `/var/log/messages` file: `Cannot parse damaged job history record 2.5285; skipping`	75980
If you run the `isi status` command, the value displayed for the sum of per-node throughput might differ from the value displayed for the sum of cluster throughput. This occurs because some data is briefly cached. The issue is temporary. Workaround: Re-run the `isi status` command. For more information, see article 88690 on the EMC Online Support site.	73554
Reconfiguring aggregate interfaces can leave active events for inactive interfaces.	72200

Events, alerts, and cluster monitoring known issues	ID
Workaround: Cancel the events manually.	
Event system databases that store historical events might fail to upgrade properly. If a database fails to upgrade, the database is replaced by an empty database with a new format and historical events are lost.	71840
A network interface that is configured as a standby without an IP address triggers an interface down event. Workaround: Quiet the event manually.	71399
Monitoring with SNMP, InsightIQ, or the `isi statistics` command can fail when a cluster is heavily loaded.	68559
The final phase of the FSAnalyze job runs on one node and can consume excessive resources on that node.	64854
While a cluster processes a heavy I/O load, graphs in the OneFS web administration interface might display the following message: `Warning: Unreliable Data` Workaround: Run the `isi statistics` command.	62736
When using Simple Network Management Protocol (SNMP) to report on aggregated interfaces—for example, LACP, LAGG, and fec—the interface speed is displayed as 100 MB instead of 2 GB. For more information, see article 89363 on the EMC Online Support site.	55247
You might receive an alert that a temporary license is expired even though a permanent license is installed. Workaround: Use the command-line interface or the web administration interface to quiet the alert.	24504

File system

File system known issues	ID
When OneFS sets the POSIX mode bits on an Alternate Data Stream (ADS) after replicating a file to another cluster over SyncIQ, the file's ADS might be inaccessible if all of the following conditions exist: • The replicated file has only POSIX permissions applied to it. • The replicated file does not have a security descriptor. • OneFS did not set the execute mode bit on the ADS. If this issue occurs, users might receive a `Permission denied` or `Access denied` error when they try to access the ADS, and applications that rely on the ADS might be adversely affected. Workaround: Reapply the POSIX permissions for the affected ADS files or add an ACL to the affected files.	164905

File system known issues	ID
A node might fail to shut down or reboot if the shutdown process is unable to stop the lwsm process in less than two minutes. If this issue occurs the following error appears in the `/var/log/messages` file: `rc.shutdown: 120 second watchdog timeout expired. Shutdown terminated.` If you encounter this issue, wait five minutes and then try to reboot the node by running the reboot command. If the node fails to reboot, contact EMC Isilon Technical Support for assistance.	152338

If a dedupe job is running on a file that is also in the process of being deleted, the workers for the job can be delayed long enough to generate a hangdump file. The dedupe job continues afterwards. If this issue is encountered, messages similar to the following appear in the `/var/log/messages` file:
 	
`isi_hangdump: Lock timeout: 720.008538 from efs.lin.lock.initiator.oldest_waiter`
`isi_hangdump: LOCK TIMEOUT AT 1421800091 UTC`
`isi_hangdump: Hangdump timeout after 0 seconds: Received HUP`
`isi_hangdump: Lock timeout: 725.018597 from efs.lin.lock.initiator.oldest_waiter`
`isi_hangdump: Lock timeout: 730.028656 from efs.lin.lock.initiator.oldest_waiter`
`isi_hangdump: Lock timeout: 735.038715 from efs.lin.lock.initiator.oldest_waiter`
`isi_hangdump: Lock timeout: 740.048774 from efs.lin.lock.initiator.oldest_waiter`
`isi_hangdump: END OF DUMP AT 1421800091 UTC`
	141028

The lwio process might fail while the cluster is under heavy load, causing clients to become disconnected. If this occurs, the following lines appear in the `/var/log/messages` file:
 	
`/lib/libc.so.7:thr_kill+0xc`
`/usr/likewise/lib/liblwiocommon.so.0:LwIoAssertionFailed+0xb6`
`/usr/likewise/lib/libiomgr.so.0:IopFmIrpStateDispatchFsdCleanupDone+0x26`
`/usr/likewise/lib/libiomgr.so.0:IoFmIrpDispatchContinue+0x36c`
`/usr/lib/libisi_cpool_rdr.so:_Z16cprdr_pre_createP21_IO_FLT_CALLBACK_DATAP23_IO_FLT_RELATED_OBJECTSPPvPPFvS0_S3_ES4_+0x646`
`/usr/lib/libisi_cpool_rdr.so:_Z19process_pre_op_itemP13_LW_WORK_ITEMPv+0x54`
`/usr/likewise/lib/liblwbase.so.0:WorkThread+0x256`
`/usr/likewise/lib/liblwbase.so.0:LwRtlThreadRoutine+0xee`
`/lib/libthr.so.3:_pthread_getprio+0x15d`
	140203

If a node in a three node cluster unexpectedly shuts down and is not re-added to the cluster, and if you then add a new node to the cluster, one of the remaining nodes might unexpectedly reboot and you might have to wait for a significant amount of time before you can successfully add the new node to the cluster. Workaround: Add the new node while no writes are being made to the cluster.	124603
LDAP user and group ownership cannot be configured through the OneFS web administration interface. Workaround: Use the command-line interface to configure LDAP user and group ownership.	103983

File system known issues	ID

An Alternate Data Stream (ADS) block-accounting error might cause the Inode Format Manager (IFM) module to fail, causing the following message to be logged to the stack trace:
 	
`kernel:isi_assert+0xde`
`kernel:isi_assert_mayhalt+0x70`
`efs.ko:ifm_compute_new_ads_summary+0x9a`
`efs.ko:ifm_update_ads_summary+0x15b`
`efs.ko:ifm_end_operation+0x11ad`
`efs.ko:txn_i_end_all_inode_ops+0x11d`
`efs.ko:txn_i_end_operations+0x5e`
`efs.ko:txn_i_end+0x3d`
`efs.ko:bam_remove+0x198`
`efs.ko:ifs_vnop_wrapremove+0x1bf`
`kernel:VOP_REMOVE_APV+0x33`
`kernel:kern_unlinkat+0x2a6`
`kernel:isi_syscall+0x49`
`kernel:syscall+0x26e`
 	
Workaround: Ignore this error message. This is a transient error that OneFS will recover from automatically.
	100118

If you view the contents of a directory that contains more than 999 files and directories through the OneFS File System Explorer, the attributes of the files and directories—for example the size and date modified—are not displayed. In addition, the **Protected** column indicates that the files and directories are not protected regardless of the actual state of the files and directories being displayed. Workaround: To ensure that the correct information about your files and folders is displayed in the File System Explorer, organize files and directories into subdirectories so that each subdirectory contains fewer than 1,000 files and directories.	99170
If you run the `isi statistics` command to view available storage for a cluster, the storage capacity of nodes that have not been provisioned into storage pools is listed in the per node storage statistics, but is not reflected in the total storage available on the cluster. The capacity of nodes that have not been provisioned into storage pools is omitted from the cluster total because you cannot write data to unprovisioned nodes.	86328
Some configuration changes cannot be made if the cluster is 99 percent full. As a result, the cluster might stop responding until sufficient free space is made available. See Best Practices Guide for Maintaining Enough Free Space on Isilon Clusters and Pools on the EMC Online Support site.	74272
When FlexProtect is run with `verify upgrade check` enabled and one or more drives are down, OneFS occasionally reports false data corruption. If this issue occurs, contact EMC Isilon Technical Support.	73276
If you run an improperly formatted shutdown command, a node might be placed into read-only mode and could fail to shutdown. In some cases the node is inaccessible through the network but is still accessible through a serial connection.For more information, see article 89544 on the EMC Online Support site.	54120

File transfer

File transfer known issues	ID
The FTP output of the `isi statistics` command might be inaccurate.	129599
In the OneFS web administration interface, on the **Diagnostics › Settings** page, if you enter an invalid address in the **HTTP host** or **FTP host** field, `Connection Verification Succeeded` is displayed when you click **Verify**.	70448

Hardware

Hardware known issues	ID
The following warning message might appear on the serial console when a node starts up: `WARNING: softAVL is missing Isilon P/N(s)for vendor ID="QLogicNetXtreme II BCM57810 1", firmware ID="7.8.51.0", hwType="io_card",hwName="Broadcom 10GigE` This error is caused by a change of the device ID string in the updated BXE driver that is included in this release of OneFS. For more information about this warning, including instructions about how to permanently remove the message, see article 466850 on the EMC Online Support site.	154538
If you reboot or shut down a node with a Broadcom 10 GbE network interface card that is configured for legacy fec aggregation, the node might stop responding until it is manually powered off.	136915
If the power supply fan in an HD400 node fails, the power supply indicator light turns yellow, but no alert is sent. If this condition is not addressed, the power supply will eventually fail and an alert will be sent for the power supply failure. Contact EMC Isilon Technical Support if you encounter this issue.	135814
Due to an issue that can prevent the `shutdown -p` command from correctly shutting down a node, the `isi firmware update` command might incorrectly report that a firmware update has failed.	133606
If you run the `isi firmware update` command to update the firmware on a different node—for example, if you run the command on node 1 to update the firmware on node 2—OneFS might incorrectly report that the firmware update failed. This issue occurs because, the first node attempts to confirm the success of the update before the second node has updated it's configuration with the new revision number.	133317
Node firmware updates will fail if HPM downloads return error code `D5` during the upgrade process. Workaround: Retry updating the node firmware. If this issue persists, contact EMC Isilon Technical Support.	132523
Chassis Management Controller (CMC) firmware update procedures might fail.	123303

Hardware known issues	ID
Workaround: Run the following command and then retry the update: `/usr/bin/isi_hwtools/isi_ipmicmc`	
An internal sensor that monitors components might not properly detect the source of a hardware component failure, such as the I2C bus. If this occurs, the wrong alert or no alert might be generated.	73050
If a node without a recognized system configuration number is added to a cluster, the node will be split from the rest of the cluster and unable to accept client connections. In addition, the node will be marked as down on the **Cluster Status** page in the OneFS web administration interface and if you run the `isi_hw_status -i` command the Class and Series are listed as unknown. Workaround: Contact EMC Isilon Technical support for assistance with this issue.	71354
A newly created cluster might not be visible to unconfigured nodes for up to three minutes. As a result, nodes will fail to join the cluster during that time period.	69503
If the `/etc/isilon_system_config` file or any `/etc VPD` file is blank, an `isi_dongle_sync -p` operation will not update the VPD EEPROM data.	67932
There are multiple issues with shutting down a node improperly that can potentially lead to data loss. Workaround: Follow the instructions for shutting down nodes exactly. For more information, see article 16529 on the EMC Online Support site.	35144

HDFS

HDFS known issues	ID
If the Hadoop datanode services are left running on Hadoop clients that are connected to a cluster, the isi_hdfs_d process will continuously log the following messages to the `/var/log/messages` file and the `/var/log/isi_hdfs_d.log` file as the process receives requests: `org.apache.hadoop.hdfs.server.protocol.DatanodeProtocol from verify_ipc_protocol (/build/mnt/src/isilon/bin/isi_hdfs_d/protoutil.c:18) from parse_connection_context (/build/mnt/src/isilon/bin/isi_hdfs_d/protoutil.c:100) from ver2_2_parse_connection_context (/build/mnt/src/isilon/bin/isi_hdfs_d/protocol_v2_2.c:388) from process_out_of_band_rpc (/build/mnt/src/isilon/bin/isi_hdfs_d/protocol_v2_2.c:1000)`	135993

While the cluster is under heavy HDFS load, the isi_hdfs_d process might restart. If this occurs, the following lines appear in the `/var/log/messages` file:
`/lib/libc.so.7:__sys_kill+0xc`
`/usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/usr/lib/libisi_hdfs.so.1:hdfs_enc_mkdirat_p+0x2b1`
`/usr/lib/libisi_hdfs.so.1:hdfs_mkdir_p+0x41`
`/usr/bin/isi_hdfs_d:config_init_directory+0x13`
	123802

iSCSI

iSCSI known issues	ID
The iSCSI protocol can log a `data digest error` message in the iSCSI log. No workaround is required; the protocol will recover and reconnect.	83537
VSS 32-bit installation succeeds on a Windows initiator, but the provider does not appear in the list of installed providers. This issue affects Windows Server 2003 only. Workaround: Install the Microsoft iSCSI Software Initiator. For more information, see article 88763 on the EMC Online Support site.	74303
In the OneFS web administration interface, the **iSCSI Summary** page sometimes loads slowly. If this occurs, the page might time out and the isi_webui_d process might be consuming a high percentage of CPU resources on one or more nodes.	73038
If you create a new target after you move iSCSI shadow clone LUNs, the OneFS web administration interface might become unresponsive.	71919

Job engine

Job engine known issues	ID
In rare instances, if a drive fails while IntegrityScan is running, the IntegrityScan job can fail. In addition, if you run the `isi job events list --job-type integrityscan` command, a message similar to the following appears on the console, where *<x>* is the job ID: `2015-02-12T15:35:31 <x> IntegrityScan 1 State change Failed` The job should automatically restart and then run to completion.	139708
In rare instances, if a drive fails while MediaScan is running, the MediaScan job can fail. In addition, if you run the `isi job events list --job-type mediascan` command, a message similar to the following appears on the console, where *<x>* is the job ID: `2015-02-12T15:35:31 <x> MediaScan 1 State change Failed` The job should automatically restart and then run to completion.	139704

The isi_job_d process might fail while a QuotaScan job is running. If this happens, the QuotaScan job will continually pause and resume, and the following lines might appear in the `/var/log/messages` file:
`/lib/libc.so.7:thr_kill+0xc`
`/usr/lib/libisi_util.so.1:isi_assert_halt+0xa0`
`/usr/lib/libisi_job.so.1:tw_opendir+0x207`
`/usr/lib/libisi_job.so.1:tw_tree_init+0x327`
`/usr/bin/isi_job_d:treewalk_task_next_item+0x150`
`/usr/bin/isi_job_d:quotascan_task_next_item+0x4c`
	134301

Job engine known issues	ID

`/usr/bin/isi_job_d:worker_process_task+0x307`
`/usr/bin/isi_job_d:worker_main+0x11cd : /lib/libthr.so.`
`3:_pthread_getprio+0x15d`
 	
Workaround: Cancel the QuotaScan job.	

If you queue multiple jobs while smartfailing drives, AutoBalance jobs might fail.	133771
The MediaScan job reports errors for drives that have been removed from the cluster. Workaround: Don't fail a drive after a MediaScan job has started, or cancel the job.	132083
If the MultiScan, Collect, or Autobalance jobs are disabled before a rolling upgrade, the jobs will automatically become enabled after the rolling upgrade completes. Workaround: If MultiScan, Collect, or Autobalance jobs are disabled before a rolling upgrade, and you want those jobs to be disabled after the rolling upgrade completes, manually disable those jobs after the rolling upgrade completes.	124744

Instead of displaying active jobs before inactive jobs, the `isi job status` command displays jobs in numerical order by running ID.
Run one of the following commands depending on the information you need:
• To see a list of running jobs:
`isi job jobs list --state running`
• To see the full list of jobs sorted by state - for example, running, paused, and waiting:
`isi job jobs list --sort state`
For more information about commands, see isi job jobs list and isi job status in the OneFS 7.2.1 CLI Administration Guide.
	114802, 114583

The `isi job reports view job` command sometimes returns reports twice.	112265

The Dedupe and DedupeAssessment jobs can only run with a job-impact level of low. If you attempt to change the Default Impact Policy to ant other level, the following message appears where ‹job type› is either Dedupe or DedupeAssessment:
`Job Types Edit Failed`
`The edit to the job type did not save due to the following error:`
`Cannot change impact policy of job type<job type>`
	110129

When you run a DomainMark job after taking a snapshot, and then run a SnapRevert job with a job impact policy set higher than LOW, the impact policy has no effect. For more information, see article 88597 on the EMC Online Support site.	93603
Job engine operations occasionally fail on heavily loaded or busy clusters. If the command fails, a message similar to the following is displayed on the console: `Unable to pause integrity scan: pause command failed: Resource temporarily unavailable.` Workaround: If an operation fails, wait a moment and then retry the operation.	72109

Migration

Migration known issues	ID
If the `/var` partition on the cluster is full or nearly full and you run the `isi_vol_copy` or `isi_vol_copy_vnx` command to migrate data to the cluster, the migration process might stop or might generate errors. Workaround: Decreasing the logging level for the isi_vol_copy or isi_vol_copy_vnx process can reduce the amount of space that is required to complete the migration process. Before you migrate data to the cluster, ensure the `/var` partition and the cluster have sufficient available space to accommodate the migration process and the migrated data. The migration process generates new log entries in the log files on the `/var` partition. For more information about maintaining sufficient free space on an Isilon cluster, see Best Practices Guide for Maintaining Enough Free Space on Isilon Clusters and Pools on the EMC Online Support site.	136679

If you start a restartable migration of data from a source array to an Isilon cluster, and you specify the source array by using the source array's IP address or hostname, if the migration process stops, you must use the same IP address or hostname to identify the source array when you restart the migration process. For example, if you run the following command to start a restartable migration process, where *‹hostname›* is the hostname for the source array, you must use *‹hostname›* to restart the migration process:
`isi_vol_copy -sa ndmp:GLfSRQmI <hostname>/vol/vol1/test_data /ifs/test1`
If you identify the source array by using an IP address when you restart a migration process that was originally started by using a hostname, or if you use a hostname to restart a migration process that was originally started by using an IP address, a message similar to the following appears on the console and the migration process does not restart:
`Found a migration 1_85559_1414612755 with conflicting destination path`
	136395

If you migrate ACLs to the cluster through the `isi_vol_copy_vnx` command and then attempt to read those ACLs over NFSv4, the read will fail with the following error message: `An NFS server error occurred`	131299
When migrating FIFO files using the `isi_vol_copy` utility, the following message displays: `Save checkpoint error: Could not match file history.`	100366
If the `isi_vol_copy` command is run twice, with different source paths but with the same target path, the second instance fails without migrating any files.	100365

Networking

Networking known issues	ID
In some cases, sbflush_internal is called when a network socket is already closed. If this issue occurs, the affected node might restart unexpectedly. If a node restarts as a result of this issue, the following message appears: `Software Watchdog failed (userspace is starved!)`	152330
After group changes occur—for example, after a node is shut down—SmartConnect is unable to resolve DNS requests from clients for up to twelve seconds. As a result, clients are unable to make new connections to the cluster for a short period of time.	134858
In clusters with a large number of nodes, after an InfiniBand switch is rebooted, the cluster might experience a high level of group change activity for approximately two hours. Because, by default, a single Device Work Thread (DWT) is handling all node transitions to the new InfiniBand connections, some requests are not handled in a timely manner. As a result, nodes might not successfully failover to the new InfiniBand connection, and, in some cases, might fail to rejoin the cluster. Workaround: To increase the number of DWT threads handling requests to failover to a new InfiniBand connection, run the following command: `sysctl efs.rbm.dwt_threads=4` For more information about viewing and setting sysctl options, see article 89232 on the EMC Online Support site. **Note** Increasing the number of DWT threads might affect CPU performance depending on the number of processors in the node.	134665
lxgbe interfaces might report a status of inactive, even if the cable and the port that the cable is plugged into are functioning correctly.	127706
If a port on an A100 node has IP addresses assigned to it, the port will reinitialize when the node is booted up.	126464
When you use SmartConnect to log in to the OneFS web administration interface, the session fails or returns you to the login page. Workaround: Connect to the cluster with a static IP address instead of a hostname. For more information, see article 88893 on the EMC Online Support site.	91640

After a group change, the dnsiq_d process might fail. If this occurs, the following message appears in the `/var/log/messages` file:
`/usr/sbin/isi_dnsiq_d:vip_configured+0x54`
`/usr/sbin/isi_dnsiq_d:vip_ifconfig_down+0x18`
`/usr/sbin/isi_dnsiq_d:apply_flx_subnet+0x7c`
`/usr/sbin/isi_dnsiq_d:gmp_group_changed+0x122`
`/usr/sbin/isi_dnsiq_d:main+0x660`
`/usr/sbin/isi_dnsiq_d:_start1+0x80`
`/usr/sbin/isi_dnsiq_d:_start+0x15`
	78588

If a node with a static IP address is Smartfailed, the IP address is assigned to another node. In some cases, the IP address that is moved might be moved to a	71687

Networking known issues	ID
node that already has an IP address assigned to it, replacing the IP address on that node.	

NFS

NFS known issues	ID
If an NFSv4 client sends a request to OPEN a file and the OneFS NFS server returns an error to the NFSv4 client, some clients might not correctly increment the sequence ID of the request before resending it to the server. If the error returned by the NFS server is transient and the NFSv4 client resends the OPEN request with a sequence ID that is not correctly incremented, an `NFS4ERR_BAD_SEQID` error is returned to the client and the client is denied access to the file. Under these circumstances, if you set the logging level for the nfs process to debug, the following message appears in the `/var/log/nfs.log` file: `opStatus: 10026 on v4 operation 18`. Workaround: Under these circumstances, the NFS server cannot resolve the error, and the affected client's connection to the cluster must be reset to clear the error and allow access to the file.	164891

On occasion, when OneFS is shutting down the uNFS server, a syscall made by the server does not return a response within the allowed 15-minute grace period. As a result, the uNFS server is forcibly shut down and lines similar to the following appear in the `/var/log/messages` file:
`/lib/libc.so.7:syscall+0xc`
`/usr/likewise/lib/lw-svcm/onefs.so:OnefsQuerySetInformationFile+0xa7`
`/usr/likewise/lib/lw-svcm/onefs.so:OnefsSetInformationFile+0x3b`
`/usr/likewise/lib/lw-svcm/onefs.so:OnefsIrpSpark+0x109`
`/usr/likewise/lib/lw-svcm/onefs.so:OnefsIrpWork+0xfa`
`/usr/likewise/lib/lw-svcm/onefs.so:OnefsAsyncStart+0x55`
`/usr/likewise/lib/lw-svcm/onefs.so:OnefsDriverDispatch+0x6f`
`/usr/likewise/lib/libiomgr.so.0:IopFmIrpStateDispatchFsdExec+0x9d`
`/usr/likewise/lib/libiomgr.so.0:IoFmIrpDispatchContinue+0x56c`
`/usr/likewise/lib/libiomgr.so.0:IopIrpDispatch+0x1d0`
`/usr/likewise/lib/libiomgr.so.0:IopQuerySetInformationFile+0x1fc`
`/usr/likewise/lib/libiomgr.so.0:IoSetInformationFile+0x44`
`/usr/likewise/lib/lw-svcm/nfs.so:Nfs4SetattrSetInfoFile+0x5a2`
`/usr/likewise/lib/lw-svcm/nfs.so:Nfs4Setattr+0x3bd`
`/usr/likewise/lib/lw-svcm/nfs.so:NfsProtoNfs4ProcSetAttr+0x178`
`/usr/likewise/lib/lw-svcm/nfs.so:NfsProtoNfs4ProcCompound+0x87e`
`/usr/likewise/lib/lw-svcm/nfs.so:NfsProtoNfs4Dispatch+0x486`
`/usr/likewise/lib/lw-svcm/nfs.so:NfsProtoNfs4CallDispatch+0x3e`
`/usr/likewise/lib/liblwbase.so.0:SparkMain+0xb7`
	136358, 135529

OneFS might report that NFS clients are still connected to the cluster after the clients have disconnected.	135376

Under certain, rare circumstances, the NFS process might unexpectedly restart, causing all NFS clients to be disconnected from the cluster. If this occurs, the following lines appear in the `/var/log/messages` file:
`/lib/libc.so.7:thr_kill+0xc`
`/lib/libc.so.7:__assert+0x35`
`/usr/likewise/lib/libiomgr.so.0:IoFileSetContext+0x32`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsStoreCCB+0x20`
	129684

NFS known issues	ID

`/usr/likewise/lib/lwio-driver/onefs.so:OnefsNfsCreateFile+0xf4b`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsCreateInternal+0x1209`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsSemlockAvailableWorker+0x92`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsAsyncUpcallCallbackWorker+0x1dd`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsAsyncUpcallCallback+0xe8`
`/usr/lib/libisi_ecs.so.1:oplocks_event_dispatcher+0xb9`
`/usr/likewise/lib/lwio-driver/onefs.so:OnefsOplockChannelRead+0x56`
`/usr/likewise/lib/liblwbase.so.0:EventThread+0x6dc`
`/usr/likewise/lib/liblwbase.so.0:LwRtlThreadRoutine+0xee`
`/lib/libthr.so.3:_pthread_getprio+0x15d`	

If an SMB client has an opportunistic lock (oplock) on a file and the file is renamed or deleted by an NFS client, the SMB client does not relinquish its oplock, and the file data on the SMB client is not updated. This issue is caused by an extremely rare race condition that might occur in OneFS 6.0 or later. For more information, see article 88591 on the EMC Online Support site.	94168
Moving files between exports in an NFSv4 overriding-exports scenario may cause unforeseen consequences. Workaround: Configure exports so that they are not exporting similar paths or mapping to two different credentials.	70616

OneFS API

OneFS API known issues	ID
The Isilon Swift API cannot be used on a cluster to which the Security Technology Implementation Guides (STIG) hardening profile is applied. This limitation exists because, in order for a OneFS cluster to comply with STIG ID: WA00520 A22, the Apache 2.2 server must not be configured as a proxy server. To prevent the Apache 2.2 server from being configured as a proxy server, OneFS does not load the following modules, which are required by Swift, on a cluster to which the STIG hardening profile is applied: • proxy_module • proxy_connect_module • proxy_http_module	147627, 146726
If a client is connected to the cluster through Swift, OneFS logs some messages as errors when they should be logged as warnings. For example, if a Swift client attempts to write to a read-only file, a message similar to the following appears in the `/var/log/lwswift.log` file: `Error message: The calling user was denied one or more of the desired access bits.`	134770
Under some circumstances, multiple isi_papi_d process threads might call the same code at the same time. If this occurs, the isi_papi_d process might unexpectedly restart.	130692
In the RESTful Access to the Namespace (RAN) API, when a file is created through the PUT operation, a temporary file of the same name with a randomly generated	104388

OneFS API known issues	ID
suffix is placed in the target directory. Under normal circumstances, the temporary file is removed after the operation succeeds or fails. However, the temporary file may remain in the target directory if the server crashes or is restarted during the PUT operation.	

OneFS web administration interface

OneFS web administration interface known issues	ID
On the **Protocols** › **Windows Sharing** › **SMB Shares**tab in the OneFS web administration interface, if you click **Reset** or **Cancel** in the **Add a User or Group** dialog box while adding or viewing an SMB Share, the **Add a User or Group** dialog becomes inoperable. Workaround: Refresh the OneFS web administration web page.	139712
If you attempt to upload cluster information through the OneFS web administration interface and the upload fails, the interface for uploading information stops responding. If you attempt to upload information again, OneFS will display a `Gather Succeeded` message, but no cluster information will be uploaded.	133974
Modifying the **File Count** value on the **Cluster Diagnostics** › **Packet Capture** tab in the OneFS web administration interface has no effect.	106299
The default SSL port (8080) for the web administration interface cannot be modified. For more information, see article 88725 on the EMC Online Support site.	94026
If a snapshot revert or replication domain exists for a directory, **File System Explorer** incorrectly displays the directory as a SmartLock directory. This is a display issue only.	81598

Security

Security known issues	ID
By default, the Very Secure FTP Daemon (vsftpd) service supports clear-text authentication, which is a possible security risk. For more information about this issue, see the *Protocols* section of the *OneFS 7.2 Security Configuration Guide*.	127738

Security profiles

Security profiles known issues	ID
When a hardened security profile is applied to a OneFS 7.2.1 and later cluster, the Federal Information Processing Standard (FIPS) is applied to OneFS security key management. As a result, SHA-1 checksum values are added to the key manager database, and the key manager database, which previously used MD5 checksum values, uses only SHA-1 checksum values for new entries. As a result, if new SEDs	147734

Security profiles known issues	ID
are added to the cluster, the key manager database is updated with SHA-1 checksum values only. If the cluster is later joined to another cluster that is running an earlier version of OneFS that relies on MD5 values, or the cluster is later downgraded to an earlier version of OneFS that relies on MD5 values, the SEDs will not be recognized. Workaround: Prior to joining the cluster to a cluster that is running an earlier version of OneFS, or prior to downgrading the cluster to an earlier version of OneFS, perform the following steps to disable FIPs: 1. Open an SSH connection on any node in the cluster and log in using the compadmin account. 2. Run the following command to open the `openssl.cnf` for editing: `edit /etc/mcp/templates/openssl.cnf` 3. Change the values assigned to the `fips_mode` entry from `yes` to `no`. 4. Save and close the `openssl.cnf` file.	
If the Security Technology Implementation Guides (STIG) hardening profile is applied to a OneFS 7.2.1 or later cluster, in order to comply with SRG item GEN005504, the SSH process (sshd) must only listen on IP addresses that are designated for management traffic. In order to ensure that the sshd process only listens on the designated addresses, a ListenAddress directive is added to the `sshd_config` file for each IP address assigned to a node and a listen socket is opened for each ListenAddress directive. Because the sshd process only supports sixteen open sockets, if you assign greater than sixteen IP addresses to nodes in a cluster to which the STIG hardening profile has been applied, the sshd process stops unexpectedly, open SSH sessions are disconnected, you will be unable to establish new SSH connections to the affected nodes, and errors similar to the following appear on the console and in the `/varl/og/messages` file: `sshd[87472]: fatal: Too many listen sockets. Enlarge MAX_LISTEN_SOCKS` Workaround: Reconfigure the affected nodes through the OneFS web administration interface so that fewer than seventeen IP addresses are assigned to each node.	143726
On hardened clusters, after you run a command, an `fc` command is run to read the last command that was logged in the ZSH command history list. As a result, after you establish an SSH connection and then log on to a cluster to which a hardened security profile has been applied, the following message is displayed: `isi_precmd:fc:1: bad history number: 1` This message is displayed because the ZSH command history list does not yet have any entries.	142220
Because some OneFS processes—for example, the MCP process—periodically update files on the system but do not explicitly set permissions on those files, after a file has been updated by one of these processes, the permissions applied to those files might not meet the requirements of a STIG hardened profile. Under normal circumstances, on a STIG hardened cluster, a regularly scheduled cron job resets the permissions on the updated files to the expected levels. However, if you revert a STIG hardened cluster before the scheduled cron job resets the	142007

Security profiles known issues	ID

permissions on the updated files, the hardening revert script might detect some files that are not set to the expected level of permissions. If this occurs, messages similar to the following appear on the console:
 	
`Found the following Issue(s) on the cluster:`
`Issue #1 (Islion Control_id:isi_GEN000000-LNX00520)`
`Node:node-1`
`1: /etc/sysctl.conf: Actual permission is 0644 but Expected permission is 0600`
`Node:node-3`
`1: /etc/sysctl.conf: Actual permission is 0644 but Expected permission is 0600`
`Node:node-2`
`1: /etc/sysctl.conf: Actual permission is 0644 but Expected permission is 0600`	

If the Security Technology Implementation Guides (STIG) hardening profile is applied to a OneFS 7.2.1 or later cluster, in order to comply with WA00565_A22, some HTTP methods are disabled. As a result, after applying a STIG hardened profile to the cluster, only the GET, POST, OPTIONS, PUT, and DELETE methods are enabled and DAV operations that rely on other methods might fail.	140794

SmartLock

SmartLock known issues	ID
If you attempt to create a hard link to a file in a WORM (Write Once Read Many) directory, the following incorrect error message displays: `Numerical argument out of domain`	73790

SmartQuotas

SmartQuotas known issues	ID
Writing files past the quota limit over NFSv4 generates an input-output error.	69816

SMB

SMB known issues	ID
If all of the following conditions were met, an `Input/Output` error sometimes appeared after you attempted to transfer a file to the cluster over NFS: • The file was stored in an NFS export or on a file system that was set to ISO-8859-1 character encoding • The Isilon cluster was set to UTF-8 encoding • The name of the file or directory contained a Latin character from the ISO-8859-1 character set	163522

SMB known issues	ID
The lwio process might fail while a node is being shut down.	135869

If you shut down a node while a cluster is under heavy load, the following lines might appear in the `/var/log/messages` file:
`/lib/libc.so.7:recvfrom+0xc`
`/usr/lib/libisi_gconfig_c_client.so.1:gconfig_connection_flush+0x375`
`/usr/lib/libisi_gconfig_c_client.so.1:gconfig_connection_read_message+0x47`
`/usr/lib/libisi_gconfig_c_client.so.1:gconfig_client_update_entries_count+0x799`
`/usr/lib/libisi_gconfig_c_client.so.1:gconfig_client_wait_for_config_change+0x274`
`/usr/likewise/lib/lwio-driver/srv.so:StoreChangesWatcherThreadRoutine+0xf3`
`/lib/libthr.so.3:_pthread_getprio+0x15d`
	134661

The lwio process might fail while running the NetWkstaGetInfo function at info level 102.	134324
If an application sends OneFS a request for alternate data streams, but specifies a buffer size that is too small to receive all of the alternate data streams, OneFS will report that the streams do not exist, instead of reporting that the buffer size was too small.	134299
Alternate data streams might be inaccessible through Windows PowerShell.	134250
The lwio process might fail while running the SrvBuildNegotiateResponse_SMB_V1_NTLM_0_12 function due to an incorrectly formed request during NTLM authentication, resulting in a temporary loss of authentication service.	134248
If you add a new Access Control Entry (ACE) that grants run-as-root permissions to an Access Control List (ACL) on an SMB share, OneFS adds a duplicate ACE if there is already an entry granting full control to the identity. The extra ACE grants no extra permissions. Workaround: Remove the extra ACE by running the `isi smb permissions` command.	72337

Upgrade and installation

Upgrade and installation known issues	ID
If you initiate a simultaneous upgrade through the OneFS web administration interface, the following incorrect message appears: `A rolling upgrade is currently in progress. Some changes to configuration may be disallowed.`	133409

Upgrade and installation known issues	ID

If you run the `sudo isi update` command to begin a OneFS upgrade on a cluster in compliance mode, the pre-upgrade check might report that unresolved critical errors similar to the following were detected:
 	
`Device ad4 located at Carrier board J4 is not a supported boot disk. Please replace with a supported boot disk immediately.`
 	
`Device ad3 located at Carrier board J3 is not a supported boot disk. Please replace with a supported boot disk immediately.`
 	
These errors can safely be ignored.
	131929

During a rolling upgrade, if you select `Confirm before rebooting nodes` in the OneFS web administration interface or specify the `--manual` option from the command-line interface to request a prompt before each node is rebooted, you must respond to each prompt to reboot a node within 15 minutes. If a prompt to reboot is not answered within the 15-minute timeout period, the rolling upgrade process will stop progressing . **Note** In the web administration interface the upgrade job appears to be in progress and there is no indication that the upgrade is halted due to the timeout. In addition, you can click the reboot button, but the node is not rebooted. Workaround: Restart the rolling upgrade from either the web administration interface or the command-line interface. The process will restart from where it left off, and should complete successfully.	126799
If a node is rebooted during a rolling upgrade, and the node fails, the upgrade process might continue to run indefinitely, even after all other nodes have been upgraded. Workaround: After the failed node has been removed, reboot another node in the cluster.	125320
If Collect or MultiScan jobs are in progress when either a rolling upgrade or a cluster reboot is initiated, the job might fail instead of being cancelled. **Note** If the Collect or MultiScan jobs continue to fail after the rolling upgrade is complete, it is unlikely that the failure was caused by this issue.	123903
During a OneFS upgrade, the scripts that upgrade the OneFS configuration information might timeout if the configuration information that is being upgraded exceeds the maximum size allowed in the configuration database. If this issue occurs, you might not be able to create or modify the configuration information for one or more OneFS components. For example, you might not be able to configure NFS exports, SMB shares, or network components after the upgrade is complete.	123584
If a FlexProtect or FlexProtectLin job is started when a rolling upgrade is in progress, OneFS might cancel the job, and the job might not complete until after the rolling upgrade is complete. If the FlexProtect job was started because a device failed during the upgrade, pause the upgrade until the FlexProtect job completes. Do not pause the FlexProtect job.	123167

Upgrade and installation known issues	ID
Workaround: Run a FlexProtect job to completion before you perform an OneFS upgrade. This is considered a best practice for all upgrades. For more information about upgrade best practices, see the Upgrade Planning and Process Guide on the EMC Online Support site.	
If an upgrade job is started on a cluster containing a node with a degraded boot drive, the upgrade engine crashes on initialization, preventing the upgrade from proceeding. For more information, see article 88746 on the EMC Online Support site.	98072
The OneFS web administration interface does not prevent multiple rolling upgrades from being started simultaneously. If multiple rolling upgrades are running simultaneously, the upgrades fail.	84376

Virtual plug-ins

Virtual plug-ins known issues	ID
If you try to register a VMware vSphere API for Storage Awareness (VASA) provider by using an IPv6 address, the registration fails. Workaround: Register the VASA provider by using a domain name that resolves to the required IP address or by using a SmartConnect zone name.	136368
Adding an Isilon vendor provider might fail when you enable VASA support. In addition, the VASA information that appears in vCenter might be incorrect. These issues can occur if you create a data store or virtual machine through the VMware vSphere PowerCLI. Workaround: You can resolve this issue by creating data stores through either the VMware vCenter graphical user interface or the VMware ESXi command-line interface.	97735

CHAPTER 8

OneFS Release Resources

Sources for information about and help with the OneFS operating system.

- OneFS information and documentation 116
- Functional areas in the OneFS release notes 117
- Where to go for support 120
- Provide feedback about this document 121

OneFS information and documentation

You can access OneFS product documentation and technical information about EMC Isilon products through the following channels.

Channel	Description
Isilon Info Hubs	Isilon info hubs organize Isilon documentation, videos, blogs, and user-contributed content into topic areas, making it easy to find content about subjects that interest you.
EMC Isilon Community Network	The EMC Isilon Community connects you to a central hub of information and experts. From this site, you can demo our products, ask questions, view technical videos, and get the latest Isilon product documentation.
EMC Isilon OneFS Product Page	The EMC Isilon OneFS product page provides access to Isilon product documentation and current software releases. Access to this page requires that you log in to the EMC Online Support site.
EMC Isilon Security Advisories	An EMC security advisory (ESA) is published in the EMC knowledgebase when an EMC product is impacted by a security vulnerability and a remedy that must be applied by customers is available from EMC. For the most up to date list of Isilon ESAs, see EMC Isilon Security Advisories (ESAs) on the EMC Isilon Community Network.
EMC Isilon Technical Advisories	An EMC technical advisory (ETA) is published in the EMC knowledgebase when EMC identifies an issue that may cause serious negative impact to a production environment. For the most up to date list of Isilon ETAs, see EMC Isilon Technical Advisories (ETAs) on the EMC Isilon Community Network.
EMC Isilon YouTube playlist	The EMC Isilon YouTube playlist on the EMC Corporate YouTube channel contains Isilon how-to videos, information about new features, information about Isilon hardware, and technical overviews.
Help on This Page	From any page in the OneFS web administration interface, select **Help on this Page** from the **Help** menu to see specific information about the controls and settings on that page. The **Help on this Page** option does not require internet connectivity.
Isi Knowledge	The Isi Knowledge blog links to Isilon support content, announces the of availability of content, offers product tips, and provides information about new ID.TV videos.
Online Help	Select **Online Help** from the **Help** menu in the OneFS web administration interface to see information from the *OneFS Web Administration Guide* and the *OneFS Event Reference*. The Online Help contains the latest available versions for these guides. The **Online Help** option requires internet connectivity.

Functional areas in the OneFS release notes

This section contains a list of the functional areas that are used to categorize content in the OneFS release notes and descriptions of the types of content that each category contains.

Antivirus

This functional area is used to categorize new features, changes, and issues that affect the way OneFS interacts with antivirus software.

Authentication

This functional area is used to categorize new features, changes, and issues that affect authentication on the cluster. This includes, but is not limited to:

- Access control lists (ACLs)
- LDAP
- NIS
- Role-based access control (RBAC)

Backup, recovery, and snapshots

This functional area is used to categorize new features, changes, and issues that affect backup, recovery, and snapshots. This includes, but is not limited to:

- NDMP
- Snapshots
- SyncIQ
- Symantec NetBackup

Cluster configuration

This functional area is used to categorize new features, changes, and issues that affect cluster configuration. This includes, but is not limited to:

- CloudPools
- Licensing
- NTP
- OneFS registry (gconfig)
- SmartPools

Command-line interface

This functional area is used to categorize new features, changes, and issues that affect the OneFS command-line interface.

Diagnostic tools

This functional area is used to categorize new features, changes, and issues that affect tools that are used to research and diagnose cluster issues. This includes, but is not limited to:

- EMC Secure Remote Services (ESRS)

- Gather info (`isi_gather_info`)
- Help in the OneFS web administration interface

Events, alerts, and cluster monitoring

This functional area is used to categorize new features, changes, and issues that affect utilities that are used to detect and record system events and utilities that are used to monitor cluster health and general statistics. This includes, but is not limited to:

- Alerts
- Protocol auditing
- Cluster event log (CELOG)
- File system analytics (FSA)
- Onsite Verification Test (OVT)
- Simple Network Management Protocol (SNMP)
- Statistics
- Status

File system

This functional area is used to categorize new features, changes, and issues that affect the OneFS file system. This includes, but is not limited to:

- Cluster group management
- File system coalescer
- File system events (not CELOG)
- FreeBSD
- L3 cache
- MCP
- Network Lock Manager (NLM)
- OneFS Kernel

File transfer

This functional area is used to categorize new features, changes, and issues that affect FTP and HTTP connections to the cluster.

Hardware

This functional area is used to categorize new features, changes, and issues that affect Isilon hardware in a OneFS cluster .

HDFS

This functional area is used to categorize new features, changes, and issues that affect the HDFS protocol.

iSCSI

This functional area is used to categorize new features, changes, and issues that affect the iSCSI protocol and iSCSI devices connected to a OneFS cluster.

Note

Support for the iSCSI protocol is deprecated in this version of OneFS.

Job engine

This functional area is used to categorize new features, changes, and issues that affect the OneFS job engine and deduplication in OneFS.

Migration

This functional area is used to categorize new features, changes, and issues that affect migration of data from a NAS array or a OneFS cluster to a OneFS cluster through the `isi_vol_copy` utility or the `isi_vol_copy_vnx` utility.

Networking

This functional area is used to categorize new features, changes, and issues that affect the OneFS external network and the OneFS back-end network. This includes, but is not limited to:

- Fibre Channel
- Flexnet
- InfiniBand
- SmartConnect
- TCP/IP

NFS

This functional area is used to categorize new features, changes, and issues that affect NFS connections to the cluster.

OneFS API

This functional area is used to categorize new features, changes, and issues that affect the OneFS Platform API and SWIFT.

OneFS web administration interface

This functional area is used to categorize new features, changes, and issues that affect the web administration interface.

Performance

This functional area is used to categorize new features, changes, and issues that affect cluster performance.

Security

This functional area is used to categorize new features, changes, and issues that are related to security fixes and vulnerabilities.

Security Profiles

This functional area is used to categorize new features, changes, and issues that affect hardened profiles such as the security technical information guides (STIG).

SmartQuotas

This functional area is used to categorize new features, changes, and issues that affect SmartQuotas.

SMB

This functional area is used to categorize new features, changes, and issues that affect SMB connections to the cluster.

Upgrade and installation

This functional area is used to categorize new features, changes, and issues that affect OneFS upgrades, installation of OneFS patches, and the reformatting and reimaging of Isilon nodes by using a USB flash drive.

Virtual plug-ins

This functional area is used to categorize new features, changes, and issues that affect virtual plug-ins. This includes, but is not limited to:

- Isilon for vCenter
- OneFS Simulator
- Storage Replication Adapter (SRA)
- vStorage APIs for Array Integration (VAAI)
- VMware vSphere API for Storage Awareness (VASA)

vOneFS

This functional area is used to categorize new features, changes, and issues that affect vOneFS.

Where to go for support

Contact EMC Isilon Technical Support for any questions about EMC Isilon products.

Online Support	Live Chat Create a Service Request
Telephone Support	United States: 1-800-SVC-4EMC (800-782-4362) Canada: 800-543-4782 Worldwide: +1-508-497-7901 For local phone numbers for a specific country, see EMC Customer Support Centers.
Help with Online Support	For questions specific to EMC Online Support registration or access, email support@emc.com.
Isilon Info Hubs	For the list of Isilon info hubs, see the Isilon Info Hubs page on the EMC Isilon Community Network. Isilon info hubs organize Isilon documentation, videos, blogs, and user-contributed content into topic areas, making it easy to find content about subjects that interest you.

Support for IsilonSD Edge

If you are running a free version of IsilonSD Edge, community support is available through the EMC Isilon Community Network. However, if you have purchased one or more licenses of IsilonSD Edge, you can contact EMC Isilon Technical Support for assistance, provided you have a valid support contract for the product.

Provide feedback about this document

We value your feedback. Please let us know how we can improve this document.

- Take the survey at http://bit.ly/isi-docfeedback.
- Send your comments or suggestions to docfeedback@isilon.com.