



IBM System Storage N series
Data ONTAP 7.1.3 Filer Release Notes

Copyright and trademark information

Copyright information

Copyright © 1994 - 2008 Network Appliance, Inc. All rights reserved.

Portions copyright © 2007, 2008 IBM Corporation. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Portions © 1998–2001 The OpenSSL Project. All rights reserved.

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.

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, copyright © Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS "AS IS" CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.

- All advertising materials mentioning features or use of this software must display the following acknowledgment:

 - This product includes software developed by the University of California, Berkeley and its contributors.

 - Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS

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.

Portions of the software were created by Netscape Communications Corp.

The contents of those portions are subject to the Netscape Public License Version 1.0 (the "License"); you may not use those portions except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/NPL/>.

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is Mozilla Communicator client code, released March 31, 1998.

The Initial Developer of the Original Code is Netscape Communications Corp. Portions created by Netscape are Copyright © 1998 Netscape Communications Corp. All rights reserved.

This software contains materials from third parties licensed to Network Appliance Inc. which is sublicensed, and not sold, and title to such material is not passed to the end user. All rights reserved by the licensors. You shall not sublicense or permit timesharing, rental, facility management or service bureau usage of the Software.

Portions developed by the Apache Software Foundation (<http://www.apache.org/>). Copyright © 1999 The Apache Software Foundation.

Portions Copyright © 1995–1998, Jean-loup Gailly and Mark Adler

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions Copyright © 1995 University of Southern California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

Portions of this product are derived from version 2.4.11 of the libxml2 library, which is copyrighted by the World Wide Web Consortium.

Network Appliance modified the libxml2 software on December 6, 2001, to enable it to compile cleanly on Windows, Solaris, and Linux. The changes have been sent to the maintainers of libxml2. The unmodified libxml2 software can be downloaded from <http://www.xmlsoft.org/>.

Copyright © 1994–2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>.

Software derived from copyrighted material of the World Wide Web Consortium is subject to the following license and disclaimer:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.

Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [Date-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>"

Notice of any changes or modifications to the W3C files, including the date changes were made. (We recommend you provide URLs to the location from which the code is derived.)

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. 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 Network Appliance.

The product described in this manual may be protected by one or more U.S. 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

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: IBM, the IBM logo, AIX, System Storage.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor, StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance 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.

Table of Contents

Chapter 1	About the Data ONTAP 7.1.3 Release	1
Chapter 2	Changes introduced in Data ONTAP 7.1.3	3
Chapter 3	New and changed features	5
	New platform and hardware support	6
	Manageability enhancements	7
	Protocol related enhancements	8
	Data protection enhancements	11
	Storage resource management enhancements	12
	Networking-related enhancements	13
	Windows-related improvements	14
	Clustering	15
	Licensing changes	16
Chapter 4	New and changed commands and options	17
Chapter 5	Requirements for running Data ONTAP 7.1.3	27
Chapter 6	Important considerations	29
Chapter 7	Known problems and limitations.	45
	Manageability issues	48
	File access protocol issues	51
	Block access protocol issues	53
	Data protection issues.	54
	Storage management issues.	58
	Cluster configuration issues	61

Data ONTAP® 7.1.3 is a maintenance release in the Data ONTAP 7.1 release family. This release contains fixes for functionality defects as well as security vulnerabilities. We believe these vulnerabilities require us to invoke our CERT notification. Data ONTAP 7.1.3 fixes the issues we uncovered and we strongly recommend that you upgrade now. Customers running prior versions of Data ONTAP 7.1 are strongly recommended to upgrade to Data ONTAP 7.1.3, as soon as feasible, if their platforms are supported in Data ONTAP 7.1.3.

Please see the “[Requirements for running Data ONTAP 7.1.3](#)” on page 27 of these release notes for platforms supported in Data ONTAP 7.1.3.

Please note that Data ONTAP 7.1.3 is the last maintenance release in the 7.1 release family and no regularly scheduled patch releases are planned at this time. Customers requiring fixes to bugs already fixed in Data ONTAP 7.2.x are recommended to upgrade to Data ONTAP 7.2.5.1, if their platforms are supported in Data ONTAP 7.2. For up to date information regarding the latest release of Data ONTAP, see the following Web page:

www-1.ibm.com/support/docview.wss?rs=1147&uid=ssg1S7001894

About these release notes

All the information in this document applies to the Data ONTAP 7.1.3 release for N series storage systems, also sometimes called *filers* or *appliances*.

Note

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

All the information in this document applies to the Data ONTAP 7.1.3 release for IBM® System Storage® N series storage systems. Most of the information in this document also applies to gateway systems running Data ONTAP 7.1.3, except where otherwise indicated. For more information regarding gateway systems running Data ONTAP 7.1.3, see the *Data ONTAP 7.1.3 Gateway Release Notes*.

Where to find more information

For additional information about Data ONTAP 7.1.3, visit the IBM NAS product web page at:

<http://www.ibm.com/storage/nas/>.

Support information, including documentation in PDF format, can be found at the following web page:

<http://www.ibm.com/storage/support/nas/>.

Upgrade information

If you are upgrading to this release, you should read the following items:

- ◆ **Important considerations**

This information helps you identify and resolve issues that might affect the availability of your systems. See “[Important considerations](#)” on page 29.

- ◆ ***Data ONTAP 7.1.2 Upgrade Guide***

This document describes how to upgrade to the Data ONTAP 7.1.2 release. It provides information that you need before upgrading your software or reverting it to an earlier version.

Note

If this is the first release you are installing in the Data ONTAP® 7.1 release family, it is not necessary to read this section. You can go directly to “[New and changed features](#)” on page 5.

SnapLock

The Data ONTAP 7.1.3 release supports SnapLock®. However, you cannot enable deduplication on a SnapLock volume.

Other new features

For information about the new features in Data ONTAP 7.1 release family, see “[New and changed features](#)” on page 5.

Problem fixes

Data ONTAP 7.1.3 includes fixes to a small set of problems reported by customers since the Data ONTAP 7.1.2 release.

This section covers features that were added or changed in the Data ONTAP® 7.1 release family. The topics that follow describe the major changes; see also “[New and changed commands and options](#)” on page 17.

- ◆ “[New platform and hardware support](#)” on page 6
- ◆ “[Manageability enhancements](#)” on page 7
- ◆ “[Protocol related enhancements](#)” on page 8
- ◆ “[Data protection enhancements](#)” on page 11
- ◆ “[Storage resource management enhancements](#)” on page 12
- ◆ “[Networking-related enhancements](#)” on page 13
- ◆ “[Windows-related improvements](#)” on page 14
- ◆ “[Clustering](#)” on page 15
- ◆ “[Licensing changes](#)” on page 16

New platform and hardware support

ESH2 firmware version 18

Data ONTAP 7.1.2.1 and later includes ESH2 firmware version 18.

N3700 support for EXN1000 storage expansion unit

Data ONTAP 7.1.1 and later releases support the ability to attach IBM System Storage N series EXN1000 (2861-001) storage expansion units to IBM System Storage N series N3700 (2863-A10) or (2863-A20) systems. This ability allows you to implement cost-effective SATA-based external storage for the Fibre Channel-based N3700 series.

A minimum of five SATA drives installed in each EXN1000 is required for this feature. These SATA drives are 250 GB, 320 GB, and 500 GB in capacity. A maximum of three EXN1000 storage expansion units can be added to any one N3700 system.

Note

You can attach either EXN1000 or EXN2000 storage expansion units to a single N3700 system. Attaching both EXN1000 and EXN2000 storage expansion units to a single N3700 is not supported.

The maximum raw capacity of 16 TB for the N3700 system remains the same, regardless of the number of shelves attached or the sizes of the drives in the shelves.

Manageability enhancements

Password aging and improved password security

Provides support for site-determined password aging policies and support for RFC 2307.

Dynamic DNS (DDNS)

Dynamic updates to the Domain Name System (DNS) allow the storage system to automatically update DNS servers with its DNS information. DNS updates have to be enabled on the storage system as well as the DNS server.

Slave NIS

You can now configure a storage system to maintain a slave copy of the Network Information System (NIS) maps. Unless the NIS maps are very large and are updated frequently, using a slave NIS database improves performance and reduces network traffic.

SNMP updates

Data ONTAP 7.1.1 includes a new Simple Network Management Protocol (SNMP) trap in the IBM custom Management Information Base (MIB) file. The new maxdirsize trap causes an alert to be sent when the 64MB directory limit is nearly reached or has been reached.

Protocol related enhancements

Blocks protocols

Blocks protocol enhancements include the following:

- ◆ Improved support for iSCSI
 - ❖ Support for multiple iSCSI connections per session

Multiple iSCSI connections per session improves high availability and bandwidth aggregation. The multiple iSCSI connection feature provides sophisticated error recovery techniques that reduce the impact of error conditions on applications.
 - ❖ Support for iSCSI error recovery levels 1 and 2

Improves data transmission error handling and recovery with initiators that also support this feature.
- ◆ Emulex PCI-X iSCSI hardware target support

Data ONTAP 7.1.1 and later releases support the following Emulex PCI-X iSCSI hardware target adapters:

 - ❖ Dual-port GbE iSCSI target adapter (copper) FC 1010
 - ❖ Dual-port GbE iSCSI target adapter (optical) FC 1011

The following features are supported with these adapters:

 - ❖ VLAN Tagged Interface
 - ❖ Jumbo frames
 - ❖ Cluster failover
 - ❖ vFiler™ functionality
 - ❖ Multi-connection Sessions (MCS)
 - ❖ Error Recovery Levels 0, 1, and 2

The following features are not supported with these adapters:

 - ❖ VLAN GVRP registration
 - ❖ VIFS/Link aggregation
- ◆ Single_image FCP port mode

Single_image FCP port mode is a new cfmode. With single_image mode, all LUNs in the cluster are exported on all FCP target ports in the cluster. The cluster has a single global FCP nodename. Any request that storage system A receives that cannot be mapped to a LUN mapped on storage system A is sent across the interconnect to storage system B. LUN maps are checked across both storage systems for conflicts.
- ◆ FCP Port sets

FCP Port sets can be used when FCP single_image mode is on, and with non-clustered systems, to provide the ability to define sets of target and initiator ports. By binding an igroup to a port set, you can control which LUNs can be seen by which ports.

- ◆ 4 Gb target HBA support

Data ONTAP 7.1.1 and later releases support the PCI-x 4 Gb PCI-x Fibre Channel target host bus adapter (HBA).

Note

Only single_image and standby cfmodes are supported with this HBA. If you upgrade from a prior version of Data ONTAP and want to use the 4Gb HBA, you must manually change the cfmode to single_image. For more information, see "Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations."

For more information about blocks protocols, see the *Data ONTAP Block Access Management Guide*.

For more information about CIFS, FTP, HTTP, LDAP, and NFS protocol support, see the *Data ONTAP File Access and Protocols Management Guide*.

LDAP

LDAP (Lightweight Directory Access Protocol) enhancements are summarized here (for more information see the *Data ONTAP File Access and Protocols Managements Guide*):

- ◆ LDAP over SSL

LDAP (Lightweight Directory Access Protocol) over SSL (Secure Sockets Layer) adds the ability to encrypt and authenticate LDAP traffic using cryptographic certificates over SSL. (SSL support already exists on Data ONTAP, but previously LDAP could not use it.) Secure LDAP is supported against all LDAP servers, including OpenLDAP and Windows Active Directory. Mutual (client side) authentication is not currently supported.

- ◆ LDAP authentication using Simple Authentication and Security Layer (SASL)

The Simple Authentication and Security Layer (SASL) is a method for adding authentication support to connection-based protocols. By using SASL, a user can be authenticated to an LDAP server without passing the user password in clear text. SASL requires a mechanism to be negotiated between the client and the server before doing the user authentication. There are several industry standard authentication mechanisms that can be used

with SASL, including GSSAPI for Kerberos V, DIGEST-MD5, and Plain and External for use with Transport Layer Security (TLS).

Data protection enhancements

SnapLock

SnapLock is supported in Data ONTAP 7.1.3, however, you cannot enable deduplication on a SnapLock volume.

LockVault

LockVault enhances existing SnapLock functionality and builds SnapVault features on top of basic and enhanced SnapLock functionality. LockVault provides a solution for sites where data must be retained for specific periods of time. For example, LockVault allows for compliance with various regulations, such as SEC 17a-4 and Sarbanes-Oxley.

Resynchronizing vFiler unit systems after disaster recovery

The `vfiler unit dr resync` command enables you to resynchronize your original vFiler unit with the currently activated disaster-recovery vFiler unit before bringing up the original vFiler unit.

VERITAS NetBackup & SnapVault

Data ONTAP 7.1.1 includes two new integration points with VERITAS NetBackup software. First, NetBackups (NBUs) can now be used to configure & manage SnapVault relationships. Second, SnapVault on NearStore now provides a space-efficient backup target for NBU-managed backups on third-party primary storage.

Support for SCSI reservations for the Symantec VERITAS NetBackup 6 Shared Storage Option

Data ONTAP 7.1.1 supports SCSI Reserve/Release commands for the Symantec VERITAS NetBackup 6 Shared Storage Option, which allows users to dynamically share tape devices with NDMP backup policies. To enable the SCSI Reserve/Release commands, set the `tape.reservations` option to `scsi`. For more information, see the `na_options(1)` man page.

Support for an LSI Logic RoHS-compliant SCSI tape card

Data ONTAP 7.1.1 supports an LSI Logic SCSI tape HBA (FC 1016). This adapter is the Reduction of Hazardous Substances (RoHS) compliant replacement for the previous FC 1016 HBA. It provides equivalent functionality and is supported on N5000 series storage systems.

Storage resource management enhancements

The storage resource management enhancements are described below

Maintenance Center and disk health management

Maintenance Center adds an improved disk health assessment test, and, if a drive does not meet the requirements of health assessment, Maintenance Center determines if the disk can still be used as a spare or if it needs to be returned to IBM.

For more information, see the *Data ONTAP Storage Management Guide*.

Performance improvements for Microsoft Exchange environments

WAFL extents are a performance improvement feature for Microsoft Exchange environments. WAFL extents are off by default and controlled through the `vol` command options `vol-name extent`.

Policy-based space management

Policy-based space management provides a variety of space reservation settings including thin provisioning or under provisioning and fractional reservations. Policy-based space management helps a storage administrator keep the amount of unutilized space on a storage system to a minimum. Policy-based space management includes these functions:

- ◆ Snapshot autodelete
- ◆ Volume autosize

Networking-related enhancements

Support for a RoHS-compliant quad-port copper (RJ-45) PCI-X Network Interface Card (NIC) (FC 1007)

Data ONTAP 7.1.2 and later releases support a RoHS-compliant quad-port copper (RJ-45) PCI-X Network Interface Card (NIC) (FC 1007).

Note

The FC 1007 NIC is supported on the N5200, N5500, N5200 gateway and N5500 gateway platforms, on which it replaces the previous level FC 1007 NIC. It is not supported on the N5300 and N5600 gateway platforms, which will continue to use the previous level FC 1007 NIC

Support for a ROHS-compliant quad-port Fibre-Channel (FC) Initiator PCI-X host bus adapter (HBA) (FC 1025)

Data ONTAP 7.1.2 and later releases support a ROHS-compliant quad-port Fibre-Channel (FC) Initiator PCI-X host bus adapter (HBA) (FC1025)

Note

The FC1025 NIC is supported on the N5200 and N5500 platforms, except those using Fabric MetroCluster connections. It is not supported on gateway and N5600 platforms.

Support for a RoHS-compliant dual-port Ethernet card

Data ONTAP 7.1.1 supports a RoHS-compliant dual-port copper 10/100/1000 Ethernet card (FC 1008).

Support for the Brocade SilkWorm 200E 8- and 16-port switches for SnapMirror

Data ONTAP 7.1.1 supports the Brocade SilkWorm 200E 8- and 16-port switches for SnapMirror.

Windows-related improvements

The Windows-related improvements are described below.

Improved GPO support

Data ONTAP 7.1.1 includes improved support for Windows Group Policy Objects (GPOs), which enable policy-based administration for computers in an Active Directory environment. Supported GPOs include a File System security policy, an Event Log, Auditing, and startup and shutdown scripts. For details, see the File Access and Protocols Management Guide.

Native file blocking

Data ONTAP 7.1.1 provides new file screening software that runs natively on the storage system, in addition to existing support for third-party file screening software that runs on client systems. Native file blocking provides simple denial of restricted file types. For more information, see the File Access and Protocol Management Guide.

Clustering

Negotiated failover for network interface failures

Negotiated failover provides a framework that allows nodes in a cluster to exchange health status information using up to four NIC interfaces. If one node is healthy and its partner is not, the healthy node takes over. If both nodes become impaired, no takeover is triggered.

Support for the Brocade SilkWorm 200E 8- and 16-port switches for Fabric Attached MetroClusters

Data ONTAP 7.1.1 supports the Brocade SilkWorm 200E 8- and 16-port switches for Fabric Attached MetroClusters.

Licensing changes

There are no licensing changes.

This section provides information about the commands, options, and configuration files that have been changed or added to the Data ONTAP 7.1 release family. These changes are described in the following topics:

- ◆ “[New commands in Data ONTAP 7.1](#)” on page 17
- ◆ “[Changed commands in Data ONTAP 7.1.1](#)” on page 19
- ◆ “[Replaced or removed commands in Data ONTAP 7.1](#)” on page 20
- ◆ “[New configuration files in Data ONTAP 7.1](#)” on page 21
- ◆ “[New options in Data ONTAP 7.1.1](#)” on page 21
- ◆ “[Replaced or Removed options in Data ONTAP 7.1.1](#)” on page 25

New commands in Data ONTAP 7.1

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

Command	Purpose	Documentation
<code>cifs gresult</code>	Displays information about the Windows Group Policy Objects (GPOs) currently applicable to the storage system and the results of applying them.	na_cifs(1) man page <i>Data ONTAP File Access and Protocols Managements Guide</i>
<code>cifs gupdate</code>	Updates GPOs on the storage system with the most current Group Policy settings available in an Active Directory domain.	na_cifs(1) man page <i>Data ONTAP File Access and Protocols Managements Guide</i>
<code>disk checksum</code> {disk_name all} [-c block zoned]	Specifies the checksum type (block or zone) for one or all ATA drives. The default is block checksums.	na_disk(1) man page
<code>iscsi alias</code>	Creates a new iSCSI alias.	na_iscsi(1) man page

Command	Purpose	Documentation
iscsi portal	Displays a list of the portals (IP address and/or TCP port number), and their portal group assignments, over which the storage system operates the iSCSI service.	na_iscsi(1) man page
iscsi tgroup	Manages the assignment of storage system network interfaces to target portal groups.	na_iscsi(1) man page
keymgr	Manages security certificates.	na_keymgr(1) man page
portset help	Lists portset commands.	na_portset(1) man page
portset add	Adds a port to a portset.	na_portset(1) man page
portset create	Creates a new portset.	na_portset(1) man page
portset destroy	Destroys portsets.	na_portset(1) man page
portset remove	Removes ports from a portset.	na_portset(1) man page
portset show	Displays portsets.	na_portset(1) man page
snap autodelete	Automatically deletes Snapshots when the flexible volume is nearly full.	na_snap(1) man page
storage show disk -x	Displays disk-specific information including serial number, vendor name, and model.	na_storage(1) man page
vol autosize	Configures a FlexVol volume to grow automatically when it is nearly full.	na_vol(1) man page <i>Data ONTAP Storage Management Guide</i>
vol create trad_volname [-L [compliance enterprise]]	Specifies type of SnapLock volume being created. Only needed if both SnapLock Compliance and SnapLock Enterprise are licensed.	na_vol(1) man page <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>

Command	Purpose	Documentation
vol options vol-name no_i2p [on off]	Disables the inode to pathname scanner.	na_vol(1) man page
vol options vol-name extent	Enables application writes to be written in the volume as a write of a larger group of related data blocks called an extent.	na_vol(1) man page
vol options vol-name try_first [volume_grow snap_delete]	Defines how Data ONTAP reclaims space when a volume is nearly full specifically, whether to automatically increase the volume first or whether to automatically delete Snapshot copies first.	na_vol(1) man page

Changed commands in Data ONTAP 7.1.1

For each command, the following table gives this information:

- ◆ The change in the command
- ◆ The location of documentation about the feature

Command	Change	Documentation	Release introduced
fcg config	Now includes the speed option, which allows you to change the speed setting for an adapter (in GB/second). Depending on the capability of your adapter, you can set the speed to 4, 2, 1, or auto (auto-negotiation, which is the default).	na_fcg(1) man page <i>Data ONTAP Block Access Management Guide</i>	7.1

Command	Change	Documentation	Release introduced
ifconfig	Now includes the <code>-nfo</code> option which specifies that negotiated failover is to be disabled for the network interface. This option applies only to storage systems in a cluster.	na_ifconfig(1) man page <i>Data ONTAP Network Management Guide</i>	7.1
iscsi stats	Now reports statistics for the entire storage system, instead of for an individual adapter.	na_iscsi(1) man page	7.1
useradmin	Now includes the options <code>-m</code> and <code>-M</code> . The <code>-m</code> option specifies the minimum allowable age of the user's password (in days) before the user can change it again. The <code>-M</code> option specifies the maximum allowable age of the user's password (in days).	na_useradmin(1) man page <i>Data ONTAP System Administration Guide</i>	7.1
vol options volname [optname optval]	Now adds the maximum retention period infinite for the options <code>snaplock_maximum_period</code> , <code>snaplock_default_period</code> , and <code>snaplock_minimum_period</code> <code>fconfig</code> .	na_vol(1) man page	7.1

Replaced or removed commands in Data ONTAP 7.1

The following commands are deprecated or removed:

- ◆ `dafs`, removed
- ◆ `iswt` interface, replaced by `iscsi` interface
- ◆ `iswt session show`, replaced by `iscsi session show`

- ◆ iswt connection show, replaced by iscsi connection show
- ◆ iscsi show initiator, replaced by iscsi initiator show
- ◆ iscsi config, removed
- ◆ iscsi show adapter, removed

New configuration files in Data ONTAP 7.1

There are no new configuration files in Data ONTAP 7.1.2.

New options in Data ONTAP 7.1.1

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 allowed values, or the type of the value, or an example value used with the option

Option	Purpose	Default value or example
cf.takeover.on_network_interface_failure	Allows negotiated takeover to be enabled when the cluster nodes detect failures in network interfaces. Only those network interfaces that have explicitly enabled negotiated failover via the ifconfig command will be monitored.	off (on off)
cf.takeover.on_network_interface_failure.policy	Determines what policy to apply for triggering negotiated failover when network interfaces fail.	all_nics (all_nics any_nics)
cf.takeover.on_short_uptime	Determines whether a cluster failover will happen if a storage system fails within sixty seconds of booting up.	on (on off)

Option	Purpose	Default value or example
<code>cifs.audit.nfs.enable</code>	Enables CIFS auditing of files in UNIX security style qtrees. When enabled, auditable events performed on these files are recorded in the log file. Auditable events are specified by the Windows System Access Control Lists (SACLs) either on the file itself, or on the file specified in the value of <code>cifs.audit.nfs.filter.filename</code> .	<code>off (on off)</code>
<code>cifs.audit.nfs.filter.filename</code>	Points to the filter file used to identify which file events get included in the CIFS log by default. When an operation is performed on UNIX security style qtree files that have no SACL set, the SACL set on the file identified by this option is used to determine which events get logged.	<code>none</code>
<code>cifs.client.dup-detection</code>	Enables/disables GPO support on storage systems. Directs Windows servers to attempt to detect duplicate sessions in order to terminate any sessions that did not terminate when a client system rebooted.	<code>name (name ip-address off)</code>
<code>cifs.gpo.enable</code>	Enables/disables GPO support on storage systems.	<code>on off (default off)</code>
<code>cifs.gpo.trace.enable</code>	Enables/disables display of GPO diagnostic information.	<code>on off (default off)</code>
<code>disk.maint_center.spares_check</code>	Enables/disables checking the number of available spares before Data ONTAP puts a disk that has exceeded the threshold of media errors into the maintenance center.	<code>on off (default on)</code>

Option	Purpose	Default value or example
iscsi.isns.rev	Determines the draft level of the iSNS specification the iSNS service on the storage system is compatible with.	22 (example iSNS specification draft level)
iscsi.max_connections_per_session	Specifies the number of iSCSI connections per session allowed by the storage system.	The default is use_system_default, which currently equals 1.
iscsi.max_error_recovery_level	Specifies the maximum iSCSI error recovery level allowed by the storage system.	The default is use_system_default, which currently equals 0.
kerberos.replay_cache.enable	<p>Prevents passive replay attacks by storing user authenticators on the filer, and by insuring that the authenticators are not reused by attackers.</p> <hr/> <p>Note Storing and comparing the user authenticators can result in a substantial performance penalty for higher workloads on the filer.</p> <hr/>	off (on off)
nfs.notify.carryover	Specifies that when set to off, the hosts present in the /etc/sm/notify file are not sent Network Status Monitor (NSM) reboot notifications after a storage system panic/reboot.	on (on off)

Option	Purpose	Default value or example
nfs.v4.acl.enable	Specifies that when enabled, Access Control Lists (ACLs) are supported for NFS version 4. The ACL option controls setting and getting NFSV4 ACLs. It does not control enforcement of these ACLs for access checking. This feature is not supported over NFS versions 2 and 3.	off (on off)
nis.netgroup.domain_search.enable	Specifies whether netgroup entry comparisons will consider the domain names in the search directive from <code>/etc/resolv.conf</code> .	on (on off)
nis.slave.enable	Enables an NIS slave on the storage system.	off (on off)
security.passwd.firstlogin.enable	Controls whether all administrators (except for root) must change their passwords upon first login. A value of on means that newly created administrators, or administrators whose passwords were changed by another administrator, may only run the passwd command until they change their password again.	off (on off)
security.passwd.lockout.numtries	Controls how many attempts an administrator can make to log in before the account is disabled. This account may be reenabled by having a different administrator change the disabled administrator's password. If the new password value is 0, then it remains unused, and failing to log in will never disable an account.	4,294,967,295

Option	Purpose	Default value or example
security.passwd.rules.history	Controls whether an administrator can reuse a previous password. A value of 5 means that the storage system will store 5 passwords, none of which an administrator can reuse. A value of 0 means that an administrator is not restricted by any previous password. Make sure that the security.passwd.rules.enable command has the value on or security.passwd.rules.history is ignored.	0
snapvault.lockvault_log_volume	Configures the LockVault Log Volume. Valid values for this option are online SnapLock volume names.	sl_vol1
tape.reservations	Enables SCSI reservations or persistent reservations for all tape drives, medium changers, bridges, and tape libraries (including those with embedded bridges) attached to the filer via Fibre Channel, including those attached through switches.	off (scsi persistent off)

Replaced or Removed options in Data ONTAP 7.1.1

The following options have been replaced or removed in this release:

- ◆ cifs.restrict_anonymous.enable (replaced by cifs.restrict_anonymous)
- ◆ iscsi.iswt.tcp_window_size (replaced by iscsi.tcp_window_size, although iscsi.iswt.tcp_window_size continues to be the command for backwards compatibility).
- ◆ tape.persistent_reservations (replaced by tape.reservations)

This section lists the storage systems and firmware you need in order to run Data ONTAP 7.1.3 on individual and active/active configuration systems.

Supported systems

The following storage systems are supported:

- ◆ N5200 and N5500 storage systems
- ◆ N3700 storage systems

Notes:

- ◆ If you want to run IBM diagnostics software on the 4 Gb Fibre Channel target HBAs, you must use diagnostics version 4.6.5 or later. Earlier versions of diagnostics software do not recognize 4 Gb Fibre Channel target HBAs.
- ◆ For information about installing 4 Gb Fibre Channel target HBAs in storage system expansion slots, see “[Only single_image and standby cfmodes are supported with 4-Gb HBAs](#)” on page 53 in these release notes, and the *Data ONTAP Introduction and Planning Guide*.

Systems you can use in clusters

Clusters are supported on the following IBM storage systems with Fibre Channel connections to disk shelves:

- ◆ N5500 (2865-A10 or 2865-A20)
- ◆ N5200 (2864-A10 or 2864-A20)
- ◆ N3700 (2863-A10 or 2863-A20)

N3700 series systems support clusters regardless of whether the shelf connected to the system is FC-AL or SATA.

The IBM System Storage EXN1000 SATA expansion unit is available for attachment to the N3700 when running Data ONTAP 7.1.1 or later.

Maximum total capacity supported

For information about your storage system model's capacity and maximum volume size, see the *Data ONTAP Introduction and Planning Guide*, located on the IBM NAS support Web site:

<http://www.ibm.com/storage/support/nas>

Storage system firmware

For information about the latest storage system firmware, see the NAS support Web site:

www.ibm.com/storage/support/nas

Note

The latest system firmware is included with Data ONTAP upgrade packages for CompactFlash®-based storage systems. For more information, see the *Data ONTAP Upgrade Guide*.

Disk firmware

For information about the latest disk firmware, see the NAS support Web site:

www.ibm.com/storage/support/nas

Note

New disk firmware is sometimes included with Data ONTAP upgrade packages. For more information, see the *Data ONTAP Upgrade Guide*.

The following items can help you identify and resolve issues that might affect the operation of your N series storage system.

- ◆ [“If you want to use SnapLock Compliance on Data ONTAP 7.1, an upgrade to DataONTAP 7.1.3 is mandatory”](#) on page 30
- ◆ [“Data ONTAP 7.1.1 is a disruptive install for customers that have been running iSCSI host sessions”](#) on page 30
- ◆ [“If you have AT-based disk shelves in an active/active configuration”](#) on page 33
- ◆ [“If you are reverting to a Data ONTAP release with a lower maximum capacity”](#) on page 34
- ◆ [“If you are using FPolicy”](#) on page 34
- ◆ [“If you have AT-based disk shelves in an active/active configuration”](#) on page 33
- ◆ [“If you are upgrading from a 2 Gb onboard port to a 4 Gb adapter in a SAN environment”](#) on page 35
- ◆ [“If you upgrade storage systems with iSCSI targets”](#) on page 36
- ◆ [“If you are running DataFabric Manager with a Business Continuance License enabled”](#) on page 36
- ◆ [“Corequisites for NDMP support with Storage Manager”](#) on page 36
- ◆ [“Replacing traditional volumes with flexible volumes”](#) on page 36
- ◆ [“Monitor error logs and spare drive counts”](#) on page 36
- ◆ [“Upgrading storage systems that contain two or more non-original disk drives”](#) on page 37
- ◆ [“Upgrading systems running SnapMirror”](#) on page 38
- ◆ [“Replacing disk drives”](#) on page 39
- ◆ [“Disabling automatic scheduling of snapshots”](#) on page 39
- ◆ [“Unexpected warnings and messages regarding flexible volumes”](#) on page 40
- ◆ [“Customers are strongly encouraged to enable AutoSupport”](#) on page 40
- ◆ [“If you are replacing one of the controllers in a N3700 system in an active/active configuration”](#) on page 42
- ◆ [“If you are updating disk firmware or you are upgrading Data ONTAP software in systems with EXN1000 or EXN2000 storage expansion units \(such as N3700 systems\)”](#) on page 43
- ◆ [“If you have a SnapLock Compliance license”](#) on page 43

If you want to use SnapLock Compliance on Data ONTAP 7.1, an upgrade to DataONTAP 7.1.3 is mandatory

If you want to use SnapLock Compliance and you are running Data ONTAP 7.1 or later, you must upgrade to Data ONTAP 7.1.3. Earlier versions of the Data ONTAP 7.1 family no longer support SnapLock Compliance.

In addition, storage systems that do not include any SnapLock Compliance volumes but are mirrored, copied, or cascaded from SnapLock Compliance volumes must also be upgraded. Volume SnapMirror®, vol copy, and aggr copy operations will fail if the source system is running a supported SnapLock release and the destination system is running a Data ONTAP release in the 7.1 family that is earlier than 7.1.3.

Also, storage systems that do not include any SnapLock Compliance volumes but are in a fail-over cluster or MetroCluster configuration with a system that includes SnapLock Compliance volumes must be upgraded. At failover, the SnapLock compliance volume will be taken offline if the original system is running a supported SnapLock release and the failover system is running a Data ONTAP release in the 7.1 family that is earlier than 7.1.3.

Data ONTAP 7.1.1 is a disruptive install for customers that have been running iSCSI host sessions

An IBM N series branded filer running Data ONTAP 7.1 may have its default iSCSI Qualified Name (iqn) start with a prefix set to: "iqn.1992-08.com.netapp". The full nodename can be displayed by running the `iscsi nodename` CLI command. An iSCSI host will use the iqn assigned to the iSCSI target to perform a login sequence in order to access LUNs hosted by the target.

A change was made to Data ONTAP 7.1.1 which sets the default iqn prefix to: "iqn.1986-03.com.ibm" on IBM branded N series filers. If the filer is upgraded from a version of Data ONTAP whose default nodename starts with the prefix, "iqn.1992-08.com.netapp", to a version of Data ONTAP whose default nodename starts with the prefix, "iqn.1986-03.com.ibm", iSCSI hosts unaware of this change may encounter failures with the iSCSI login phase because it will use the old default nodename to attempt to login. This can lead to application failure due to the inability to access data.

iSCSI sessions will need to be stopped and restarted after hosts are configured with the correct iSCSI target nodename.

To allow hosts to access LUN's after an iqn name change, the following general steps should be taken:

Procedure for making iSCSI Configuration Changes after an upgrade to Data ONTAP 7.1.1 in a dual cluster configuration:

1. On the first controller, download Data ONTAP 7.1.1.

2. On the first controller, disabled cluster takeover capability (issue `cf disable` command).
3. Reboot the first controller.
4. On the second controller, download Data ONTAP 7.1.1.
5. On the second controller, disable cluster takeover capability (issue `cf disable` command).
6. Reboot the second controller.
7. Use **Host Procedures to Regain Access to LUNs after iSCSI Configuration Changes** (see below).
8. Enable takeover capability of both controllers in the cluster (issue `cf enable` command on either controller).
9. Restart I/O.

Note

For more detailed upgrade procedures consult the *Data ONTAP 7.1.1 Upgrade Guide*.

Procedure for making iSCSI Configuration Changes after an upgrade to Data ONTAP 7.1.1 in a single cluster configuration:

10. Download Data ONTAP 7.1.1.
11. Reboot the controller.
12. Use **Host Procedures to Regain Access to LUNs after iSCSI Configuration Changes** (see below).
13. Restart I/O.

Note

For more detailed upgrade procedures consult the *Data ONTAP 7.1.1 Upgrade Guide*.

Host Procedures to Regain Access to LUNs after iSCSI Configuration Changes:**Procedure for AIX Hosts**

1. Unmount file systems using `iscsi luns`.
2. Issue `varyoffvg` (vary off the `iscsi lun volume group(s)`).

3. Issue `exportvg` (export the volume group).
4. Issue `rmdev -dl hdiskX` (remove disk device(s)).
5. Edit `/etc/iscsi/targets` to change iqn name.
6. Issue `cfdmgrp` (rescan the bus for luns).
7. Issue `importvg` (import the volume group(s)).
8. Issue `varyonvg` (vary on the volume group(s)).
9. Mount filesystems.

Note

Consult the specific iSCSI host support kit for more detailed information at:
<http://www-03.ibm.com/servers/storage/support/nas/iSCSIHost/installing.html>

Procedure for Unix Hosts

Note

These are generic procedures, make sure to consult your specific iSCSI Host Support Kit documentation prior to performing this upgrade

1. Unmount file systems.
2. Issue `iscsi stop` on host.
3. Change iqn name on the controller(s).

Note

Make sure to consult the specific iSCSI Host Support Kit for your installation.

4. Issue `iscsi start`.
5. Scrub `/etc/fstab` to update device paths using the current target iqn (use name changed in Step 3).

Note

This filename may vary on the host you are upgrading, make sure to consult the specific iSCSI Host Support Kit for your installation.

6. Mount file systems

Note

Consult the specific iSCSI host support kit for more detailed information at:
<http://www-03.ibm.com/servers/storage/support/nas/iSCSIHost/installing.html>

Procedure for Windows Hosts

1. Run the Microsoft iSCSI Initiator applet
2. Refresh iSCSI initiator targets
3. Click the new nodenames
4. Log onto the new nodenames

Note

Refer to the following documentation for more detailed information:

iSCSI initiator documentation	
C:\WINDOWS\iSCSI\uguide.doc	User's guide, includes CLI reference.
C:\WINDOWS\iSCSI\readme.txt	Initiator configuration information.
C:\WINDOWS\iSCSI\relnotes.txt	Initiator and host configuration and troubleshooting information.

If you have AT-based disk shelves in an active/active configuration

During disk shelf firmware updates for AT-FC, AT-FC2, and AT-FCX modules, the data on the disk shelf cannot be accessed until the firmware update is complete. If you are planning a nondisruptive upgrade for a system with any of these shelf modules attached, you must first determine whether the disk shelf module firmware for your system is up-to-date.

Attention:

You cannot use the nondisruptive method to upgrade Data ONTAP if you have AT-FC, AT-FC2, or AT-FCX-based disk shelves in an active/active configuration, and if the module firmware for these shelves is outdated.

Data ONTAP upgrade packages include the latest firmware versions for a number of system components, including disk shelf modules. During a Data ONTAP upgrade, the disk shelf module firmware is updated automatically if the version in the download package is later than the version installed in your system.

Client services might encounter delays accessing data when disk shelf firmware is updated on AT-FC, AT-FC2, or AT-FCX modules.

Systems with the following disk shelf modules are affected:

- ◆ AT-FCX
- ◆ AT-FC2
- ◆ AT-FC

If you are reverting to a Data ONTAP release with a lower maximum capacity

When you revert to an earlier Data ONTAP release, your storage system must conform to the maximum capacity limitations of the earlier release. If you upgraded your system to a release that supports greater capacities and you configured storage to utilize the new capacities, you must reconfigure your system to the lower capacity limits before you revert. If you don't reconfigure in this way, the storage system will not boot up following the revert process until the excess capacity has been disconnected.

To access any storage in excess of system limits, you must disconnect it from the system you are reverting. Excess capacity must be disconnected for relocation to a different storage system so that aggregates remain intact.

The system to which storage is relocated must meet the following requirements:

- ◆ It has spare capacity room to accommodate the relocated storage.
- ◆ It is running the same Data ONTAP release to which the previous system is being reverted.
- ◆ It is running a Data ONTAP release that supports the relocated disks.

For more information about physically moving aggregates, see the *Data ONTAP Storage Management Guide*. For more information about maximum capacity limits for a given Data ONTAP release, see the *Data ONTAP System Configuration Guide* entries for that release.

If you are using FPolicy

If you use FPolicy (for file screening and other supported features), do not upgrade to the Data ONTAP 7.1 release family. FPolicy is not supported in the Data ONTAP 7.1 release family. We recommend you upgrade to Data ONTAP 7.2.2 or later releases.

If you are upgrading a system that includes FlexVol volumes

If your storage system includes FlexVol volumes and you are planning to upgrade to Data ONTAP 7.1.1 from an earlier release family, confirm that no FlexVol volumes on your system are marked "waf1 inconsistent." If you find any inconsistent FlexVol volumes, contact technical support immediately and correct the problem before beginning the upgrade process.

Attention: If your storage system contains a waf1 inconsistent FlexVol volume, do not upgrade to Data ONTAP 7.1.1 from an earlier release family. If you do, the inconsistent volume will no longer be accessible and the data it contains will effectively be lost. Reverting to an earlier release is not possible in this case.

If you are upgrading from a 2 Gb onboard port to a 4 Gb adapter in a SAN environment

If you have a N5000 series storage system in a SAN environment and you are upgrading from a 2 Gb onboard port to a 4 Gb adapter, you must stop the FCP service and reconfigure the onboard port before installing the 4 Gb adapter. Failure to do so might result in a storage system panic and lost data.

Note

When you upgrade from the onboard ports to the 4 Gb HBAs, the World Wide Port Names (WWPNs) will change, which requires that you reconfigure any UNIX hosts with the newly assigned WWPNs

To reconfigure the 2 Gb onboard port, complete the following steps:

1. Stop the FCP service by entering the following command:

```
fcv stop
```

Result: The FCP service is stopped and all target adapters are taken offline.

2. Set the embedded FC ports to unconfigured:

```
fcadmin config -t unconfig <ports>
```

Result: The embedded FC ports are unconfigured.

Example:

```
fcadmin config -t unconfig 0b 0d
```

3. Ensure the cfmode is set to `single_image` or `standby`.
4. Shut down the storage system.
5. Install the 4 Gb adapter according to the instructions provided with the adapter.
6. Power on the storage system.

If you upgrade storage systems with iSCSI targets

In Data ONTAP 7.1, the iSCSI target portal group tags have changed from the values used in previous releases. You must reconfigure any Linux or HP-UX iSCSI hosts that use these tags or they will not be able identify the iSCSI target provided by the IBM storage system. For more information, see the *Data ONTAP 7.1.1 Upgrade Guide*, located on the IBM NAS support Web site at:

<http://www.ibm.com/storage/support/nas>

If you are running DataFabric Manager with a Business Continuance License enabled

If you are running DataFabric Manager with a Business Continuance License enabled, you must upgrade to DataFabric Manager 3.0.1R1 or later. Not doing so will cause SnapVault relationship updates to fail and display the following error message:

"destination requested snapshot that does not exist on the source"

Corequisites for NDMP support with Storage Manager

NDMP support with IBM Tivoli® Storage Manager has the following corequisites:

- ❖ Tivoli Storage Manager Version 5.3.2.2 or later

When configuring TSM for NDMP operations on the N series storage system, the NETAPDUMP data format must be used for the data mover and storage pool definitions.

Replacing traditional volumes with flexible volumes

Flexible volumes have different best practices, optimal configurations, and performance characteristics than traditional volumes. Users are encouraged to ensure that they understand these differences by referring to the FlexVol™ volume documentation and deploy the configuration that is optimal for your environment.

For detailed information about configuring flexible volumes and aggregates, see the *Data ONTAP Storage Management Guide*.

Monitor error logs and spare drive counts

Users are encouraged to periodically monitor error logs and spare drive counts to ensure all drives are available. Remove and replace failed drives as soon as possible to ensure adequate spares.

Upgrading storage systems that contain two or more non-original disk drives

Under certain conditions, it is possible to inadvertently enable software-based ownership on a storage system that previously used a hardware-based ownership scheme. When this occurs, volumes on disks that have no software-based ownership information will be unavailable, which may prevent Data ONTAP from starting up. Though unlikely, there is also a slight potential for data to be lost, as described later in this discussion of disk ownership during upgrades.

To learn more about this and to ensure that you do not have problems at startup and do not risk data loss, read and use the information in the following sections.

About software-based disk ownership:

Data ONTAP provide the ability for you to designate disk ownership, rather than letting Data ONTAP automatically assign ownership based on the location of the disk drive. This ownership information gets stored on the disk. When Data ONTAP starts, it determines ownership by retrieving any ownership information on the disk. If no software ownership information is present, hardware disk ownership is assigned based on disk location.

Note

When a disk drive is added to a hardware-based ownership system while Data ONTAP is running, Data ONTAP automatically removes any software-based disk ownership information from the disk; ownership information can also be removed manually by booting the storage system in maintenance mode and using the `disk remove_ownership` command.

N3700 uses software-based ownership. N5000 uses hardware-based ownership.

How software-based disk ownership can be enabled:

Accidental enabling of software-based disk ownership happens only either of the following conditions are met:

- ◆ A storage system that uses hardware-based disk ownership has two or more disk drives replaced while Data ONTAP is not running. (If the disk drives are replaced while Data ONTAP is running, the software-based ownership information is automatically deleted from the disks.)
- ◆ The replacement disk drives have not had ownership information removed after they were installed in the software-based ownership environment.

When either of these conditions are met, the storage system converts to using a software-based ownership scheme. If this happens, only disks with software-based ownership information on them will be available at startup. Data ONTAP might fail to boot, and there is a small possibility that data could be lost.

Preventing and correcting accidental enabling of software-based disk ownership:

To avoid accidental conversion to software-based disk ownership, complete the following steps the first time you start the storage system after the upgrade process.

1. Watch the console as you boot the storage system for the first time following the upgrade process.
2. If you see the message "software ownership has been enabled for this system," type Ctrl-C to display the special boot menu. If you don't see this message, the storage system is still using hardware-based ownership; discontinue this procedure and boot the storage system in normal mode.
3. Select **Maintenance mode boot** (option 5).
4. Use the `disk remove_ownership` command to remove ownership information from the replacement disks.

This problem can also occur after the upgrade, if you replace disk drives while the storage system is not running.

Note

If you have already experienced an accidental conversion to software-based ownership, use this same procedure to start the storage system in maintenance mode and remove ownership information from the replacement disks.

Upgrading systems running SnapMirror

If you are upgrading Data ONTAP on storage systems that are running the SnapMirror software and are using volume SnapMirror replication, it is important that you upgrade storage systems that have SnapMirror destination volumes before you upgrade storage systems that have SnapMirror source volumes. This is necessary because when using SnapMirror volume replication, the destination volume must run under a version of Data ONTAP equal to, or later than, that of the SnapMirror source volume.

Also, be aware that synchronous SnapMirror is not allowed to achieve synchronous mode (In-Sync) when the versions of Data ONTAP do not match. If the destination is upgraded first, SnapMirror will automatically do per-minute updates instead of synchronous updates until the source is upgraded to the same version of Data ONTAP. Once both source and destination are using the same version of Data ONTAP, synchronous mode is achieved automatically.

For further important information about updating storage systems that are running SnapMirror, see the *Data ONTAP 7.1.1 Upgrade Guide*.

Replacing disk drives

Disk drives should only be replaced with drives supported by IBM for the N series product. Other disk drives may not function properly.

Disabling automatic scheduling of snapshots

When a volume or aggregate is created on the N series storage system, snapshots are automatically scheduled for that volume or aggregate. If automatic snapshots are not desired, you can disable them following volume creation from the graphical FilerView™ interface or following volume or aggregate creation from the command line.

From FilerView, do the following:

1. Click **Volumes > Snapshots > Manage**.
2. Click on the volume name.
3. Clear the **Scheduled Snapshots** checkbox and click **Apply**.

From the command line, type the following to disable snapshots for volumes:

```
vol options <volume name> nosnap on
```

From the command line, type the following to disable snapshots for aggregates:

```
aggregate options <aggr name> nosnap on
```

To recover the space reserved for the snapshots, issue the following commands:

For an aggregate:

```
snap reserve -A <aggr> 0
```

For a volume:

```
snap reserve <vol> 0
```

The automatic snapshot schedules can also be modified to fit your needs. For information, refer to the *Online Backup and Recovery Guide*.

Unexpected warnings and messages regarding flexible volumes

If an aggregate is unexpectedly taken offline because of a hardware failure, for example if you have an aggregate on a loop and the loop fails, you might see unexpected warnings and messages. Data ONTAP assigns a unique file system identifier (FSID) to every flexible volume on a system. Data ONTAP cannot check the FSIDs of offline flexible volumes. Data ONTAP does detect the possibility of FSID conflict and issues warnings and messages alerting you to the possibility of duplicate FSID creation.

During the time the aggregate is offline you may see warnings and messages during operations such as these:

- ◆ Creating a new flexible volume, including by cloning

In this case you will see a warning. You can override the warning, but if you do, you might see additional problems.

- ◆ Bringing a flexible volume online

Before the flexible volume is brought online, Data ONTAP checks the FSID. The error message you will see is:

```
n3700-1> vol create temp1 aggr0 2g
Creation of volume 'temp1' with size 2g on containing aggregate
'aggr0' has completed.
n3700-1> Thu Apr 6 17:31:03 GMT [n3700-1:
volaggr.offline:CRITICAL]: Some aggregates are offline. Volume
creation could cause duplicate FSIDs.
```

If there is a duplicate FSID you see the following message:

```
vol online: FSID marked invalid because it already exists.
Unable to bring volume 'volname' online
```

The volume does not come online. You cannot override this warning. Knowledge Base article KB4361 provides additional information.

- ◆ Accessing of a volume by a vFiler unit

You might see a message from the vFiler unit about a volume being in an unexpected state. The volume reported might not be an offline volume. NFS clients will not be able to access the reported vFiler volume. You can fix this problem by first taking the volume offline: `vol offline vol1`. Then bring the volume back online: `vol online vol1`.

Customers are strongly encouraged to enable AutoSupport

With the IBM System Storage N series storage controllers, the Autosupport feature was disabled in the following Data ONTAP releases.

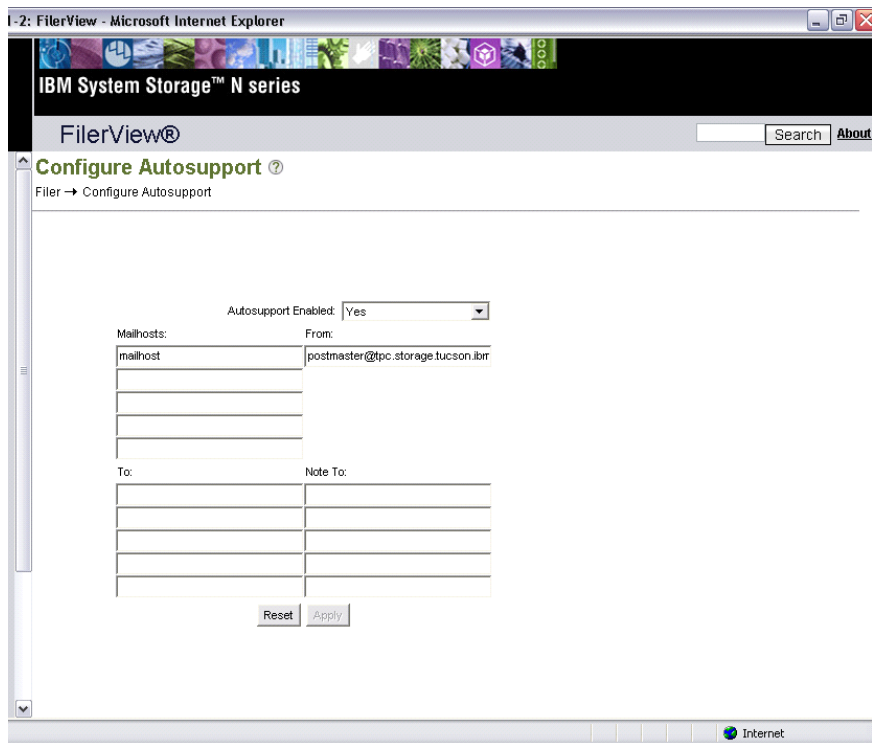
- ◆ Data ONTAP 7.1h2
- ◆ Data ONTAP 7.1.0.1p1

◆ Data ONTAP 7.1.0.1p2

IBM strongly encourages customers to enable Autosupport. Autosupport enables the IBM System Storage N series machines to send diagnostic information back to IBM when error conditions are encountered within the machines. Autosupport can be enabled via the command line interface via the following command:

```
options autosupport.enable on
```

Autosupport can also be enabled via the FilerView web-browser user interface by selecting **Filer** in the left-hand navigation frame and then selecting **Configure autosupport**, as shown in the following screen capture:



If you have ever run a Data ONTAP release on an N series machine where the Autosupport is off by default you should check to make sure that the option is enabled (on). Installing a new Data ONTAP release where the default is different does NOT automatically change the option (it does not change the current setting). Autosupport will be enabled (default setting is on) in the following Data ONTAP release:

◆ Data ONTAP 7.2

So, installing Data ONTAP 7.2 (where the Autosupport default is on) on an existing N series machine currently running Data ONTAP 7.1.0.1p2 (where the Autosupport default is off) does not change the current setting to on (enabled).

In addition to enabling Autosupport, IBM customers should not change the Autosupport transport protocol default of HTTPS to any other value (HTTP or SMTP). The only Autosupport transport protocol supported by IBM is HTTPS. Changing the Autosupport transport protocol, to any other value besides the default of HTTPS, will result in no Autosupport messages being delivered to IBM.

IBM Customers are reminded to work with their IBM sales, marketing or account representative to make sure that their customer information in IBM's RETAIN Common Customer Profile Facility (CCPF) is accurate. By doing so, it will enable:

- ◆ the Autosupport messages to be received by IBM and
- ◆ a problem record (PMR) to be opened in IBM's problem reporting database (RETAIN) and
- ◆ the problem record to be queued to the appropriate IBM support team and
- ◆ provide the necessary information for the IBM support team to contact the customer regarding the problem.

If you are replacing one of the controllers in a N3700 system in an active/active configuration

If you perform a hot plug-in of one of the controllers in an N3700 storage system in an active/active configuration, a system panic will occur. This situation occurs if you disable the active/active configuration, halt both of the nodes, physically remove the controller you are replacing, boot the other controller, and then hot plug-in the new controller.

The panic message that displays on the console is as follows:

```
PANIC: ic_do_read(): Out of desc xid's. in process cf_firmware
```

To avoid this situation, make sure both controllers are installed before booting either one of the controllers. If you rebooted one of the controllers while the other one was not installed, it is necessary to halt the node, install the missing controller, and boot both controllers.

If you are updating disk firmware or you are upgrading Data ONTAP software in systems with EXN1000 or EXN2000 storage expansion units (such as N3700 systems)

Disk firmware is automatically updated during the Data ONTAP upgrade whenever the disk drives have a firmware version older than the firmware version that is included with the Data ONTAP release. As disk firmware is updated, it is possible (although unlikely) that all disk drives on a shelf will receive the update at the same time. This simultaneous update might cause the firmware update process to fail, rendering the disks unusable and causing loss of data if the failure occurs to multiple disks in the same RAID group.

If you have AutoSupport enabled, you can use AutoSupport to verify your disk firmware version in advance and find out whether your system is vulnerable. If AutoSupport indicates that all disks on a disk shelf need firmware updates, failure to follow these instructions could cause disk drive damage and data loss.

For instructions on how to use AutoSupport to detect outdated disk firmware, see the *Data ONTAP Upgrade Guide*.

If you have a SnapLock Compliance license

There are two types of Snaplock Licenses in Data ONTAP: SnapLock Compliance and SnapLock Enterprise. When you display licenses for your storage system, the Snaplock Compliance license is listed as snaplock and the Snaplock Enterprise license is listed as snaplock_enterprise. This information applies to Snaplock Compliance licenses only.

If you upgrade to any Data ONTAP version released on or after March 7, 2005, you must obtain a unique updated SnapLock Compliance license code from the Protocol Licenses page. Previous SnapLock Compliance license codes have been replaced with the updated SnapLock Compliance license codes on this page.

Important:

If you upgrade an IBM N series storage system to any Data ONTAP version released on or after March 7, 2005, without changing the license code, the attributes and SnapLock protection on already existing SnapLock files (and SnapLock volumes) will not be affected in any way. But because the old license is no longer functional, *you will not be able to commit any new files to SnapLock protected status until the license code is updated*, nor will you be able to create new SnapLock volumes.

This section describes potential problems and limitations identified with the Data ONTAP® 7.1 release family.

- ◆ Storage management issues:
 - ❖ [“NetBackup Managed SnapVault does not support SnapLock Compliance \(WORM\) volumes”](#) on page 59
 - ❖ [“Change in `sysstat -x` command output”](#) on page 58
 - ❖ [“Data ONTAP can silently select varying disk sizes when enlarging or creating an aggregate”](#) on page 58
 - ❖ [“Default quota requirement”](#) on page 58
 - ❖ [“Partition alignment problems cause an iSCSI LUN performance problem”](#) on page 58
 - ❖ [“Qtree path not shown in quota report for new qtrees”](#) on page 58
 - ❖ [“Qtree quotas may prevent explicit quotas from overriding the default”](#) on page 59
 - ❖ [“ASUPs for bypassed disks may be sent erroneously”](#) on page 60
- ◆ Block access protocol issues:
 - ❖ [“Only `single_image` and `standby` cfmodes are supported with 4-Gb HBAs”](#) on page 53
 - ❖ [“Ongoing cfmode support”](#) on page 53
 - ❖ [“iSCSI target HBA support”](#) on page 53
- ◆ Manageability issues:
 - ❖ [“Mapping SNMP traps to EMS”](#) on page 48
 - ❖ [“4 Gb HBA is not recognized by earlier diagnostics software”](#) on page 48
 - ❖ [“Synchronous SnapMirror and flexible volumes”](#) on page 48
 - ❖ [“SSH Tectia client does not work properly unless forwarding is disabled”](#) on page 48
 - ❖ [“Maximum number of SSH sessions”](#) on page 49
 - ❖ [“Long-running commands invoked using SSH may be terminated”](#) on page 48
 - ❖ [“SSH authentication agent forwarding is not supported”](#) on page 48
 - ❖ [“Pseudo-terminals are not supported for SSH”](#) on page 49
 - ❖ [“PuTTY command line editing does not work properly”](#) on page 49
 - ❖ [“Non-local users cannot use login protocols”](#) on page 49

- ❖ “FilerView does not support roles and capabilities for all tasks” on page 49
- ❖ “Reversion cannot take place during inode-to-pathname initialization” on page 49
- ◆ File access protocol issues:
 - ❖ “FPolicy names must be shorter than 80 characters” on page 51
 - ❖ “NFSv4 client compatibility” on page 51
 - ❖ “NFS configuration issue for Hummingbird with Kerberos” on page 51
 - ❖ “Client notification messages in Windows domains require NetBIOS” on page 51
 - ❖ “Unsupported Windows features in the file serving environment” on page 52
 - ❖ “Do not use vlan delete -q with ipsec” on page 52
 - ❖ “iSNS service and Microsoft iSNS Server v.3.0” on page 52
 - ❖ “Enable caching with NIS lookup” on page 52
 - ❖ “Group Policy Objects (GPOs) are applied differently on storage systems and Windows systems” on page 52
- ◆ Data protection issues:
 - ❖ “SnapLock Qtree SnapMirror (QSM) resynchronization restrictions” on page 54
 - ❖ “SnapLock ComplianceClock restrictions” on page 54
 - ❖ “Restore incorrectly estimates amount of space” on page 55
 - ❖ “Incremental restore from tape might not work” on page 55
 - ❖ “Stagger Snapshot schedules to avoid transfer conflicts” on page 55
 - ❖ “Synchronous SnapMirror and storage system performance” on page 56
 - ❖ “The vscan feature does not support NIS authentication” on page 56
 - ❖ “Compatibility with SnapDrive and Snap Manager” on page 56
 - ❖ “FlexVol volumes cannot be migrated using SnapMover” on page 56
 - ❖ “Potentially reduced performance for IPsec traffic” on page 56
- ◆ Cluster configuration issues:
 - ❖ “CFO partner node's SSL certificate is used after takeover and giveback” on page 61
 - ❖ “Possible I/O errors during cluster takeover or giveback on heavily loaded systems” on page 61
 - ❖ “Complete setup through the console” on page 61
 - ❖ “Failed disks can cause giveback to fail” on page 61

- ❖ “SnapMover cannot be used on MetroClusters” on page 61

Manageability issues

Mapping SNMP traps to EMS

The Syslog Translator contains information about Event Management System (EMS) messages and the associated SNMP traps. It does not provide a mapping from SNMP traps to EMS messages.

4 Gb HBA is not recognized by earlier diagnostics software

If you want to run IBM diagnostics software on the 4 Gb target HBAs, you must use diagnostics version 4.6.5 or later. If you run an earlier version of diagnostics software on a N5000 series system running Data ONTAP 7.1.1, you will see an error message indicating that the 4 Gb HBA is an unrecognized device.

Synchronous SnapMirror and flexible volumes

A filer cannot be both the source of a synchronous SnapMirror relationship using flexible volumes and the destination of another synchronous SnapMirror relationship using flexible volumes.

Long-running commands invoked using SSH may be terminated

If you invoke a command that takes longer than one minute to complete using the non-interactive SSH protocol, it may be terminated before it completes. To avoid this issue, use the interactive SSH protocol or the less secure non-interactive **rsh** protocol.

SSH Tectia client does not work properly unless forwarding is disabled

If you use the SSH Tectia client from SSH Communications Security to access your storage system, you must disable all forwarding. If you do not disable forwarding, in non-interactive mode the client authenticates but does not run the specified command, and the SSH session remains in use until the storage system is rebooted.

SSH authentication agent forwarding is not supported

Data ONTAP does not support SSH authentication agent forwarding.

Maximum number of SSH sessions

Data ONTAP supports a maximum of 12 SSH sessions.

Pseudo-terminals are not supported for SSH

SSH pseudo-terminals such as `pty` are not supported by Data ONTAP. This may cause some clients to produce a message stating that "Server refused to allow `pty`". Though generally benign, this limitation might cause certain client shortcuts to fail.

PuTTY command line editing does not work properly

If you use PuTTY to access your storage system, and you are using SSH2, you might see the message "Server refused to allow `pty`" and command line editing features might not work properly. If this happens, use SSH1. To use SSH1, set the storage system `ssh2.enable` option to `Off` and set the `ssh1.enable` option to `On`.

Non-local users cannot use login protocols

Non-local users, or users added with the `useradmin domainuser add` command, cannot log in using the login protocols (`telnet`, `console`, `rsh`, `ssh`, or `http-admin`) to a system running Data ONTAP. Non-local users can log in only by using the Windows RPC mechanism, and can use only the API capability once they are logged in. This is true regardless of the capabilities associated with any N series system group the non-local user has been placed in.

FilerView does not support roles and capabilities for all tasks

For some tasks, the FilerView® administration tool accesses the storage system as `root`, regardless of what user name you used to log in to FilerView. This means that when you use FilerView, you might not be limited by which capabilities are associated with your user name. To prevent a user from using FilerView, assign that user only to groups that do not have the `login-http-admin` capability.

Reversion cannot take place during inode-to-pathname initialization

When you upgrade to Data ONTAP 7.1, inode-to-pathname information is written to each existing volume. You cannot revert to an earlier release while the system is being scanned to initialize the inode-to-pathname information. Before proceeding with the reversion, you must do one of the following:

- ◆ Wait until the scan for inode-to-pathname information completes.
 - This can take many hours, depending on the number of inodes and volumes on your system.
- ◆ Abort the scan.

To abort the inode-to-pathname scan, complete the following step.

Enter the following command:

```
vol options no_i2p on
```

Result: Inode-to-pathname translation is disabled immediately and the reversion can proceed.

Note

Because any inode-to-pathname information is removed during a reversion, reverting to an earlier release can take several hours, depending on the number of inodes and volumes on your system, and whether the scan for inode-to-pathname information has been completed.

File access protocol issues

FPolicy names must be shorter than 80 characters

The Fpolicy feature (Data ONTAP file screening policy) allows you to create policy names with a maximum of 80 characters. However, if a policy name has exactly 80 characters, it cannot be enabled, displayed, modified, or deleted. Until this problem is resolved in a later Data ONTAP release, you can work around this problem by ensuring that policy names are no longer than 79 characters.

NFSv4 client compatibility

If you use NFSv4: When your NFSv4 clients are in a different domain than your N series device is in, 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 that needs to access a storage system using NFSv4, ensure that the Data ONTAP option `nfs.v4.enable` is set to `on`. In new installations, this option is set to `off` by default.

If you do not use NFSv4: Ensure that the Data ONTAP option `nfs.v4.enable` is set to `off`.

For more information about NFSv4, see the *File Access and Protocols Management Guide*.

NFS configuration issue for Hummingbird with Kerberos

The Data ONTAP option `nfs.mount_rootonly` should be set to `off` for implementations that require support for Hummingbird software using Kerberos v5 authentication.

Client notification messages in Windows domains require NetBIOS

The Windows client notification feature used for client messaging and shutdown notices requires NetBIOS over TCP to be enabled in Data ONTAP. On Windows clients, NetBIOS over TCP must be enabled and the Windows Messenger service must be running. (Windows 2003 and Windows XP SP2 clients have messaging disabled by default.)

Unsupported Windows features in the file serving environment

Data ONTAP 7.1.1 does not support every available Windows 2000 and Windows XP feature. For example, this release does not support the following:

- ◆ Encrypting File System
- ◆ Logging of NT File System events in the change journal
- ◆ File Replication Service
- ◆ Microsoft Windows Indexing Service
- ◆ Remote storage through Hierarchical Storage Management
- ◆ Local user account creation from the User Manager or Microsoft Manager Console
- ◆ Quota management from Windows 2000 clients
- ◆ Windows 2000 quota semantics
- ◆ The LMHOSTS file
- ◆ NT File System native compression

Do not use `vlan delete -q` with ipsec

It is recommended that you delete VLAN interfaces one by one (that is, do not use `vlan delete -q`) when ipsec is enabled.

iSNS service and Microsoft iSNS Server v.3.0

By default, Data ONTAP 7.1.1 and later works with Microsoft iSNS server releases earlier than v.3.0. If you want to use v.3.0 of the Microsoft iSNS server, set the `iscsi.isns.rev` option to 22:

```
options iscsi.isns.rev 22
```

Group Policy Objects (GPOs) are applied differently on storage systems and Windows systems

Data ONTAP 7.1 and later releases support Group Policy Objects (GPOs) that are relevant to storage system management. However, Event Log and Audit (Local Policies) GPOs are applied differently on IBM storage systems than on Windows systems. For more information, see the *Data ONTAP File Access & Protocols Management Guide*.

Enable caching with NIS lookup

If you use NIS for group lookup services, disabling NIS group caching can cause severe degradation in performance. Whenever you enable NIS lookups (`nis.enable`), it is recommended that you also enable caching (`nis.group_update.enable`). Failure to enable these two options together could lead to timeouts as CIFS clients attempt authentication.

Block access protocol issues

Only single_image and standby cfmodes are supported with 4-Gb HBAs

Only single_image and standby cfmodes are supported with the new 4-Gb Fibre Channel Host Bus Adapters (HBAs) on N5000 storage systems. If you are upgrading from a 2-Gb HBA to a 4-Gb HBA, ensure you change the cfmode to single_image or standby cfmode before upgrading.

If you attempt to run a N5000 system with a 4-Gb HBA in an unsupported cfmode, the 4-Gb HBA is set to offline and an error message is displayed.

In addition, Data ONTAP does not allow changing from a supported cfmode to an unsupported cfmode with the 4-Gb HBA installed on these systems.

Ongoing cfmode support

Starting with the 4Gb support on the N5000 series storage systems, only single_image and standby cfmodes are supported. The remaining cfmodes will continue to be supported on legacy systems, as documented in the *Data ONTAP Block Access Management Guide*.

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 do so, an error message is displayed.

Data protection issues

SnapLock Qtree SnapMirror (QSM) resynchronization restrictions

Beginning with Data ONTAP 7.1.3 in the 7.1 release family, Qtree SnapMirror resynchronization is not supported for mirroring volumes from a less strict to a SnapLock Compliance qtree. When the source is a non-SnapLock or SnapLock Enterprise qtree, and the destination is a SnapLock Compliance qtree, resynchronization is not allowed. The following table shows all the possible combinations and if resynchronization is allowed:

	Destination	SnapLock Compliance	SnapLock Enterprise	Non-SnapLock
Source	SnapLock Compliance	Yes	Yes	Yes
	SnapLock Enterprise	No	Yes	Yes
	Non-SnapLock	No	Yes	Yes

In addition, to allow a QSM resynchronization between the SnapLock Compliance qtrees, you need to make at least one transfer from the SnapLock Compliance qtree source to the SnapLock Compliance qtree destination after the upgrade to Data ONTAP 7.1.3.

SnapLock ComplianceClock restrictions

The ComplianceClock on a SnapLock qtree may be transferred with a Qtree SnapMirror copy to either another SnapLock Compliance qtree or a SnapLock Enterprise qtree, but not to a non-SnapLock qtree. After the qtree is copied, the ComplianceClock is stored on the destination qtree. At a Qtree SnapMirror break, the qtree ComplianceClock is checked against the Data ONTAP ComplianceClock to ensure that the Data ONTAP ComplianceClock is not ahead of the qtree clock. If the Data ONTAP ComplianceClock is ahead of the qtree ComplianceClock, the Data ONTAP compliance clock is drifted back in time until the clocks are synchronized.

Normally, the maximum amount of time that the ComplianceClock can fall behind is bounded by the length of the longest SnapLock Qtree SnapMirror transfer to that system. However there are two cases that the amount of time that the ComplianceClock can fall behind is unbounded. Those two cases are:

- ◆ when the source and destination SnapLock qtree are on the same system

- ◆ when there is a cycle of SnapLock Qtree SnapMirror transfers across the storage systems, such as snap mirroring of one SnapLock qtree from filer A to filer B, and snap mirroring of another SnapLock qtree from filer B to filer A

You can work around these situations by changing your SnapMirror configuration so that the source and destination SnapLock qtrees of every transfer are on different storage systems. Also, change the configurations so that SnapLock Qtree SnapMirror operations are all going "one-way," such as always from storage system A to storage system B and never from storage system B back to storage system A.

Restore incorrectly estimates amount of space

The restoration process using the `restore` command can incorrectly estimate the amount of space needed in a volume to restore the requested data from tape. If the space on the volume is not sufficient, however, the restore, continues and aborts at a later time with a message stating that there is no more space left on the device.

Incremental restore from tape might not work

An incremental restore from tape might not work because tape positioning after a reboot might be different from what is assumed. If you position the tape using the `mt` command or specify the file number by using the `restore` command with the `-s` option, the restore will work.

Stagger Snapshot schedules to avoid transfer conflicts

Data ONTAP provides a default Snapshot schedule for each volume that can be modified using the `snap sched` command. The `snap sched` command configures a per-volume schedule that creates, rotates, and deletes hourly, nightly, and weekly Snapshot copies. If either the SnapMirror or SnapVault feature is scheduled to perform Snapshot management at the same time as a `snap sched` activity, then the Snapshot management operations scheduled using the `snap sched` command might fail with syslog messages, "Skipping creation of hourly snapshot" and "Snapshot already exists".

To avoid this condition, stagger your Snapshot update schedules so that SnapMirror activity does not begin or end at the exact time a `snap sched` operation attempts to create a Snapshot. Additionally, if `snap sched` Snapshot copies conflict with SnapVault activity, use the `snapvault snap sched` command to configure equivalent schedules.

For more information about the SnapMirror and SnapVault features, and for a detailed description of how to use the `snap sched` command, see the *Data ONTAP Online Backup and Recovery Guide*.

Synchronous SnapMirror and storage system performance

The synchronous mode of SnapMirror will have an impact on the performance of a storage system. This impact is dependent on the power of your storage system, its load, and the number of simultaneous synchronous SnapMirror relationships. To keep this impact under control, no more than eight simultaneous synchronous SnapMirror relationships are supported.

Please engage your professional services consultant for help with best-practice configuration for synchronous SnapMirror.

The vscan feature does not support NIS authentication

The storage system anti-virus vscan feature requires NT LAN Manager or Kerberos authentication; it does not support Network Information Service authentication. The storage system validates any vscan server which connects to the storage system, and it requires the vscan server to connect as a user who is in the storage system's Backup Operators group.

Compatibility with SnapDrive and Snap Manager

Data ONTAP 7.1.1 is not supported by all versions of IBM software for Windows clients. Before installing or upgrading to Data ONTAP 7.1.1, make sure that the SnapManager® and SnapDrive™ versions you are running or plan to install are compatible with this release.

FlexVol volumes cannot be migrated using SnapMover

Data ONTAP 7.1 and later releases do not support migrating FlexVol volumes or aggregates containing these volumes using SnapMover software. You can use SnapMover to migrate only traditional volumes between clustered nodes.

Potentially reduced performance for IPsec traffic

Data ONTAP 7.1 and later releases do not support hardware-assisted IPsec encryption and instead use software encryption. If your storage system includes hardware-assisted encryption, and you use IPsec encryption for high volumes of traffic, you might see reduced throughput for the IPsec encrypted traffic.

Storage management issues

Default quota requirement

Users are encouraged to create default user, group, and tree quotas for any volume with quotas enabled. Without these quotas, you might see very high CPU utilization.

Partition alignment problems cause an iSCSI LUN performance problem

Under some circumstances you will get poor performance on iSCSI LUNs used by Linux hosts because of partition alignment problems. Refer to tech tip titled “Partiton alignment problems cause an iSCSI LUN performance problem” for more information.

Qtree path not shown in quota report for new qtrees

In the quota specifier field of quota reports, the fully qualified path to qtrees created in Data ONTAP 7.1 and later is not displayed.

Change in `sysstat -x` command output

Change in `sysstat -x` command output in Data ONTAP 7.1.1 and later releases, the DAFS column is no longer displayed in the `sysstat -x` command output. Note that any scripts that currently depend on `sysstat -x` command output might need to be revised.

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

When you create a new aggregate or add disks to an existing aggregate using automatic disk selection, Data ONTAP selects disks based on various criteria including size, speed, and checksum type. If you have disks of more than one size in your storage system, automatic disk selection can give unexpected results.

For best results if disks of varying size are in use, always specify the desired disk size when creating or adding to an aggregate.

For example, to add eight 68-GB disks to `aggr1`, you could enter the following command:

```
aggr add aggr1 8@68G
```

To find out what disks Data ONTAP will automatically select, you can use the `-n` option for the `aggr create` or `aggr add` command. The `-n` option lists, without performing the creation or addition, the disks that would be selected automatically if you created an aggregate or added to an aggregate. 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.

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

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

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

NetBackup Managed SnapVault does not support SnapLock Compliance (WORM) volumes

NetBackup Managed SnapVault replication preserves all backups in Snapshot copies. By definition, WORM volumes have retention periods set. The Snapshot copies taken by NetBackup Managed SnapVault replication do not have retention periods set and can be deleted by any privileged user. This ability to delete Snapshot copies makes the destination volume noncompliant even if it was originally set to be compliant; therefore, NetBackup Managed SnapVault replication cannot be used for compliance purposes.

Neither the Data ONTAP operating system nor the VERITAS NetBackup application prevent you from using a WORM volume as a destination volume for NetBackup Managed SnapVault replication. You must ensure that you do not select a WORM volume as a destination volume for NetBackup Managed SnapVault replication.

Qtree quotas may prevent explicit quotas from overriding the default

If you have a default user or group quota on a volume, and you wish to override that default quota for a particular user or group, then if you also have one or more qtree quotas on that volume, the explicit quota may not take precedence over the default quota within those qtrees as expected.

To avoid this issue, you can override the default quota on the qtrees that have a quota.

For example, suppose you have the following quota file:

```
* user@/vol/vol1 16K
14717 user@/vol/vol1 1M
/vol/vol1/mytree tree 10M
```

The presence of the qtree quota causes all users, including user 14717, to be limited to 16K of space in the mytree qtree. However, the addition of the following line enables the override:

```
14717 user@/vol/vol1/mytree 1M
```

Now, user 14717 is able to use 1M of space anywhere in vol1.

**ASUPs for
bypassed disks
may be sent
erroneously**

When a disk port is bypassed, Data ONTAP sends an AutoSupport message similar to this: IN/OUT ports of the ESH or ESH2 disk module have been bypassed, even if there is no problem with any of the disks. If you receive this AutoSupport message, run the 'storage show hub' command on the storage system in question. If the results show that no disk ports are bypassed, and one or more shelf ports (shown as IN or OUT) are bypassed, that is probably the cause of the AutoSupport message. Replace any modules with bypassed ports, and ensure that all OUT ports are properly terminated.

Note

If no disks or shelf ports are shown as bypassed, no further action is necessary. This message can be caused by a temporary condition that resolves itself without intervention.

Cluster configuration issues

CFO partner node's SSL certificate is used after takeover and giveback

If you have two clustered nodes, each with its own certificate, and one node takes over the other node, then after you issue the `cf giveback` command, both nodes may present the certificate for the takeover node. If this happens, restart SSL on the node that was taken over.

Possible I/O errors during cluster takeover or giveback on heavily loaded systems

In certain environments (such as an AIX or Windows hosts accessing the storage system using the iSCSI protocol) if an event occurs that results in a cluster takeover or giveback (such as a hardware failure) while the storage system is heavily loaded, the takeover or giveback may take an extended period of time. This delay can cause I/O errors on the AIX or Windows hosts as a result of I/O timeouts being triggered.

In these environments, users are encouraged to monitor the CPU load on a storage system to ensure that the storage system is not overloaded. The performance can be monitored using Filer At-A-Glance from the FilerView interface or by using the `sysstat` command on the storage system. IBM suggests keeping storage systems CPU loads at or below 70%.

Complete setup through the console

When performing the initial setup from the console, complete the setup through the step of setting the gateway TCP/IP address before attempting to use the GUI. If you exit the setup command earlier, you might not be able to connect to your filer through the GUI because the TCP/IP configuration has not been completed.

Failed disks can cause giveback to fail

In rare instances, if a storage system in an active/active configuration has failed disks, giveback can fail. To avoid this issue, which is related to disk reservations, remove any failed drives before entering the giveback command. If you have problems during the giveback process, contact technical support.

SnapMover cannot be used on MetroClusters

If you have the `cluster_remote` license, which is required for MetroCluster configurations, you cannot also license the SnapMover functionality.

This section corrects or adds to information published in the Data ONTAP 7.1.1 manuals.

Changes to the Upgrade Guide

Upgrade from Data ONTAP 7.1 could take longer than usual: Replace the existing topic of the same name on page 22 with this topic: When you upgrade from a release earlier than Data ONTAP 7.1, Inode to Pathname (I2P) information is written to each existing volume. Depending on the number of volumes on your system and the number of files in each volume, this could cause the upgrade to take more time than previous Data ONTAP upgrades. However, the I2P updates take place in the background and have no significant effect on system availability or performance.

After upgrading to Data ONTAP 7.1 or later from an earlier release, new I2P (Inode to Pathname) mappings are created for all files and directories in the volume. This is done by a WAFL scan utility that updates small amounts of metadata for all files and directories in the volume and then creates a new metadata file for volumes with files that have hard links. The I2P updates are done in the background; the scan has minimal impact on the your systems performance. The time to complete the I2P scan on a volume is dependent on the volume size and the number of files in the volume.

Note

Data ONTAP 7.1 quota initialization is terminated when you halt, reboot, or upgrade your storage system. For more information, see the *Data ONTAP Storage Management Guide*.

Slow data replication and backup transfers after upgrade: Insert this new topic on page 22 after the topic "Upgrade from Data ONTAP 7.1 and earlier could take longer than usual ": After upgrading to Data ONTAP 7.1 or later from an earlier release, data replication and backup features will require significantly more time because all updated I2P data needs to be transferred. After the initial transfer with I2P mappings, subsequent transfers will be accomplished within normal time periods.

The metadata that is updated and added by the I2P scan affects the data replication and backup features of Data ONTAP, including volume SnapMirror, qtree SnapMirror, and SnapVault operations. This is because the extra data that is updated and added on the source system is sent to the destination or backup system. Any SnapMirror updates than occur while the I2P scan is running on the

source system, and the first SnapMirror update that occurs after the I2P scans are completed will result in the transfer of a larger amount of data and therefore might require significantly more time.

After upgrading Data ONTAP on the source system, these large data transfers will occur only until all the I2P data is sent to the destination system. Any subsequent SnapMirror updates will be of regular size, and will therefore be accomplished in a normal time period.

About disk firmware upgrades: Replace the second bullet item on page 66 with the following text:

- ◆ Data ONTAP detects disk firmware updates in the `/etc/disk_fw` directory.
- ◆ Data ONTAP scans the `/etc/disk_fw` directory for new disk firmware every two minutes.

Updating shelf firmware: The following text should be inserted after page 72:

When you upgrade Data ONTAP, shelf firmware is updated automatically if the firmware on the shelves is older than the firmware that is bundled with the Data ONTAP system files. You can also update shelf firmware by downloading the most recent firmware for your shelves from the IBM support Web site and installing the files.

Attention:

During the firmware update for the AT-FC, AT-FC2, and AT-FCX modules, the data on the disk shelf cannot be accessed until the firmware update is complete. The disruption lasts about 70 seconds per module. That is, updates for the A module cause a disruption of about 70 seconds, then several minutes later another 70-second disruption occurs while updates for the B module are taking place.

It is strongly recommended that you do not place firmware files for AT-FCX modules in the `/etc/shelf_fw` directory unless you intend to update shelf firmware immediately. Several events in Data ONTAP can trigger an automatic shelf firmware update if there is a shelf firmware file in the `/etc/shelf_fw` directory that has a higher revision number than the current firmware on the shelf module.

Complete the following steps to download and update disk shelf firmware:

1. Find and download the most recent firmware for your shelves from the IBM support Web site located at:
`http://www.ibm.com/storage/support/nas`
2. Extract your firmware files to the `/etc/shelf_fw` directory in the root volume of your storage system.

3. Enter the following command at the storage system console to access the advanced administrative commands:

```
priv set advanced
```

The prompt now displays an asterisk (*) after the storage system name to indicate that you are in the advanced mode:

```
filename*>
```

4. Depending on your upgrade scenario, enter one of the following commands to upgrade the shelf firmware.

If you want to upgrade the shelf firmware on...	Then enter the following command at the storage system console:
All the disk shelves in your storage system	storage download shelf
The shelves attached to a specific adapter	storage download shelf <i>adapter_name</i>

5. To confirm that you want to upgrade the firmware, enter y for yes.
6. Enter the following command to verify the new shelf firmware:
7. Enter the following command to return to the standard administrative console prompt:

```
sysconfig -v
```

```
priv set
```

The prompt returns to the standard console prompt:

```
filename>
```

Updating disk shelf firmware

Insert this new section on page 72 following the existing section "Upgrading disk firmware."

Make sure that you update to the latest disk shelf firmware version when you upgrade Data ONTAP. In some upgrade scenarios, disk shelf firmware updates are mandatory.

About disk shelf firmware updates: When you upgrade Data ONTAP, disk shelf firmware (firmware for modules on disk shelves) is updated automatically if the firmware on the shelves is older than the firmware that is bundled with the

Data ONTAP system files. You can also update disk shelf firmware by downloading the most recent firmware for your shelf modules from the NOW site and installing the files.

The module (AT series, ESH series, LRC, or SAS) in a disk shelf maintains the integrity of the loop when disks are swapped and provides signal retiming for enhanced loop stability. In AT- and ESH-based shelves, there are two modules in the middle of the rear of the disk shelf, one for Channel A and one for Channel B. SAS modules are internal components in N3300 and N3600 storage systems. Updated firmware for these modules is made available periodically.

Each storage system is shipped with an `/etc/shelf_fw` directory that contains the latest disk shelf firmware versions available at that time.

Disk shelf firmware updates can be added to this directory at the following times:

- ◆ After a Data ONTAP upgrade
Disk shelf firmware updates are often included in Data ONTAP upgrade packages. If the version in `/etc/shelf_fw` is higher than the installed version, the new version will be downloaded and installed during the reboot or `cf giveback` phase as part of the Data ONTAP upgrade process.
- ◆ During a manual firmware update
You might need to download a disk shelf firmware update from the NOW site if you plan to perform a nondisruptive upgrade of Data ONTAP software, or if you receive a notice from IBM.

Data ONTAP scans the `/etc/shelf_fw` directory for new firmware every two minutes (on systems with software-based disk ownership). If new disk shelf firmware is detected -- that is, if there is a disk shelf firmware file in the `/etc/shelf_fw` directory that has a higher revision number than the current firmware on the shelf module -- the new firmware is automatically downloaded to the disk shelf module.

The following events in Data ONTAP can also trigger an automatic disk shelf firmware update when there is new firmware in the `/etc/shelf_fw` directory:

- ◆ The `reboot` command is issued.
- ◆ The `cf giveback` command is issued.
- ◆ New disk drives are inserted.
- ◆ New shelf modules are inserted.
- ◆ NetApp Health Trigger (NHT) AutoSupport messages are sent

Note

If your system does not use software-based disk ownership, Data ONTAP does not scan the the /etc/shelf_fw directory for new disk shelf firmware. However, the other trigger events are still applicable if software-based disk ownership is not used. For more information about software-based disk ownership, see the *Data ONTAP Storage Management Guide*.

For more information about disk shelves and disk shelf modules, see the *Data ONTAP Active/Active Configuration Guide* and the Hardware and Service Guide for your shelves.

Service availability during disk shelf firmware updates

When you upgrade to the current Data ONTAP release, the availability of storage system services during a disk shelf firmware update depends on the type of shelf modules.

During firmware updates to disk shelves controlled by ESH series or LRC modules, the data on the disk shelf remains accessible.

During firmware updates for disk shelves controlled by AT-FCX, AT-FC, AT-FC2, or SAS modules, the data on the disk shelf cannot be accessed until the firmware update is complete.

The following table summarizes Data ONTAP service availability during disk shelf firmware updates for these modules:

Module	Disk shelf model	Disruption
AT-FCX	EXN1000	about 70 seconds per module *
AT-FC2		
AT-FC		
ESH4	EXN4000 or EXN2000	none
ESH2	EXN1000	
ESH	EXN1000	
LRC	EXN1000	none
SAS	N3300 and N3600 series internal shelves only	about 40 seconds per module *

* "Per module" means that updates for the A module cause a disruption, then several minutes later another disruption occurs while updates for the B module are taking place.

Attention:

You cannot use the nondisruptive method to upgrade Data ONTAP if you have AT-FCX, AT-FC, AT-FC2, or SAS-based disk shelves in an active/active configuration, and if the firmware for these modules is outdated.

Detecting outdated disk shelf firmware

If you want to perform a nondisruptive upgrade of Data ONTAP software when there are AT- or SAS-based disk shelves attached to your system, or if you are directed to update disk shelf firmware, you must find out what firmware is installed on disk shelves attached to your system.

Steps

1. At the storage system command line, enter the following command:

```
sysconfig -v
```

2. Locate the shelf information in the `sysconfig -v` output.

```
Shelf 1: AT-FCX Firmware rev. AT-FCX A: 35 AT-FCX B: 35  
Shelf 2: AT-FCX Firmware rev. AT-FCX A: 35 AT-FCX B: 35
```

3. Go to the disk shelf firmware information on the IBM support Web site and determine the most recent firmware version for your shelves.
4. Take the appropriate action.

If the disk shelf firmware version in the `sysconfig -v` output is the same as the most recent version on the Web site, no disk shelf firmware update is required at this time.

If the disk shelf firmware version in the `sysconfig -v` output is earlier than the most recent version on the Web site, update your disk shelf firmware manually.

Command for updating system firmware

On page 61, the following passage replaces the existing text under "Command for updating system firmware":

Either the `update_flash` command or the `update-flash` command (with a hyphen instead of an underscore) loads the system firmware from media (CompactFlash card or firmware diskette) into your storage systems flash Programmable ROM (PROM).

The `update_flash` command must be run from the boot environment prompt (CFE prompt).

CIFS must be terminated before a giveback operation in a nondisruptive upgrade

The following step should be inserted after the existing Step 12 on page 105:

Choose the option that describes your configuration.

If, in System A, CIFS is...	Then...
Not in use	Go to step 13.
In use	Enter the following command: <pre>cifs terminate -t nn</pre> where <code>nn</code> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to step 13.

Changes to the Cluster Administration and Installation Guide

In Chapter 6 of the *Cluster Installation and Administration Guide*, the section "Hot-swapping a disk shelf module" should include the following caution prior to the start of the procedure:

Caution

If there is newer firmware in the `/etc/shelf_fw` directory than that on the replacement module, the system automatically runs a firmware update, causing a service interruption.

Changes to the Storage Management Guide

The FlexCache technology described in Chapter 6, "Volume Management," of the *Data ONTAP Storage Management Guide* is not available in this Data ONTAP 7.1.1 release.

Changes to the Block Access Management Guide for iSCSI and FCP

In Chapter 5 of the *Data ONTAP Block Access Management Guide*, under "Restrictions on resizing a LUN", it states that Data ONTAP imposes a maximum increase to 2 TBs for a LUN. In reality, the upper limit depends on the host operating system.

Refer to the following table for the upper limits on LUN sizes:

Operating System	Maximum LUN size
Windows	<ul style="list-style-type: none"> ◆ 12 TB for GPT-based disks on Windows 2003 SP1 and later ◆ 2 TB for MBR-based disks on Windows 2003 SP1 and later, and all other versions of Windows
Linux	<ul style="list-style-type: none"> ◆ 2 TB
HP-UX	<ul style="list-style-type: none"> ◆ 2047 GB with HP-UX 11i v1 ◆ 2 TB with all other versions
Solaris	<ul style="list-style-type: none"> ◆ 16 TB with Solaris 10 ◆ 2 TB with Solaris 9, VxVM and appropriate patches ◆ 1023 GB for all other versions
AIX	<ul style="list-style-type: none"> ◆ 2 TB with AIX 5.2ML7 and later and AIX 5.3ML3 and later ◆ 1 TB with Veritas and all other versions of AIX
VMware	<ul style="list-style-type: none"> ◆ 2 TB

The following passage should be added to Chapter 13, "Managing the Fibre Channel SAN."

Changing the adapter speed

Use the `fcg config` command to change the FC adapter speed. These are the available speeds:

- ◆ Autonegotiate (default)

- ◆ 1 Gb
- ◆ 2 Gb
- ◆ 4 Gb

Follow these steps to change the speed.

1. Set the adapter to down (to be offline) using the following command:

```
fcv config adapter down
```

Note

This might temporarily interrupt FCP service on the adapter.

Example

```
: device1> fcv config 2a down
: Wed Jun 15 14:04:47 GMT [device1:
: scsitarget.ispfct.offlineStart:notice]:
: Offlining Fibre Channel target adapter 2a.
: Wed Jun 15 14:04:47 GMT [device1:
: scsitarget.ispfct.offlineComplete:notice]: Fibre Channel
: target adapter
: 2a offlined.
```

2. Enter the following command:

```
fcv config adapter speed [auto|1|2|4]
```

Example

```
: device1> fcv config 2a speed 4
```

3. Enter the following command:

```
fcv config adapter up
```

Result

The new adapter speed is set and the adapter is online again.

Example

```
:device1> fcv config 2a up
:Wed Jun 15 14:05:04 GMT [device1:
scsitarget.ispfct.onlining:notice]:
: Onlining Fibre Channel target adapter 2a.
: device1> fcv config
: 2a: ONLINE Loop No Fabric
: host address 0000da
: portname 50:0a:09:81:96:97:a7:f3 nodename
: 50:0a:09:80:86:97:a7:f3
mediatype auto speed 4Gb
```

Note

Although the `fcv config` command displays the current adapter speed setting, it does not necessarily display the actual speed at which the adapter is running. For example, if the speed is set to `auto`, the actual speed might be 1 Gb, 2 Gb, or 4 Gb. To view the actual speed at which the adapter is running, use the `show adapter -v` command and examine the **Data Link Rate** value (shown in bold below).

Example

```
device1> fcv show adapter -v
Slot: 5a
Description: Fibre Channel Target Adapter 5a (Dual-channel, QLogic
2312 (2352) rev. 2)
Status: ONLINE
Host Port Address: 010200
Firmware Rev: 4.0.18
PCI Bus Width: 64-bit
PCI Clock Speed: 33 MHz
FC Nodename: 50:0a:09:80:87:69:27:ff (500a0980876927ff)
FC Portname: 50:0a:09:83:87:69:27:ff (500a0983876927ff)
Cacheline Size: 16
FC Packet Size: 2048
SRAM Parity: Yes
External GBIC: No
Data Link Rate: 4 GBit
Adapter Type: Local
Fabric Established: Yes
Connection Established: PTP
Mediatype: auto
Partner Adapter: NoneStandby: No
Target Port ID: 0x1
Slot: 5b
Description: Fibre Channel Target Adapter 5b (Dual-channel, QLogic
2312 (2352) rev. 2)
```

**Changes to the
Tape Backup and
Recovery Guide**

The following passage should be added to Chapter 2, "Tape Drive Management."

**Support for SCSI reservations for the Symantec
VERITAS NetBackup 6 Share Storage Option**

Data ONTAP 7.1.1 supports SCSI Reserve/Release commands for the Symantec VERITAS NetBackup 6 Shared Storage Option, which allows users to dynamically share tape devices with NDMP backup policies. To enable the SCSI Reserve/Release commands, set the `tape.reservations` option to `scsi`. For more information, see `na_options(1)` man page.

If you are changing tape reservations from SCSI persistent commands to SCSI reserve and release commands, you must ensure that all persistent reservation keys have been released before changing. To change from SCSI persistent to SCSI-2 tape reservation, complete the following steps.

1. Stop all backup software and tape status checking daemons.
2. From one storage system, enter the following command:

```
options tape.reservations persistent
```
3. From the same storage system, for each attached device exhibiting the reservation problem, enter the following command:

```
storage release device_type device
```
4. After releasing all devices, enter the following command:

```
options tape.reservations scsi
```

Note

Do not issue any commands from any host or storage system that access tapes or changers until after setting the `tape.reservation` option to `scsi`. Doing so might reestablish an unwanted persistent reservation key on one or more devices.

5. Set the `tape.reservation` option to `scsi` on all storage systems.

Note

You can perform this step and the following two steps from one storage system.

6. Ensure that all backup daemons are set to use SCSI-2 reservations, not persistent reservations.
7. Restart all backup daemons.

Changes to the File Access and Protocols Management Guide

The FPolicy feature (which enables file screening policies) is not supported in the Data ONTAP 7.1 release family. The text on pages 295-316 of the *Data ONTAP File Access and Protocols Management Guide* is not currently valid for Data ONTAP 7.1.x releases.

The following sentence should be added to the section "About NetBIOS over TCP" on page 165 in the *Data ONTAP File Access and Protocols Management Guide*:

To verify the status of NetBIOS over TCP on your storage system, use the `nbtstat` command as described in the `nbtstat(1)` man page.

The **Attention** notice under "Tracing CIFS logins" on p. 343 should be updated to read as follows:

Attention

Use CIFS login tracing carefully because it reports every CIFS login. Persistent use can result in excessive console and log messages, which can affect system performance. The `cifs.trace_login` option should be turned on for diagnostic purposes only; it is recommended that it be kept off at other times (the default is off).

Changes to the MultiStore Management Guide

The note on page 14 of the *Data ONTAP MultiStore Management Guide* should be replaced with the following note:

You can move only traditional volumes, not FlexVol® volumes.

Readers' Comments — We'd Like to Hear from You

IBM System Storage N series
Data ONTAP 7.1.3 Filer Release Notes

Publication No. GC26-7862-05

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

E-mail address



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Dept. GZW
9000 South Rita Road
Tuscon, AZ
U.S.A. 85744-0001



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



NA 210-04136_A0, Printed in USA

GC26-7862-05

