

IBM System Storage N series



Data ONTAP 8.0.4 7-Mode Release Notes

Contents

Data ONTAP 8.0.4 7-Mode Release Notes	1
Changes introduced since Data ONTAP 8.0	3
Changes introduced in Data ONTAP 8.0.4 7-Mode release	3
Changes in the Data ONTAP 8.0.3 7-Mode release	5
Changes in the Data ONTAP 8.0.2 7-Mode release	9
Changes in the Data ONTAP 8.0.1 7-Mode release	13
New and changed features	17
New and changed platform and hardware support	19
Support for 1-TB Flash Cache modules	19
Support for N6200 series storage systems	19
Support for the N7x50T series storage systems	20
Support for 512-GB and 256-GB Performance Acceleration Module (PAM II)	21
Support for SSDs	21
Support for 900-GB SAS disks	21
Supported tape drives	21
Support for 3-TB SATA disks	21
Support for 2-TB SATA disks	21
Changes in storage expansion unit support	23
Support for EXN3500 storage expansion units	23
Support for EXN3000 storage expansion units	23
Manageability enhancements	24
Reversion and downgrade support	24
SSH server version for Data ONTAP	24
Supported OpenSSH client versions	24
Support for a 64-bit root volume	25
The Service Processor for the N6200 and N7x50T series storage systems	25
Enhancements that allow you to specify SSL versions	26
Root aggregate Snapshot reserve values reduced	26
SSL security improvements	27
Default security settings for manageability protocols	27
Requirement to set root account password	28
Uses of the systemshell and the diagnostic account	28
Boot menu enhancements and changes	28
Read reallocation of data	29
Extents for FlexVol volumes	29
Partner recipients for AutoSupport email notifications	29
AutoSupport message file attachments	30
New disk drive AutoSupport message sent weekly	30
Storage resource management enhancements	31
Support for 64-bit aggregates	31
File space utilization report	31
Upgrading to Data ONTAP 8.0.1 increases volume capacity after deduplication	32
Maximum simultaneous FlexClone file or FlexClone LUN operations per storage system	32

Support for FlexClone files and FlexClone LUNs on vFiler units	32
Increased maximum RAID group size for SATA disks	32
Ability to decrease maxfiles.	32
High-availability pair (formerly active/active configuration) enhancements.	33
High-availability pair terminology changes	33
Networking and Security protocols enhancements	34
Ability to block data traffic on the e0M interface	34
Support for the SMB 2.0 protocol	35
Support for CDP	35
Port-based load-balancing option for multimode interface groups	35
TSO in NICs.	35
64-bit SNMP object identifiers (OIDs) are supported	36
Configuration for receiving untagged traffic on VLAN tagged interfaces.	36
Vifs are now known as interface groups	36
File access protocol enhancements	37
Maximum number of auxiliary UNIX groups supported for a user	37
Support for multiple open instances of the SMB named pipe on an FPolicy server	37
Block protocols enhancements	38
iSCSI RADIUS support	38
Support for Data Motion for Volumes	38
VMware VAAI features for ESX hosts support	39
igroup scalability limits	39
Data protection enhancements.	40
SnapVault update schedules for volume SnapMirror destination backups	40
Support for 840 disk drives on fabric-attached MetroCluster configuration	41
Support for out-of-order frame delivery for SnapMirror over Fibre Channel	41
Support for SnapMirror relationship between two FlexClone volumes	41
Support for SnapMirror network compression	41
Support for designated range of ports for NDMP data connections	42
If volumes with deduplication exceed the maximum supported size	42
NearStore license is not required for deduplication	42
Maximum concurrent deduplication operations supported	42
Prerequisite for backing up data to tape using NDMP services	42
Snapshot copy schedules	43
SnapMirror support for 64-bit volumes	43
SnapVault support for 64-bit volumes	43
MultiStore enhancements	43
New option for the vfiler dr configure command	44
Changes to MultiStore support	45
Data Motion is not supported in the Data ONTAP 8.0 7-Mode release family	45
New and changed commands and options.	47
New commands in Data ONTAP 8.0 7-Mode	48
Changed commands in Data ONTAP 8.0 7-Mode	52
Replaced or removed commands in Data ONTAP 8.0 7-Mode	56
New options in Data ONTAP 8.0 7-Mode	57
Changed options in Data ONTAP 8.0 7-Mode	59
Requirements for running Data ONTAP 8.0 7-Mode.	61
Firmware.	62

Gateway requirements and support information	63
Important cautions	65
Unsupported features	66
Upgrade and revert cautions	67
CIFS startup might take several minutes after cluster failover giveback	67
Change in nondisruptive upgrade procedure	68
Upgrade process changes with Data ONTAP 7-Mode software images	69
iSCSI traffic not supported on e0M and e0P interfaces	70
New FlexVol volume space requirements when upgrading from a release earlier than Data ONTAP 7.3	70
Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later	70
If your root FlexVol volume does not meet the new size requirements	71
Nondisruptive upgrades on systems with VMware ESX Server hosts	71
Requirements for nondisruptive upgrades on systems with deduplicated volumes	72
New reversion and downgrade procedures	72
Manageability cautions	73
Certificates signed with MD5 could cause loss of connectivity to secure clients	73
If you need to boot the storage system from a Data ONTAP image stored on a remote server	73
TLS is not supported in the Data ONTAP 8.0 release family	74
RLM over IPv6 not supported in the Data ONTAP 8.0 release family	74
Storage management cautions	74
Disk failures might increase for some disk models after upgrading to this version of Data ONTAP	74
Network protocol cautions	76
System might panic if the network configuration entries in the /etc/rc file are not in the recommended order	76
Reconfiguring the losk interface results in system outage	76
IPv6 is not supported	76
IPsec is not supported	76
TOE is not supported	76
File access and protocols cautions	77
CIFS setup using ONTAPI library might lead to the removal of already existing user-created CIFS shares	77
Storage system might panic when not specifying an export path	77
FTPS is not supported	77
SFTP is not supported	78
Block access protocol cautions	79
Non-ASCII characters in the comment string of the lun setup command result in system panic	79
Identical WWPNs can occur on both nodes in a cluster	79
If you use Solaris 10 and the emlxs Fibre Channel Adapter driver	79
cfmode support change in Data ONTAP 8.0	79
Data protection cautions	81
SnapValidator for Oracle not supported	81
SnapLock is not supported	81
New RLM firmware and its upgrade requirements	85
Known problems and limitations	87
Manageability issues	88
Cached data in a Flash Cache module is cleared during system reboot and takeover and giveback	88

The wrfile command might fail to capture all input during a cut-and-paste operation performed through the RLM or SP	88
NTP version used to communicate with a newly configured NTP server defaults to v3.	89
RLM configuration is not saved if the RLM's Ethernet cable is not connected	89
If your system uses rtc or rdate as the protocol for time synchronization	89
Obsolete timed options	89
Issues regarding the timed log.	90
Shutdown delay in N3400 systems when critical thresholds are crossed	90
Storage resource management issues.	91
Writes to a LUN fail when the combined size of the LUN and the change log is greater than the volume size	91
Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance.	91
After changing the port assigned to ACP, a reboot is required to regain use of the original port	91
Moving disks to a system running an earlier version of Data ONTAP could cause disk identification issues	91
Data ONTAP can silently select varying disk sizes or array LUN sizes when creating an aggregate	92
Adding disks when varying sized disks are in use	92
Discovering what disks Data ONTAP will automatically select	93
SSD Spare blocks consumed limit reported as N/A.	93
High-availability pair issues	94
Storage system in an HA pair might get rebooted during a controller or NVRAM replacement	94
If takeover is enabled for network interface failure, automatic giveback should be disabled	94
Network protocol issues.	95
Partner node might be unavailable during controller failover when an interface name is used as an alternate host name in the /etc/hosts file	95
Change in routing configuration maximum value	95
NIS group file entries should not exceed 1024 characters per line	95
File access protocol issues	96
Value of the option nfs.ifc.rcv.high is unexpectedly reset	97
Enabling new EMS messages to help troubleshoot Kerberos issues	97
Microsoft Windows 7 clients fail to rename files and directories	97
File copy operation on a widelink path inside a CIFS share might fail when using SMB 2.x	98
Domain controller responds with an access denied error to an SMB 2.x tree connect request	98
Controlling NFS requests from nonreserved ports	98
Enabling or disabling NFSv2	99
Error NFS4_BADOWNER when NFSv4 client passes UID as string	99
Domain user unable to execute CLI commands.	100
CIFS shares with comment ending with backslash disappear	100
Removing stale entries from the access cache	100
CIFS authentication failure with Windows Server 2008 R2 read-only domain controllers	100
Too many CIFS audit files can cause system performance degradation	101
CIFS client authentication fails if host name exceeds 16 characters	101
Failed CIFS connections due to signature mismatch	101
Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance	101
Excessive pending CIFS authentication requests can cause service disruption.	102
Client notification messages in Windows domains require NetBIOS	102
Configuration issue for clients that mount NFS shares using a non-reserved port	102
Widelinks are not accessible from Mac clients	103

Empty CIFS/SMB2.x change notifications	103
Block access protocol issues	104
vStorage priority setting is not persistent after a reboot	104
Disks offline in Windows 2008 after a standard upgrade.	104
iSCSI Target HBA support.	104
If an iSCSI target HBA fails	105
If you use the automatic sizing feature on thinly provisioned LUNs when snap reserve is set to a non-zero value	105
Data protection issues	106
Tape devices behind DataFort are not detected unless a known SCSI device is mapped to LUN 0	106
Tape backup and volume move operations must not be performed simultaneously.	107
SMTape restore to an offline volume overwrites the data on that volume	107
Deleted files occupy additional space on SMTape restored volumes	107
Limitations when using SnapVault with non-default vFiler units	107
Restore engine does not honor any mandatory or advisory locks on files	108
Restore operation to a non-existent volume restores data to the root volume	109
Suboptimal space savings with deduplication when 16-GB Performance Acceleration Module is installed.	109
If you use DataFabric Manager server to manage SnapMirror or SnapVault relationships between vFiler units.	109
Changes to published documentation	111
Changes to the Data ONTAP Upgrade and Revert/Downgrade Guide	112
Correction to the information in the Planning your upgrade with Upgrade Advisor topic	112
Updates to software image installation examples	112
Updated NDU requirements	118
Additional preparation for nondisruptive upgrades	118
Additional SAN upgrade requirement.	120
Additional requirement for AT-FCX storage expansion unit firmware updates	120
AT-FC and AT-FC2 shelves no longer supported	121
Changes to the System Administration Guide	122
Changes to BMC command information	122
Changes to information about the /etc/rc file	122
Corrections to the timed options.	123
Updated password requirements	123
Corrections for the location of the secureadmin.pem file.	123
Correction to the description of the autosupport.to option	123
Change in the frequency of the performance AutoSupport message	124
Use FlexShare with HDDs only	124
Correction to the free space required for reallocation scans	124
Changes to the supported OpenSSH client versions	124
Changes to storage system configuration backup and restore	124
Changes to the instructions on editing the /etc/rc file	125
Terminology changes for PAM and WAFL extended cache	126
Corrections for root FlexVol volume size requirement.	126
Changes to the root aggregate and root volume Snapshot reserve values on new systems	126
Changes to the Storage Management Guide	128
Corrections to the data compression commands	128
Correction to the availability of FlexClone files and FlexClone LUNs in vfiler contexts	129
Capacity information for 3-TB disks	129

Use FlexShare with HDDs only	129
Incorrect cross-reference to maximum volume size information	129
Maximum volume size for deduplication	129
Maximum number of mirrored aggregates per system	129
ACP example output shows incorrect IP addresses	129
Using disk sanitization to remove data from disks.	130
Changes to the Software Setup Guide	132
Correction to the setup password rules	132
Changes to the High-Availability Configuration Guide	132
Hardware-assisted takeover support on systems that use Service Processor	133
Corrections to the example in "Monitoring HA pair status" topic	133
Changes to the Network Management Guide	133
Changes to information about routed daemon	133
Supported SNMP versions in Data ONTAP	134
TLS not supported in Data ONTAP 8.0	134
Changes to the File Access and Protocols Management Guide	135
Enabling or disabling SFTP log files	135
Specifying the maximum size of the current SFTP log files	135
Correction to specifying the FTP authentication style	136
Incorrect reference to nonexistent SFTP log files	136
Updated compatibility information available online	136
Corrections to the FPolicy commands to monitor directory operations	136
Unsupported Windows features in the file serving environment	137
NFSv4 client compatibility.	137
FTP server does not support Unicode characters	137
Changes to the Block Access Management Guide for iSCSI and FC	138
Correction to the DCB setting example	138
Correction to links in the "Unified Ethernet network management" and "What data center bridging is" topics	138
Setting volume options for the autodelete configuration	138
Correction to configure onboard adapters for target mode	139
Changes to the Data Protection Online Backup and Recovery Guide.	141
Corrections to information about deployment of SnapMirror over Fibre Channel	141
Connection name is required to enable SnapMirror compression	141
Change to the Snapshot copy reserve value	141
Single file reversion to or from symbolic links using SnapRestore not supported	142
Changes to the Data Protection Tape Backup and Recovery Guide	142
Support for SMTape backup of compressed volumes	142
Designating range of ports for NDMP data connections	142
Changes to the MultiStore Management Guide	144
Correction to the effects of storage system reboot on a vFiler unit	144
Requirement for backup and replication operations on vFiler units	144
Naming guidelines for vFiler units	144
SnapMirror relationship schedules are not modified for existing volumes	145
Corrections to the List of RSH and SSH commands	145
Corrections to the SnapVault support on vFiler units	145
Changes to the Storage Efficiency Management Guide	147
Corrections to deduplication and volume SnapMirror information	147
Changes to the Gateway Implementation Guide	147
Changes to the Gateway Implementation Guide for Native Storage Expansion Units	147

Format change for the Gateway Implementation Guides. 148
Changes to the Gateway MetroCluster Guide 148

Websites 149

Copyright and trademark information 151
Trademark information. 152

Notices. 155

Data ONTAP 8.0.4 7-Mode Release Notes

These *Release Notes* describe new features, enhancements, and known issues for Data ONTAP® 8.0.4 7-Mode, as well as additional information about running this release on specific storage systems.

Note: The terms *flexible volumes* and *FlexVol volumes* are used interchangeably in Data ONTAP documentation.

About the Data ONTAP 8.0.4 7-Mode release

Data ONTAP 8.0.4 7-Mode is a General Availability (GA) release in the Data ONTAP 8.0 7-Mode release family. For more information about Data ONTAP releases, see the following publication:

IBM System Storage N series Data ONTAP Release Model

For up to date information regarding the latest release of Data ONTAP, refer to the following publication:

NEWS: Recommended Release for IBM System Storage N series Data ONTAP

The Data ONTAP 8.0.4 7-Mode release includes support for the 1-TB PCI-e-attached Flash Cache module that uses Flash Technology, support for the 3-TB SATA disks, and support for the OpenSSH client version 4.4p1 and 64-bit SNMP object identifiers (OIDs). This release also includes reversion and Data ONTAP downgrade support.

The Data ONTAP 8.0.4 7-Mode release also includes the ability to block data traffic on the e0M interface. New storage systems that are shipped with Data ONTAP 8.0.4 7-Mode will have reduced reserve values for root aggregates and root volume Snapshot copies.

This release also includes additional enhancements to command-line executables, new cautions, update to known limitations, and documentation enhancements.

For a complete description of these and other new features, see “New and changed features” on page 17.

Attention: If you are upgrading to Data ONTAP 8.0.4 7-Mode or reverting from this release to an earlier release, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Data ONTAP 8.0.4 7-Mode supports all the features in the earlier versions of Data ONTAP 7.x releases with some exceptions, including the following:

- SnapValidator for Oracle
- SnapLock
- IPv6

Do not upgrade to Data ONTAP 8.0.4 7-Mode, if you have configured any of these features.

For more information, see “Unsupported features” on page 66.

Changes introduced since Data ONTAP 8.0

If you are already familiar with an earlier release in the Data ONTAP 8.0 7-Mode family you can see changes introduced with each release in the Data ONTAP 8.0 7-Mode family.

Note: If you are new to the Data ONTAP 8.0 family, it is not necessary to read this section; you can go directly to “New and changed features” on page 17.

- “Changes introduced in Data ONTAP 8.0.4 7-Mode release”
- “Changes in the Data ONTAP 8.0.3 7-Mode release” on page 5
- “Changes in the Data ONTAP 8.0.2 7-Mode release” on page 9
- “Changes in the Data ONTAP 8.0.1 7-Mode release” on page 13

Changes introduced in Data ONTAP 8.0.4 7-Mode release

The Data ONTAP 8.0.4 7-Mode release includes hardware and software enhancements, and problem fixes.

New or changed features in Data ONTAP 8.0.4 7-Mode

Data ONTAP 8.0.4 7-Mode includes the following new feature:

- “Support for 900-GB SAS disks” on page 21

For information about all new and changed features in the Data ONTAP 8.0 7-Mode release family, see “New and changed features” on page 17.

New commands in Data ONTAP 8.0.4 7-Mode

Data ONTAP 8.0.4 7-Mode introduces the following new command:

- `bmc battery-info`

For more information, see “New commands in Data ONTAP 8.0 7-Mode” on page 48.

Commands removed in Data ONTAP 8.0.4 7-Mode

No commands are removed in Data ONTAP 8.0.4 7-Mode.

For more information, see “Replaced or removed commands in Data ONTAP 8.0 7-Mode” on page 56

New option in Data ONTAP 8.0.4 7-Mode

Data ONTAP 8.0.4 7-Mode includes the following new option:

- `nfs.rpcsec.trace`

For more information, see “New options in Data ONTAP 8.0 7-Mode” on page 57.

New cautions in Data ONTAP 8.0.4 7-Mode

Data ONTAP 8.0.4 7-Mode includes the following new cautions:

- “Disk failures might increase for some disk models after upgrading to this version of Data ONTAP” on page 74
- “CIFS startup might take several minutes after cluster failover giveback” on page 67
- “CIFS setup using ONTAPI library might lead to the removal of already existing user-created CIFS shares” on page 77

For more information about all the important cautions in the Data ONTAP 8.0 release family, see “Important cautions” on page 65.

Cautions removed in Data ONTAP 8.0.4 7-Mode

The following caution that was included in Data ONTAP 8.0.x Release Notes has been removed in Data ONTAP 8.0.4 7-Mode:

- Oplock breaks might cause panic when using durable handles

New known issues and limitations in Data ONTAP 8.0.4 7-Mode

Data ONTAP 8.0.4 7-Mode includes the following new known issues and limitations:

- “Cached data in a Flash Cache module is cleared during system reboot and takeover and giveback” on page 88
- “Writes to a LUN fail when the combined size of the LUN and the change log is greater than the volume size” on page 91
- “Partner node might be unavailable during controller failover when an interface name is used as an alternate host name in the `/etc/hosts` file” on page 95
- “Value of the option `nfs.ifc.rcv.high` is unexpectedly reset” on page 97
- “Enabling new EMS messages to help troubleshoot Kerberos issues” on page 97

- “Microsoft Windows 7 clients fail to rename files and directories” on page 97
- “File copy operation on a widelink path inside a CIFS share might fail when using SMB 2.x” on page 98
- “Domain controller responds with an access denied error to an SMB 2.x tree connect request” on page 98
- “Controlling NFS requests from nonreserved ports” on page 98
- “Enabling or disabling NFSv2” on page 99
- “Error NFS4_BADOWNER when NFSv4 client passes UID as string” on page 99
- “vStorage priority setting is not persistent after a reboot” on page 104

For more information about all the limitations in the Data ONTAP 8.0 release family, see “Known problems and limitations” on page 87.

Limitations removed in Data ONTAP 8.0.4 7-Mode

The following limitations are removed in Data ONTAP 8.0.4 7-Mode:

- User mapping fails for users in a trusted domain
- SnapVault tries count not persistent across reboots for manually created SnapVault schedules
- Reboot required after updating the krb5.conf file
- Specification of the default keytab file
- File screening requests are not accepted after FPolicy times out

Documentation changes for Data ONTAP 8.0.4 7-Mode

The following new or changed information supplements the documentation you received for Data ONTAP 8.0.4 7-Mode:

- “Changes to the System Administration Guide” on page 122
- “Changes to the Software Setup Guide” on page 132
- “Changes to the File Access and Protocols Management Guide” on page 135

For more information about these and other important documentation changes, see “Changes to published documentation” on page 111.

Changes in the Data ONTAP 8.0.3 7-Mode release

The Data ONTAP 8.0.3 7-Mode includes support for software enhancements, and problem fixes.

New and changed features in Data ONTAP 8.0.3 7-Mode

The Data ONTAP 8.0.3 7-Mode release introduces the following new and changed features:

- “SSH server version for Data ONTAP” on page 24
- “New option for the vfiler dr configure command” on page 44

For information about all new and changed features in the Data ONTAP 8.0 7-Mode release family, see “New and changed features” on page 17.

New commands in Data ONTAP 8.0.3 7-Mode

No new commands are added in Data ONTAP 8.0.3 7-Mode.

For more information, see “New commands in Data ONTAP 8.0 7-Mode” on page 48.

Changed commands in Data ONTAP 8.0.3 7-Mode

The following command is changed in Data ONTAP 8.0.3 7-Mode:

- flexcache fstat

For more information, see “Changed commands in Data ONTAP 8.0 7-Mode” on page 52.

Commands removed in Data ONTAP 8.0.3 7-Mode

No commands are removed in Data ONTAP 8.0.3 7-Mode.

For more information, see “Replaced or removed commands in Data ONTAP 8.0 7-Mode” on page 56.

New options in Data ONTAP 8.0.3 7-Mode

Data ONTAP 8.0.3 7-Mode includes the following new option:

- lun_ic_alua_changed

For more information, see “New options in Data ONTAP 8.0 7-Mode” on page 57.

Changed options in Data ONTAP 8.0.3 7-Mode

No options are changed in Data ONTAP 8.0.3 7-Mode.

For more information, see “Requirements for running Data ONTAP 8.0 7-Mode” on page 61.

New cautions in Data ONTAP 8.0.3 7-Mode

The following new cautions are introduced in Data ONTAP 8.0.3 7-Mode:

- “System might panic if the network configuration entries in the `/etc/rc` file are not in the recommended order” on page 76
- “Storage system might panic when not specifying an export path” on page 77
- “Non-ASCII characters in the comment string of the `lun setup` command result in system panic” on page 79

For more information about all important cautions in the Data ONTAP 8.0 release family, see “Important cautions” on page 65.

Cautions removed in Data ONTAP 8.0.3 7-Mode

The following caution that was included in Data ONTAP 8.0.x Release Notes has been removed in Data ONTAP 8.0.3 7-Mode:

- Do not remove the special system files that are for restoring LUNs in snapshot copies

New known issues and limitations in Data ONTAP 8.0.3 7-Mode

Data ONTAP 8.0.3 7-Mode includes the following new known issues and limitations:

- “The `wrfile` command might fail to capture all input during a cut-and-paste operation performed through the RLM or SP” on page 88
- “Storage system in an HA pair might get rebooted during a controller or NVRAM replacement” on page 94
- “Change in routing configuration maximum value” on page 95
- “Domain user unable to execute CLI commands” on page 100
- “CIFS shares with comment ending with backslash disappear” on page 100
- “Removing stale entries from the access cache” on page 100
- “CIFS authentication failure with Windows Server 2008 R2 read-only domain controllers” on page 100
- “Too many CIFS audit files can cause system performance degradation” on page 101
- “CIFS client authentication fails if host name exceeds 16 characters” on page 101

- “Tape devices behind DataFort are not detected unless a known SCSI device is mapped to LUN 0” on page 106

For more information about all the limitations in the Data ONTAP 8.0 release family, see “Known problems and limitations” on page 87.

Known limitations changed in Data ONTAP 8.0.3 7-Mode

There are no known limitations that have been changed in Data ONTAP 8.0.3 7-Mode.

For more information, see “Known problems and limitations” on page 87.

Limitations removed in Data ONTAP 8.0.3 7-Mode

The following limitations are removed in Data ONTAP 8.0.3 7-Mode:

- Exportfs commands via non-interactive ssh fail intermittently
- Potential issue after removing disk ownership in Maintenance mode
- FPolicy quota and user reports show incorrect file size
- CIFS termination required before issuing a lock break
- Data transfer speed reduces when LSI Logic 1030 Ultra320 SCSI HBA is used with tape drives
- Shutting down a storage system when SMTape operations are active might result in a panic

Documentation changes for Data ONTAP 8.0.3 7-Mode

The following new or changed information supplements the documentation you received for Data ONTAP 8.0.3 7-Mode:

- “Changes to the Data ONTAP Upgrade and Revert/Downgrade Guide” on page 112
- “Changes to the System Administration Guide” on page 122
- “Changes to the Storage Management Guide” on page 128
- “Changes to the High-Availability Configuration Guide” on page 132
- “Changes to the Network Management Guide” on page 133
- “Changes to the Block Access Management Guide for iSCSI and FC” on page 138
- “Changes to the Data Protection Online Backup and Recovery Guide” on page 141
- “Changes to the MultiStore Management Guide” on page 144
- “Changes to the Gateway MetroCluster Guide” on page 148

For more information about these and other important documentation changes, see “Changes to published documentation” on page 111.

Changes in the Data ONTAP 8.0.2 7-Mode release

Data ONTAP 8.0.2 7-Mode includes support for new hardware, software enhancements, and problem fixes.

New and changed features in Data ONTAP 8.0.2 7-Mode

The Data ONTAP 8.0.2 7-Mode release introduces the following new and changed features:

- “Support for 1-TB Flash Cache modules” on page 19
- “Support for 3-TB SATA disks” on page 21
- “Reversion and downgrade support” on page 24
- “Supported OpenSSH client versions” on page 24
- “Root aggregate Snapshot reserve values reduced” on page 26
- “Ability to block data traffic on the e0M interface” on page 34
- “TSO in NICs” on page 35
- “SnapVault update schedules for volume SnapMirror destination backups” on page 40

For information on all new and changed features in the Data ONTAP 8.0 7-Mode release family, see “New and changed features” on page 17.

New commands in Data ONTAP 8.0.2 7-Mode

No new commands were added in Data ONTAP 8.0.2 7-Mode.

For more information, see “New commands in Data ONTAP 8.0 7-Mode” on page 48.

Changed commands in Data ONTAP 8.0.2 7-Mode

The following commands have been changed in Data ONTAP 8.0.2 7-Mode:

- `sysconfig -v`

For more information, see “Changed commands in Data ONTAP 8.0 7-Mode” on page 52.

Commands removed in Data ONTAP 8.0.2 7-Mode

No commands were removed in Data ONTAP 8.0.2 7-Mode.

For more information, see “Replaced or removed commands in Data ONTAP 8.0 7-Mode” on page 56.

New options in Data ONTAP 8.0.2 7-Mode

Data ONTAP 8.0.2 7-Mode includes the following options:

- `snapvault.snapshot_for_dr_backup`
- `interface.blocked.mgmt_data_traffic`
- `nfs.always.deny.truncate`

For more information, see “New options in Data ONTAP 8.0 7-Mode” on page 57.

Changed options in Data ONTAP 8.0.2 7-Mode

Data ONTAP 8.0.2 7-Mode includes the following changed option:

- `autosupport.minimal.subject.id`
- `cifs.max_mpx`

For information about all new and changed command and options in the Data ONTAP 8.0 7-Mode release family, see “New and changed commands and options” on page 47.

Requirements for running Data ONTAP 8.0.

The following requirement that was included in earlier Data ONTAP 8.0.x Release Notes has been removed from Data ONTAP 8.0.2 7-Mode:

- Storage systems that support Multipath Storage

For more information, see “Changed options in Data ONTAP 8.0 7-Mode” on page 59.

New cautions in Data ONTAP 8.0.2 7-Mode

The following new cautions were introduced in Data ONTAP 8.0.2 7-Mode:

- “iSCSI traffic not supported on e0M and e0P interfaces” on page 70
- “New reversion and downgrade procedures” on page 72
- “Certificates signed with MD5 could cause loss of connectivity to secure clients” on page 73

For more information, see “Important cautions” on page 65.

Cautions removed in Data ONTAP 8.0.2 7-Mode

- If you deploy ALUA on storage systems with FCP or iSCSI
- If you have been using the flex_cache license
- Upgrading N6000 series systems
- If you are upgrading to Data ONTAP 8.0 or later from Data ONTAP 7.x and you use quotas
- If you are reverting to a Data ONTAP release with a lower maximum capacity *
- If you are reverting to a previous release and you use deduplication *

The topic If you are upgrading an N5500 storage system, was incorrectly included in the Data ONTAP 8.0.1 Release Notes and has been removed from Data ONTAP 8.0.2 7-Mode Release Notes.

* Reversion and downgrade topics are now covered in the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

New known limitations in Data ONTAP 8.0.2 7-Mode

The following new limitations were introduced in Data ONTAP 8.0.2 7-Mode:

- “NTP version used to communicate with a newly configured NTP server defaults to v3” on page 89
- “Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance” on page 91
- “After changing the port assigned to ACP, a reboot is required to regain use of the original port” on page 91
- “Failed CIFS connections due to signature mismatch” on page 101
- “Excessive pending CIFS authentication requests can cause service disruption” on page 102
- “Empty CIFS/SMB2.x change notifications” on page 103

For information on all known limitations in the Data ONTAP 8.0 release family, see “Known problems and limitations” on page 87.

Known issues and limitations changed in Data ONTAP 8.0.2 7-Mode

There were no known issues and limitations changed in Data ONTAP 8.0.2 7-Mode.

For information on all known limitations in the Data ONTAP 8.0 release family, see “Known problems and limitations” on page 87.

Limitations removed in Data ONTAP 8.0.2 7-Mode

The following limitations were removed from Data ONTAP 8.0.2 7-Mode:

- The `ifconfig` command and the N3400 e0P port
- Creating more than 64 mirrored aggregates might impair disaster recovery
- Forcing a FlexVol volume online must be done in diag mode
- Do not disable the IP fastpath option on an N6200 series system in a High Availability (HA) pair
- Incorrect warning of termination in spite of a successful giveback
- Brocade 200E switch speed must be set manually
- Enable caching with NIS lookup
- Storage system panic when adding preferred domain controllers

Documentation changes for Data ONTAP 8.0.2 7-Mode

The following new or changed information supplements the documentation you received for the Data ONTAP 8.0.2 7-Mode release:

- “Changes to the System Administration Guide” on page 122
- “Changes to the Storage Management Guide” on page 128
- “Changes to the Network Management Guide” on page 133
- “Changes to the File Access and Protocols Management Guide” on page 135
- “Changes to the Block Access Management Guide for iSCSI and FC” on page 138
- “Changes to the Data Protection Online Backup and Recovery Guide” on page 141
- “Changes to the Data Protection Tape Backup and Recovery Guide” on page 142
- “Changes to the MultiStore Management Guide” on page 144
- “Changes to the Storage Efficiency Management Guide” on page 147
- “Changes to the Gateway Implementation Guide” on page 147

For more information about important documentation changes, see “Changes to published documentation” on page 111.

Changes in the Data ONTAP 8.0.1 7-Mode release

The Data ONTAP 8.0.1 7-Mode release includes support for new hardware, software enhancements, and problem fixes.

New and changed features in Data ONTAP 8.0.1 7-Mode

Data ONTAP 8.0.1 7-Mode release includes the following new features:

- “Support for N6200 series storage systems” on page 19
- “Support for the N7x50T series storage systems” on page 20
- “Support for 512-GB and 256-GB Performance Acceleration Module (PAM II)” on page 21
- “Support for SSDs” on page 21
- “Supported tape drives” on page 21
- “Support for EXN3500 storage expansion units” on page 23
- “Support for a 64-bit root volume” on page 25
- “The Service Processor for the N6200 and N7x50T series storage systems” on page 25
- “Root aggregate Snapshot reserve values reduced” on page 26
- “Enhancements that allow you to specify SSL versions” on page 26
- “File space utilization report” on page 31
- “Upgrading to Data ONTAP 8.0.1 increases volume capacity after deduplication” on page 32
- “Support for FlexClone files and FlexClone LUNs on vFiler units” on page 32
- “Increased maximum RAID group size for SATA disks” on page 32
- “Support for the SMB 2.0 protocol” on page 35
- “TSO in NICs” on page 35
- “Support for CDP” on page 35
- “Port-based load-balancing option for multimode interface groups” on page 35
- “Support for Data Motion for Volumes” on page 38
- “VMware VAAI features for ESX hosts support” on page 39
- “igroup scalability limits” on page 39
- “Support for 840 disk drives on fabric-attached MetroCluster configuration” on page 41
- “Support for out-of-order frame delivery for SnapMirror over Fibre Channel” on page 41
- “Support for SnapMirror relationship between two FlexClone volumes” on page 41

- “Support for SnapMirror network compression” on page 41
- “Support for designated range of ports for NDMP data connections” on page 42

For information on all new and changed features in the Data ONTAP 8.0 7-mode release family, see “New and changed features” on page 17.

New commands in Data ONTAP 8.0.1 7-Mode

The following commands were introduced in Data ONTAP 8.0.1 7-Mode:

- `cdpd`

For more information, see “New commands in Data ONTAP 8.0 7-Mode” on page 48.

Changed commands in Data ONTAP 8.0.1 7-Mode

The following commands have been changed in Data ONTAP 8.0.1 7-Mode:

- `ifgrp create multi`
- `software update`
- `storage show acp`
- `storage show fault`

For more information, see “Changed commands in Data ONTAP 8.0 7-Mode” on page 52.

New options in Data ONTAP 8.0.1 7-Mode

Data ONTAP 8.0.1 7-Mode includes the following options:

- `autosupport.periodic.tx_window time`
- `ndmpd.data_port_range`
- `nfs.v3.snapshot.active.fsid.enable`
- `nfs.v4.snapshot.active.fsid.enable`
- `ssl.v2.enable`
- `ssl.v3.enable`

For more information, see “New options in Data ONTAP 8.0 7-Mode” on page 57.

Changed options in Data ONTAP 8.0.1 7-mode

Data ONTAP 8.0.1 includes the following changed option:

- `cifs.max_mpx`

For information about all new and changed command and options in the Data ONTAP 8.0 7-Mode release family, see “New and changed commands and options” on page 47.

New or changed cautions in Data ONTAP 8.0.1 7-Mode

The following new cautions were introduced in Data ONTAP 8.0.1 7-Mode:

- “Change in nondisruptive upgrade procedure” on page 68
- “SnapValidator for Oracle not supported” on page 81

The following caution changed in Data ONTAP 8.0.1 7-Mode RC3:

- “Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later” on page 70
- “SnapLock is not supported” on page 81

For more information, see “Important cautions” on page 65.

Cautions removed in Data ONTAP 8.0.1 7-Mode

- If you are upgrading to Data ONTAP 8.0 or later from Data ONTAP 7 and you use quotas
- If you have been using the flex_cache license
- If you are updating a system with AT-FCX disk shelves attached
- If you are upgrading a system that includes FlexVol volumes
- If you revert from Data ONTAP 8.0 7-Mode and attempt to access LUN data in a Snapshot copy
- Compression for SnapMirror transfers not supported
- SMB 2.0 is not supported
- Using ACP can cause system panic
- 512-GB and 256-GB Performance Acceleration modules not supported

Known limitations in Data ONTAP 8.0.1 7-Mode

The following new limitations were introduced in Data ONTAP 8.0.1 7-Mode:

- “Shutdown delay in N3400 systems when critical thresholds are crossed” on page 90
- “SSD Spare blocks consumed limit reported as N/A” on page 93
- “Failed CIFS connections due to signature mismatch” on page 101
- “Excessive pending CIFS authentication requests can cause service disruption” on page 102
- “Widelinks are not accessible from Mac clients” on page 103

- “Tape backup and volume move operations must not be performed simultaneously” on page 107
- “SMTape restore to an offline volume overwrites the data on that volume” on page 107
- “Deleted files occupy additional space on SMTape restored volumes” on page 107

For information on all known limitations in the Data ONTAP 8.0 release family, see “Known problems and limitations” on page 87.

Known issues and limitations removed in Data ONTAP 8.0.1 7-Mode

The following known issues and limitations were removed from Data ONTAP 8.0.1 7-Mode:

- The `-m` option for the `disk replace start` command is not recognized
- Attempting to assign more than 500 disks or array LUNs with a single `disk assign` command causes a system panic
- Qtree quotas might prevent explicit quotas from overriding the default
- If you establish a volume SnapMirror relationship between a source system with Data ONTAP 7.x and a destination system with Data ONTAP 8.0
- Down nodes in a cluster containing a FlexCache origin volume can prevent the FlexCache volume from connecting to the origin

Documentation changes for Data ONTAP 8.0.1 7-Mode

The following new or changed information supplements the documentation you received for the Data ONTAP 8.0.1 7-Mode release:

- “Changes to the System Administration Guide” on page 122
- “Changes to the File Access and Protocols Management Guide” on page 135
- “Changes to the Block Access Management Guide for iSCSI and FC” on page 138
- “Changes to the Data Protection Online Backup and Recovery Guide” on page 141
- “Changes to the Data Protection Tape Backup and Recovery Guide” on page 142
- “Changes to the MultiStore Management Guide” on page 144
- “Changes to the Storage Efficiency Management Guide” on page 147

For more information about important documentation changes, see “Changes to published documentation” on page 111.

New and changed features

This section covers features that were added or changed in the Data ONTAP 8.0 7-Mode release family.

Note: Some new and changed features in this release might also be introduced in a maintenance release of an earlier 7G release family. Before upgrading, be sure to consult with your IBM representative about new Data ONTAP functionality to determine the best solution for your business needs.

Data ONTAP 8.0 7-Mode and later Data ONTAP 8.0.x 7-Mode releases provide a new aggregate type, 64-bit aggregates, with a larger maximum aggregate and FlexVol volume sizes, as well as improved performance, resiliency, and management capabilities for storage resources.

For more information about 64-bit aggregates, see “Support for 64-bit aggregates” on page 31.

Data ONTAP 8.0 7-Mode and later Data ONTAP 8.0.x 7-Mode releases also provide new platform and hardware support as described in “New and changed platform and hardware support” on page 19, additional management capabilities such as SSL security improvements, additional capabilities using iSCSI and FC protocols for N series SAN environments, and enhanced data protection technologies such as SnapMirror and SnapVault.

Additional major changes made in Data ONTAP 8.0.4 7-Mode are described here. For other important information in these release notes about changes to the product made by these releases, see “New and changed commands and options” on page 47.

- “New and changed platform and hardware support” on page 19
- “Changes in storage expansion unit support” on page 23
- “Manageability enhancements” on page 24
- “Storage resource management enhancements” on page 31
- “Networking and Security protocols enhancements” on page 34
- “File access protocol enhancements” on page 37
- “Block protocols enhancements” on page 38
- “Data protection enhancements” on page 40
- “High-availability pair (formerly active/active configuration) enhancements” on page 33
- “MultiStore enhancements” on page 43

- “Changes to MultiStore support” on page 45

New and changed platform and hardware support

This Data ONTAP release supports new and changed hardware as listed and provides an overview of the new and changed hardware.

For information about supported systems, see “Supported systems” on page 61.

- “Support for 1-TB Flash Cache modules”
- “Support for N6200 series storage systems”
- “Support for the N7x50T series storage systems” on page 20
- “Support for 512-GB and 256-GB Performance Acceleration Module (PAM II)” on page 21
- “Support for SSDs” on page 21
- “Supported tape drives” on page 21
- “Support for 3-TB SATA disks” on page 21
- “Support for 2-TB SATA disks” on page 21

Support for 1-TB Flash Cache modules

Data ONTAP 8.0.2 and later releases support the 1-TB PCIe-attached Flash Cache module that uses Flash technology.

The addition of Flash Cache optimizes the performance of random read-intensive workloads such as file services, messaging, virtual infrastructure, and OLTP databases without using more high-performance disk drives.

For more information about Flash Cache modules supported with your storage system models, see the *IBM System Storage N series Introduction and Planning Guide*.

Support for N6200 series storage systems

Data ONTAP 8.0.1 and later releases support the N6200 series storage system.

The N6200 series storage systems support the following features:

- One or two controllers in a chassis
 - One-controller systems can be configured with a filler blank or an I/O expansion module, which provides four full-length PCIe slots.
- Two Gigabit Ethernet ports per controller
- Two HA ports per controller
- Two 3-Gb or 6-Gb SAS ports per controller
- Two 4-Gb Fibre Channel ports per controller
- One USB port per controller
- One private management (ACP) port per controller

- One remote management port per controller
- One full-length PCIe slot and one 3/4-length PCIe slot per controller
- Service Processor
- 4 GB, 8 GB, or 16 GB main memory per controller
- Support for EXN1000, EXN2000, EXN3000, EXN3500 and EXN4000 storage expansion units
- Maximum FlexVol volumes supported: N6210: 200; N6240 and N6270: 500
- Two on-board 10-GbE ports that are not available for configuration

Support for the N7x50T series storage systems

Data ONTAP 8.0.1 and later Data ONTAP releases support N7x50T series storage systems.

Architectural features of the N7x50T series platform include:

- PCIe v2.0 (Gen 2) x8 architecture
- Much higher slot count on the N7x50T series storage systems
 - Use as many Flash Cache cards, network adapters, and storage adapters as the solution requires with slots to spare.
- Latest 64-bit processing architecture
- Faster DDR3 memory and significant increase in the amount of memory on N7x50T
 - Available memory per controller is 96-GB.
- High-performance onboard 10-GbE and 8-Gb FC ports
 - 4 onboard 10-GbE ports with stateless offload per controller.
 - 4 onboard 8/4/2-Gb FC ports configurable as targets or initiators per controller.
- Reliability, availability, serviceability, and manageability (RASM) improvements
 - Persistent write log destages NVRAM to CompactFlash to safely protect write data not yet committed to disk during extended outages.
 - Service Processor (SP), the next-generation RLM technology, adds more RASM features and capabilities.
- I/O expansion module (IOXM)
 - Provides 10 additional slots to those in the controller.
 - Provides configuration flexibility for HA solutions in 6U and 12U footprints.

Support for 512-GB and 256-GB Performance Acceleration Module (PAM II)

Data ONTAP 8.0.1 7-Mode and later Data ONTAP 8.0.x 7-Mode releases support a PCIe-attached Performance Acceleration Module (PAM II also known as Flash Cache module) that uses Flash technology.

The addition of PAM II optimizes the performance of random read-intensive workloads such as file services, messaging, virtual infrastructure, and OLTP databases without using more high-performance disk drives.

For more information about PAM II modules supported with your storage system models, see *IBM System Storage N series Introduction and Planning Guide*.

Support for SSDs

Data ONTAP 8.0.1 7-Mode and later Data ONTAP 8.0.x releases support solid-state disks (SSDs). SSDs are flash memory-based storage devices that provide better overall performance than hard disk drives, or HDDs, which are mechanical devices using rotating media.

For more information about SSDs, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*. For information about which storage system models support SSDs, see *IBM System Storage N series Introduction and Planning Guide*.

Support for 900-GB SAS disks

In Data ONTAP 8.0.2 and later releases, 900-GB disks are supported for EXN3500 SAS storage expansion units. These are performance disks (rather than capacity disks) and their speed is 10K.

Supported tape drives

Data ONTAP 8.0.1 and later releases support Quantum LTO-4, Quantum LTO-5, Hewlett-Packard LTO-4, Hewlett-Packard LTO-5, and IBM LTO-5 tape drives.

Support for 3-TB SATA disks

In Data ONTAP 8.0.2 and later releases, 3-TB SATA disks are supported for the SAS disk connection type.

The 3-TB SATA disks have a right-sized capacity of 2,538 GB (where GB = 1,000*1,024*1,024), and 5,198,943,744 available blocks (sectors).

Support for 2-TB SATA disks

In Data ONTAP 8.0 and later releases, 2-TB SATA disks are supported for both the SAS and FC disk connection types. With these larger disks and 64-bit aggregates, you can create larger aggregates before being affected by spindle and aggregate size limits.

For more information about 64-bit aggregates, see “Support for 64-bit aggregates” on page 31.

Changes in storage expansion unit support

Data ONTAP 8.0 7-Mode release family includes some changes to storage expansion unit support.

- “Support for EXN3500 storage expansion units”
- “Support for EXN3000 storage expansion units”

Support for EXN3500 storage expansion units

The EXN3500 storage expansion unit contains external SAS disk drives and 6-Gb I/O modules.

EXN3500 storage expansion units are supported on the following Data ONTAP releases:

- Data ONTAP 8.x releases: Data ONTAP 8.0P1 or later; except N3400 storage systems.
- Data ONTAP 8.x releases: Data ONTAP 8.0.1 or later.

For more information about which disk drives and storage system models are supported with the EXN3500 storage expansion unit, see the *IBM System Storage N series Introduction and Planning Guide*. For more information about the EXN3500 storage expansion unit, see the *EXN3500 Storage Expansion Unit Installation and Service Guide*.

Support for EXN3000 storage expansion units

Data ONTAP 8.0.2P3 and later releases support EXN3000 storage expansion units. The EXN3000 storage expansion unit is also supported on Data ONTAP 7.3.6 and later releases.

For more information about which disk drives and storage system models are supported with the EXN3000 storage expansion unit, see the *IBM System Storage N series Introduction and Planning Guide*. For more information about the EXN3000 storage expansion unit, see the *EXN3000 Storage Expansion Unit Installation and Service Guide*.

Manageability enhancements

Data ONTAP releases provide additional management capabilities using MultiStore, FilerView, and other tools. This section provides an overview of these management capabilities.

- “Reversion and downgrade support”
- “SSH server version for Data ONTAP”
- “Supported OpenSSH client versions”
- “Support for a 64-bit root volume” on page 25
- “The Service Processor for the N6200 and N7x50T series storage systems” on page 25
- “Enhancements that allow you to specify SSL versions” on page 26
- “Root aggregate Snapshot reserve values reduced” on page 26
- “SSL security improvements” on page 27
- “Default security settings for manageability protocols” on page 27
- “Requirement to set root account password” on page 28
- “Uses of the systemshell and the diagnostic account” on page 28
- “Boot menu enhancements and changes” on page 28
- “Read reallocation of data” on page 29
- “Extents for FlexVol volumes” on page 29
- “Partner recipients for AutoSupport email notifications” on page 29
- “AutoSupport message file attachments” on page 30
- “New disk drive AutoSupport message sent weekly” on page 30

Reversion and downgrade support

Beginning in Data ONTAP 8.0.2 7-Mode, you can revert to a Data ONTAP release in an earlier release family or downgrade to an earlier Data ONTAP release in the same family. You can perform these operations without assistance from technical support when transitioning new or test systems to earlier releases.

However, you must call technical support if you have problems during or after upgrading. If you need to transition to an earlier release, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide* to help determine the correct action in your environment.

SSH server version for Data ONTAP

The SSH server version running on Data ONTAP is Data ONTAP SSH version 1.0.

Supported OpenSSH client versions

Data ONTAP 8.0.2 7-Mode and later releases support OpenSSH client version 4.4p1 only. To enhance security, OpenSSH client version 3.8p1 is no longer supported because it does not contain the latest security fix.

Support for a 64-bit root volume

Although your storage system is shipped with the root volume in a 32-bit aggregate, you can designate a different volume to be the new root volume. Starting in Data ONTAP 8.0.1, you can use a volume in a 64-bit aggregate for the root volume.

For more information about the root volume, see the *IBM System Storage N series Introduction and Planning Guide*. For more information about 64-bit aggregates and volumes, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

The Service Processor for the N6200 and N7x50T series storage systems

The Service Processor (SP) is a remote management device that is included in all system models except for the N3300, N3400, N3600, N5000 series, N6000 series, and N7000 series storage systems. The SP enables you to access, monitor, and troubleshoot the storage system remotely.

The SP provides the following capabilities:

- The SP enables you to access the storage system remotely to diagnose, shut down, power-cycle, or reboot the system, regardless of the state of the storage controller.

The SP is powered by a standby voltage, which is available as long as the system has input power to at least one of the system's power supplies.

The SP is connected to the system through the serial console. You can log in to the SP by using a Secure Shell client application from an administration host. You can then use the SP CLI to monitor and troubleshoot the system remotely. In addition, you can use the SP to access the system console and run Data ONTAP commands remotely.

You can access the SP from the system console or access the system console from the SP. The SP allows you to open both an SP CLI session and a separate system console session simultaneously.

- The SP monitors environmental sensors and logs system events to help you take timely and effective service actions in the event that a system problem occurs.

The SP monitors the system temperatures, voltages, currents, and fan speeds. When the SP detects that an environmental sensor has reached an abnormal condition, it logs the abnormal readings, notifies Data ONTAP of the issue, and takes proactive actions as necessary to send alerts and "down system" notifications through an AutoSupport message.

If SNMP is enabled for the SP, the SP generates SNMP traps to configured trap hosts for all "down system" events.

The SP also logs system events such as boot progress, Field Replaceable Unit (FRU) changes, Data ONTAP-generated events, and SP command history.

- Hardware-assisted takeover is available on systems that support the SP and have the SP configured.

For more information about the SP, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

For more information about hardware-assisted takeover, see the *Data ONTAP 8.0 7-Mode High-Availability Configuration Guide*.

You must ensure that your system has the latest SP firmware version. You can use the `sp status` command to display the SP firmware version on your system, and, if necessary, update the SP to the latest version that is available for download on the IBM N series support web site, which is accessed and navigated as described in “Websites” on page 149.

For instructions on how to download and update the SP firmware, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Enhancements that allow you to specify SSL versions

If your storage system has the SSL protocol enabled, you can further specify the SSL version to use by using the `ssl.v2.enable` and the `ssl.v3.enable` options.

Root aggregate Snapshot reserve values reduced

The root aggregate and root volume Snapshot reserve values for new storage systems have been reduced starting with Data ONTAP 8.0.1.

To increase the amount of usable space on the storage system while still retaining all the benefits of volume Snapshot copies, IBM ships all new systems running Data ONTAP 8.0.1 or later with the following new Snapshot reserve values:

- Snapshot reserve for the root aggregate is set to 0%.
The previous value was 5%.
- Snapshot reserve for the root volume is set to 5%.
The previous value was 20%.

No other aggregates or volumes are affected, nor is the default Snapshot reserve value for newly created aggregates or volumes.

Existing systems, and systems configured at the factory to run MetroCluster, are not affected by this change.

Attention: If you later configure MetroCluster, or use RAID SyncMirror or the aggregate copy capability on the root aggregate, the root aggregate Snapshot reserve *must* be increased to 5% by using the snap reserve command.

For more information about using the snap reserve command, see the `na_snap(1)` man page.

SSL security improvements

As a precautionary measure due to security vulnerability CVE-2009-3555, the SSL renegotiation feature is disabled in Data ONTAP.

Default security settings for manageability protocols

On storage systems shipped with Data ONTAP 8.0 7-Mode or later, secure protocols are enabled and non-secure protocols are disabled by default.

SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 7-Mode or later. For these systems, the following are the default security settings:

- Secure protocols (including SSH, SSL, and HTTPS) are enabled by default.
- Non-secure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default.

On storage systems shipped with Data ONTAP 8.0 7-Mode or later, the following are the default option settings for SSH and SSL:

- `options ssh.enable on`
- `options ssh2.enable on`
- `options ssh1.enable off`
- `options ssh.passwd_auth.enable on`
- `options ssh.pubkey_auth.enable on`
- `options httpd.admin.ssl.enable on`

Also on storage systems shipped with Data ONTAP 8.0 7-Mode or later, the following are the default option settings for the non-secure protocols:

- `options ftpd.enable off`
- `options httpd.admin.enable off`
- `options httpd.enable off`
- `options rsh.enable off`
- `options telnet.distinct.enable on`
- `options telnet.enable off`

Note: These default settings apply only to storage systems shipped with Data ONTAP 8.0 7-Mode or later. For storage systems upgraded from an earlier

version to Data ONTAP 8.0 7-Mode or later, the above default settings do not apply. Instead, for those upgraded systems, the settings remain unchanged after the upgrade. Also, if you make security setting modifications after upgrading to Data ONTAP 8.0 7-Mode or later, the modifications are preserved even if the system reverts back to the previous Data ONTAP version.

For more information about the default security settings, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Requirement to set root account password

During the initial setup of a storage system shipped with Data ONTAP 8.0 7-Mode or later, you are prompted to set up a password for the root account.

The following are the default password rules for all accounts when `security.passwd.rules.enable` is set to on (the default):

- The password must be at least eight characters long.
- The password must contain at least one number.
- The password must contain at least two alphabetic characters.
- The password must not contain the Ctrl-c or Ctrl-d key combination or the two-character string `^D`.

Note: Subsequent invocations of the setup command do not prompt you to set up a password for the root account.

For more information about setting up the storage system, see the *Data ONTAP 8.0 7-Mode Software Setup Guide*.

Uses of the systemshell and the diagnostic account

A diagnostic account, named “diag,” is provided with your storage system running Data ONTAP 8.0 7-Mode or later. You can enable the diagnostic account to perform troubleshooting tasks in the systemshell. The diagnostic account and the systemshell are intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

The diagnostic account is the only account that can be used to access the systemshell, through the advanced command `systemshell`. The diagnostic account is disabled by default. You must enable the account and set up its password before using it. Neither the diagnostic account nor the systemshell is intended for general administrative purposes.

Boot menu enhancements and changes

Starting with Data ONTAP 8.0, the boot menu includes new menu options that allow you to restore the configuration information from the root volume to the boot device, to install new software, or to reboot the storage system.

Data ONTAP 8.0 and later releases allow you to create a new root FlexVol volume but not a new traditional root volume from the boot menu. However, preexisting traditional root volumes are still supported.

For information about the boot menu, see the `na_floppyboot(1)` man page and the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Read reallocation of data

For workloads that perform a mixture of random writes and large and multiple sequential reads, the read reallocation function improves the file's layout and sequential read performance.

When you enable read reallocation, by using the `vol options read_realloc` command, Data ONTAP analyzes the parts of the file that are read sequentially. If the associated blocks are not already largely contiguous, Data ONTAP updates the file's layout by rewriting those blocks to another location on disk. The rewrite improves the file's layout, thus improving the sequential read performance the next time that section of the file is read.

For more information about the `vol options read_realloc` command, see the `na_vol(1)` man page and the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Extents for FlexVol volumes

Enabling extents, by setting the `vol options extent` command to `on`, might improve performance of Exchange database validation. However, to enable extents when storage space is a concern, you should set the `vol options extent` command to `space_optimized` instead of `on`.

The `space_optimized` option conserves space. However, it results in degraded read performance through the Snapshot copies. Therefore, if fast read performance through Snapshot copies is a higher priority to you than storage space, do not use the `space_optimized` option.

For more information about the `vol options extent` command, see the `na_vol(1)` man page and the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Partner recipients for AutoSupport email notifications

The new option `autosupport.partner.to` allows you to define the list of partner recipients for AutoSupport email notifications.

Whereas the `autosupport.support.to` option is read-only and displays where email-based AutoSupport notifications are sent to technical support when the `autosupport.support.enable` option is set to `on` (the default), the `autosupport.partner.to` option allows you to define a list of partner

recipients who will receive all AutoSupport notifications regardless of the severity level or the setting of the `autosupport.support.enable` option.

Using the `autosupport.partner.to` option, you can specify up to five email addresses for partner recipients to receive AutoSupport notifications. The email addresses should be entered as a comma-separated list with no spaces in between. By default, the `autosupport.partner.to` option is not defined and has an empty list.

To receive any email-based AutoSupport notifications, you must also define a mail host by using the `autosupport.mailhost` option.

AutoSupport message file attachments

In Data ONTAP 8.0 7-Mode and later releases, the contents of `/etc/messages` and `/etc/log/ems` are no longer sent in the body of AutoSupport messages. Instead, they are sent as `.gz` attachments to AutoSupport messages.

New disk drive AutoSupport message sent weekly

A weekly health trigger AutoSupport message provides information about any failed disk drives.

The message is sent each Sunday between 12:00 a.m. and 1:00 a.m. local time. If no drives failed during the past week, no weekly drive health test message is sent.

By default, the health test message is sent only to technical support.

Storage resource management enhancements

This Data ONTAP release provides improved performance, resiliency, and management capabilities for storage resources.

- “Support for 64-bit aggregates”
- “File space utilization report”
- “Upgrading to Data ONTAP 8.0.1 increases volume capacity after deduplication” on page 32
- “Maximum simultaneous FlexClone file or FlexClone LUN operations per storage system” on page 32
- “Support for FlexClone files and FlexClone LUNs on vFiler units” on page 32
- “Increased maximum RAID group size for SATA disks” on page 32
- “Ability to decrease maxfiles” on page 32

Support for 64-bit aggregates

In Data ONTAP 8.0 7-Mode and later releases, Data ONTAP supports a new type of aggregate, the 64-bit aggregate. 64-bit aggregates have a larger maximum size than 32-bit aggregates. In addition, volumes contained by 64-bit aggregates have a larger maximum size than volumes contained by 32-bit aggregates.

The maximum size of a 64-bit aggregate and the volumes it contains depends on the model of your storage system. To determine the maximum 64-bit aggregate size for your storage system model, see the *IBM System Storage N series Introduction and Planning Guide*. The maximum size of a volume contained by a 64-bit aggregate corresponds to the maximum size of its containing aggregate.

The maximum size for 32-bit aggregates remains 16 TB.

For more information, see the *Data ONTAP 8.0 7-Mode Storage Management Guide* and the `na_aggr(1)` man page.

File space utilization report

The file space utilization report enables you to see the files and the amount of space that they occupy in a deduplicated volume. You can choose to either move or delete the files to reclaim the space.

This report provides a view of the total number of blocks in a file and the number of blocks that are shared by non-deduplicated or non-cloned files.

Note: Total blocks refer to the number of blocks in a file, including blocks that are required for storing the file metadata.

Upgrading to Data ONTAP 8.0.1 increases volume capacity after deduplication

After you upgrade to Data ONTAP 8.0.1 and later from an earlier release of Data ONTAP, you can increase the deduplicated volume size up to 16 TB on all supported platform models that support Data ONTAP 8.0.1.

Maximum simultaneous FlexClone file or FlexClone LUN operations per storage system

Starting with Data ONTAP 8.0 7-Mode and later, you can simultaneously run a maximum of 500 FlexClone file or FlexClone LUN operations on a storage system.

For more information about FlexClone files and FlexClone LUNs, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

Support for FlexClone files and FlexClone LUNs on vFiler units

In Data ONTAP 8.0.1 and later release of the 8.0 release family, the FlexClone file and FlexClone LUN commands are available in the both the default vFiler context.

You can use FlexClone files and FlexClone LUNs feature to create writable, space-efficient clones of parent files and parent LUNs within a vFiler unit. Storage owned by a vFiler unit cannot be accessed or discovered from other vFiler units by using the FlexClone file or FlexClone LUN commands.

For more information about FlexClone files and FlexClone LUNs, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

Increased maximum RAID group size for SATA disks

Starting in Data ONTAP 8.0.1, the maximum RAID group size allowed for ATA, BSAS, and SATA disks has increased from 16 to 20 disks. The default size remains the same, at 14 disks.

Ability to decrease maxfiles

Starting in Data ONTAP 8.0, you can decrease the value of the maxfiles parameter for a volume, as long as you do not decrease it below the current number of files stored in the volume.

In previous versions of Data ONTAP, you could increase the value of the maxfiles parameter, but not decrease it.

High-availability pair (formerly active/active configuration) enhancements

This section provides an overview of High-availability pair related changes.

For more information about HA features, see the *Data ONTAP 8.0 7-Mode High-Availability Configuration Guide*.

- “High-availability pair terminology changes”

High-availability pair terminology changes

The high-availability (HA) functionality referred to as an HA pair in Data ONTAP 8.0 was called an active/active configuration in the Data ONTAP 7.2 and 7.3 release families. In the Data ONTAP 7.1 release family and earlier releases, an HA pair was called a cluster. HA functionality is a distinct, separate feature from Data ONTAP 8.0 Cluster-Mode clustering. For more information about these features, see the *Data ONTAP 8.0 7-Mode High-Availability Configuration Guide*.

Networking and Security protocols enhancements

This Data ONTAP release includes a number of new features and enhancements for networking and security protocol enhancements. This section provides an overview of these features and enhancements.

For more information about these features, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

- “Ability to block data traffic on the e0M interface”
- “Support for the SMB 2.0 protocol” on page 35
- “Support for CDP” on page 35
- “Port-based load-balancing option for multimode interface groups” on page 35
- “TSO in NICs” on page 35
- “64-bit SNMP object identifiers (OIDs) are supported” on page 36
- “Configuration for receiving untagged traffic on VLAN tagged interfaces” on page 36
- “Vifs are now known as interface groups” on page 36

Ability to block data traffic on the e0M interface

In Data ONTAP 8.0.2 and later releases, you can block certain types of traffic from the e0M interface, including SnapMirror transfers, SnapVault transfers, and data transfers that use the CIFS, NFS, and NDMP protocols. This feature enables you to optimize system performance.

Data ONTAP 8.0.2 introduces a new option `interface.blocked.mgmt_data_traffic` to control the blocking of data transfer on the e0M interface. New systems that ship with Data ONTAP 8.0.2 and later releases have the default value of this option set to `on`, which prevents data transfer on the e0M interface. Systems upgraded to Data ONTAP 8.0.2, 7.3.6, and later releases have the default value of the option set to `off`, which allows data transfer on the e0M interface.

Note: Data transfers that use the iSCSI protocol are always blocked on systems upgraded to Data ONTAP 8.0 and later. The data transfers are not affected by the option `interface.blocked.mgmt_data_traffic` option setting.

The e0M port is a low-bandwidth interface that should be used only for management traffic that uses SSH and other management protocols. Configuring e0M to serve data traffic can cause performance degradation and routing problems. Therefore, e0M should be configured on a dedicated management LAN or it should be configured down. If an e0M interface is serving management traffic, it should be partnered with another e0M interface.

It is a best practice to set the `interface.blocked.mgmt_data_traffic` option to on, and to use the e0M interface only for management traffic. For more information about the e0M interface, see the *Data ONTAP 8.0 7-Mode System Administration Guide* and the *Data ONTAP 8.0 7-Mode Software Setup Guide*.

Support for the SMB 2.0 protocol

Data ONTAP 8.0.1 and later support the SMB 2.0 protocol, which is more suitable than the original SMB protocol in environments requiring an increased level of scalability and data integrity.

For more information, see the *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide*.

Support for CDP

Cisco Discovery Protocol (CDP) is supported in Data ONTAP 7.3.3 and later releases of the 7.x release family, and in Data ONTAP 8.0.1 and later. CDP enables you to automatically discover and view information about directly connected CDP-enabled devices in a network.

For more information about CDP, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

Port-based load-balancing option for multimode interface groups

Port-based load balancing is supported for multimode interface groups in Data ONTAP 7.3.2 and later releases of the 7.x release family, and in Data ONTAP 8.0.1 and later. On a multimode interface group, you can uniformly distribute outgoing traffic based on the transport layer (TCP or UDP) ports and network layer addresses by using the port-based load-balancing method.

The port-based load-balancing method uses a fast hashing algorithm on the source and destination IP addresses along with the transport layer (TCP or UDP) port number.

For more information about port-based load balancing, see the *Data ONTAP 8.0 7-Mode Network Management Guide*

TSO in NICs

TCP segmentation offload (TSO) is a hardware feature supported on NICs to increase the host CPU's efficiency.

Starting from Data ONTAP 8.0.1, the network interface cards that support TSO are as follows:

- Dual 10G Ethernet Controller T320E-SFP+ and T320-XFP
- 10 Gigabit Ethernet Controller IX1-SFP+

Starting in Data ONTAP 8.0.2, the following network interface card also supports TSO:

- Dual 10G Ethernet Controller CNA-SFP+

For more information about TSO and the statistics displayed by the NICs, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

64-bit SNMP object identifiers (OIDs) are supported

In Data ONTAP 7.2 and 7.3 release families, some OIDs in the `netapp.mib` file have low and high 32-bit values. In Data ONTAP 8.0 7-Mode and later, these OIDs are replaced with 64-bit values. You should use SNMPv2 or SNMPv3 to query your storage system by using the 64-bit OIDs.

For example, the `miscHighNfsOps` and `miscLowNfsOps` OIDs have been replaced with the 64-bit OID `misc64NfsOps`. In the `netapp.mib` file, `miscHighNfsOps` and `miscLowNfsOps` are marked as deprecated.

Configuration for receiving untagged traffic on VLAN tagged interfaces

In Data ONTAP 8.0 7-Mode and later, you can configure an IP address for a network interface with VLANs. Any untagged traffic goes to the base interface, and the tagged traffic goes to the respective VLAN.

For more information about this feature, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

For information about reverting with a configuration for receiving tagged and untagged frames on the same network interface, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Vifs are now known as interface groups

Starting with Data ONTAP 8.0 7-Mode and later, vifs are known as interface groups. An interface group refers to a single virtual interface that is created by grouping together multiple physical interfaces. In the Data ONTAP 7.2 and 7.3 release families, this functionality is referred to as a vif.

You can use the `ifgrp` command to create and configure interface groups. For more information about interface groups and the `ifgrp` command, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

File access protocol enhancements

Data ONTAP 8.0 7-Mode and later releases provide new features of NFSv4 and FPolicy, new commands for debugging CIFS file access control problems, a new option for specifying the minimum levels of authentication and session security that clients must use, and compatibility with Windows Vista and Windows Server 2008.

For more information about these features, see the *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide*.

- “Maximum number of auxiliary UNIX groups supported for a user”
- “Support for multiple open instances of the SMB named pipe on an FPolicy server”

Maximum number of auxiliary UNIX groups supported for a user

If you use Kerberos V5 authentication, the maximum number of auxiliary UNIX groups that a user can be a member of is 32 by default. You can increase the maximum to 256 groups by setting the `nfs.max_num_aux_groups` option to 256.

If you do not use Kerberos V5 authentication, the maximum number of auxiliary UNIX groups that a user can be a member of is 16.

For more information about the `nfs.max_num_aux_groups` option, see the `na_options(1)` man page.

Support for multiple open instances of the SMB named pipe on an FPolicy server

In Data ONTAP 8.0 7-Mode, you can enable multiple open instances of the SMB named pipe on an FPolicy server by using the `fpolicy.multiple_pipes` option.

For more information, see the *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide*.

Block protocols enhancements

This Data ONTAP release includes a number of new features and enhancements for block protocol enhancements. This section provides an overview of these features and enhancements.

For more information about these features, see the *IBM System Storage N series Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

- “iSCSI RADIUS support”
- “Support for Data Motion for Volumes”
- “VMware VAAI features for ESX hosts support” on page 39
- “igroup scalability limits” on page 39

iSCSI RADIUS support

Data ONTAP 8.0 7-Mode introduces support for iSCSI RADIUS, which enables you to centrally manage iSCSI initiator authentication.

RADIUS still uses CHAP to authenticate iSCSI initiators, but it enables you to manage the authentication process from a central RADIUS server, rather than manage them manually on each storage system. In larger SAN environments, this can greatly simplify iSCSI initiator management and provide added security.

In addition to simplifying CHAP password management, RADIUS also reduces the load on your storage system. The RADIUS client service runs on the storage system and communicates with the RADIUS server and initiators, but most of the authentication processing is handled by the RADIUS server.

For more information, refer to the *Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

Support for Data Motion for Volumes

Data ONTAP 8.0.1 7-Mode or later supports IBM N series Data Motion for Volumes, which enables you to non-disruptively move a volume from one aggregate to another within the same controller for capacity utilization, improved performance, and to satisfy service-level agreements. In a SAN environment, FlexVol volumes and the LUNs in the volumes are moved non-disruptively from one aggregate to another.

The volume move occurs in three phases: setup phase, data copy phase, and cutover phase.

For more information about the volume move, see *Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

VMware VAAI features for ESX hosts support

Data ONTAP 8.0.1 and later supports certain VMware vStorage APIs for Array Integration (VAAI) features when the ESX host is running ESX 4.1 or later. These features help offload operations from the ESX host to the storage system and increase the network throughput.

The ESX host enables the features automatically in the correct environment. You can determine the extent to which your system is using the VAAI features by checking the statistics contained in the VAAI counters.

For more information about the VMware VAAI features for ESX hosts, see *Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

igroup scalability limits

Data ONTAP 8.0.1 introduces support to expand the host side scalability by expanding the number of igroups supported on N6000 series and N7000 series systems from 512 per controller to 1024 per controller.

Data protection enhancements

This Data ONTAP release includes a number of new features and enhancements for data protection enhancements. This section provides an overview of these features and enhancements.

For more information about these features, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide* and the *Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide*.

- “SnapVault update schedules for volume SnapMirror destination backups”
- “Support for 840 disk drives on fabric-attached MetroCluster configuration” on page 41
- “Support for out-of-order frame delivery for SnapMirror over Fibre Channel” on page 41
- “Support for SnapMirror relationship between two FlexClone volumes” on page 41
- “Support for SnapMirror network compression” on page 41
- “Support for designated range of ports for NDMP data connections” on page 42
- “If volumes with deduplication exceed the maximum supported size” on page 42
- “NearStore license is not required for deduplication” on page 42
- “Maximum concurrent deduplication operations supported” on page 42
- “Prerequisite for backing up data to tape using NDMP services” on page 42
- “Snapshot copy schedules” on page 43
- “SnapMirror support for 64-bit volumes” on page 43
- “SnapVault support for 64-bit volumes” on page 43

SnapVault update schedules for volume SnapMirror destination backups

When backing up a volume SnapMirror destination using SnapVault in Data ONTAP 8.0.2, the `snapvault.snapshot_for_dr_backup` option allows you to choose the SnapVault behavior for updating the destination system.

You can set the following values for the `snapvault.snapshot_for_dr_backup` option:

- `named_snapshot_only`: Schedules the SnapVault updates from the most recent Snapshot copy with the scheduled base name for scheduled backup.
- `named_snapshot_preferred`: Schedules the SnapVault updates from the most recent Snapshot copy created by volume SnapMirror if the scheduled Snapshot copy is not available.
- `vsm_base_only`: Schedules the SnapVault updates from the most recent Snapshot copy created by volume SnapMirror. This is the default value.

In Data ONTAP 8.0 and Data ONTAP 8.0.1 7-Mode, SnapVault uses a named Snapshot copy and not the Snapshot copy created by volume SnapMirror to update the destination system.

Support for 840 disk drives on fabric-attached MetroCluster configuration

In Data ONTAP 7.3.5 and later releases of 7.x release family, Data ONTAP 8.0.1 and later releases of 8.x release family, fabric-attached MetroCluster configurations can support up to 840 disk drives.

For more information, see the IBM N series interoperability matrix.

Related information

IBM N series interoperability matrix: www.ibm.com/systems/storage/network/interophome.html

Support for out-of-order frame delivery for SnapMirror over Fibre Channel

Data ONTAP 8.0.1 and later support out-of-order frame delivery for SnapMirror over Fibre Channel.

When enabled, out-of-order frame delivery for SnapMirror over Fibre Channel eliminates the requirement to deliver all the frames in the same sequence in which the frames were transmitted. This feature ensures the uninterrupted SnapMirror transfers regardless of the order in which the frames are delivered.

For more information about out-of-order frame delivery, see *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

Support for SnapMirror relationship between two FlexClone volumes

Data ONTAP 8.0.1 and later support SnapMirror relationship between two FlexClone volumes.

The SnapMirror relationship between two FlexClone volumes that have the common base Snapshot copy helps in establishing a SnapMirror relationship without transferring the common Snapshot data again to the destination system.

For more information about SnapMirror relationship between two FlexClone volumes, see *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

Support for SnapMirror network compression

Data ONTAP 7.3.2 and later releases of the 7.x release family, and Data ONTAP 8.0.1 and later releases of 8.x release family support SnapMirror network compression. This feature is supported only for asynchronous volume SnapMirror.

Note: The SnapMirror destination system should be using a Data ONTAP release that supports SnapMirror network compression.

SnapMirror network compression compresses the data stream on the source system, transfers the compressed data stream over the network, and uncompresses the data stream on the destination system before writing it to disk. It helps in optimizing network bandwidth utilization between SnapMirror source and destination systems. This feature can be especially useful for connections that have relatively low bandwidth, such as WAN connections.

For more information about SnapMirror network compression, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

Support for designated range of ports for NDMP data connections

Data ONTAP 8.0.1 and later releases supports designated range of ports that can be used for NDMP data connections in response to NDMP_DATA_LISTEN and NDMP_MOVER_LISTEN operations. Therefore, you can perform data migration by using the ndmcopy command and three-way tape backups even in environments where the source and destination networks are separated by a firewall.

If volumes with deduplication exceed the maximum supported size

Volumes continue to be online when they exceed the maximum supported size. For more information about the maximum volume size allowed for different storage systems, with and without deduplication, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

NearStore license is not required for deduplication

Starting with Data ONTAP 8.0.1 7-Mode, you do not have to enable the NearStore personality license to use deduplication.

Maximum concurrent deduplication operations supported

You can run a maximum of eight concurrent deduplication operations on all storage systems.

Prerequisite for backing up data to tape using NDMP services

Starting from Data ONTAP 8.0 7-Mode, NDMP users must have the login-ndmp capability for successfully authenticating NDMP sessions.

A predefined role named backup, by default, has the login-ndmp capability. To provide a user with the login-ndmp capability, the backup role can be assigned to the group to which the user belongs. However, when a group is assigned the backup role, all users within the group get the login-ndmp capability. Therefore, it is best to group all NDMP users in a single group that has the backup role.

Snapshot copy schedules

Scheduled Snapshot copy creation might fail for various reasons, such as a volume being offline. When the volume becomes online again, Data ONTAP detects the scheduled Snapshot copy creation that has failed and automatically creates a Snapshot copy for that schedule. This feature is available in Data ONTAP 8.0 7-Mode and later releases in the Data ONTAP 8.0 7-Mode release family.

For more information about this feature, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

SnapMirror support for 64-bit volumes

Data ONTAP 8.0 7-Mode and later releases support 64-bit volumes in addition to 32-bit volumes. SnapMirror supports replication between 64-bit volumes and also between 32-bit volumes. However, in Data ONTAP 8.0, there are certain considerations when replicating volumes and qtrees using SnapMirror.

When replicating volumes by using SnapMirror, synchronously or asynchronously, the source and destination volumes must be of the same type. Both the source and destination volumes must be 64-bit volumes or they must both be 32-bit volumes.

When replicating qtrees by using SnapMirror, the source and destination volumes can be of different types. Either the source and or destination volume can be 64-bit or 32-bit.

For more information about SnapMirror support for 64-bit volumes and see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

SnapVault support for 64-bit volumes

Data ONTAP 8.0 7-Mode and later releases in this release family support 64-bit volumes. SnapVault supports replication between 64-bit volumes, in addition to supporting replication between 32-bit volumes. When replicating qtrees by using SnapVault, the source and destination volumes can be of different types. The source and destination volumes can be either 64-bit or 32-bit.

For more information about SnapVault replication, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

MultiStore enhancements

This Data ONTAP release includes a number of new features and enhancements for MultiStore enhancements. This section provides an overview of these features and enhancements.

- “New option for the vfiler dr configure command” on page 44

New option for the vfiler dr configure command

Data ONTAP 8.0.3 and later provide a new option `-u` for the `vfiler dr configure` command, which can be used to set up a disaster recovery vFiler unit without reinitializing the existing SnapMirror relationship between the source and destination storage systems.

If there is no SnapMirror relationship set between the source and destination storage systems, and you run the `vfiler dr configure` command with the `-u` option, then the SnapMirror relationship is in the uninitialized state. If you run the `vfiler dr activate` command when the SnapMirror relationship is in the uninitialized state, the activation of the disaster recovery vFiler unit fails. To initialize the SnapMirror relationship between the source and destination storage systems, you must run the `vfiler dr resync` command on the destination storage system.

For more information about using this option, see the `na_vfiler(1)` man page.

Changes to MultiStore support

Data ONTAP provides MultiStore capabilities and changes.

- “Data Motion is not supported in the Data ONTAP 8.0 7-Mode release family”

Data Motion is not supported in the Data ONTAP 8.0 7-Mode release family

The Data Motion feature that enables users to migrate vFiler units between storage systems was introduced in a previous release, but it is not supported in Data ONTAP 8.0 7-Mode release family.

The Data Motion feature integrates virtual storage, mirroring, and provisioning software technologies so that you can perform data migrations non-disruptively in both physical and virtual environments.

New and changed commands and options

This section provides information about the commands, options, and configuration files that have been changed or added to the Data ONTAP 8.0 7-Mode release family. These listings are not exhaustive. For detailed information about specific commands and options, see the corresponding man pages.

These changes are described in the following topics:

- “New commands in Data ONTAP 8.0 7-Mode” on page 48
- “Changed commands in Data ONTAP 8.0 7-Mode” on page 52
- “Replaced or removed commands in Data ONTAP 8.0 7-Mode” on page 56
- “New options in Data ONTAP 8.0 7-Mode” on page 57
- “Changed options in Data ONTAP 8.0 7-Mode” on page 59

New commands in Data ONTAP 8.0 7-Mode

Many new commands have been added in the Data ONTAP 8.0 7-Mode release family.

For each command family and each command, the following table gives this information:

- The purpose of the command
- The location of documentation about the feature
- The Data ONTAP 8.0 7-Mode release in which the command was introduced

Command	Purpose	Documentation	Release introduced
acpadmin configure	Configures the ACP subsystem.	<i>Data ONTAP 8.0 7-Mode Storage Management Guide</i> na_acpadmin(1) man page	8.0
acpadmin list_all	Displays information about the ACP (Alternate Control Path) storage subsystem.	<i>Data ONTAP 8.0 7-Mode Storage Management Guide</i> na_acpadmin(1) man page	8.0
bmc battery-info	Displays battery information, for example, battery name, charging status, serial number, part number, firmware and hardware versions, chemistry type, manufacturer, and date of manufacture.	na_bmc(1) man page	8.0.4
cdpd	Displays information about devices that advertise themselves by using the CDPv1 protocol.	<i>Data ONTAP 8.0 7-Mode Network Management Guide</i>	8.0.1
du	Displays the files and amount of space that they occupy in a deduplicated volume.	<i>Data ONTAP 8.0 7-Mode Storage Management Guide</i>	8.0.1
ifgrp	Manages interface groups on your storage system. This command replaces the vif command. However, the options remain the same.	<i>Data ONTAP 8.0 7-Mode Network Management Guide</i>	8.0

Command	Purpose	Documentation	Release introduced
revert_to	Transitions Data ONTAP to a release in an earlier release family	<i>Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide</i> revert_to(1) man page	8.0
smtape abort	Aborts a backup or restore job initiated by the SMTape engine.	<i>Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide</i> na_smtape man page	8.0
smtape backup	Backs up blocks of data to tape.	<i>Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide</i> na_smtape man pages	8.0
smtape continue	Continues an SMTape-initiated backup or restore operation after it has reached the end of current tape and is in the wait state to write output to or accept input from a new tape.	<i>Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide</i> na_smtape man page	8.0
smtape restore	Restores blocks of data from tape.	<i>Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide</i> na_smtape man page	8.0
smtape status	Display the status of SMTape-initiated backup and restore operations	<i>Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide</i> na_smtape man page	8.0

Command	Purpose	Documentation	Release introduced
sp help	Displays the Data ONTAP sp commands that you can enter at the storage system prompt.	<i>Data ONTAP 8.0 7-Mode System Administration Guide</i> na_sp(1) man page	8.0.1
sp reboot	Reboots the SP and causes the SP to perform a self-test.	<i>Data ONTAP 8.0 7-Mode System Administration Guide</i> na_sp(1) man page	8.0.1
sp setup	Initiates the interactive SP setup script.	<i>Data ONTAP 8.0 7-Mode System Administration Guide</i> na_sp(1) man page	8.0.1
sp status	Displays the current status and the network configuration of the SP.	<i>Data ONTAP 8.0 7-Mode System Administration Guide</i> na_sp(1) man page	8.0.1
sp test autosupport	Sends a test email to all recipients specified with the autosupport.to option.	<i>Data ONTAP 8.0 7-Mode System Administration Guide</i> na_sp(1) man page	8.0.1
sp test snmp	Performs SNMP test on the SP, forcing the SP to send a test SNMP trap to all trap hosts specified in the snmp traphost command.	<i>Data ONTAP 8.0 7-Mode System Administration Guide</i> na_sp(1) man page	8.0.1
sp update	Updates the SP firmware.	<i>Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide</i> na_sp(1) man page	8.0.1
storage array modify	Modifies attributes of array profile records.	na_storage(1) man page	8.0
storage array purge-database	Removes all records from the controller's array profile database.	na_storage(1) man page	8.0
storage array remove	Removes records for the specified array from the controller's array profile database.	na_storage(1) man page	8.0

Command	Purpose	Documentation	Release introduced
storage array remove-port	Removes ports associated with an array profile.	na_storage(1) man page	8.0
storage array show	Lists all array profile records known to the controller.	na_storage(1) man page	8.0
storage array show-config	Provides a summary of the connectivity to SAN-attached arrays connected to the gateway system's FC initiator ports.	<i>Gateway Installation Requirements and Reference Guide</i> na_storage(1) man page	8.0
storage array show-luns array-name	Lists all array LUNs exported from a named storage array.	na_storage(1) man page	8.0
storage array show-ports	Lists all target ports for all storage arrays known to the gateway system.	na_storage(1) man page	8.0
storage show acp	Displays information about the ACP (Alternate Control Path) storage subsystem.	<i>Data ONTAP 8.0 7-Mode Storage Management Guide</i> na_acpadmin(1) man page	8.0

Changed commands in Data ONTAP 8.0 7-Mode

This table lists the commands that have been changed in the Data ONTAP 8.0 7-Mode release family.

For each command family and each command, the following table gives this information:

- The change in the command
- The location of documentation about the feature
- The Data ONTAP 8.0 7-Mode release in which the change was introduced

Command	Change	Documentation	Release command changed in
aggr create	The -B option has been added. This option is used to create a 64-bit aggregate.	na_aggr(1) man page <i>Data ONTAP 8.0 7-Mode Storage Management Guide</i>	8.0
df	The -x option has been added. This option suppresses the display of the .snapshot lines.	na_df(1) man page	8.0
disk assign	You can no longer use this command to change the ownership of a disk or array LUN that is already owned unless you are running the command on the system that owns the disk already.	na_disk(1) man page	8.0
environment shelf	The Complex Programmable Logic Device (CPLD) version has been added to the output.	na_environment(1) man page	8.0
flexcache fstat	The -inode-file option has been added. This option displays the number of blocks used by the FlexCache volume's inode file.	na_flexcache(1) man page	8.0.3

Command	Change	Documentation	Release command changed in
halt	The -s option no longer powers off the controller. Instead, it halts at the LOADER prompt. For systems that have FRU fault LEDs on the motherboard, Data ONTAP also clears the LEDs when subsequently booted.	na_halt(1) man page	8.0.4
ifgrp create multi	The port value (port-based load-balancing) for the -b option has been added.	<i>Data ONTAP 8.0 7-Mode Network Management Guide</i>	8.0.1
software update	The -r option, which suppresses automatic reboot, is now the default. To request an automatic reboot, you now must specify the -R option.	na_software(1) man page <i>Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide</i>	8.0.1
storage show acp	A column for <i>Module type</i> has been added to the output. This shows the type of the I/O module.	na_storage(1) man page	8.0.1
storage show expander	The system interconnect link (SIL) ports (the last four rows of output) are now shown as [SIL0] through [SIL3]. Previously they were shown as [DIS0] through [DIS3].	na_storage(1) man page	8.0
storage show expander	The EXN3000 and EXN3500 storage expansion unit ports are now listed as [SQR0]-[SQR3] for the square ports and [CIR4]-[CIR7] for the circle ports. Previously these ports were listed as [IO 0]-[IO7].	na_storage(1) man page	8.0

Command	Change	Documentation	Release command changed in
storage show fault	You can now specify a storage expansion unit name for this command to see information about a specific storage expansion unit.	na_storage(1) man page	8.0.1
sysconfig	When the -r option is used to display disk information for SAS-connected disks, the CHAN (channel) column now displays the port as "A" or "B", just as it does for FC-connected disks. Previously, "1" and "2" were used.	na_sysconfig(1) man page	8.0
sysconfig	When the -a option is used, the output includes status and power information for any IOXM modules present.	na_sysconfig(1) man page	8.0.1
sysconfig	When the -v option is used, any disabled ports show "N/A" for the data link rate.	na_sysconfig(1) man page	8.0.2
vif	This command has been replaced by the ifgrp command.	<i>Data ONTAP 8.0 7-Mode Network Management Guide</i>	8.0
vol options	The extent option enables extents but can only be used on FlexVol volumes.	na_vol(1) man page and the <i>Data ONTAP 8.0 7-Mode System Administration Guide</i>	8.0
vol options	The read_realloc option enables read reallocation but can only be used on FlexVol volumes.	na_vol(1) man page	8.0
vol status	When the -r option is used to display disk information for SAS-connected disks, the CHAN (channel) column now displays the port as A or B, just as it does for FC-connected disks. Previously, 1 and 2 were used.	na_vol(1) man page	8.0

Replaced or removed commands in Data ONTAP 8.0 7-Mode

The following information lists replaced or removed commands in Data ONTAP 8.0 7-Mode.

Support for the following commands will be discontinued in a future version of Data ONTAP. To perform the equivalent operation, use the `aggr` command.

- `vol add`
- `vol media_scrub`
- `vol migrate`
- `vol mirror`
- `vol scrub`
- `vol split`
- `vol verify`
- `vol wafiron`

These commands are removed in Data ONTAP 8.0 7-Mode.

- `snapmirror store`
- `snapmirror use`
- `snapmirror retrieve`

In Data ONTAP 8.0 7-Mode the `vif` command has been replaced by the `ifgrp` command. For more information, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

New options in Data ONTAP 8.0 7-Mode

Options that can be used with the `options` command are included in this table.

For each new option that can be used with the `options` command, the following table gives this information:

- A description of the option's purpose
- The default value or an example value used with the option
- The Data ONTAP 8.0 7-Mode release in which the option was introduced

For more information about the `options` command and individual options, see the `na_options(1)` man page.

Option	Purpose	Default value	Release introduced
<code>autosupport.performance_data.doit</code> <i>any_string</i>	Triggers a performance snapshot AutoSupport message when any string is added.	N/A	8.0
<code>autosupport.periodic.tx_window</code> <i>time</i>	Specifies the randomized delay window for periodic AutoSupport messages.	60 minutes	8.0.1
<code>cf.takeover.on_panic</code>	Triggers a takeover if the partner mode panics.	0n	8.0
<code>cf.takeover.on_reboot</code>	Triggers a takeover if the partner node reboots.	The default is on, unless FCP or iSCSI is licensed, in which case the default is off.	8.0
<code>cf.giveback.auto.delay.seconds</code>	Adjusts the giveback delay time for automatic giveback.	300 seconds	8.0
<code>lun_ic_alua_changed</code>	Disables ALUA State Change Unit Attention on interconnect up or down events.	off	8.0.1 and 8.0.3 Note: This option is not available in Data ONTAP 8.0.2.
<code>interface.blocked.mgmt_data_traffic</code>	Blocks or allows data traffic on the management interface, e0M	off for systems upgraded from an earlier release; on for new systems	8.0.2
<code>ndmpd.data_port_range</code>	Specifies a port range on which the NDMP server can listen for data connections.	all	8.0.1

Option	Purpose	Default value	Release introduced
nfs.always.deny.truncate	Controls whether NFSv2 and NFSv3 clients can truncate files in UNIX qtrees that are also opened from a CIFS client with DENY write permissions.	on	8.0.2
nfs.max_num_aux_groups	Specifies the maximum number of auxiliary UNIX groups that a user can be a member of.	32	8.0
nfc.rpcsec.trace	Enables EMS messages for troubleshooting Kerberos authentication and connectivity issues.	off	8.0.4
nfs.v3.snapshot.active.fsid.enable	Determines whether the FSID of objects in a Snapshot copy matches the FSID of the active file system for NFSv3.	on	8.0.1
nfs.v4.snapshot.active.fsid.enable	Determines whether the FSID of objects in a Snapshot copy matches the FSID of the active file system for NFSv4.	off	8.0.1
fpolicy.multiple_pipes	Enables multiple open instances of the SMB named pipe on an FPolicy ser	on	8.0
snapvault.snapshot_for_dr_backup	Enables you to specify the Snapshot copy to use for updating the destination system when backing up a volume SnapMirror destination using SnapVault.	vsm_base_only	8.0.2
sp.setup	Displays whether the SP has been configured.	(N/A)	8.0.1
ssl.v2.enable	Enables or disables SSLv2.	on	8.0.1
ssl.v3.enable	Enables or disables SSLv3.	on	8.0.1

Changed options in Data ONTAP 8.0 7-Mode

This table lists options that have changed or become obsolete.

For each option, the following table gives this information:

- The name of the changed option
- The nature of the change
- The Data ONTAP 8.0 7-Mode release in which the change was introduced

For more information about the `options` command and individual options, see the `na_options(1)` man page.

Command	Change	Release introduced
<code>autosupport.minimal.subject.id</code>	The default for this option has been changed to <i>systemid</i> for systems shipped with Data ONTAP 8.0.2 7-Mode or later.	8.0.2
<code>cifs.max_mpx</code>	The default for this option has been changed to 253.	8.0.2
<code>flexscale.max_io_qdepth</code>	This option has become obsolete.	8.0
<code>httpd.admin.enable</code>	The default for this option has been changed to <code>off</code> for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
<code>httpd.admin.ssl.enable</code>	The default for this option has been changed to <code>on</code> for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
<code>rsh.enable</code>	The default for this option has been changed to <code>off</code> for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
<code>ssh.enable</code>	The default for this option has been changed to <code>on</code> for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
<code>ssh2.enable</code>	The default for this option has been changed to <code>on</code> for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0

Command	Change	Release introduced
security.passwd.rules.everyone	The default for this option has been changed to on for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
security.passwd.rules.history	The default for this option has been changed to 6 for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
telnet.distinct.enable	The default for this option has been changed to on for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
telnet.enable	The default for this option has been changed to off for systems shipped with Data ONTAP 8.0 7-Mode or later.	8.0
timed.log	This option has become obsolete.	8.0
timed.max_skew	This option has been deprecated.	8.0
timed.proto	The rtc and the rdate parameters of this option have been deprecated. The default has been changed to ntp.	8.0
timed.sched	This option has been deprecated.	8.0
timed.window	This option has been deprecated.	8.0

Requirements for running Data ONTAP 8.0 7-Mode

You can verify whether you have the storage systems and firmware needed to run the software in the Data ONTAP 8.0 7-Mode release family.

To find out your storage system model's capacity and maximum volume size, see the *IBM System Storage N series Introduction and Planning Guide*.

Supported systems

You need one of the supported storage systems listed in this section to run a release in the Data ONTAP 8.0 7-Mode family.

The following models of storage systems are supported:

- N5300 and N5600 storage systems
- N3400 storage systems
- N6000 series systems
- N6200 series systems
- N7000 series systems
- N7x50T

All systems supported by Data ONTAP 8.0 7-Mode release family can be configured as a high-availability pair.

Unsupported systems

There are some storage systems that are not supported in Data ONTAP 8.0 7-Mode.

The following N series storage systems models are not supported in Data ONTAP 8.0 7-Mode and later Data ONTAP 8.0.x 7-Mode releases:

- N3220, N3240, N3300, and N3600 storage systems
- N5200 and N5500 storage systems

Maximum total capacity supported

For information about each storage system model's capacity and maximum volume size, see the *IBM System Storage N series Introduction and Planning Guide*.

Firmware

You need to confirm that you have the latest firmware for your storage system, disks, and storage expansion units.

The following storage system components have firmware that sometimes requires upgrading:

- Motherboard (also known as system or storage system firmware)
- Disk drives
- Disk storage expansion unit

Storage system firmware

It is best to upgrade to the latest version of system firmware for your storage system. For the latest firmware, you can go to the IBM N series support website, which is accessed and navigated as described in “Websites” on page 149.

Note: The latest system firmware is included with Data ONTAP upgrade packages. For more information, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Disk firmware

For information about the latest disk firmware, see the IBM N series support website, which is accessed and navigated as described in “Websites” on page 149.

Note: New disk firmware is sometimes included with Data ONTAP upgrade packages. For more information, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Disk storage expansion unit

For information about the latest disk storage expansion unit and ESH (Embedded Switched Hub) firmware, see the IBM N series support website, which is accessed and navigated as described in “Websites” on page 149.

Gateway requirements and support information

Not all Data ONTAP releases support the same features, configurations, storage system models, and storage array models for gateway systems. The Gateway Support Matrix is the final authority on supported configurations, storage array firmware and microcode versions, switches, and so on.

See the Gateway Support Matrix, located on the IBM N series support website, which is accessed and navigated as described in “Websites” on page 149, for complete information about supported gateway models and storage arrays, as well as supported microcode, license code, and storage array firmware versions. The Gateway Support Matrix also identifies switch connectivity and topology that can be used with gateway systems.

Important cautions

Before upgrading to this release of Data ONTAP, make sure that you read the following items to identify and resolve issues that might affect the operation of your storage systems.

- “Unsupported features” on page 66
- “Upgrade and revert cautions” on page 67
- “Block access protocol cautions” on page 79
- “Data protection cautions” on page 81
- “Manageability cautions” on page 73
- “Storage management cautions” on page 74
- “Network protocol cautions” on page 76
- “File access and protocols cautions” on page 77
- “New RLM firmware and its upgrade requirements” on page 85

Unsupported features

Some Data ONTAP features are not supported in the Data ONTAP 8.0 7-Mode release family.

Features not available in Data ONTAP 8.0 7-Mode release family

If you have configured any of the following features, do not upgrade to any release in the Data ONTAP 8.0 7-Mode release family:

- “IPv6 is not supported” on page 76
- “SnapValidator for Oracle not supported” on page 81
- “SnapLock is not supported” on page 81

Doing so might cause your system to halt or make storage data unavailable to your clients. If you require these features in your business environment, contact IBM N series support team to evaluate the best upgrade path.

The following features are not supported in the Data ONTAP 8.0. 7-Mode release family:

- “Data Motion is not supported in the Data ONTAP 8.0 7-Mode release family” on page 45
- “If you need to boot the storage system from a Data ONTAP image stored on a remote server” on page 73
- “TLS is not supported in the Data ONTAP 8.0 release family” on page 74
- “FTPS is not supported” on page 77
- “SFTP is not supported” on page 78
- “IPsec is not supported” on page 76
- “TOE is not supported” on page 76

The storage system will continue to function but without the services provided by these features.

Features discontinued in Data ONTAP 8.0 7-Mode release family

Beginning in Data ONTAP 8.0 7-Mode, the following features are discontinued:

- “Upgrade process changes with Data ONTAP 7-Mode software images” on page 69
- Hardware-based disk ownership
- Media kits
- SnapVault for NetBackup

Upgrade and revert cautions

If you are upgrading to this Data ONTAP release or reverting from this release to an earlier release, you should review these issues, and if any apply in your environment, take appropriate action.

For more information about procedures and planning, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

- “CIFS startup might take several minutes after cluster failover giveback”
- “Change in nondisruptive upgrade procedure” on page 68
- “Upgrade process changes with Data ONTAP 7-Mode software images” on page 69
- “iSCSI traffic not supported on e0M and e0P interfaces” on page 70
- “New FlexVol volume space requirements when upgrading from a release earlier than Data ONTAP 7.3” on page 70
 - “Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later” on page 70
- “If your root FlexVol volume does not meet the new size requirements” on page 71
- “Nondisruptive upgrades on systems with VMware ESX Server hosts” on page 71
- “Requirements for nondisruptive upgrades on systems with deduplicated volumes” on page 72
- “New reversion and downgrade procedures” on page 72

CIFS startup might take several minutes after cluster failover giveback

After a cluster failover giveback, the CIFS protocol might take several minutes to start up. If vFiler units are present, they might not be initialized for the duration of the CIFS startup time. The delay in CIFS startup is typically due to network infrastructure issues related to domain controllers and LDAP servers.

About this task

To ensure that the domain controllers, DNS servers, and LDAP servers are accessible, it is recommended that you run the following tests on the partner node for which giveback is to be initiated before initiating a giveback.

Note: If the storage system contains vFiler units, you must complete the procedure on all vFiler units running CIFS.

Procedure

1. Test domain controller connectivity and responsiveness by entering the following command: `cifs testdc`

If the command fails, reconnect to a domain controller by using the `cifs resetdc` command. The `cifs resetdc` command terminates existing domain controller connections and then discovers available domain controllers and establishes new connections.

2. Verify that the `cifs resetdc` command succeeded by entering the following command: `cifs domaininfo`
If the command fails, identify and fix the domain connectivity issue before proceeding.
3. Test Active Directory LDAP connectivity by entering the following command: `cifs adupdate`
If the command fails, identify and fix the Active Directory LDAP issues before proceeding.
4. Look up the SID of a domain user for the domain to which the CIFS server belongs by entering the following command: `wcc -s your_domain\user_account`
If the command fails, identify and fix the issue before proceeding.
5. Optional: If LDAP is configured for user mapping, run LDAP queries by entering the following commands:
`priv set advanced`
`getXXbyYY getpwbyname_r user_name`
`getXXbyYY getpwbyuid_r uid`
If the commands fail, identify and fix the LDAP issues before initiating giveback.

Results

After you verify domain controller, DNS server, and LDAP server accessibility, you can proceed with giveback.

Note: Successful execution of these commands does not guarantee that CIFS startup during giveback will occur quickly. The servers that are responding now might fail later to respond when giveback occurs.

Change in nondisruptive upgrade procedure

When you upgrade to this Data ONTAP release, you should be aware of changes in the software update command syntax and the requirement for boot environment access.

Beginning with 8.0.1, system firmware is updated automatically. It is no longer necessary to access the boot environment for upgrades from a Data ONTAP 8.0.x release to a later release.

Note: This change only refers to minor nondisruptive upgrades (within a release family). When upgrading to a Data ONTAP 8.0.x release from an earlier release family, you must enter the boot environment during the takeover phase to ensure that the Data ONTAP 8.0.x release is running the correct firmware release.

Beginning with 8.0.1, the `-r` option (no automatic reboot) is the default for the software update command.

Note: However, when you first upgrade to Data ONTAP 8.0.1 or a later release in the 8.0 family, you must specify the `-r` option for nondisruptive upgrades. In subsequent major and minor nondisruptive upgrades, it will no longer be required to specify the option.

Upgrade process changes with Data ONTAP 7-Mode software images

When you upgrade to Data ONTAP 8.0 and later releases, you must obtain Data ONTAP software images from the N series support website (accessed and navigated as described in “Websites” on page 149) and copy them to your storage system. Media kits, including CD-ROMs with Data ONTAP software images, are no longer available, and the process of installing upgrade images from UNIX or Windows clients is no longer supported.

Beginning with Data ONTAP 8.0, media kits including CD-ROMs are no longer available. You must obtain software images from the IBM N series support website (accessed and navigated as described in “Websites” on page 149) and copy them to your storage system or to an HTTP server in your environment.

Beginning with Data ONTAP 8.0, the following processes are no longer supported for extracting and installing Data ONTAP upgrade images:

- Using the `tar` command from UNIX clients
- Using the `setup.exe` file and WinZip from Windows clients

For the upgrade to Data ONTAP 8.0 and later releases, `.exe` images are no longer available. You must use one of the following image types depending on the upgrade you are performing:

- `.zip` images, for upgrades from an earlier release family to Data ONTAP 8.0
- `.tgz` images, for upgrades from any Data ONTAP 8.0 release to a later release

After you have upgraded to Data ONTAP 8.0 or later, you can only use `.tgz` images for further upgrades.

After you have copied the `.zip` or `.tgz` image to your storage system, you can install it with the software update command. Alternatively, you can make the

.zip or .tgz image available from an HTTP server in your environment, then use the software update command to copy the upgrade image from the HTTP server, extract and install the system files, and download the files to the boot device with one command.

For more information about upgrade images, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

iSCSI traffic not supported on e0M and e0P interfaces

If you are running iSCSI traffic and you are upgrading from Data ONTAP version 7.3 or earlier to version 8.0, you will no longer be able to run iSCSI traffic on the e0M and e0P interfaces.

In addition, if you revert from Data ONTAP version 8.0, the e0M and e0P interfaces remains disabled for iSCSI traffic by default, as follows:

- Target portal group configuration related to e0M and e0P may change.
- e0M and e0P may be removed from the iSCSI access list configuration.

New FlexVol volume space requirements when upgrading from a release earlier than Data ONTAP 7.3

Upgrading to Data ONTAP 8.0 or later from a release earlier than Data ONTAP 7.3 will cause existing FlexVol volumes to require more free space from their containing aggregate.

Data ONTAP 7.3 includes an improvement to free space accounting. As a result, existing FlexVol volumes reserve additional space, resulting in a loss of 0.5% of free space.

Attention: You must ensure that all systems upgrading to Data ONTAP 7.3 or later have at most 96 percent used space in their aggregates. If there is insufficient free space in an aggregate to satisfy the increased requirement from its FlexVol volumes, the space guarantee for one or more volumes in that aggregate could be disabled.

For more information, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later

If you suspect that your system has almost used all of its free space, or if you use thin provisioning, you should check the amount of space in use by each aggregate. If any aggregate is 97 percent full or more, *do not* proceed with the upgrade until you have used `aggrSpaceCheck` tools to determine your system capacity and plan your upgrade.

About this task

1. Check your system's capacity by entering the following command:
`df -A`

If the capacity field shows...	Then...
96% or less for all aggregates	You can proceed with your upgrade to Data ONTAP 7.3; no further action is required.
97% or more for any aggregate	Continue to plan your upgrade with the <code>aggrSpaceCheck</code> tool listed in Step 2.

2. Download the Aggregate Space Checker (`aggrSpaceCheck`) from the IBM N series support website, which is accessed and navigated as described in "Websites" on page 149.

After downloading the Aggregate Space Checker, use the tool's Reference Guide to install and use it to assess the free space requirements for your system in Data ONTAP 7.3 and later releases. If you do not have sufficient free space, the tool will recommend a course of action to ensure a successful upgrade.

After using these tools and completing the upgrade, make sure that your space guarantees are configured according to your requirements.

If your root FlexVol volume does not meet the new size requirements

The minimum required size for root FlexVol volumes has been increased for every system running Data ONTAP 8.0 release family or later. If your root FlexVol volume does not meet the new requirements, you should increase its size as soon as you complete the upgrade procedure.

The root volume must have enough space to contain system files, log files, and core files. If a system problem occurs, these files are needed to provide technical support.

For more information about root volumes and the root FlexVol volume requirements for your platform, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Nondisruptive upgrades on systems with VMware ESX Server hosts

Before performing a nondisruptive upgrade on storage systems exporting data over NFS to VMware ESX Server hosts, you should verify that your clients' NAS components are correctly configured. This verification will help ensure that VMware guest operating systems do not experience service disruption during the upgrade.

In particular, the following parameters should be set correctly:

- NFS Heartbeat parameters on the ESX Server
- Timeout values for SCSI disks on guest operating systems

It is also highly recommended that the file systems using virtual machine disk format (VMDK) on Windows be correctly aligned with the storage systems' WAFL file system.

For more information about verifying and updating these configurations, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

Requirements for nondisruptive upgrades on systems with deduplicated volumes

You can perform major and minor nondisruptive upgrades when deduplication is enabled provided that no more than 300 FlexVol volumes (or 200 FlexVol volumes on N3400 storage systems) have deduplication enabled, and that no deduplication operations are running during the Data ONTAP upgrade.

For information about deduplication, see the *Data ONTAP 8.0 7-Mode Storage Management Guide* and the `sis(1)` man page.

New reversion and downgrade procedures

If you need to transition to an earlier Data ONTAP release, be sure to check the new procedures in the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide* to help determine the correct action in your environment.

Manageability cautions

If you are a storage system administrator, you should familiarize yourself with these manageability issues.

For more information about these cautions, see the *Data ONTAP 8.0 7-Mode Systems Administration Guide*.

- “Certificates signed with MD5 could cause loss of connectivity to secure clients”
- “If you need to boot the storage system from a Data ONTAP image stored on a remote server”
- “TLS is not supported in the Data ONTAP 8.0 release family” on page 74
- “RLM over IPv6 not supported in the Data ONTAP 8.0 release family” on page 74

Certificates signed with MD5 could cause loss of connectivity to secure clients

To enhance security, starting with Data ONTAP 8.0.2, Data ONTAP uses the SHA256 message-digest algorithm to sign the contents of digital certificates (including certificate signing requests (CSRs) and root certificates) on the storage system. Use of the MD5 message-digest algorithm, which was used to sign CSRs and root certificates, is no longer officially supported.

Depending on the certificate depth verification, clients might need to use SHA256 to verify digital certificates presented by Data ONTAP 8.0.2 and later.

Data ONTAP 8.0 and 8.0.1 use the MD5 message-digest algorithm to sign digital certificates. Due to the CVE-2004-2761-IETF X.509 certificate MD5 signature collision vulnerability, and to minimize security risks when using a certificate signed with MD5, you should have the CSRs further signed by a certificate authority (CA) using SHA256 or SHA1.

If you need to boot the storage system from a Data ONTAP image stored on a remote server

In Data ONTAP 8.0 and later release families, netboot is not a supported function, unless you are restoring the Data ONTAP image on the boot device, such as a PC CompactFlash card. If you need to boot the storage system from a Data ONTAP image stored on a remote server, contact technical support.

For information about how to replace a PC CompactFlash card or boot device, or how to restore the Data ONTAP image on the card, see the *Hardware and Service Guide* that is applicable to your storage system model.

TLS is not supported in the Data ONTAP 8.0 release family

Transport Layer Security version 1.0 (TLSv1.0) was introduced in a previous Data ONTAP release, but it is not a supported function in the Data ONTAP 8.0 release family.

If you enabled TLS in a previous Data ONTAP release family, it is no longer available when you run the Data ONTAP 8.0 release family. The system uses SSL instead of TLS for communication.

RLM over IPv6 not supported in the Data ONTAP 8.0 release family

If you upgrade from the Data ONTAP 7.3 release family and your RLM is configured with IPv6, existing IPv6 configuration is automatically removed during the upgrade because IPv6 is not supported in the Data ONTAP 8.0 release family.

If your RLM configuration does not include static IPv4 or DHCP IPv4, you need to reconfigure the RLM by using the `rlm setup` command after upgrading to the Data ONTAP 8.0 release family.

For information about the RLM, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Storage management cautions

You should review these issues and take appropriate action before upgrading or reinstalling.

- “Disk failures might increase for some disk models after upgrading to this version of Data ONTAP”

Disk failures might increase for some disk models after upgrading to this version of Data ONTAP

Starting with Data ONTAP 8.0.4, the software install or upgrade includes new storage expansion unit firmware versions for EXN1000 and EXN3000 storage expansion units that provide enhanced disk error detection and prediction capabilities. Therefore, after upgrading to this version of Data ONTAP, you might experience an increase in the number of disk failures for certain storage expansion units and disk models.

Increased disk failures after upgrading to this version of Data ONTAP is expected behavior. You should follow standard best practices for spares when you upgrade, to ensure that sufficient spares are always available.

The following table shows the storage expansion unit and disk models that might show an increased failure rate after upgrading:

Storage expansion unit	Output from sysconfig -a	Capacity
EXN1000	X269_SMOOS01TSSX	1 TB
	X268_SMOOST75SSX	750 GB
	X267_SMOOST50SSX	500 GB
	X262_SMOOST25SSX	250 GB
EXN3000	X302_SMOOS01TSSM	1 TB
	X310_SMOOST50SSM	500 GB

Network protocol cautions

This information addresses network protocols cautions. Review the following section for important information regarding network management.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

- “System might panic if the network configuration entries in the `/etc/rc` file are not in the recommended order”
- “Reconfiguring the `losk` interface results in system outage”
- “IPv6 is not supported”
- “IPsec is not supported”
- “TOE is not supported”

System might panic if the network configuration entries in the `/etc/rc` file are not in the recommended order

You should ensure that the entries in the `/etc/rc` file are in the recommended order. The incorrect order can cause prolonged giveback or system reboot time, which might result in a system panic.

For more information about the ordering of entries in the `/etc/rc` file, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Reconfiguring the `losk` interface results in system outage

If you modify the default configuration of the `losk` interface, it can result in a system outage. If any IP address or routing configurations related to the `losk` interface exists in the `/etc/rc` file, you must remove these configurations to avoid undesirable behavior of the storage system.

IPv6 is not supported

IPv6 is not supported in Data ONTAP 8.0 and later. If you want to upgrade from a Data ONTAP version that supports IPv6 to Data ONTAP 8.0 or later, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

IPsec is not supported

IPsec is not supported in Data ONTAP 8.0 and later. If IPsec is enabled on your storage system, be aware that it will be disabled automatically during the upgrade to the Data ONTAP 8.0 release family.

TOE is not supported

In Data ONTAP 8.0 7-Mode and later Data ONTAP 8.0.x 7-Mode releases, the TOE functionality is not supported. There is no adverse effect on the performance of your storage system because TOE is disabled.

File access and protocols cautions

This information addresses file access and protocols cautions. Review the following section for important information regarding file access and protocols management.

For more information about these cautions, see the *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide*.

- “CIFS setup using ONTAPI library might lead to the removal of already existing user-created CIFS shares”
- “Storage system might panic when not specifying an export path”
- “FTPS is not supported”
- “SFTP is not supported” on page 78

CIFS setup using ONTAPI library might lead to the removal of already existing user-created CIFS shares

Due to a software bug in Data ONTAP, using the ONTAPI library to execute CIFS setup might lead to the removal of existing user-created shares. The default CIFS shares are not affected.

System Manager 2.x uses the ONTAPI library. To work around this issue, use the CLI to set up CIFS.

Storage system might panic when not specifying an export path

If you enter the CLI command `exportfs -uf` to remove all entries from the access cache and unexport file system paths, the storage system might panic if you do not specify an export path. It is valid to use the `-f` option without specifying an export path, but not the `-u` option.

To avoid this issue, do not use the `-u` and `-f` options at the same time. Use the `exportfs -f` command to remove all entries from the access cache. Use the `exportfs -ua` command to unexport all file system paths. Use the `exportfs -u path` command to unexport a specific file system path.

FTPS is not supported

File Transfer Protocol over Secure Socket Layer (FTPS) was introduced in a previous Data ONTAP release, but it is not supported in the Data ONTAP 8.0 release family.

If you enabled FTPS in a previous Data ONTAP release family, it is no longer available as an option for file access when you run the Data ONTAP 8.0 release family.

SFTP is not supported

Secure File Transfer Protocol (SFTP) was introduced in a previous Data ONTAP release, but it is not supported in the Data ONTAP 8.0 release family.

If you enabled SFTP in a previous Data ONTAP release family, it is no longer available as an option for file access when you run the Data ONTAP 8.0 release family.

Block access protocol cautions

If your storage system is deployed in a SAN environment, you should review these issues and take appropriate action before upgrading or reinstalling.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

- “Non-ASCII characters in the comment string of the lun setup command result in system panic”
- “Identical WWPNs can occur on both nodes in a cluster”
- “If you use Solaris 10 and the emlxs Fibre Channel Adapter driver”
- “cfmode support change in Data ONTAP 8.0”

Non-ASCII characters in the comment string of the lun setup command result in system panic

If you enter non-ASCII characters in the comment string while running the lun setup command, the system panics.

To avoid system panic, you must enter only ASCII characters in the comment string of the lun setup command.

Identical WWPNs can occur on both nodes in a cluster

If a system and its disks are disconnected from its partner and then connected to a different partner, it is possible (though rare) that both systems can end up with the same WWPN.

If this situation occurs, the systems can behave as if they are both primary systems or both secondary systems, causing each system to compute identical WWPNs for its target ports. As a result, hosts will experience problems trying to access these duplicate ports in the fabric.

If you use Solaris 10 and the emlxs Fibre Channel Adapter driver

If you are using Solaris 10 and the emlxs Fibre Channel Adapter driver, you need to upgrade to the 2.30j version driver revision or later before you upgrade to Data ONTAP 7.3 or later.

cfmode support change in Data ONTAP 8.0

For Fibre Channel storage area network (FC SAN) configurations in Data ONTAP 8.0 and later releases, only `single_image` cfmode (cluster failover mode) is supported. If you are upgrading high-availability FC SAN systems from an earlier release and they are configured for any other cfmode, you must migrate them to `single_image` mode *before* upgrading to Data ONTAP 8.0 or later.

For detailed instructions about migrating to `single_image` mode, see the [Changing the Cluster `cfmode` Setting in Fibre Channel SAN Configuration](#) article.

Data protection cautions

If your storage system is configured with licenses for data protection technologies, you should review these issues and take appropriate action before upgrading or reinstalling.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

- “SnapValidator for Oracle not supported”
- “SnapLock is not supported”

SnapValidator for Oracle not supported

SnapValidator for Oracle is not supported in Data ONTAP 8.0 7-Mode and later Data ONTAP 8.0.x 7-Mode releases. For more information, contact IBM technical support.

SnapLock is not supported

SnapLock is not supported in Data ONTAP 8.0 and in later releases in the 8.0.x family. Therefore, if you are running SnapLock on your current release of Data ONTAP, you should not upgrade to Data ONTAP 8.0 release family.

If you run the following commands with options related to SnapLock, they will not work and you will see a snaplock license is required message:

- `vol create -L`
- `aggr create -L`
- `date -c`
- `date -c initialize`

Download of Data ONTAP 8.0 or later releases in the 8.0.x family on storage systems with SnapLock volume

In some versions of Data ONTAP 7.2.x and Data ONTAP 7.3.x with SnapLock volumes, you might be able to download Data ONTAP 8.0 or later. However, Data ONTAP 8.0 does not support the SnapLock feature, therefore, your storage system will halt when you reboot it with Data ONTAP 8.0 release family.

Following is an example of an error message displayed on the storage system console:

```
Found SnapLock disk:v5.29 Use fcadmin device_map for shelf and
slot info This release does not support SnapLock.Remove the
SnapLock disks from this system.
Found SnapLock disk: v5.32 Use fcadmin device_map for shelf and slot info
This release does not support SnapLock.Remove the SnapLock disks from
this system This release does not support SnapLock.
Halting the system!!!
To recover-boot with a release that supports SnapLock or unplug the
SnapLock disks.
```

In such a scenario, contact IBM technical support immediately.

If you upgrade a storage system with the SnapLock license to Data ONTAP 8.0 or later releases in the 8.0.x family

You cannot add new SnapLock licenses in a storage system running on Data ONTAP 8.0 or later. Previously installed SnapLock licenses are retained, but are disabled. You can view the SnapLock license using the `license` command; however, all operations that require a SnapLock license fail. Therefore, you cannot create new SnapLock volumes or aggregates in Data ONTAP 8.0 or later.

A storage system might have the SnapLock license enabled, without any SnapLock volumes or aggregates. In such a case, if you upgrade the storage system to Data ONTAP 8.0 or later, the SnapLock license is disabled automatically and an error message `snaplock.unsupported.version` is raised while the storage system is booting up. The SnapLock license is enabled automatically when you boot with a Data ONTAP release that supports the SnapLock feature.

If you upgrade a storage system with SnapLock volumes and aggregates to Data ONTAP 8.0 or later releases in the 8.0.x family

If you upgrade a storage system containing SnapLock volumes and aggregates to Data ONTAP 8.0 or later, the storage system halts.

The storage system console displays a list of disks that contain the SnapLock volumes and aggregates. Following is an example of an error message on the storage system console:

```
Found SnapLock disk : v5.29 Use fcdm device_map for shelf and
slot info This release does not support SnapLock. Remove the
SnapLock disks from this system.
Found SnapLock disk : v5.32 Use fcdm device_map for shelf and slot info
This release does not support SnapLock. Remove the SnapLock disks from
this system This release does not support SnapLock.
Halting the system !!!
To recover - boot with a release that supports SnapLock or unplug the
SnapLock disks
```

If you reboot a storage system running on Data ONTAP 8.0 or later with SnapLock disks and non-SnapLock disks (from an earlier release), the storage system will halt.

If you connect disks that contains SnapLock aggregates to a storage system with Data ONTAP 8.0 or later releases in the 8.0.x family

If you upgrade the storage system to Data ONTAP 8.0 or later with disks containing SnapLock aggregates, the storage system halts in the early boot process.

If you connect disks containing SnapLock aggregates to a storage system running Data ONTAP 8.0 or later, the storage system remains online. However, the storage system lists these disks in the broken disk pool and displays the `snaplock.disk.on.unsupported.version` error message on the storage system console. The SnapLock disks in the broken disk pool are protected from data corruption.

Note: The SnapLock disks cannot be used in any other aggregate. When the SnapLock disks are moved into the broken disk pool, the ComplianceClock associated with these volumes will not get updated. This might result in ComplianceClock skew when you reattach the SnapLock disks to a storage system running a Data ONTAP release that supports the SnapLock feature.

If you reboot the storage system running on Data ONTAP 8.0 or later releases in the 8.0.x family after connecting disks with SnapLock aggregates

If you reboot the storage system running on Data ONTAP 8.0 or later after connecting disks containing SnapLock aggregates, the reboot is successful. However, you are not able to use these disks. The storage system lists these disks in the broken disk pool and displays the `snaplock.disk.on.unsupported.version` error message on the storage system console.

Recovering from a halt caused by SnapLock

If the storage system halts because you have SnapLock volumes on the system and you attempted to upgrade to Data ONTAP 8.0 or later, contact technical support immediately.

New RLM firmware and its upgrade requirements

RLM firmware versions 4.0 and later require a different layout on flash media. You must ensure that you are running the latest 3.1.x RLM firmware to enable the transition to the new layout, and then update to the 4.0 or later firmware.

Your RLM must be running the latest 3.1.x firmware to update to 4.0. If you are running a firmware version earlier than 3.1, you must first perform an intermediate update to the latest 3.1.x firmware, and then update from 3.1 to 4.0 in a separate operation.

Attention: Regardless of whether you update RLM firmware from the Data ONTAP CLI or the RLM CLI, *do not* update directly from a firmware version earlier than 3.1 to 4.0 or later. Doing so corrupts the RLM flash device.

If you are updating to version 4.0 or later from either the Data ONTAP CLI or the RLM CLI, you must run the `r1m update` command with the `-f` option for a full image update. Further updates do not require the `-f` option.

If you are updating RLM firmware from the RLM CLI, you can use the normal procedure.

For information about installing and updating the RLM firmware, see the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

For information about configuring the RLM, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Known problems and limitations

Some unexpected and potentially undesired behaviors after upgrading to this release and in some cases, workarounds to avoid these behaviors, have been identified.

- “Manageability issues” on page 88
- “Storage resource management issues” on page 91
- “High-availability pair issues” on page 94
- “Network protocol issues” on page 95
- “File access protocol issues” on page 96
- “Block access protocol issues” on page 104
- “Data protection issues” on page 106

Manageability issues

This section describes known issues and limitations with the management interface that affect your ability to manage the storage system. These issues might include problems with command behavior, command output, or error messages presented in the Data ONTAP CLI or web interface and problems with UNIX or operating system commands used to interface with your storage system.

- “Cached data in a Flash Cache module is cleared during system reboot and takeover and giveback”
- “The `wrfile` command might fail to capture all input during a cut-and-paste operation performed through the RLM or SP”
- “NTP version used to communicate with a newly configured NTP server defaults to v3” on page 89
- “RLM configuration is not saved if the RLM's Ethernet cable is not connected” on page 89
- “If your system uses `rtc` or `rdate` as the protocol for time synchronization” on page 89
- “Obsolete timed options” on page 89
- “Issues regarding the timed log” on page 90
- “Shutdown delay in N3400 systems when critical thresholds are crossed” on page 90

Cached data in a Flash Cache module is cleared during system reboot and takeover and giveback

Your storage system clears the cached data in a Flash Cache module during a system reboot and takeover and giveback. For example, the system clears the cache when you reboot to upgrade Data ONTAP.

Because there is no cached data in the Flash Cache module, the system serves initial read requests from disk, which results in decreased read performance during this period. The system repopulates the cache as it serves read requests.

The `wrfile` command might fail to capture all input during a cut-and-paste operation performed through the RLM or SP

If you access the system console from the RLM or SP and you perform a cut-and-paste operation on graphic input that involves using the `wrfile` command, the command might fail to capture all input.

This behavior is a result of the unusually high sustained data rate in the cut-and-paste operation. You can avoid this issue by using an SSH console session instead of accessing the system console through the RLM or SP.

NTP version used to communicate with a newly configured NTP server defaults to v3

Starting with Data ONTAP 8.0.2, the Network Time Protocol (NTP) version to be used for communicating with a newly configured NTP server defaults to v3 instead of v4. This change is to address situations where certain time servers support only NTP v3.

Releases prior to Data ONTAP 8.0 use Simple NTP v3 (SNTPv3) by default, while Data ONTAP 8.0 and 8.0.1 use NTP v4 by default. For Data ONTAP 8.0.2 and later releases, the NTP daemon continues to use the highest supported version (v4 in this case) to communicate with the time servers that were configured prior to Data ONTAP 8.0.2.

To reset the time servers to use NTP v3, you can use the options `timed.servers` command to reconfigure the list of servers.

RLM configuration is not saved if the RLM's Ethernet cable is not connected

If the RLM's Ethernet cable is not connected when you use the `rlm setup` command to configure the RLM, after connecting the cable you must rerun `rlm setup`.

If your system uses `rtc` or `rdate` as the protocol for time synchronization

Starting with Data ONTAP 8.0, the Network Time Protocol (NTP) protocol is the only supported protocol for time synchronization. The `rtc` and the `rdate` protocols of the `timed.proto` option are obsolete and no longer take effect after you upgrade to Data ONTAP 8.0 or later.

If your system does not already use NTP as the time-synchronization protocol, it will not keep accurate time after you upgrade it to Data ONTAP 8.0 or later. Problems can occur when the storage system clock is inaccurate.

If your system does not use NTP as the protocol for time synchronization, immediately after upgrading to Data ONTAP 8.0 or later, you must set `timed.proto` to `ntp`, set `timed.servers` to use time servers that support NTP, and ensure that `timed.enable` is set to `on`.

Note: After you set `timed.proto` to `ntp`, the setting remains in effect even if you revert back to a release prior to Data ONTAP 8.0.

For information about how to synchronize the system time, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Obsolete `timed` options

Starting with Data ONTAP 8.0, several `timed` options are obsolete although they remain visible in the CLI and can be modified.

The following `timed` options have no effect in Data ONTAP 8.0 or later:

- The `timed.log` option
- The `timed.max_skew` option
- The `timed.min_skew` option
- The `timed.sched` option
- The `timed.window` option
- The `rtc` and the `rdate` values of the `timed.proto` option

If you attempt to set these options when the system is running Data ONTAP 8.0 or later, they have no effect. However, these settings do take effect if the system is reverted back to a release prior to Data ONTAP 8.0.

Issues regarding the `timed log`

Starting with Data ONTAP 8.0, the NTP daemon automatically adjusts the storage system time to keep the system clock synchronized with the specified NTP server. However, time adjustments made by the NTP daemon are not logged even when the `timed.log` option is set to on.

Shutdown delay in N3400 systems when critical thresholds are crossed

In N3400 systems running Data ONTAP 8.0.1 and later releases in the 8.0.x family, after a critical temperature or voltage threshold has been crossed, there is a two-minute delay before the system shuts down. The purpose of the delay is to provide enough time for AutoSupport messages to be delivered.

There is no shutdown delay if three or more fans fail.

If the system is running a heavy load or if the network is slow, the AutoSupport message might not be delivered. However, EMS messages report that thresholds have been crossed, regardless of whether an AutoSupport message is delivered.

Storage resource management issues

If you have more than one size of disk installed in your storage system, if you have configured quotas or if you use FlexCache volumes you should familiarize yourself with these storage resource management issues.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

- “Writes to a LUN fail when the combined size of the LUN and the change log is greater than the volume size”
- “Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance”
- “After changing the port assigned to ACP, a reboot is required to regain use of the original port”
- “Moving disks to a system running an earlier version of Data ONTAP could cause disk identification issues”
- “Data ONTAP can silently select varying disk sizes or array LUN sizes when creating an aggregate” on page 92
 - “Adding disks when varying sized disks are in use” on page 92
 - “Discovering what disks Data ONTAP will automatically select” on page 93
- “SSD Spare blocks consumed limit reported as N/A” on page 93

Writes to a LUN fail when the combined size of the LUN and the change log is greater than the volume size

If the size of a LUN and the size of the change log when combined together is greater than the volume itself, then subsequent writes to the LUN might fail due to insufficient free space in the volume.

Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance

If you configure user mapping for a default user quota on a storage system for which CIFS is licensed but not configured, a later attempt to delete a qtree could result in significantly degraded performance. If you license CIFS, you must also configure it completely to avoid this issue.

After changing the port assigned to ACP, a reboot is required to regain use of the original port

If you reconfigure ACP to use a new port, the port that was configured for ACP before becomes available for use by another subsystem only after you reboot the storage system.

Moving disks to a system running an earlier version of Data ONTAP could cause disk identification issues

If you remove a disk from a storage system running Data ONTAP 7.3 or later and place it into a storage system running an earlier version of Data ONTAP,

the software disk ownership information will be incorrect for the earlier version and might prevent Data ONTAP from identifying or correctly assigning the disk.

The disk identification issues cannot be solved except by returning the disk to the original storage system and removing the software disk ownership information using the more recent version of Data ONTAP.

To avoid causing disk identification issues, always use the recommended procedure for removing a data disk, which includes removing the software disk ownership information from the disk before removing it.

For more information, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

Note: Gateway systems support disks starting with Data ONTAP 7.3. Therefore, you cannot move a disk from a gateway system to a gateway system running a version of Data ONTAP earlier than 7.3.

Data ONTAP can silently select varying disk sizes or array LUN sizes when creating an aggregate

If disks or array LUNs of varying size are present in your storage system, relying on automatic disk selection can result in unexpected sizes being selected by Data ONTAP.

When you create a new aggregate using automatic disk or array LUN selection, Data ONTAP selects disks or array LUNs based on various criteria including size, speed, and checksum type. To ensure that the disks or array LUNs selected are the size you want, specify the disk or array LUN size when creating an aggregate.

For more information, see the `na_aggr(1)` man page.

Adding disks when varying sized disks are in use

When you have disks of different sizes available, always specify the disk size when creating or adding to an aggregate.

To specify the disk size when you add a disk to an aggregate, enter the following command:

```
aggr add aggr_name num_disks@size
```

Example

For example, to add four 635-GB disks to `aggr1`, enter the following command:

```
aggr add aggr1 4@635G
```

Discovering what disks Data ONTAP will automatically select

To list which disks would be selected, use the `-n` option for the `aggr create` or `aggr add` command.

When you use the `-n` option, the disks that would be selected automatically, if you created an aggregate or added to an aggregate, are displayed. The operation is not performed. If the selected disks are not what you intended, you can specify a disk list when you enter the `aggr create` or `aggr add` command.

To list the disks that would be selected for an `aggr create` command, enter the following command:

```
aggr create aggr_name -n disk_list
```

Example

For example, to list the disks that would be used for the creation of the `newaggr` aggregate using eight automatically selected disks, enter the following command:

```
aggr create newaggr -n 8
```

SSD Spare blocks consumed limit reported as N/A

When Data ONTAP reports status information for SSDs, the Spare blocks consumed limit value is displayed as N/A. The limit for currently shipping SSDs is 90%.

High-availability pair issues

If you are using HA pairs, you might need to familiarize yourself with these issues.

For more information about these issues, see the *IBM System Storage N series Data ONTAP 8.0 7-Mode High-availability Configuration Guide*.

- “Storage system in an HA pair might get rebooted during a controller or NVRAM replacement”
- “If takeover is enabled for network interface failure, automatic giveback should be disabled”

Storage system in an HA pair might get rebooted during a controller or NVRAM replacement

During a controller or NVRAM replacement, if the ha-config chassis and ha-config controller values are set to non-ha, the storage system reboots after displaying the following error message: 0 disk found.

If you encounter this error message, you need to boot the storage system into Maintenance mode and set the ha-config chassis and ha-config controller values to ha by performing the following steps:

1. Run the following command to check the ha-config values:
ha-config show
2. Run the following command to change the value of the controller:
ha-config modify controller ha
3. Run the following command to change the chassis value:
ha-config modify chassis ha
4. Run the following command again to confirm that the ha-config values are set to ha:
ha-config show

If takeover is enabled for network interface failure, automatic giveback should be disabled

If you set the `cf.takeover.on_network_interface_failure` option to on to enable automatic failover on network interface failures, ensure that you have the `cf.giveback.auto.enable` option set to off.

If the `cf.giveback.auto.enable` option is set to on, a network interface failure could result in a continuous series of takeovers and givebacks until the network interface problem is resolved.

Network protocol issues

You should familiarize yourself with these network protocol issues in this release.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

- “Partner node might be unavailable during controller failover when an interface name is used as an alternate host name in the `/etc/hosts` file”
- “Change in routing configuration maximum value”
- “NIS group file entries should not exceed 1024 characters per line”

Partner node might be unavailable during controller failover when an interface name is used as an alternate host name in the `/etc/hosts` file

When the interface name used for configuring the partner interface is present in the `/etc/hosts` file, the IP address of the local node is substituted instead of the IP address of the partner node. This results in the partner node being unavailable during takeover.

To avoid this issue, you can use one of the following workarounds:

- Do not use interface names as alternate host names in the `/etc/hosts` file.
- Use the IP address of the partner node instead of the interface name while configuring the partner interface.

Note: You can use the IP address of the partner node while configuring the partner interface only when the interface belongs to the default IPspace.

Change in routing configuration maximum value

When you run the `tracert` command with the `-m` option, you must ensure that the maximum value of TTL does not exceed 255.

NIS group file entries should not exceed 1024 characters per line

Data ONTAP supports a maximum of 1024 characters per line in the `/etc/group` file on the NIS server. NIS group file entries with more than 1024 characters can result in login denials.

To avoid login denials, you must reorganize NIS group entries such that each line does not exceed 1024 characters.

File access protocol issues

If your storage systems provide CIFS, NFS, or FTP client services, you might need to familiarize yourself with file access protocol issues.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide*.

- “Value of the option `nfs.ifc.rcv.high` is unexpectedly reset” on page 97
- “Enabling new EMS messages to help troubleshoot Kerberos issues” on page 97
- “Microsoft Windows 7 clients fail to rename files and directories” on page 97
- “File copy operation on a widelink path inside a CIFS share might fail when using SMB 2.x” on page 98
- “Domain controller responds with an access denied error to an SMB 2.x tree connect request” on page 98
- “Controlling NFS requests from nonreserved ports” on page 98
- “Enabling or disabling NFSv2” on page 99
- “Error `NFS4_BADOWNER` when NFSv4 client passes UID as string” on page 99
- “Domain user unable to execute CLI commands” on page 100
- “CIFS shares with comment ending with backslash disappear” on page 100
- “Removing stale entries from the access cache” on page 100
- “CIFS authentication failure with Windows Server 2008 R2 read-only domain controllers” on page 100
- “Too many CIFS audit files can cause system performance degradation” on page 101
- “CIFS client authentication fails if host name exceeds 16 characters” on page 101
- “Failed CIFS connections due to signature mismatch” on page 101
- “Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance” on page 91
- “Excessive pending CIFS authentication requests can cause service disruption” on page 102
- “Client notification messages in Windows domains require NetBIOS” on page 102
- “Configuration issue for clients that mount NFS shares using a non-reserved port” on page 102
- “Widelinks are not accessible from Mac clients” on page 103
- “Empty CIFS/SMB2.x change notifications” on page 103

Value of the option `nfs.ifc.rcv.high` is unexpectedly reset

The option `nfs.ifc.rcv.high` value might unexpectedly get reset to a previous value if it is changed by the option `nfs.tcp.recvwindowsize`.

The option `nfs.ifc.rcv.high` controls the high-water mark after which NFS level flow control takes effect. This option is also controlled by the option `nfs.tcp.recvwindowsize`. Changing the option `nfs.tcp.recvwindowsize` automatically changes the value of the option `nfs.ifc.rcv.high`.

However, there are some scenarios in which the option `nfs.ifc.rcv.high` value gets reset to a previous value if it is changed by the option `nfs.tcp.recvwindowsize`. To avoid the issue, first configure the option `nfs.tcp.recvwindowsize`, then manually configure the option `nfs.ifc.rcv.high`.

Note: The value for the option `nfs.ifc.rcv.high` should be 1.5 times the value of the option `nfs.tcp.recvwindowsize`.

Enabling new EMS messages to help troubleshoot Kerberos issues

New EMS messages have been added to Data ONTAP to help troubleshoot Kerberos authentication and connectivity issues.

Enable the new EMS messages for Kerberos by entering the following command:

```
options nfs.rpcsec.trace on
```

After you finish troubleshooting, disable the EMS message for Kerberos by entering the following command:

```
options nfs.rpcsec.trace off
```

Microsoft Windows 7 clients fail to rename files and directories

When a Microsoft Windows 7 client accesses an SMB share with widelinks enabled, it can fail intermittently to rename files and directories on the mapped share. This can lead to a situation where attempts by Microsoft Office applications to save documents on the share fail.

This issue is caused by a bug in the Windows 7 client for which a hot fix is available from Microsoft. The bug is described and information about how to obtain the hot fix is provided in the Microsoft KB 2703610. This KB article also provides information about a workaround for this issue if you do not want to apply the hot fix.

File copy operation on a widelink path inside a CIFS share might fail when using SMB 2.x

If you try to copy a large file to a folder that is within a widelink path inside the root of the CIFS share and the size of the file being copied is greater than the available space at the root of share, the copy operation fails.

This issue occurs if you are using Windows Explorer for the copy operation on an SMB 2.x capable client with the SMB 2.x protocol enabled.

This happens because, when calculating whether there is available space for the copy operation to succeed, the Windows client refers to the root of the share when calculating available space. This leads to a condition where the operation fails if there is insufficient space at the root of the share, even if the folder inside the root of the share resides within a widelink to a different volume that has enough available space for the copy operation to succeed.

The following are possible workarounds for this issue:

- Copy the file using the DOS command prompt.
- Copy the file using the Windows PowerShell interface.
- Create a CIFS share mounted on the widelink point, map the share on the Windows client, and perform the copy operation.

Note: This issue does not occur if SMB 1.0 is used for the file copy operation.

Domain controller responds with an access denied error to an SMB 2.x tree connect request

A domain controller can respond with a 0xc0000022 error (Access Denied) to an SMB 2.x tree connect request from the storage server on IPC\$.

This issue is most likely due to a race condition between changing the `smb2.signing.required` option and creating a session setup message.

To avoid this issue, disable SMB 2.x on the storage system by entering the following command:

```
options cifs.smb2.client.enable off
```

Controlling NFS requests from nonreserved ports

You can reject NFS mount requests from nonreserved ports by enabling the `nfs.mount_rootonly` option. To reject all NFS requests from nonreserved ports, you can enable the `nfs.nfs_rootonly` option.

About this task

By default, the option `nfs.mount_rootonly` is on.

By default, the option `nfs.nfs_rootonly` is off.

These options do not apply to the NULL procedure.

Procedure

Perform one of the following actions:

If you want to...	Enter the command...
Allow NFS mount requests from nonreserved ports	<code>options nfs.mount_rootonly off</code>
Reject NFS mount requests from nonreserved ports	<code>options nfs.mount_rootonly on</code>
Allow all NFS requests from nonreserved ports	<code>options nfs.nfs_rootonly off</code>
Reject all NFS requests from nonreserved ports	<code>options nfs.nfs_rootonly on</code>

Enabling or disabling NFSv2

You can enable or disable NFSv2 by modifying the `nfs.v2.enable` option. This allows file access for clients using the NFSv2 protocol. By default, NFSv2 is enabled.

Procedure

Perform one of the following actions:

If you want to...	Enter the command...
Enable NFSv2	<code>options nfs.v2.enable on</code>
Disable NFSv2	<code>options nfs.v2.enable off</code>

Error NFS4_BADOWNER when NFSv4 client passes UID as string

NFSv4 specifies users and groups as strings instead of the 32-bit numeric values used by NFSv2 and v3. However, when an NFSv4 client passes the UID as a string instead of the format `user_name@domain_name`, the NFSv4 server returns the error NFSV4_BADOWNER when setting file attributes and user nobody (65534) when obtaining file attributes.

Domain user unable to execute CLI commands

If you log in to the storage system CLI as a domain user and enter the command `cifs gpupdate`, all subsequently entered CLI commands fail with the following error message:

```
[useradmin.unauthorized.user:warning] User domain_name\user_name denied access - missing required capability: 'cli-options'.
```

If you encounter this issue, log out, then log back in to the CLI.

CIFS shares with comment ending with backslash disappear

If you create a CIFS share with a comment that ends with a backslash (\), the CIFS share might disappear after rebooting the storage system or restarting CIFS.

To prevent this issue from occurring, do not end comments for CIFS shares with a backslash.

To recover CIFS shares that disappeared due to this issue, follow these steps:

1. Open the file `/etc/cifsconfig_share.cfg`.
2. Locate the command that originally created the CIFS shares that disappeared.
3. Execute the command again but either remove the trailing backslash or add a character such as a space after the backslash.

Removing stale entries from the access cache

If you are using the `exportfs -f` command to flush the access cache due to a changed reverse lookup DNS entry for a client, you must use the `-n` parameter to also flush the reverse lookup DNS cache. This prevents issues due to stale reverse lookup DNS cache entries.

For more information, see the `na_exportfs(1)` man page.

CIFS authentication failure with Windows Server 2008 R2 read-only domain controllers

The storage system might lose CIFS functionality when attempting to authenticate with a read-only domain controller (RODC) and the writable domain controller is not configured to replicate passwords for the storage system.

To avoid this, you must maintain at least one writable domain controller in the CIFS domain and perform one of the following steps:

- Configure the writable domain controller to allow the RODC to replicate passwords for the storage system.

For more information, see the article about Windows Server 2008 Password Replication Policy Administration at <http://technet.microsoft.com>

- Use the `cifs prefdc add` command to add the writable domain controller to the list of preferred domain controllers on the storage system.
For information about the `cifs prefdc add` command, see the *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide*.

Too many CIFS audit files can cause system performance degradation

If the option `cifs.audit.autosave.file.limit` is set to 0, it allows an unlimited number of log files to be saved in the `/etc/log` directory. However, too many log files in this directory can cause system performance degradation.

If you set this option to 0, you should regularly monitor the `/etc/log` directory and remove unnecessary log files to prevent performance issues.

CIFS client authentication fails if host name exceeds 16 characters

If a CIFS client with a NetBIOS name of more than 16 characters attempts to connect to the storage system, the authentication fails.

The storage system log records the following error message:

```
AUTH:SPNEGO-Could not unpack NTLMSSP Authenticate token.
```

The connecting client displays the following error message:

```
NT_STATUS_MORE_PROCESSING_REQUIRED.
```

To avoid this issue, you must ensure that the CIFS client name does not exceed 16 characters so that it complies with NetBIOS specifications.

Failed CIFS connections due to signature mismatch

If you enable SMB signing on an alive and active CIFS server, existing CIFS connections might fail due to signature mismatch.

To prevent this, use the `cifs terminate` command to shut down the CIFS server and ensure that all existing CIFS connections are terminated before enabling SMB signing. After enabling SMB signing, use the `cifs restart` command to restart the CIFS server.

Configuring user mapping for a default user quota with CIFS licensed but not configured could severely degrade performance

If you configure user mapping for a default user quota on a storage system for which CIFS is licensed but not configured, a later attempt to delete a qtree could result in significantly degraded performance. If you license CIFS, you must also configure it completely to avoid this issue.

Excessive pending CIFS authentication requests can cause service disruption

If domain controllers become unavailable to a vFiler unit, pending CIFS authentication requests accumulate on the storage system. If the issue is not resolved, this can prevent other vFiler units on the storage system from successfully completing CIFS authentication requests.

Because a vFiler unit shares system resources with other vFiler units that are on the storage system, pending CIFS authentication requests for a vFiler unit could affect resources available to other vFiler units.

Data ONTAP enables you to monitor pending CIFS authentication requests and take corrective action if needed. It generates SNMP traps in the following situations:

- The amount of pending CIFS authentication requests reaches 50 percent of total possible requests
- The pending authentication requests drop down under 10 percent, meaning within a normal range again

For information about SNMP, see the *Data ONTAP 8.0 7-Mode Network Management Guide*.

Client notification messages in Windows domains require NetBIOS

The Windows client notification feature used for client messaging, shutdown notices, and vscan alerts requires NetBIOS over TCP to be enabled in Data ONTAP.

Similarly, NetBIOS over TCP must be enabled on Windows clients and the Windows Messenger service must be running.

By default, the Windows Messenger service is disabled on Windows 2003 and Windows XP SP2 clients.

Configuration issue for clients that mount NFS shares using a non-reserved port

The `nfs.mount_rootonly` option must be set to off on a storage system that must support clients that mount NFS shares using a non-reserved port even when the user is logged in as root. Such clients include Hummingbird clients and Solaris NFS/IPv6 clients.

If the `nfs.mount_rootonly` option is set to on, Data ONTAP does not allow NFS clients that use nonreserved ports (that is, ports with numbers higher than 1023) to mount the NFS shares.

Widelinks are not accessible from Mac clients

When a user attempts to connect to a share using widelinks from a Mac OS X client, the attempt fails. Widelinks are not accessible from Mac OS X clients.

Empty CIFS/SMB2.x change notifications

CIFS/SMB2.x clients connected to the storage system might request to be notified about changes to a particular directory. In certain situations, the notifications that Data ONTAP returns are empty.

Data ONTAP collects all notifications that occur within 500 milliseconds after an event is generated to avoid subsequent change notify requests for every single such event. If the maximum capacity of the buffer that is allocated to hold these notifications is reached, Data ONTAP responds with an empty notification.

To work around this issue, you can decrease the time period for accumulating events. Decreasing the time period reduces the chance of maxing out the buffer and receiving empty change notifications.

To decrease the time period, enter the following command:

```
setflag smb_boxcar_expire_ms time
```

time is the length of the time period in milliseconds.

By reducing this time period, Data ONTAP returns change notification responses quicker and with fewer events per response. While this might not completely eliminate empty notifications, a properly reduced time period reduces them significantly.

Due to the reduced number of events returned with each change notification response, clients have to send more change notify requests to obtain all events.

Block access protocol issues

If your storage systems are part of a SAN environment, you might need to familiarize yourself with block access protocol issues in this release.

For more information about these issues, see the *IBM System Storage N series Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

- “vStorage priority setting is not persistent after a reboot”
- “Disks offline in Windows 2008 after a standard upgrade”
- “iSCSI Target HBA support”
- “If an iSCSI target HBA fails” on page 105
- “If you use the automatic sizing feature on thinly provisioned LUNs when snap reserve is set to a non-zero value” on page 105

vStorage priority setting is not persistent after a reboot

When you reboot your vStorage, the vstorage priority setting might change.

To verify the setting, run the following command in advanced mode: `vstorage show`.

If you need to change the vstorage priority setting use the following command: `vstorage set_priority [LOWER|EQUAL|AUTO]`.

Disks offline in Windows 2008 after a standard upgrade

During a standard upgrade to Data ONTAP 7.3.3 and later releases, LUNs are assigned new revision numbers. Windows Server 2008 software interprets the LUNs with new revision numbers as new disks and sets them offline; this status is shown in Windows 2008 management interfaces after the upgrade. Windows Server 2003 ignores the LUN revision number.

You can work around this problem using the nondisruptive upgrade method, which allows the LUNs to maintain their revision numbers. You can also bring the disks online after the upgrade using Windows disk management tools or SnapDrive functionality.

iSCSI Target HBA support

Some storage systems support the use of an iSCSI target HBA, which contains special network interfaces that offload part of the iSCSI protocol processing. You cannot combine these iSCSI hardware-accelerated interfaces with standard iSCSI storage system interfaces in the same target portal group.

If you attempt to combine these interfaces, an error message is displayed.

If an iSCSI target HBA fails

If your iSCSI target host bus adapter (HBA) fails, the iSCSI initiators connected to this HBA might time out, prompting the following error message on the storage system:

```
iSNAP.fw.Crashed
```

The timeout on the iSCSI initiators might be caused by HBA firmware defects, hardware failure, or device driver defects.

Note: If multipath network I/O (MPIO) is configured on the hosts, I/O should be re-routed as expected, thereby preventing any disruption in service.

If you use the automatic sizing feature on thinly provisioned LUNs when snap reserve is set to a non-zero value

Generally, before you thinly provision LUNs, you should set snap reserve to zero. However, there are rare exceptions that require you to set snap reserve to a value other than zero. In these instances, you must use the automatic sizing feature for thinly provisioned LUNs in FlexVol volumes to work properly.

Using the automatic sizing feature is required because the space from deleted Snapshot copies can only be used to fulfill Snapshot space requests. Furthermore, the automatic deletion process will not begin until the snap reserve value is exceeded.

Step

Enter the following command:

```
vol autosize vol_name [-m size] [-I size] on
```

-m *size* is the maximum size to which the volume will grow. Specify a size in k (KB), m (MB), g (GB), or t (TB).

-I *size* is the increment by which the volume's size increases. Specify a size in k (KB), m (MB), g (GB) or t (TB).

If the specified FlexVol volume is about to run out of free space and is smaller than its maximum size, and if there is space available in its containing aggregate, the FlexVol volume's size will increase by the specified increment.

Data protection issues

If you use data protection products that include Snapshot technology (such as SnapRestore, SnapVault, SnapMirror, and SnapManager), you might have to familiarize yourself with relevant data protection issues.

For more information about these issues, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide* and the *Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide*.

- “Tape devices behind DataFort are not detected unless a known SCSI device is mapped to LUN 0”
- “Tape backup and volume move operations must not be performed simultaneously” on page 107
- “SMTape restore to an offline volume overwrites the data on that volume” on page 107
- “Deleted files occupy additional space on SMTape restored volumes” on page 107
- “Limitations when using SnapVault with non-default vFiler units” on page 107
- “Restore engine does not honor any mandatory or advisory locks on files” on page 108
- “Restore operation to a non-existent volume restores data to the root volume” on page 109
- “Suboptimal space savings with deduplication when 16-GB Performance Acceleration Module is installed” on page 109
- “If you use DataFabric Manager server to manage SnapMirror or SnapVault relationships between vFiler units” on page 109

Tape devices behind DataFort are not detected unless a known SCSI device is mapped to LUN 0

If a known SCSI device is not mapped to LUN 0 when using a DataFort appliance, the tape drives and media changers behind the DataFort appliance are not detected.

To avoid this issue, you must ensure that the following actions have been performed:

- When configuring a tape library, you must assign LUN 0 to SCSI devices such as, tape drives, media changers, and SCC devices for them to be detected and accessed by the storage system.
- When connecting a tape library through a DataFort appliance, the tape drive, media changer, or SCC device assigned to LUN 0 must also be mapped to a DataFort Cryptainer. This enables the storage system to detect LUN 0 and also the SCSI devices attached to it.

Tape backup and volume move operations must not be performed simultaneously

You must not perform a tape backup or restore operation while a volume move operation is in cutover phase. Similarly, you must not perform a volume move operation while a tape backup or restore is in progress. You should wait until one of the operations is complete before initiating the other.

For more information about volume move operation, see *Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

SMTape restore to an offline volume overwrites the data on that volume

If you perform an SMTape restore to an offline volume, the data on tape overwrites the data on that offline volume.

Deleted files occupy additional space on SMTape restored volumes

On an SMTape restored volume, a Snapshot copy might lock the data blocks of a deleted file, thus occupying additional space on the restored volume.

When you delete a large file from a volume, the data blocks of that file are not freed immediately. If you perform an SMTape backup immediately after deleting the file, some data blocks of the deleted file are also backed up. After restore, those data blocks reside on the volume and are locked in the Snapshot copy.

To free the space on the restored volume, delete the Snapshot copy that locked the data blocks of the deleted file.

Limitations when using SnapVault with non-default vFiler units

There are certain limitations when using SnapVault with non-default vFiler units. These limitations apply to all Data ONTAP releases that support non-default vFiler units.

The management of SnapVault secondary (creation or modification of SnapVault relationships and schedules at the SnapVault secondary) is only supported from the default vFiler unit (vfiler0). The management of SnapVault secondary is not supported from a non-default vFiler context. If the volume containing the SnapVault destination qtree is owned by a non-default vFiler unit, the SnapVault secondary needs to be managed through the default vFiler unit (vfiler0).

Table 1. vFiler unit support with SnapVault secondary volumes for different combinations

Management of SnapVault secondary volume	Ownership of SnapVault secondary volume	
	Default vFiler unit (vfiler0)	Non-default vFiler unit

Table 1. vFiler unit support with SnapVault secondary volumes for different combinations (continued)

Management of SnapVault secondary volume	Ownership of SnapVault secondary volume	
	Default vFiler unit (vfiler0)	Yes
Non-default vFiler unit	No	No

The management of SnapVault primary in a vfiler context is supported.

Table 2. vFiler unit support with SnapVault primary volumes for different combinations

Management of SnapVault primary volume	Ownership of SnapVault primary volume	
	Default vFiler unit (vFiler0)	Non-default vFiler unit
Default vFiler unit (vFiler0)	Yes	Yes
Non-default vFiler unit	No	Yes (From a non-default vFiler context, you can only manage volumes owned by that non-default vFiler unit.)

DataFabric Manager server support for the management of SnapVault relationships

DataFabric Manager server supports the management of SnapVault relationships for volumes through the default vFiler (vFiler0) context only. When using DataFabric Manager server, the following limitations apply for SnapVault relationships involving non-default vFiler units.

- You can only view SnapVault relationships configured through the default vFiler unit (vfiler0). You cannot view any SnapVault relationships configured through non-default vFiler units.
- You can configure new SnapVault relationships for a volume only through the default vFiler unit (vfiler0), even if the volume belongs to a non-default vFiler unit.

Restore engine does not honor any mandatory or advisory locks on files

The destination volume for a restore operation might have files with mandatory or advisory locks. When you initiate a restore operation to such a volume, the restore engine does not honor these locks. It ignores these locks and overwrites the files.

Restore operation to a non-existent volume restores data to the root volume

If the destination of your restore operation is a non-existent volume, data is restored to the root volume.

Suboptimal space savings with deduplication when 16-GB Performance Acceleration Module is installed

When you run a deduplication scan by using the `sis start -s` command on a volume with existing data and with the 16-GB Performance Acceleration Module installed, you obtain suboptimal space savings.

If you use DataFabric Manager server to manage SnapMirror or SnapVault relationships between vFiler units

For Data ONTAP 7.2 and later, SnapMirror and SnapVault relationships can be created using vFiler units. However, DataFabric Manager server 3.4 and earlier releases cannot use vFiler units for managing SnapMirror and SnapVault relationships.

As a result, you might encounter the following issues:

- If the `snapvault.access` and `snapmirror.access` options on the source system allow access only to the destination vFiler unit, then the relationship creation, scheduled backups, on-demand backups, SnapMirror updates, and SnapMirror resync processes from DataFabric Manager server fail, and you receive an error message: request denied by source filer, check access permissions on source

Workaround: To allow access to the destination hosting storage system, set the `snapmirror.access` and `snapvault.access` options on the source system.

- If the `ndmpd.preferred_interfaces` option is not specified on the source hosting system, then the backups from DataFabric Manager server might not use the correct network interface.

Workaround: Enable the `ndmpd.preferred_interfaces` option on the source hosting system.

- The backups and SnapMirror updates from DataFabric Manager server fail and you receive the error message source unknown. This occurs when both of these conditions are met:
 - A relationship between two vFiler units is imported into DataFabric Manager server by auto-discovery or is added manually.
 - The destination hosting system is not able to contact the source vFiler IP address.

Workaround: Ensure that the host name or IP address of the source system that is used to create relationships can be reached from the destination hosting system.

Changes to published documentation

Some information about this release has become available after the set of guides provided with this release were published. The information should be used in conjunction with the guides provided with this release of Data ONTAP.

- “Changes to the Data ONTAP Upgrade and Revert/Downgrade Guide” on page 112
- “Changes to the System Administration Guide” on page 122
- “Changes to the Storage Management Guide” on page 128
- “Changes to the Software Setup Guide” on page 132
- “Changes to the High-Availability Configuration Guide” on page 132
- “Changes to the Network Management Guide” on page 133
- “Changes to the File Access and Protocols Management Guide” on page 135
- “Changes to the Block Access Management Guide for iSCSI and FC” on page 138
- “Changes to the Data Protection Online Backup and Recovery Guide” on page 141
- “Changes to the Data Protection Tape Backup and Recovery Guide” on page 142
- “Changes to the MultiStore Management Guide” on page 144
- “Changes to the Storage Efficiency Management Guide” on page 147
- “Changes to the Gateway Implementation Guide” on page 147
- “Changes to the Gateway MetroCluster Guide” on page 148

Changes to the Data ONTAP Upgrade and Revert/Downgrade Guide

New information about upgrade procedures has become available since the previous revision of the *Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide*.

- “Correction to the information in the Planning your upgrade with Upgrade Advisor topic”
- “Updates to software image installation examples”
- “Updated NDU requirements” on page 118
- “Additional preparation for nondisruptive upgrades” on page 118
- “Additional SAN upgrade requirement” on page 120
- “Additional requirement for AT-FCX storage expansion unit firmware updates” on page 120
- “AT-FC and AT-FC2 shelves no longer supported” on page 121

Correction to the information in the Planning your upgrade with Upgrade Advisor topic

The note in the "Planning your upgrade" topic in the Upgrade and Revert Guide contains misleading information.

The note has been corrected as follows:

It is a best practice to use Upgrade Advisor to plan your upgrade. Nonetheless, you might find useful detail and related information in the Upgrade and Revert guide that complements your Upgrade Advisor plan.

If you are not able to use Upgrade Advisor, you should create your upgrade plan manually, by using the guidelines provided in the Upgrade and Revert guide.

Updates to software image installation examples

The following topics provide corrected sample commands for installing images on Data ONTAP 7.3.x systems.

Installing software images from an HTTP server

To install software images, you must know the URL of an HTTP server in your environment that is configured to serve software images.

Step

1. From the storage system prompt, enter the following command: `software update url options`
 - *url* is the URL of the HTTP server and subdirectory.
 - *options* is one or more of the following:

- The `-d` option prevents the **download** command from being run automatically after the system files are installed.
- The `-f` option overwrites the existing image in the `/etc/software` directory.
- The `-r` option prevents the system from rebooting automatically after the **download** command has finished (default).
- The `-R` option causes the system to reboot automatically after the **download** command has finished.

Attention: Beginning in Data ONTAP 8.0.1 or later in the 8.0.x release family, the software update options have changed; the `-r` option (no automatic reboot) is the default, and the `-R` option must be specified to override the `-r` option. However, if you are upgrading from any release earlier than Data ONTAP 8.0.1, you must include the `-r` option to prevent automatic reboot if you are performing a nondisruptive upgrade or if you are upgrading firmware. For more information, see the `software(1)` man page for the Data ONTAP version currently running on your system.

Example

You can use the following commands to copy and install the Data ONTAP software image:

If you are running Data ONTAP..	And you want to...	Then you can enter...
7.3.5 or later 7.3.x release	Copy and install the image from your HTTP server	software update http://www.example.com/downloads/x86-64/my_new_setup_i.zip -d
	Copy from your HTTP server and overwrite an existing image	software update http://www.example.com/downloads/x86-64/my_new_setup_i.zip -d -f
	Copy and install the image from your HTTP server, then download the new system files to the boot device immediately after installing them	software update http://www.example.com/downloads/x86-64/my_new_setup_i.zip
	Copy and install the image from your HTTP server to a single system, then download the new system files and reboot immediately	software update http://www.example.com/downloads/x86-64/my_new_setup_i.zip -R

If you are running Data ONTAP..	And you want to...	Then you can enter...
7.3.4 or earlier 7.3.x release	Copy and install the image from your HTTP server	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.zip -d -r
	Copy from your HTTP server and overwrite an existing image	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.zip -d -r -f
	Copy and install the image from your HTTP server, then download the new system files to the boot device immediately after installing them	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.zip -r
	Copy and install the image from your HTTP server to a single system, then download the new system files and reboot immediately	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.zip
8.0.1 or later	Copy and install the image from your HTTP server	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.tgz -d
	Copy from your HTTP server and overwrite an existing image	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.tgz -d -f
	Copy and install the image from your HTTP server, then download the new system files to the boot device immediately after installing them	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.tgz
	Copy and install the image from your HTTP server to a single system, then download the new system files and reboot immediately	software update http://www.example.com/ downloads/x86-64/ my_new_setup_i.tgz -R

If you are running Data ONTAP..	And you want to..	Then you can enter..
8.0	Copy and install the image from your HTTP server	software update http://www.example.com/downloads/x86-64/my_new_setup_i.tgz -d -r
	Copy from your HTTP server and overwrite an existing image	software update http://www.example.com/downloads/x86-64/my_new_setup_i.tgz -d -r -f
	Copy and install the image from your HTTP server, then download the new system files to the boot device immediately after installing them	software update http://www.example.com/downloads/x86-64/my_new_setup_i.tgz -r
	Copy and install the image from your HTTP server to a single system, then download the new system files and reboot immediately	software update http://www.example.com/downloads/x86-64/my_new_setup_i.tgz

When you use the **software update** command without the options, a message similar to the following appears on your storage system console:

```
software: You can cancel this operation by hitting Ctrl-C in the next 6 seconds.
software: Depending on system load, it might take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to &lt;filename>;
software: 100% file read from location.
software: /etc/software/&lt;filename>; has been copied.
software: installing software, this could take a few minutes...
software: Data ONTAP Package Manager Verifier 1
software: Validating metadata entries in /etc/boot/NPM_METADATA.txt
software: Checking sha1 checksum of file checksum file: /etc/boot/NPM_FCSUM-pc.sha1.asc
software: Checking sha1 file checksums in /etc/boot/NPM_FCSUM-pc.sha1.asc
software: installation of <filename> completed.
Mon Oct 2 13:26:17 PDT [filer: rc:info]: software: installation of <filename> completed.
```

```
software: Reminder: You might need to upgrade Volume SnapMirror destination
software: filers associated with this filer. Volume SnapMirror can not mirror
software: if the version of ONTAP on the source filer is newer than that on
software: the destination filer.
Mon Oct 2 13:26:17 PDT [filer: download.request:notice]
```

Complete the installation by downloading to HA pairs or single systems.

Installing software images from the /etc/software directory

To install software images, the new software image must be present in the /etc/software directory on your storage system.

Step

1. From the storage system prompt, enter the following command: `software update file options`
 - *file* is the name of the software image you copied to the `/etc/software` directory.
 - *options* is one or more of the following:
 - The `-d` option prevents the **download** command from being run automatically after the system files are installed.
 - The `-f` option overwrites the existing image in the `/etc/software` directory.
 - The `-r` option prevents the system from rebooting automatically after the **download** command has finished (default).
 - The `-R` option causes the system to reboot automatically after the **download** command has finished.

Attention: Beginning in Data ONTAP 8.0.1 or later in the 8.0.x release family, the software update options have changed; the `-r` option (no automatic reboot) is the default, and the `-R` option must be specified to override the `-r` option. However, if you are upgrading from any release earlier than Data ONTAP 8.0.1, you must include the `-r` option to prevent automatic reboot if you are performing a nondisruptive upgrade or if you are upgrading firmware. For more information, see the `software(1)` man page for the Data ONTAP version currently running on your system.

Example

Use the following commands to copy and install the Data ONTAP software image:

If you are running Data ONTAP..	And you want to...	Then you can enter...
7.3.5 or later 7.3.x release	Install the new system files from the <code>/etc/software</code> directory	<code>software update my_new_setup_i.zip -d</code>
	Download the new system files to the boot device immediately after installing them	<code>software update my_new_setup_i.zip</code>
	Copy and install the image from your HTTP server	<code>software update http:// www.example.com/downloads/x86-64/ my_new_setup_i.zip</code>
	Copy from your HTTP server and overwrite an existing image	<code>software update http:// www.example.com/downloads/x86-64/ my_new_setup_i.zip -f</code>

If you are running Data ONTAP..	And you want to...	Then you can enter...
7.3.4 or earlier 7.3.x release	Install the new system files from the /etc/software directory	software update my_new_setup_i.zip -d -r
	Download the new system files to the boot device immediately after installing them	software update my_new_setup_i.zip -r
	Perform an upgrade on a single system and reboot immediately	software update my_new_setup_i.zip
8.0.1 or later	Install the new system files from the /etc/software directory	software update my_new_setup_i.tgz -d
	Download the new system files to the boot device immediately after installing them	software update my_new_setup_i.tgz
	Perform an upgrade on a single system and reboot immediately	software update -R my_new_setup_i.tgz
8.0	Install the new system files from the /etc/software directory	software update my_new_setup_i.tgz -d -r
	Download the new system files to the boot device immediately after installing them	software update my_new_setup_i.tgz -r
	Perform an upgrade on a single system and reboot immediately	software update my_new_setup_i.tgz

When you use the **software update** command without the options, a message similar to the following appears on your storage system console:

```
software: You can cancel this
operation by hitting Ctrl-C in the next 6 seconds.
software: Depending on system load, it might take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to &lt;filename>;
software: 100% file read from location.
software: /etc/software/&lt;filename>; has been copied.
software: installing software, this could take a few minutes...
software: Data ONTAP Package Manager Verifier 1
software: Validating metadata entries in /etc/boot/NPM_METADATA.txt
software: Checking sha1 checksum of file checksum file: /etc/boot/NPM_FCSUM-pc.sha1.asc
software: Checking sha1 file checksums in /etc/boot/NPM_FCSUM-pc.sha1.asc
software: installation of <filename> completed.
Mon Oct 2 13:26:17 PDT [filer: rc:info]: software: installation of <filename>
completed.
```

```

software: Reminder: You might need
to upgrade Volume SnapMirror destination
software: filers associated with this filer. Volume SnapMirror can
not mirror
software: if the version of ONTAP on the source filer is newer than
that on
software: the destination filer.
Mon Oct 2 13:26:17 PDT [filer: download.request:notice]

```

Complete the installation by downloading to HA pairs or single systems.

Updated NDU requirements

Modified requirements apply for nondisruptive upgrades to Data ONTAP 8.0.3 and later 8.0.x releases.

You should avoid exceeding maximum values for the following system elements (per storage controller) on all platforms:

Element	Value
FlexVol volumes	500 The limit for N3400 storage system is 200 FlexVol volumes.
FlexVol volumes enabled for deduplication	300
Snapshot copies	Major NDU: <ul style="list-style-type: none"> • 5,000, when upgrading from 7.3.2 and earlier releases • 12,000, when upgrading from 7.3.3 and later releases on systems with FC or SAS drives • 4,000, when upgrading from 7.3.3 and later releases on systems with SATA drives Minor NDU: <ul style="list-style-type: none"> • 20,000, when upgrading from any 8.0.x release, regardless of drive type The only exception is upgrades from systems earlier than 8.0.1 with SATA drives, where the limit is 4,000 Snapshot copies.
CPU utilization	No greater than 50%
Disk utilization	No greater than 50%

Additional preparation for nondisruptive upgrades

The procedures in "Preparing for nondisruptive upgrades" have been supplemented with four new steps.

About this task

The following steps should be inserted after the existing Step 1 ("Ensure that your HA pair is optimally configured and functioning correctly").

Procedure

1. Ensure that network ports are up and functioning correctly by entering the following command:

```
ifconfig -a
```

Example

For each interface, you see a display similar to the following:

```
e0a: flags=0x2f4c867<UP,BROADCAST,RUNNING,MULTICAST,TCPCSUM,LINK_UP> mtu 1500
  inet 192.9.200.41 netmask 0xffffffff broadcast 192.9.200.255
  partner e0a 192.9.200.42
  ether 00:0c:29:56:54:7e (auto-1000t-fd-up) flowcontrol full
```

For each interface that serves data traffic, you must ensure that each of the following is true:

- The interface has a partner that also serves data; that is, the partner is not an e0M or e0P interface.
- The link to the local interface is up.
- The mtu parameter settings are the same for both partners.
- Partnered interfaces are on the same LAN (the same broadcast domain).

For example, an interface named e0a-10 should be partnered only with another VLAN with tag 10, such as e0b-10.

- Partnered interfaces have matching settings for the `interface.blocked.protocol` option.

For example, if CIFS is blocked on e0a and e0a is partnered with e0b, CIFS should also be blocked on e0b.

If your system includes multiple interface groups, you might also want to confirm their activity with the `ifgrp status` command.

2. If you have edited the `/etc/rc` file, ensure that entries are listed in the following order:

```
hostname system_name
ifgrp [commands]
vlan [commands]
ifconfig [commands]
vfiler [commands]
route [commands]
[any other commands]
```

3. If your systems include e0M management interfaces, ensure that they are serving only management traffic on a dedicated management LAN or that

they are configured down. If an e0M interface is serving management traffic, it should be partnered with another e0M interface.

For more information about e0M configuration, see the *Data ONTAP 8.0 7-Mode System Administration Guide*.

4. If your systems include e0P interfaces for controlling SAS storage expansion units, ensure that they are connected only to a private ACP network or that they are configured down. e0P interfaces should not be partnered.

For more information about ACP configuration, see the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

Additional SAN upgrade requirement

If your storage system is in a SAN environment, you must ensure that your SAN configuration is fully supported. All SAN components—including the target Data ONTAP software version, host OS and patches, required Host Utilities software, and adapter drivers and firmware—should be listed in the IBM N series interoperability matrix, which is accessed and navigated as described in “Websites” on page 149.

Additional requirement for AT-FCX storage expansion unit firmware updates

The topic "Service availability during storage expansion unit firmware updates" currently lists two requirements for updating AT-FCX storage expansion unit firmware non-disruptively: minimum firmware version and Multipath Storage configuration. There is an additional requirement that Data ONTAP 7.3.2 or later be running.

The first row of the table should read as follows:

Module	Storage expansion unit model	System downtime required?
AT-FCX	EXN1000	With Multipath Storage, firmware version 37, and Data ONTAP 7.3.2 or later: No Without Multipath Storage: Yes

After the table, the AT-FCX information should also be updated to state that you cannot perform a Data ONTAP nondisruptive upgrade (NDU) under the following circumstances:

- AT-FCX storage expansion units are attached to your system and one or more of the following is true:
 - You have not verified that the latest AT-FCX firmware version is running.

- A Data ONTAP release earlier than 7.3.2 is running.
- AT-FCX version 36 or earlier is running.
- Multipath Storage is not configured.

AT-FC and AT-FC2 shelves no longer supported

Beginning with Data ONTAP 8.x, AT-FC and AT-FC2 storage expansion units are no longer supported. You can disregard all references to these devices in the *Data ONTAP 8.0 7-Mode Upgrade Guide*.

Changes to the System Administration Guide

Additional information has become available since the last revision of the *Data ONTAP 8.0 7-Mode System Administration Guide*.

- “Changes to BMC command information”
- “Changes to information about the `/etc/rc` file”
- “Corrections to the timed options” on page 123
- “Updated password requirements” on page 123
- “Corrections for the location of the `secureadmin.pem` file” on page 123
- “Correction to the description of the `autosupport.to` option” on page 123
- “Change in the frequency of the performance AutoSupport message” on page 124
- “Use FlexShare with HDDs only” on page 124
- “Correction to the free space required for reallocation scans” on page 124
- “Changes to the supported OpenSSH client versions” on page 124
- “Changes to storage system configuration backup and restore” on page 124
- “Changes to the instructions on editing the `/etc/rc` file” on page 125
- “Terminology changes for PAM and WAFL extended cache” on page 126
- “Corrections for root FlexVol volume size requirement” on page 126
- “Changes to the root aggregate and root volume Snapshot reserve values on new systems” on page 126

Changes to BMC command information

The *Data ONTAP 7-Mode System Administration Guide* is amended with information about the `bmc battery-info` command.

In the *Data ONTAP commands for managing the BMC* section, the command table has been amended with the following information:

If you want to ...	Use this Data ONTAP command ...
Display battery information, for example, battery name, charging status, serial number, part number, firmware and hardware versions, chemistry type, manufacturer, and date of manufacture	<code>bmc battery-info</code>

Changes to information about the `/etc/rc` file

The *About the `/etc/rc` file* section of the *Data ONTAP 7-Mode System Administration Guide* has updated information about the commands that the `/etc/rc` file must not contain.

The `/etc/rc` file must not contain the following types of commands:

- Commands that are executed by subsystems that are not yet available when the file is executed
- Commands that are interactive and would wait for input during the boot process

For example, you must not include the `iscsi` commands or the `wrfile` command in the `/etc/rc` file. Doing so prevents your storage system from booting successfully.

Corrections to the timed options

The `timed.log` option has been removed from The *timed options* section of the *Data ONTAP 7- Mode System Administration Guide*. The option has become obsolete as of Data ONTAP 8.0.

Updated password requirements

The *How to manage passwords for security* section of the *Data ONTAP 7-Mode System Administration Guide* is amended with an additional requirement that the password must not contain the Ctrl-c or Ctrl-d key combination or the two-character string `^D`.

The following are the default password rules for all accounts when `security.passwd.rules.enable` is set to `on` (the default):

- The password must be at least eight characters long.
- The password must contain at least one number.
- The password must contain at least two alphabetic characters.
- The password must not contain the Ctrl-c or Ctrl-d key combination or the two-character string `^D`.

Corrections for the location of the `secureadmin.pem` file

The *Installing a certificate-authority-signed certificate* section of the *Data ONTAP 7-Mode System Administration Guide* contains incorrect information about the location of the `secureadmin.pem` file. The correct location of the file is `/etc/keymgr/csr`.

Correction to the description of the `autosupport.to` option

The section titled *AutoSupport options* in the *Data ONTAP 7-Mode System Administration Guide* is amended with information about the `autosupport.to` option.

The description of the `autosupport.to` option states that recipients receive only critical AutoSupport email notifications. This is incorrect. Email notifications are not restricted to AutoSupport messages with specific severity types. Recipients defined in `autosupport.to` receive key AutoSupport messages, as defined in factory-default settings.

Change in the frequency of the performance AutoSupport message

Starting from Data ONTAP 8.0.3, an AutoSupport message is sent when the `cm_hourly_stats` file reaches the threshold limit. Therefore, it is possible that the performance AutoSupport message might be sent more than once a week.

Use FlexShare with HDDs only

Use FlexShare with storage systems that have hard disk drives (HDDs) only. FlexShare is not designed for use with storage systems that have solid-state drives (SSDs). Enabling FlexShare on a storage system that has SSDs can result in decreased throughput to SSD-based volumes.

Correction to the free space required for reallocation scans

The topic "What a reallocation scan is" includes incorrect information.

The following content is incorrect:

Reallocation scans will not run if there is less than 10 percent free space (excluding the Snapshot reserve) in the active file system on the target volume or aggregate.

The content is corrected as follows:

Reallocation scans will not run if there is less than five percent free space (excluding the Snapshot reserve) in the active file system on the target volume or aggregate.

Changes to the supported OpenSSH client versions

The SSH protocol section of the *Data ONTAP 8.0 7-Mode System Administration Guide* has updated information for OpenSSH client versions that Data ONTAP supports.

Data ONTAP supports the following SSH clients:

- OpenSSH client version 4.4p1 on UNIX platforms
- SSH Communications Security client (SSH Tectia client) version 6.0.0 on Windows platforms
- Vandyke SecureCRT version 6.0.1 on Windows platforms
- PuTTY version 0.6.0 on Windows platforms
- F-Secure SSH client version 7.0.0 on UNIX platforms

To enhance security, OpenSSH client version 3.8p1 is no longer supported because it does not contain the latest security fix.

Changes to storage system configuration backup and restore

The *Data ONTAP 8.0 7-Mode System Administration Guide* has updated information about the backup and restore of a storage system configuration.

- The section about storage system configuration backup and cloning is amended with the following information:

When you back up a storage system configuration, the following files are backed up for the storage system and the default vFiler unit (vfiler0):

- System-specific configuration files, for example, `/etc/rc`
- System-specific registry options
- Volume configuration
- vfiler0-specific configuration, for example, `/etc/quotas`, `/etc/hosts`, `/etc/usermap.cfg`, `/etc/nsswitch.conf`, and `/etc/hosts.equiv`
- vfiler0-specific registry options, for example, NFS, CIFS, ndmpd, and NIS

If you have configured vFiler units, when you back up the configuration of a vFiler unit, the following files in the vFiler unit are backed up:

- vFiler-specific configuration files, for example, `/etc/quotas`, `/etc/hosts`, `/etc/usermap.cfg`, `/etc/nsswitch.conf`, and `/etc/hosts.equiv`
- vFiler-specific registry options, for example, NFS, CIFS, ndmpd, and NIS

vFiler configuration is backed up or restored only for the vFiler unit on which the `config dump` or `config restore` command is run.

- The section about restoring a storage system configuration is amended with the following information as Step 3:

1. Enter the following command:
`config restore [-v] config_file`

`-v` enables you to restore volume-specific configuration files, as well as storage system-specific configuration files.

2. Reboot the system to run commands in the `/etc/rc` file.
3. If you use quotas for any volumes owned by a non-default vFiler unit (a vFiler unit other than vfiler0), ensure that the quotas are in the desired state (on or off) for those volumes.

The quotas state for volumes owned by a non-default vFiler is not restored when you restore a system configuration.

Changes to the instructions on editing the `/etc/rc` file

The instructions on editing the `/etc/rc` file in the *Data ONTAP 8.0 7-Mode System Administration Guide* have updated information.

The procedure for editing the `/etc/rc` file is amended with information in Step 3 below:

1. Make a backup copy of the `/etc/rc` file.
2. Edit the `/etc/rc` file.

Note: Do not add CIFS commands to `/etc/rc`. Doing so can cause problems when the storage system boots if CIFS is not fully initialized or the commands cause deadlocks between the `/etc/rc` file and CIFS.

3. Ensure that entries in the `/etc/rc` file are listed in the following order:

```
hostname system_name
ifgrp commands
vlan commands
ifconfig commands
vfiler commands
route commands
[any other commands]
```

4. Save the edited file.
5. Reboot the storage system to test the new configuration. If the new configuration does not work as you want, repeat Step 2 through Step 4.

Terminology changes for PAM and WAFL extended cache

In the *Data ONTAP 8.0 7-Mode System Administration Guide*, mentions of "Performance Acceleration Module (PAM) family" should be amended to read "Performance Acceleration Modules (PAM) and Flash Cache Modules", and mentions of "WAFL extended cache" should be amended to read "WAFL external cache".

Corrections for root FlexVol volume size requirement

The size requirement for root FlexVol volumes on some systems has changed.

In the section titled Size requirement for root FlexVol volumes in the *Data ONTAP 8.0 7-Mode System Administration Guide*, the table is amended with the following information:

Storage system model	Minimum root FlexVol volume size
N6210	151 GB
N6240	205 GB

Changes to the root aggregate and root volume Snapshot reserve values on new systems

Starting with Data ONTAP 8.0.1, new systems are shipped with Snapshot reserve for the root aggregate set to 0 percent and Snapshot reserve for the root volume set to 5 percent.

The default reserve value for other aggregates or volumes remains the same. Existing systems are also not affected by this change.

The following note no longer applies and is considered removed from the section titled Aggregate Snapshot Reserve in the *Data ONTAP 8.0 7-Mode System Administration Guide*.

Note: If you have automatic aggregate Snapshot copy creation enabled, you should not decrease the size of the aggregate Snapshot reserve below the default of 5 percent. If you need to reclaim the space being used for the aggregate Snapshot reserve, disable automatic aggregate Snapshot copy creation.

Changes to the Storage Management Guide

New information has become available since the previous revision of the *Data ONTAP 8.0 7-Mode Storage Management Guide*.

- “Corrections to the data compression commands”
- “Correction to the availability of FlexClone files and FlexClone LUNs in vfiler contexts” on page 129
- “Capacity information for 3-TB disks” on page 129
- “Use FlexShare with HDDs only” on page 124
- “Maximum volume size for deduplication” on page 129
- “Incorrect cross-reference to maximum volume size information” on page 129
- “Maximum number of mirrored aggregates per system” on page 129
- “ACP example output shows incorrect IP addresses” on page 129
- “Using disk sanitization to remove data from disks” on page 130

Corrections to the data compression commands

The sections *Decompressing the compressed data* and *Reverting compressed volumes* include incorrect command syntax for the data compression commands.

In the section *Decompressing the compressed data*, the existing command syntax to decompress the compressed data is incorrect. The corrected command syntax is as follows:

```
vol decompress start volume_name
```

In the section *Reverting compressed volumes*, the existing command syntax to revert a volume and the example are incorrect. The corrected command syntax and example are as follows:

```
vol decompress revert volume_name
```

volume_name is the name of the volume.

Example

The following command reverts the volume VolA, which contains compressed data:

```
vol decompress revert VolA
```

Correction to the availability of FlexClone files and FlexClone LUNs in vfiler contexts

The topic "How MultiStore works with FlexClone files and FlexClone LUNs" incorrectly states that FlexClone files and LUN commands are available in default and non-default vfiler contexts.

The corrected information is as follows: MultiStore supports FlexClone files and FlexClone LUNs only in vfiler0 context. You cannot run FlexClone file and LUN commands in any other vfiler context.

Capacity information for 3-TB disks

The topic titled "Available disk capacity by disk size" does not include information for 3-TB SATA disks.

The 3-TB SATA disks have a right-sized capacity of 2,538 GB (where GB = 1,000*1,024*1,024), and 5,198,943,744 available blocks (sectors). They are used in SAS storage expansion units.

Use FlexShare with HDDs only

Use FlexShare with storage systems that have hard disk drives (HDDs) only. FlexShare is not designed for use with storage systems that have solid-state drives (SSDs). Enabling FlexShare on a storage system that has SSDs can result in decreased throughput to SSD-based volumes.

Incorrect cross-reference to maximum volume size information

The cross-reference in the section "Deduplication and volume SnapMirror" incorrectly references the section "Maximum volume size with deduplication."

Maximum volume size for deduplication

Starting in Data ONTAP 8.0.1 and in later releases, you can enable and run deduplication on volumes with a maximum size of up to 16 TB for all supported platforms. This information is missing in *Data ONTAP 8.0 7-Mode Storage Management Guide*.

Maximum number of mirrored aggregates per system

You should not create more than 64 mirrored aggregates per storage system. Data ONTAP does not prevent you from exceeding this limit, but doing so could cause plex synchronization problems after some types of failures.

ACP example output shows incorrect IP addresses

The IP addresses shown in the example showing the output of the storage show acp command shows ACP configured to use public IP addresses. This is incorrect; you should use private IP addresses when configuring ACP.

The example should read as follows:

Example

For example, if you select e0P as the interface for ACP traffic, 192.168.0.0 as the ACP domain, and 255.255.252.0 as the network mask for the ACP subnet, the storage show acp command output looks similar to the following example:

```
my-sys-1> storage show acp
Alternate Control Path: enabled
Ethernet Interface: e0P
ACP Status: Active
ACP IP address: 192.168.2.61
ACP domain: 192.168.0.0
ACP netmask: 255.255.252.0
ACP Connectivity Status: Full Connectivity
Shelf Module Reset Cnt IP address FW Version Module
Type Status
-----
7a.001.A 002 192.168.0.145 01.05
IOM6 active
7a.001.B 003 192.168.0.146 01.05
IOM6 active
7c.002.A 000 192.168.0.206 01.05
IOM6 active
7c.002.B 001 192.168.0.204 01.05
IOM6 active
```

Using disk sanitization to remove data from disks

The procedure for disk sanitization in the *Data ONTAP 8.0 7-Mode Storage Management Guide* and *Data Protection Management Guide* does not include the steps for returning a sanitized disk to the spare pool. Use this version instead.

Before you begin

- You must have installed the disk sanitization license.
Attention: After the license for disk sanitization is installed on a storage system, it is permanent, and it prevents certain Data ONTAP commands from being run.
- You must have assigned ownership for the disks that you want to sanitize. You cannot sanitize unowned disks.
- The disks that you want to sanitize must be spare disks.

Procedure

1. Verify that the disks that you want to sanitize do not belong to a RAID group in any existing aggregate by entering the following command:
sysconfig -r The disks that you want to sanitize should be listed as spare disks.
2. Sanitize the specified disk or disks by entering the following command: disk sanitize start [-p *pattern1*|-r [-p *pattern2*|-r [-p *pattern3*|-r]]] [-c *cycle_count*] *disk_list*

Attention: Do not turn off the storage system, disrupt the storage connectivity, or remove target disks while sanitizing. If sanitizing is interrupted during the formatting phase, the formatting phase must be restarted and allowed to finish before the disks are sanitized and ready to be returned to the spare pool.

If you need to abort the sanitization process, you can do so by using the `disk sanitize abort` command. If the specified disks are undergoing the formatting phase of sanitization, the abort will not occur until the disk formatting phase is complete. At that time, Data ONTAP displays a message telling you that the sanitization process was stopped.

`-p pattern1 -p pattern2 -p pattern3` specifies a cycle of one to three user-defined hex byte overwrite patterns that can be applied in succession to the disks being sanitized. The default pattern is three passes, using 0x55 for the first pass, 0xaa for the second pass, and 0x3c for the third pass.

`-r` replaces a patterned overwrite with a random overwrite for any or all of the passes.

`-c cycle_count` specifies the number of times the specified overwrite patterns will be applied. The default value is one cycle. The maximum value is seven cycles.

`disk_list` specifies a space-separated list of the IDs of the spare disks to be sanitized.

3. If you want to check the status of the disk sanitization process, enter the following command: `disk sanitize status [disk_list]`
4. After the sanitization process is complete, enter advanced privilege mode by entering the following command: `priv set advanced`
5. Assign the disks to the appropriate storage system by entering the following command: `disk assign disk_list -o system_name`
6. Return the disks to the spare pool by entering the following command: `disk unfail -s disk_list` Data ONTAP designates the specified disks as hot spares.
7. Return to administrative mode by entering the following command: `priv set`
8. Optional: You can monitor the status of the sanitization process, or verify that all disks were successfully sanitized, by using the `/etc/sanitized_disks` and `/etc/sanitization.log` files.

Status for the sanitization process is written to the `/etc/sanitization.log` file every 15 minutes.

The `/etc/sanitized_disks` file contains the serial numbers of all drives that have been successfully sanitized.

Results

The specified disks are sanitized and designated as hot spares. The serial numbers of the sanitized disks are written to `/etc/sanitized_disks`.

Examples

The following command applies the default three disk sanitization overwrite patterns for one cycle (for a total of three overwrites) to the specified disks, 8a.6, 8a.7, and 8a.8:

```
disk sanitize start 8a.6 8a.7 8a.8
```

The following command would result in three disk sanitization overwrite patterns for six cycles (for a total of 18 overwrites) to the specified disks:

```
disk sanitize start -c 6 8a.6 8a.7 8a.8
```

Changes to the Software Setup Guide

Additional information has become available since the last revision of the *Data ONTAP 7-Mode Software Setup Guide*.

- “Correction to the setup password rules”

Correction to the setup password rules

The password rules in the tables in the “Required storage system information” and the “CIFS protocol information” sections have been updated.

The following password rules are for initial setup for all accounts when `security.passwd.rules.enable` is set to `on`, which is the default value:

- The password must be at least eight characters long.
- The password must contain at least one number.
- The password must contain at least two alphabetic characters.
- The password must not contain the Ctrl-c or Ctrl-d key combination or the two-character string `^D`.

Changes to the High-Availability Configuration Guide

There are some changes to the existing content of the *Data ONTAP 8.0 7-Mode High-Availability Configuration Guide*.

- “Hardware-assisted takeover support on systems that use Service Processor” on page 133
- “Corrections to the example in “Monitoring HA pair status” topic” on page 133

Hardware-assisted takeover support on systems that use Service Processor

Although the Data ONTAP 8.0 7-Mode High-Availability Configuration Guide only mentions support for the hardware-assisted takeover feature on systems with an RLM card, the feature is also supported on systems that use a Service Processor for remote management.

Corrections to the example in "Monitoring HA pair status" topic

The `cf status` command output in the "Monitoring HA pair status" topic does not provide information about the two links that constitute the HA interconnect.

The example should read as follows:

Example

```
node1>cf status
Cluster enabled, node2 is up
RDMA Interconnect is up (Link 0 down, Link 1 up)
```

Note: Depending on the storage system model, the output might display either RDMA interconnect or VIA interconnect in the last line.

Note: If the output shows that one link is down, you must configure the link so that it is up while the other link is still active.

Changes to the Network Management Guide

New information has become available since the previous revision of the *Data ONTAP 8.0 7-Mode Network Management Guide*.

- "Changes to information about routed daemon"
- "Supported SNMP versions in Data ONTAP" on page 134
- "TLS not supported in Data ONTAP 8.0" on page 134

Changes to information about routed daemon

The topic titled "When the routed daemon should be turned off" contains incorrect information. In addition, the topic titled "The routed daemon" is updated for information about RIPv2 authentication if routed daemon is on.

The information in the topic "When the routed daemon should be turned off" is corrected as follows:

Unless you have dynamic routing requirements and understand network routing configuration, you are advised to always keep the routed daemon off. If you have dynamic routing requirements, you are advised to turn the

routed daemon on. Turning the routed daemon off might cause unexpected routing behavior if dynamic routing is used.

The topic "The routed daemon" is updated for the following information:

If you require dynamic routing in untrusted environments, you must use RIPv2 with authentication. This is because routed daemon poses a security risk if RIPv1 is used in untrusted environments.

Supported SNMP versions in Data ONTAP

The "How to monitor your storage system with SNMP" topic has updated information about support for SNMP version 2. For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP versions SNMPv1, SNMPv2c, and SNMPv3.

TLS not supported in Data ONTAP 8.0

TLS is not supported in Data ONTAP 8.0 7-Mode. The SSL section of the "IP port usage on a storage system" chapter incorrectly includes information on TLS.

Changes to the File Access and Protocols Management Guide

The *Data ONTAP 8.0 7-Mode File Access and Protocols Management Guide* should have included the following topics.

- “Enabling or disabling SFTP log files”
- “Specifying the maximum size of the current SFTP log files”
- “Correction to specifying the FTP authentication style” on page 136
- “Incorrect reference to nonexistent SFTP log files” on page 136
- “Updated compatibility information available online” on page 136
- “Corrections to the FPolicy commands to monitor directory operations” on page 136
- “Unsupported Windows features in the file serving environment” on page 137
- “NFSv4 client compatibility” on page 137
- “FTP server does not support Unicode characters” on page 137

Enabling or disabling SFTP log files

You can enable or disable SFTP log files by setting the `sftp.log_enable` option to on or off, respectively. This enables you to enable or disable SFTP event logging. By default, this option is on.

About this task

When this option is enabled, Data ONTAP logs SFTP commands and data transfer operations to the `/etc/log/sftp.cmd.*` log files.

Step

1. Perform one of the following actions:

If you want SFTP log files to be...	Enter the command...
Enabled	<code>options sftp.log_enable on</code>
Disabled	<code>options sftp.log_enable off</code>

Specifying the maximum size of the current SFTP log files

To specify the maximum size of the current `/etc/log/sftp.cmd.*` SFTP log files, you can set the `sftp.log_filesize` option. By default, the maximum size of the current SFTP log files is 512 KB.

Step

1. Enter the following command:
`options sftp.log_filesize filesize`

filesize is the maximum size of the current SFTP files expressed as a value from 1K to 4 GB. You can specify the value in gigabytes (G), megabytes (M), kilobytes (K), or bytes (blank). For more information, see the `na_options(1)` man page.

The following example sets the maximum size of the current SFTP log files to 1 GB:

```
options sftpd.log_filesize 1G
```

Correction to specifying the FTP authentication style

The description of the mixed authentication style in the topic *Specifying the FTP authentication style* is incorrect.

The correct description is as follows:

When you specify the mixed authentication style, the FTP server uses the NTLM authentication style for users with names containing a backslash (\) or "@" character; it uses the UNIX authentication style for all other users.

Incorrect reference to nonexistent SFTP log files

The *Data ONTAP 7-Mode File Access and Protocols Management Guide* contains incorrect references to SFTP log files that do not exist on the storage system.

The following two topics mention `/etc/log/sftp.xfer` log files:

- *Enabling or disabling SFTP log files*
- *Specifying the maximum size of the current SFTP log files*

However, these log files do not exist on the storage system. These references will be removed in future documentation updates.

Updated compatibility information available online

The information found in the *Data ONTAP 7-Mode File Access and Protocols Management Guide* about compatibility with client and server operating systems and versions might not be current.

For the latest information about which client and server operating systems and versions are supported, see the IBM N series interoperability matrix, which is accessed and navigated as described in "Websites" on page 149.

Corrections to the FPolicy commands to monitor directory operations

The *Data ONTAP 7-Mode File Access and Protocols Management Guide* includes incorrect FPolicy commands for monitoring directory operations.

The correct FPolicy commands to monitor directory operations are as follows:

If you want to monitor directory operations for...	Enter the command...
Creating	fpolicy monitor add policy_name create_dir
Renaming	fpolicy monitor add policy_name rename_dir
Deleting	fpolicy monitor add policy_name delete_dir

Unsupported Windows features in the file serving environment

This release does not support the following Windows features:

For example, this release does not support the following Windows features:

- Encrypted File System (EFS)
- Logging of NT File System (NTFS) events in the change journal
- Microsoft File Replication Service (FRS)
- Microsoft Windows Indexing Service
- Remote storage through Hierarchical Storage Management (HSM)
- Local user account creation from the User Manager or Microsoft Manager Console
- Quota management from Windows clients
- Windows quota semantics
- The LMHOSTS file
- NT File System (NTFS) native compression

NFSv4 client compatibility

When your NFSv4 clients are in a different domain than your device, you might need to enter the client domain name as the value for the Data ONTAP option `nfs.v4.id.domain` in order to provide mapping for file ownership and group membership.

For more information about mapping options, see RFC 3530 and your client operating system documentation.

If you have any client using NFSv4 that needs to access a storage system, ensure that the Data ONTAP option `nfs.v4.enable` is set to `on`. In new installations, this option is set to `off` by default.

FTP server does not support Unicode characters

The FTP server does not support Unicode characters; thus, file names containing Unicode characters, including Japanese characters, are displayed in FTP clients with alternate characters used in place of Unicode characters.

Changes to the Block Access Management Guide for iSCSI and FC

New information has become available since the previous revision of the *Data ONTAP 8.0 7-Mode Block Access Management Guide for iSCSI and FC*.

- "Correction to the DCB setting example"
- "Correction to links in the "Unified Ethernet network management" and "What data center bridging is" topics"
- "Setting volume options for the autodelete configuration"
- "Correction to configure onboard adapters for target mode" on page 139

Correction to the DCB setting example

The result in the topic *Displaying DCB settings* has an incorrect example. The Applications column incorrectly lists IP instead of unassigned.

The corrected example is as follows:

Example

```
system1> dcb priority show e2b
Interface Priority Applications Flow Control PGID
-----
e2b
      0      unassigned   enabled    0
      1      unassigned   disabled   1
      2      unassigned   disabled   1
      3      FCoE         enabled    2
      4      unassigned   disabled   1
      5      unassigned   disabled   1
      6      unassigned   disabled   1
      7      unassigned   disabled   1
```

Correction to links in the "Unified Ethernet network management" and "What data center bridging is" topics

The topics "Unified Ethernet network management" and "What data center bridging is" provide incorrect links.

The correct links are:

- Technical Report: Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide - media.netapp.com/documents/TR-3800.pdf
- Data Center Bridging task group - www.ieee802.org/1/pages/dcbbridges.html

Setting volume options for the autodelete configuration

The final step is missing in the procedure "Setting volume options for the autodelete configuration" topic of "Configuring volumes and LUNs when using autodelete" section.

Perform the following steps for "Setting volume options for the autodelete configuration":

1. Set the space guarantee on the volumes by entering the following command:

```
vol options vol_name guarantee volume
```
2. Ensure that autosize is disabled by entering the following command:

```
vol autosize disable vol_name
```

Note: This option is disabled by default.
3. Set fractional reserve to zero percent, if it is not already, by entering the following command:

```
vol options vol_name fractional_reserve 0
```
4. Set the Snapshot copy reserve to zero percent by entering the following command:

```
snap reserve vol_name 0
```

The Snapshot copy space and application data is now combined into one large storage pool.
5. Configure Snapshot copies to begin being automatically deleted when the volume reaches the capacity threshold percentage by entering the following command:

```
snap autodelete vol_name trigger volume
```

Note: The capacity threshold percentage is based on the size of the volume. For more details, see the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.
6. Set the try_first option to snap_delete by entering the following command:

```
vol options vol_name try_first snap_delete
```

This enables Data ONTAP to begin deleting Snapshot copies, starting with the oldest first, to free up space for application data.
7. Activate the snap autodelete settings by entering the following command:

```
snap autodelete vol_name on
```

When finished, create your space-reserved LUNs.

Correction to configure onboard adapters for target mode

The section "Configuring onboard adapters for target mode" has some incorrect information about onboard adapters and target expansion cards.

Changes to "About this task"

The existing content is incorrect: If you are installing target expansion adapters, or if you exceed the allowed number of adapter ports, you must set the onboard adapters to unconfigured before installing the expansion adapters.

The corrected content is: If you exceed the allowed number of adapter ports, you must set the onboard adapters to initiator or unconfigured before installing the expansion adapters.

Changes to the first step in the topic "Configuring onboard adapters for target mode"

The following existing content is incorrect:

1. If you have already connected the port to a switch or fabric, take it offline by entering the following command:

```
fcv config adapter down
```

adapter is the port number. You can specify more than one port.

Example

```
fcv config 0c 0d down
```

Ports 0c and 0d are taken offline.

Note: If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

The corrected content is as follows:

1. Verify that the FC ports are not already configured as target ports by entering the following command:

```
fcadmin config
```

Example

```
fcadmin config
```

```
Local
```

```
Adapter Type State Status
```

```
-----  
0a initiator CONFIGURED online
```

```
0b initiator CONFIGURED online
```

```
0c target CONFIGURED online
```

```
0d target CONFIGURED online
```

```
The preceding output displays two ports for host access.  
-----
```

Changes to the Data Protection Online Backup and Recovery Guide

New information has become available since the previous revision of the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

- “Corrections to information about deployment of SnapMirror over Fibre Channel”
- “Connection name is required to enable SnapMirror compression”
- “Change to the Snapshot copy reserve value”
- “Single file reversion to or from symbolic links using SnapRestore not supported” on page 142

Corrections to information about deployment of SnapMirror over Fibre Channel

The sections "Requirements for deploying SnapMirror over Fibre Channel" and "Functionality supported by SnapMirror over Fibre Channel" have some incorrect information about fabric configuration.

Requirements for deploying SnapMirror over Fibre Channel

The following existing content is incorrect: The SAN must be enabled for in-order frame delivery (IOD) in some configurations.

This content is corrected as follows: The SAN must be enabled for in-order frame delivery (IOD) if there is more than one path between the source and destination systems.

Functionality supported by SnapMirror over Fibre Channel

The table that describes the fabric configuration options has incorrect information about in-order delivery setting on SAN fabric. The column titled "In-order delivery setting on SAN fabric must be" should be removed.

Connection name is required to enable SnapMirror compression

To enable SnapMirror network compression, you must specify a name for the connection between a SnapMirror source and destination system pair using the *connection_name* entry in the *snapmirror.conf* file. This information is missing in the "Enabling SnapMirror network compression" topic in the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

Change to the Snapshot copy reserve value

To increase the amount of usable space on the storage system while still retaining all the benefits of volume Snapshot copies, IBM ships all new systems running Data ONTAP 8.0.1 GA with a smaller volume Snapshot reserve value.

The Snapshot copy reserve value is 5 percent of the volume size instead of the 20 percent stated in the "What the Snapshot copy reserve is" topic.

Single file reversion to or from symbolic links using SnapRestore not supported

You cannot perform single file reversion to or from symbolic links using SnapRestore. This information is not included in the "Reverting a file to a selected Snapshot copy" topic in the *Data ONTAP 8.0 7-Mode Data Protection Online Backup and Recovery Guide*.

Changes to the Data Protection Tape Backup and Recovery Guide

Some additional information about tape backup features has become available since the previous revision of the *Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide*.

- "Support for SMTape backup of compressed volumes"
- "Designating range of ports for NDMP data connections"

Support for SMTape backup of compressed volumes

Starting with Data ONTAP 8.0.1, SMTape supports backup and restore of compressed volumes.

The information about the lack of support for compressed volumes in the *Data ONTAP 8.0 7-Mode Data Protection Tape Backup and Recovery Guide* is no longer applicable.

Designating range of ports for NDMP data connections

Data ONTAP supports a designated range of TCP/IP ports that can be used for NDMP data connections in response to NDMP_DATA_LISTEN and NDMP_MOVER_LISTEN operations.

Data ONTAP 8.0 7-Mode and earlier versions do not support data migration by using `ndmcopy` command and three-way tape backups in environments where the source and destination networks are separated by a firewall. This is because the data or mover port that is used in a data transfer is unpredictable.

Starting with Data ONTAP 8.0.1, administrators can designate range of ports that can be used for NDMP data connections in response to NDMP_DATA_LISTEN and NDMP_MOVER_LISTEN operations. Therefore, Data ONTAP enables you to perform data migration by using `ndmcopy` command and three-way tape backups even in environments where the source and destination networks are separated by a firewall.

1. To enable the data port range, enter the following command:

```
options ndmpd.data_port_range {start_port-end_port}
```


The `ndmpd.data_port_range` option allows administrators to specify a port range on which the NDMP server can listen for data connections.

The *start_port* and *end_port* indicate the range of ports designated for data connection and can have values between 1024 and 65535; *start_port* must be less than or equal to *end_port*.

If a valid range is specified, NDMP uses a port within that range to listen for incoming data connections. A listen request fails if no ports in the specified range are free.

The default value for `ndmpd.data_port_range` option is `all`. The `all` implies that any available port can be used to listen for data connections.

Note: The `ndmpd.data_port_range` option is persistent across reboots.

Example:

```
Filer1> options ndmpd.data_port_range 1024-2048
```

Changes to the MultiStore Management Guide

Some additional information has become available since the last revision of *Data ONTAP 8.0 7-Mode MultiStore Management Guide* and there are some changes to the existing content.

- “Correction to the effects of storage system reboot on a vFiler unit”
- “Requirement for backup and replication operations on vFiler units”
- “Naming guidelines for vFiler units”
- “SnapMirror relationship schedules are not modified for existing volumes” on page 145
- “Corrections to the List of RSH and SSH commands” on page 145
- “Corrections to the SnapVault support on vFiler units” on page 145

Correction to the effects of storage system reboot on a vFiler unit

The topic "Effects of storage system reboot on a vFiler unit" includes an incorrect word.

The following sentence includes an incorrect word: When you stop a vFiler unit and then reboot the storage system, the stopped iremen starts running again after the reboot.

The corrected sentence is as follows: When you stop a vFiler unit and then reboot the storage system, the stopped vFiler unit starts running again after the reboot.

Requirement for backup and replication operations on vFiler units

When performing SnapVault, ndmpd, and SnapMirror operations on a vFiler unit, you must ensure that the root of the vFiler unit is a volume. This information is missing in the *Data ONTAP 8.0 7-Mode MultiStore Management Guide*.

Naming guidelines for vFiler units

Information about the guidelines you require when naming a vFiler unit is missing in the *Data ONTAP 8.0 7-Mode MultiStore Management Guide*.

To create or rename a vFiler unit successfully, you must be aware of the following guidelines when naming a vFiler unit:

- The name can contain up to 31 alphanumeric ASCII characters.
- The name can contain the dash (-) and the underscore (_) characters.
However, the name should not begin with a dash.
- The name is case-insensitive.
- The name must be unique.

You can include the name of the hosting storage system as part of the vFiler unit name so that you can easily determine the storage system that contains the vFiler unit—for example, mycompanyss1_vfiler1.

- The name vFiler0 should not be used because it is the name of the default vFiler unit.

SnapMirror relationship schedules are not modified for existing volumes

In Data ONTAP 8.0.2 and later releases, if you edit the SnapMirror relationship schedules in the `etc/snapmirror.conf` file after executing the `vfiler dr configure` command, the `vfiler dr resync` command does not modify the SnapMirror relationship schedules for the existing volumes.

For example, if you change the SnapMirror relationship schedule from asynchronous to semi-synchronous after executing the `vfiler dr configure` command, the `vfiler dr resync` command does not modify the SnapMirror relationship schedule from semi-synchronous to asynchronous during resynchronization of the vFiler unit.

Corrections to the List of RSH and SSH commands

"List of RSH and SSH commands" topic lists some commands that are not supported in Data ONTAP 8.0 and later releases.

The following commands, which are listed in the *Data ONTAP 8.0 7-Mode MultiStore Management Guide*, are not supported in Data ONTAP 8.0 and later releases:

cifs	snap
config	snapmirror
df	snapvault
dns	vol
echo	vscan
fpolicy	wcc
nbtstat	yocat
nfs	ypgroup
nis	ypmatch
route	ypwhich
sectrace	

Corrections to the SnapVault support on vFiler units

The sections "Where to enter SnapVault commands" and "Features and limitations of the snapvault command" have some incorrect information about interoperability between SnapVault and MultiStore.

Where to enter SnapVault commands

The following existing content is incorrect: Commands entered on a nondefault vFiler unit makes changes on or displays information only about that specific vFiler unit.

This content is corrected as follows: Some commands entered on a nondefault vFiler unit make changes on or display information only about that specific vFiler unit--for example, `snapvault status` and `snapvault snap sched`.

Features and limitations of the snapvault command

The following existing content is incorrect:

The features of the `snapvault` command when used in a MultiStore context are as follows:

- Additional SnapVault licenses are not required. vFiler units use the same source and destination licenses as the physical storage systems.
- The SnapVault feature can be turned on and off independently on each vFiler unit.
- The `snapvault.access` and `snapvault.enable` options can be changed independently on each vFiler unit.
- Each vFiler unit has its own `snapvault.conf` file in the `/etc` directory.
- SnapVault relationships established between vFiler units are maintained across vFiler unit migration.

This content is corrected as follows:

The features of the `snapvault` command when used in a MultiStore context are as follows:

- Additional SnapVault licenses are not required. vFiler units use the same source and destination licenses as the physical storage systems.
- The `snapvault.access` and `snapvault.enable` options can be changed independently on each vFiler unit.

Changes to the Storage Efficiency Management Guide

There are some changes to the existing content of the *Data ONTAP 8.0 7-Mode Storage Efficiency Management Guide*.

- “Corrections to deduplication and volume SnapMirror information”

Corrections to deduplication and volume SnapMirror information

The "Deduplication and volume SnapMirror" section includes some incorrect content.

The following existing content is incorrect: For example, an N5300 storage system supports a maximum volume size with deduplication of 4 TB, and an N5600 storage system supports 16 TB. When establishing a volume SnapMirror relationship between the N5300 and N5600 storage systems, you should ensure that the volume SnapMirror relationship is established for 4 TB. After a failover, if the volume size on the N5600 storage system is increased to more than 4 TB, you will not be able to perform a new baseline transfer or resynchronize the storage systems, because the N5300 storage system has a maximum volume size with deduplication of 4 TB.

The content is incorrect because starting with Data ONTAP 8.0.1, deduplication can be enabled on volumes up to 16 TB for all platform models.

Changes to the Gateway Implementation Guide

There are some changes to the existing content of the *Gateway Implementation Guide*.

- “Changes to the Gateway Implementation Guide for Native Storage Expansion Units”
- “Format change for the Gateway Implementation Guides” on page 148

Changes to the Gateway Implementation Guide for Native Storage Expansion Units

New information has become available since the previous revision of the *Gateway Implementation Guide for Native Storage Expansion Units*

In the *Gateway Implementation Guide for Native Storage Expansion Units*, Appendix B, "Setting up and Managing Multipath Storage" says that to use Multipath Storage with gateway systems you must configure the `fc-non-array-adapter-list` environment variable. The information in this chapter does not specify that using this variable is required only for ports on FC adapters; it is not required for ports on SAS adapters.

Format change for the Gateway Implementation Guides

Details about deploying all storage arrays with Gateway systems are now provided in a single implementation guide.

Previously, details about deploying gateway systems with storage arrays were provided in separate vendor-specific implementation Guides (for example, the *Gateway Implementation Guide for IBM Storage*). Details about all vendors' storage arrays are now covered in a single implementation guide called the *Gateway Implementation Guide for Third-Party Storage*. The guide contains a separate section for the information specific to each vendor.

Changes to the Gateway MetroCluster Guide

In the Inter-Switch Links (ISL) section of the *Gateway MetroCluster Guide* it is incorrectly stated that only one ISL is supported per fabric. Data ONTAP supports using one or two ISLs, depending on the configuration. The Gateway Support Matrix contains information about configuration if you have one or two ISLs.

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
<http://publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp>

Copyright and trademark information

This section includes copyright and trademark information, and important notices.

Copyright information

Copyright ©1994 - 2012 Net App, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2012 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp is a licensee of the CompactFlash and CF Logo trademarks.

NetApp NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.



NA 210-05892_A0, Printed in USA

GA32-0723-15

