



IBM System Storage N series Gateway Installation Requirements and Reference Guide

Copyright and trademark information

Copyright information

Copyright ©1994 - 2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2011 IBM Corporation. All rights reserved.

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

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, FAServer, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, RAID-DP, ReplicatorX, SANscreen, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, and Web Filer are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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

NetApp is a licensee of the CompactFlash and CF Logo trademarks.

NetApp NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Table of Contents

	Preface	xi
Chapter 1	Gateway Overview	1
	How a gateway uses storage	3
	RAID types supported by gateways	5
	Supported methods for connecting to a storage array	6
	Stages of gateway implementation	9
Chapter 2	Major Planning Tasks for a Gateway	11
Chapter 3	Guidelines for Setting Up Access to Array LUNs	17
	LUN security	18
	Requirements for paths to array LUNs	20
	Common errors when setting up access to array LUNs	23
	Valid path setup examples	24
	Invalid path setup examples	29
	How automatic port selection and monitoring works	33
	What happens when a link failure occurs.	36
	How paths are reflected in array LUN names	39
	When to check the path setup to array LUNs.	42
	Using the storage show disk -p command to check paths	43
	Using the sysconfig -v command to check paths	45
	Using the storage array show-config command to check paths	46
Chapter 4	Guidelines for a Fabric-Attached Configuration	49
	General requirements for fabric setup	50
	Switch zoning.	51
	Examples of zoning a gateway with storage arrays	52

	Guidelines for Brocade switches	56
	Guidelines for Cisco fabrics	57
	Guidelines for McDATA switches and directors.	58
	Guidelines for switches with a MetroCluster.	59
Chapter 5	Guidelines for Provisioning Array LUNs	61
	Optimizing storage array RAID groups for LUNs for Data ONTAP	62
	Types of LUNs you can map to gateways	64
	Guidelines for determining the number and size of array LUNs.	65
	Ensuring that spare core LUNs are available for core dumps	70
	How the checksum type affects usable space in an array LUN	71
Chapter 6	Requirements for a Gateway to Use Storage	79
	Requirement for a gateway to use array LUNs.	80
	Requirement for a gateway to be able to use a native disk.	82
Chapter 7	Using Neighborhoods for Storage Management	83
	Overview of a neighborhood	84
	Neighborhood requirements	87
	How a neighborhood is established.	89
	Limits to consider when planning for a neighborhood	91
	Factors that impact the neighborhood maximum LUN limit.	93
	Factor 1—Mixing gateway platforms in a neighborhood	94
	Factor 2—Storage arrays specific restrictions	95
	Factor 3—Using both native disks and array LUNs	96
Chapter 8	Guidelines for Assigning Storage to Aggregates	99
	Rules about mixing storage in aggregates	100
	Rules for aggregates if multiple storage arrays are deployed	102

Chapter 9	Guidelines for Gateway FC Initiator Port Usage	107
	Identifying the ports on your gateway	108
	Guidelines for gateway port usage	112
Chapter 10	Connecting Your System to a Storage Array	115
	When to connect your gateway to third-party storage	116
	Planning a port-to-port connectivity scheme	118
	Connection locations on a gateway.	119
	Connecting a stand-alone gateway	125
	Connecting a gateway HA pair	128
	Connecting a MetroCluster	133
	Settings for Connecting to an ASCII Terminal Console	134
Chapter 11	Limiting the Target Queue Depth	137
Appendix A	Installation Quick Start: Third-Party Storage Only	139
	Example configuration for this quick start	140
	Installation quick start	141
Appendix B	Netboot Quick Start: Third-Party Storage Only	149
Appendix C	Obtaining WWNs Manually	153
Appendix D	Configuration Requirements by Vendor.	155
	Gateway requirements for EMC CLARiiON storage arrays	156
	Gateway requirements for EMC Symmetrix storage arrays	158
	Gateway requirements for Fujitsu ETERNUS storage arrays	162
	Gateway requirements for Hitachi storage arrays	166
	Gateway requirements for HP storage arrays.	168
	Gateway requirements for IBM storage arrays.	169

	Gateway requirements for 3PAR storage arrays	172
Appendix E	Storage Array Model Equivalents	173
Appendix F	Terminology Comparison Between Vendors	175
	Glossary	179
	Index	191

Preface

About this guide

Note

This guide applies to systems, including systems with gateway functionality, running Data ONTAP 8.x 7-Mode. In the Data ONTAP 8.x 7-Mode product name, the term *7-Mode* signifies that the 8.x release has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

Note

In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website, which is accessed and navigated as described in “[Websites](#)” on page v.

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- ◆ A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- ◆ The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To

understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

- ◆ IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

- ◆ For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call: Before you call, make sure you have taken these steps to try to solve the problem yourself:

- ◆ Check all cables to make sure they are connected.
- ◆ Check the power switches to make sure the system is turned on.
- ◆ Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- ◆ Refer to the IBM N series support website for information on known problems and limitations.

Using the documentation: The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the IBM N series support website, which is accessed and navigated as described in “Websites” on page v.

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the IBM N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support: You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates: IBM N series product firmware is embedded in Data ONTAP. As with all devices, it is recommended that you run the latest level of firmware. Any firmware updates are posted to the IBM N series support website, which is accessed and navigated as described in “Websites” on page v.

Note

If you do not see new firmware updates on the IBM N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by e-mail to starpubs@us.ibm.com.

Be sure to include the following:

- ◆ Exact publication title
- ◆ Publication form number (for example, GC26-1234-02)
- ◆ Page, table, or illustration numbers
- ◆ A detailed description of any information that should be changed

About this chapter

This chapter describes how a gateway virtualizes different types of storage. It also discusses how a gateway can connect to third-party storage and summarizes the implementation phases.

Note

IBM gateway models and filers with gateway functionality are both referred to as gateways.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“How a gateway uses storage”](#) on page 3
- ◆ [“RAID types supported by gateways”](#) on page 5
- ◆ [“Supported methods for connecting to a storage array”](#) on page 6
- ◆ [“Stages of gateway implementation”](#) on page 9

How to use this guide

This guide contains general guidelines for setting up the storage array to work with the gateways. When planning your deployment, first read this guide, then read the gateway *Implementation Guide* for your storage array type. The gateway implementation guides provide vendor-specific information about storage arrays.



When you see this icon in a section of this guide, you should go to the gateway *Implementation Guide* for your vendor to get additional details that are specific to your vendor.

If your gateway uses both array LUNs and native disk shelves for storage, see also the gateway *Implementation Guide for Native Disk Shelves* for information about setting up storage on the native disk shelves connected to the gateway.

Note

The gateway supports native disk shelves starting in Data ONTAP 7.3.

Additional gateway requirements and support information

Not all Data ONTAP releases support the same features, configurations, storage array models, and gateways. See the gateway *Interoperability Matrix* at www.ibm.com/storage/support/nseries/ for complete information about what the

gateway supports for particular releases. The gateway *Interoperability Matrix* also includes information about supported storage array firmware and switches that can be used with gateways.

Attention

The gateway *Interoperability Matrix* is the final authority on supported configurations, storage arrays, storage array firmware and microcode versions, switches, and so on.

Terminology

array LUN: This guide uses the term *array LUN* (logical unit number) to describe an area on the storage array that is available for a gateway or a non gateway host to read data from or write data to. You might be accustomed to hearing a different term to describe this area; the term varies among vendors and sometimes among platforms for the same vendor. See the gateway *Implementation Guide* for your storage array type for the specific term used for your platforms.

disks, disk shelves: The terms *disk* and *disk shelf* in this document refer to native storage connected to the gateway. These terms do not refer to disks or disk shelves on a third-party storage array.

Note

Gateways support native disk shelves starting in Data ONTAP 7.3.

HA pair: Two storage systems (nodes) whose controllers are connected to each other either directly or through switches. In some versions of Data ONTAP, this configuration is referred to as an *active/active configuration*.

LUN or front-end LUN: The gateway can virtualize the storage provided by third-party storage arrays and serve it up as *LUNs* to applications and customers outside the gateway (for example, through iSCSI). Clients are unaware of where such a LUN is stored.

Data ONTAP software product manuals refer to this type of storage as a LUN. You can think about this type of storage as a *front-end LUN* or *Data ONTAP-served LUN*. It is not the same as an array LUN.

native disks, native disk shelves: These terms refer to disks and disk shelves that are sold as local storage for systems that run Data ONTAP software.

How a gateway uses storage


Storage types supported

A gateway is an open storage controller that virtualizes storage from third-party storage array vendors, native disks, or both into a single heterogeneous storage pool. The Data ONTAP software provides a unified storage software platform that simplifies managing both native disk shelves and LUNs on storage arrays. You can add storage when and where you need it, without disruption.

Note

Starting in Data ONTAP 7.3, the gateway supports native disk shelves in addition to third-party storage. See the *Gateway Implementation Guide for Native Disk Shelves* for details.

The following illustration shows a gateway that pools storage from third-party storage arrays and serves data to Windows and UNIX hosts and clients.



How storage is presented to clients

A gateway presents storage to clients either in the form of Data ONTAP file system volumes, which you manage on the gateway using Data ONTAP management features, or as a SCSI target that creates LUNs for use by clients. In both cases (file system clients and LUN clients), on the gateway you combine the array LUNs into one or more Data ONTAP volumes for presentation to the clients as files or as LUNs served by Data ONTAP. Likewise, you can combine native disks into one or more Data ONTAP volumes for presentation to clients.

RAID types supported by gateways

RAID protection for array LUNs

For array LUNs, third-party storage arrays provide the RAID protection for the LUNs.

Data ONTAP supports a variety of RAID types used by storage arrays, but restricts storage arrays from using RAID0 for the LUNs that they make available to Data ONTAP. The reason is that Data ONTAP uses RAID0 to stripe across the array LUNs. Data ONTAP use of RAID0 maximizes performance because more disk spindles are used.

RAID0 provides no data protection. Therefore, when creating “RAID groups” on storage arrays, follow the best practices of the storage array vendor to ensure that there is an adequate level of protection on the storage array so that disk failure on the storage array does not result in data loss or loss of access to data.

Note

A “RAID group” on a storage array is the arrangement of disks that together form the defined RAID level. Each RAID group supports only one RAID type. The number of disks that you select for a RAID group determines the RAID types that a particular RAID group supports. Different storage array vendors use different terms to describe this entity—RAID groups, parity groups, disk groups, Parity RAID groups, and other terms.



See the gateway *Implementation Guide* for your vendor to determine whether there are specific requirements or limitations about RAID types in configurations with gateways.

RAID protection for native disks

Starting in Data ONTAP 7.3, gateways support native disk shelves as well as third-party storage. Data ONTAP supports RAID4 and RAID-DP on the native disk shelves connected to a gateway (just as it does for disks that provide storage for IBM N series storage systems). Data ONTAP does not support RAID4 and RAID-DP with array LUNs.

Supported methods for connecting to a storage array

The following methods are supported for connecting gateways to storage arrays:

- ◆ Direct-attached
- ◆ Fabric-attached
- ◆ MetroCluster



Not all methods are supported for every vendor or for every storage array that the gateway supports. Check the gateway *Implementation Guide* for your vendor to see whether there are restrictions as to how you can connect your storage array to a gateway.

Advantages and disadvantages of direct-attached configurations

Direct-attached configurations are supported with most storage arrays for both stand-alone systems and HA pairs.

The advantages of a direct-attached configuration are as follows:

- ◆ If the storage array has enough ports to connect to the gateway, a direct-attached configuration is more cost-effective because it is not necessary to purchase switches.
- ◆ You do not have to configure and manage a Fibre Channel SAN.

The disadvantages of a direct-attached configuration are as follows:

- ◆ More storage array ports are used among all hosts.
- ◆ You cannot share ports.



Check the gateway *Implementation Guide* for your vendor to determine whether a direct-attached configuration is supported for your vendor's storage arrays.

Advantages and disadvantages of fabric-attached configurations

Fabric-attached configurations are supported for both stand-alone systems and HA pairs.

The advantages of a fabric-attached configuration are as follows:

- ◆ Switches enable you to minimize the use of ports on the storage arrays. In this type of configuration, the gateway connects to the switches and the switches connect to the storage array. Therefore, an independent connection is not required for each gateway. Ports can be very expensive on high-end storage arrays in particular.
- ◆ You do not need end-to-end cabling for the entire path, which makes it easier to change the hardware configuration when you need to.



Check the gateway *Implementation Guide* for your vendor to confirm that a fabric-attached configuration is supported for your vendor's storage arrays.

MetroCluster support with gateways

MetroClusters are supported with gateways, but not with all storage arrays. In a MetroCluster, the aggregates are mirrored. Failover is possible after loss of entire node, including storage; this is not the case for standard or mirrored HA pairs.

The two types of MetroClusters are as follows:

- ❖ A *fabric-attached MetroCluster* is an HA pair running the syncmirror_local and cluster_remote licenses, where the nodes are attached to two pairs of Fibre Channel switches, and they are separated by longer distances than is possible with a stretch MetroCluster.
- ❖ A *stretch MetroCluster* is an HA pair running the syncmirror_local and cluster_remote licenses, where the nodes are separated by more limited distances than a fabric-attached MetroCluster and no switches are used between the nodes. This configuration is sometimes called a nonswitched MetroCluster.

Note

Gateways on which native disk shelves are installed cannot be deployed in a MetroCluster configuration.



Check the gateway *Implementation Guide* for your vendor to determine whether MetroCluster is supported for your vendor's storage arrays.

Number of storage arrays supported behind a gateway

A stand-alone gateway or the nodes in an HA pair can connect to multiple storage arrays. There is no limit to the number of storage arrays you can deploy behind a gateway. The storage arrays can be from the same vendor—all from the same family or from different families. Or the storage arrays can be from different vendors.

Note

Storage arrays in the same *family* share the same performance and failover characteristics. For example, members of the same family all perform active-active failover, or they all perform active-passive failover. Storage arrays with 4 GB HBAs are not considered to be in the same family as storage arrays with 2 GB HBAs.

Different rules for assigning array LUNs to aggregates apply depending on whether the storage arrays are from the same vendor and same family, the same vendor but different family, or from different vendors. See Chapter 8, “[Guidelines for Assigning Storage to Aggregates](#),” on page 99 for more information.



Typically a storage array provides storage for both gateways and other hosts. Check the gateway *Implementation Guide* for your vendor to see whether your vendor’s storage array must be dedicated to the gateways.

Attention

If your configuration is large, contact a sales engineer for help with planning for deploying storage arrays behind the same gateway. Factors such as which storage arrays you are mixing behind a gateway and the load that is being handled (for example, for archive or database processing) might make a difference in your deployment.

Stages of gateway implementation

You can think of a gateway implementation as a front-end implementation and a back-end implementation.

Summary of back-end implementation

Typically gateways use third-party storage, although use of third-party storage is not required. Setting up the back-end implementation includes all tasks that are required to set up the gateway with a storage array, up to a point where you can install Data ONTAP software. Back-end implementation tasks include the following:

- ◆ Creating and formatting array LUNs
- ◆ Assigning ports
- ◆ Cabling
- ◆ Zoning switches (if applicable)
- ◆ In Data ONTAP, assigning specific array LUNs to a gateway
- ◆ In Data ONTAP, providing information to set up a gateway on the network.
- ◆ Installing Data ONTAP software

Note

Starting in Data ONTAP 7.3, if a gateway is ordered with disk shelves, the Data ONTAP software is installed by the factory. In such a configuration, you do not need to create the root volume and install licenses and Data ONTAP software.

Summary of front-end implementation

Setting up the front-end implementation includes the following:

- ◆ Configuring the gateway for all protocols (NAS or FCP)
- ◆ Implementing the SNAP* suite of products (Snapshot™, SnapVault®, and so on)
- ◆ Creating volumes and aggregates
- ◆ Setting up data protection, including NDMP dumps to tapes
- ◆ Setting up native disks (if your system uses native disks for storage)

About this chapter

This chapter provides an overview of the planning you need to do for a configuration with a gateway. Your gateway might use third-party storage only, native disks only, or both native disks and third-party storage.

Summary of major planning tasks

Step	Task	Where to find information
1	Determine how much storage space is needed by the hosts and clients that you expect to connect to the gateway.	—
If you are using native disk shelves with your gateways:		
2	Determine the following: ◆ Gateway FC initiator port usage ◆ If you are using both disks and array LUNs, what should go on disks and what should go on array LUNs ◆ If you have an HA pair, whether to use the Multipath Storage feature	◆ Chapter 9, “Guidelines for Gateway FC Initiator Port Usage,” on page 107 ◆ <i>Gateway Implementation Guide for Native Disk Shelves</i>
If you are using third-party storage with your gateways:		
3	Determine the requirements for setting up your storage array to work with the gateway, including configuration settings on the storage array that are required for the gateway to work with the third-party storage array.	◆ Appendix D, “Configuration Requirements by Vendor,” on page 155 ◆ <i>Gateway Implementation Guides</i>
4	Determine which supported configuration for your storage array that you want to use.	◆ “Supported methods for connecting to a storage array” on page 6 ◆ <i>Gateway Implementation Guide for your vendor</i>

Step	Task	Where to find information
5	Ensure that your environment meets gateway requirements. For example, confirm your storage array, storage array firmware, and switch firmware are supported.	◆ Gateway <i>Interoperability Matrix</i>
6	Determine how to provide access to array LUNs from gateways.	◆ Chapter 3, “ Guidelines for Setting Up Access to Array LUNs ,” on page 17.
7	Determine how you are going to use Data ONTAP aggregates and volumes to partition the space on the storage arrays that you allocated for gateways. How you use aggregates and volumes directly impacts the number and size of array LUNs you need for Data ONTAP.	◆ Chapter 8, “ Guidelines for Assigning Storage to Aggregates ,” on page 99 ◆ Data ONTAP <i>Storage Management Guide</i> (for details about aggregates and volumes)
8	Determine the number and size of LUNs on the storage array that you need for Data ONTAP.	◆ Chapter 3, “ Guidelines for Setting Up Access to Array LUNs ,” on page 17 ◆ Chapter 5, “ Guidelines for Provisioning Array LUNs ,” on page 61. ◆ Gateway <i>Interoperability Matrix</i> (for information about the maximum number of array LUNs supported per platform) ◆ Your storage array documentation
9	Plan for LUN security This task includes setting access controls on the storage array and, if switches are deployed, setting zoning on switches	◆ Chapter 3, “ Guidelines for Setting Up Access to Array LUNs ,” on page 17

Step	Task	Where to find information
10	Determine your port-to-port connectivity scheme between the gateways and the storage array, which involves considering the following: ❖ Supported configurations for your vendor ❖ Gateway FC initiator port usage ❖ Cabling of redundant paths between the gateway and storage array, either directly or through switches. ❖ Zoning switches (if applicable) Map (export) array LUNs to the ports to which the gateways are connected so that the LUNs are available to those gateways.	◆ Gateway <i>Implementation Guide</i> for your vendor (for supported configurations for your vendor) ◆ Chapter 9, “Guidelines for Gateway FC Initiator Port Usage,” on page 107 ◆ Chapter 10, “Connecting Your System to a Storage Array,” on page 115 ◆ Chapter 4, “Guidelines for a Fabric-Attached Configuration,” on page 49
11	Determine which gateway you will assign in Data ONTAP as the <i>owner</i> of the array LUN. If your system is also using native disks, Data ONTAP can automatically assign ownership of the disks	Chapter 6, “Requirements for a Gateway to Use Storage,” on page 79
If you are deploying a MetroCluster:		

Step	Task	Where to find information
12	Plan for a MetroCluster. This includes ensuring that your configuration conforms to the following: ◆ Hardware and software requirements to set up a MetroCluster ◆ Requirements for setting up SyncMirror with array LUNs To set up SyncMirror with a gateway, you must assign array LUNs to SyncMirror pools correctly. Plan for SyncMirror before you assign LUNs to your gateways. Otherwise, you must unassign array LUNs you want to be in the remote location and reassign them to the correct SyncMirror pool. Note _____ MetroCluster configurations are not supported for gateways using native disk shelves.	◆ Gateway <i>MetroCluster Guide</i> ◆ Data ONTAP <i>Data Protection and Online Backup Guide</i>
If you are using Data ONTAP data protection features:		
13	Plan for any of the data protection features you want to use that are available for your Data ONTAP release, including the following: ◆ The SNAP* suite of products (Snapshot copies, SnapVault, and so on) ◆ NDMP dumps to tape ◆ SyncMirror	◆ Data ONTAP <i>Data Protection and Online Backup Guide</i> ◆ Data ONTAP <i>Data Protection and Tape Backup Guide</i>
If you are using other Data ONTAP features:		
14	Determine which of the other features to simplify storage management that you want to use.	◆ Data ONTAP <i>Storage Management Guide</i>

Gateway capacity limits

See the gateway *Interoperability Matrix* for information about the maximum and minimum system capacity limits for specific gateways.

The capacity values shown in the gateway *Interoperability Matrix* are for each node as a whole. The combination of the size of the volumes and aggregates on the system cannot be less than the minimum system capacity or exceed the maximum system capacity.

Note

The gateway *Interoperability Matrix* is the final authority on system capacity information for each gateway platform.

About this chapter

This chapter provides information to help you plan how to set up the gateways so they can access storage on the storage arrays. If your configuration is fabric-attached, also see Chapter 4, “[Guidelines for a Fabric-Attached Configuration](#),” on page 49 for information about setting up your switches.



For any additional details specific to setting up access to array LUNs for the gateway on your storage array model, check the gateway *Implementation Guide* for your vendor.

Note

If your systems use both third-party storage and native disk shelves, see the gateway *Implementation Guide for Native Disk Shelves* for information about planning for access to native disk shelves. If your gateway does not use third-party storage, skip this chapter.

Topics in this chapter

This chapter contains the following topics:

- ◆ “[LUN security](#)” on page 18
- ◆ “[Requirements for paths to array LUNs](#)” on page 20
- ◆ “[Common errors when setting up access to array LUNs](#)” on page 23
- ◆ “[Valid path setup examples](#)” on page 24
- ◆ “[Invalid path setup examples](#)” on page 29
- ◆ “[How automatic port selection and monitoring works](#)” on page 33
- ◆ “[What happens when a link failure occurs](#)” on page 36
- ◆ “[How paths are reflected in array LUN names](#)” on page 39
- ◆ “[When to check the path setup to array LUNs](#)” on page 42
- ◆ “[Using the storage show disk -p command to check paths](#)” on page 43
- ◆ “[Using the sysconfig -v command to check paths](#)” on page 45
- ◆ “[Using the storage array show-config command to check paths](#)” on page 46

LUN security

What LUN security is

LUN security is used to isolate which hosts can access which array LUNs. (To the storage array, the gateway is a host.) LUN security is similar to switch zoning in concept, but it is performed on the storage array. *LUN security* and *LUN masking* are equivalent terms to describe this functionality.

Note

The Data ONTAP disk ownership scheme prevents one gateway from overwriting an array LUN owned by another gateway. However, it does not prevent a gateway from overwriting an array LUN accessible by a non gateway host. Likewise, without a method of preventing overwriting, a non gateway host could overwrite an array LUN used by a gateway.

Available LUN security methods

To eliminate the possibility of a non gateway overwriting array LUNs owned by a gateway, or the reverse, you must use one of the following LUN security methods.

Method 1—Use LUN security: With this method, you use a LUN security product to control which of the hosts that are zoned to the same port can see specific array LUNs over that port. You prevent other hosts from accessing those same array LUNs by hiding (or masking) those LUNs from the other hosts.

Method 2—Use port-level LUN security: With port-level LUN security, you present only the array LUNs for a particular host on a port, which enables you to present only the gateway LUNs on a particular port. Therefore, the port is dedicated to a host.

Note

Not all storage arrays support port-level security. Some storage arrays present all LUNs on all ports by default, and they do not provide a way to restrict the visibility of LUNs to particular hosts. For storage arrays that present all LUNs on all ports and that do not enable you to restrict visibility of LUNs on a port, you must use either Method 1 or Method 3.

Method 3—Dedicate the storage array to the gateway: With this method, no hosts other than the gateway are connected to the storage array.

Attention

For added protection and redundancy for the gateway, it is recommended that you do both zoning and LUN security. For example, if you do not have LUN security configured and you have to replace a SAN switch, the gateway could panic before you are able to configure the zoning on the new switch because the switch is wide open.



Check the gateway *Implementation Guide* for your vendor for any additional details regarding LUN security for your vendor's storage arrays.

Requirements for paths to array LUNs

Redundant setup of all components in the path

Gateways must connect to the storage array through a redundant Fibre Channel network. Two Fibre Channel networks or fabric zones are required so that fabric ports or switches can be taken offline for upgrades and replacements without impacting the gateways.

Setting up a redundant path to array LUNs involves making sure you have set up redundancy of all devices in the path. Be sure to set up redundancy of the devices in the path as follows:

- ◆ Each gateway FC initiator port on the same gateway is on a different bus
- ◆ You use redundant switches
- ◆ Each array port is on a different (redundant) controller on the storage array

Required number of paths to an array LUN

Data ONTAP expects and requires that a storage array provide access to a specific array LUN on only two of the storage array ports; that is, through only two redundant paths. One port is the primary path to the LUN and the other port is the alternate (secondary) path. A given array LUN is accessed through only one port at a time.


Attention

You cannot set up more than two paths to an array LUN. This is a Data ONTAP restriction.

Array port selection to ensure redundant paths to an array LUN

To provide redundant paths to a gateway array LUN, ensure that the two ports on the storage array that you select to access a given array LUN are from different components that could represent a SPOF, for example, from alternate controllers or clusters. (The terminology used for the hardware component on which host adapters and ports are located varies on different storage array models, even those from the same vendor.) You want to ensure that both ports cannot fail at the same time. A given array LUN is accessed through only one port at a time.

The following illustration shows correct and incorrect storage array port selection. The path setup in the example on the left is correct because the paths to the storage array are redundant—each connection is to a port on a different (redundant) controller on the storage array.



LUN numbers used when mapping LUNs to ports

Different LUN numbers for the same array LUN can cause problems for the gateway.

Pathing when multiple gateway FC initiator port pairs are used

Although there cannot be more than two paths to a given array LUN per gateway FC initiator port pair, this does not mean that the gateway must see all array LUNs for the gateway through one FC initiator port pair. Some gateways, for example, the N7600 gateways, have a large number of FC initiator ports available. By using multiple FC initiator port pairs, each accessing a different group of array LUNs, you can partition the load of array LUN traffic over the gateway connections to optimize performance.




Data ONTAP does not support accessing array LUNs in multiple LUN groups for all storage arrays. See the gateway *Implementation Guide* for your vendor and the gateway *Interoperability Matrix* to determine whether a configuration using multiple LUN groups is supported for your storage array.

Note

A LUN group is set of logical devices that a gateway accesses over the same paths. Each gateway must have exactly two paths to a LUN group. The storage array administrator configures a set of logical devices as a group in order to define which host WWPNs can access them.

The following illustration shows how one gateway FC initiator port pair (0c and 0f) is used to access one set of LUNs (a LUN group) over one storage array port pair and a second FC initiator port pair (0a and 0h) is used to access a second set of LUNs (a second LUN group) on the same storage array over a different storage array port pair.



Common errors when setting up access to array LUNs

The following errors are commonly made when setting up access to array LUNs that are allocated to gateways:

- ◆ Too many connections are configured to a gateway array LUN.
- ◆ Redundant paths are not configured to a gateway array LUN.
- ◆ A switch is misconfigured, which prevents all the gateways in the neighborhood from seeing the gateway array LUNs.


Make sure that the switches are configured correctly so that each gateway in the neighborhood can see all the same array LUNs as the other neighborhood members. See Chapter 7, “[Using Neighborhoods for Storage Management](#),” on page 83 for information about neighborhoods.

See the examples in “[Invalid path setup examples](#)” on page 29.

Valid path setup examples

Pathing in a configuration with one 2-port array LUN group

The following illustration shows a configuration with one 2-port array LUN group.



Applicability of this configuration: This configuration, with one 2-port array LUN group, works with all storage arrays for all Data ONTAP releases.

Configuration description: In this configuration, each of the two target ports on the storage array is accessed by two gateway FC initiator ports, one from each node in the HA pair. (Two gateway FC initiator ports, one from each node, “share” the same target port.) To ensure availability, use a redundant FC initiator port pair on each node in the HA pair. Then, if one path from a node fails, the other path from the node is used; gateway controller takeover does not occur.

Note


Both this configuration and the one 4-port array LUN group in the following example are best practice configuration recommendations. However, failover is not as good in this configuration as when using the one 4-port array LUN group in “[Pathing in a configuration with one 4-port array LUN group](#)” on page 25. The reason is that if a switch on one fabric fails, all traffic from both gateways goes through a single port on the storage array.

Note

Different storage arrays, even those from the same vendor, might label the ports differently from those shown in the example. On your storage array, ensure that the ports you select are alternate controllers.

Pathing in a configuration with one 4-port array LUN group

The following illustration shows a configuration with one 4-port array LUN group.



Applicability of this configuration: This configuration, with one 4-port array LUN group, works with all storage arrays for all Data ONTAP releases.


Configuration description: In this configuration, each target port is accessed by a single gateway FC initiator port from the HA pair. You must use zoning to ensure that there are only two paths to a specific array LUN.

Note

Both this configuration and the configuration in “[Pathing in a configuration with one 2-port array LUN group](#)” on page 24 are best practice configuration recommendations. Although the one 2-port array LUN group configuration uses fewer ports, this configuration has better failover. The reason is that if a switch fails, the array LUNs are still served from two array ports, one array port going to vs1 and the other array port going to vs2.

Pathing in an HA pair with two 2-port array LUN groups with fan-in

The following configuration shows an HA pair with two 2-port array LUN groups with fan-in.



Applicability of this configuration: This configuration of multiple LUN groups on the same storage array is not supported for all storage arrays and with all Data ONTAP releases. Prior to Data ONTAP 7.3, this configuration was supported only for storage arrays with unique WWNNs per port (for example, Hitachi and Hitachi-based storage arrays). Starting in Data ONTAP 7.3, this configuration is also supported with *some* storage arrays that have several WWPNs per WWNN. For example, it is supported with IBM and EMC CLARiiON storage arrays also.



See the gateway *Implementation Guide* for your vendor to determine whether this configuration is supported for your vendor and for the release of Data ONTAP you are running.

Configuration description: In this configuration, each gateway FC initiator port accesses multiple LUN groups. Two or more LUN groups can be fanned into the same gateway port. However, each gateway controller can see a LUN group over only two paths.

There is a redundant path to each LUN group, one path through Controller 1 and the other path through Controller 2. The port pairs configured on the storage array are redundant in the following ways:

- ◆ Port A on Controller 1 and Port A on Controller 2 are a port pair. Because they are alternate paths, both ports can access the array LUNs in LUN group 1.
- ◆ Port B on Controller 1 and Port B on Controller 2 are a port pair. Because they are alternate paths, both ports can access the array LUNs in LUN group 2.

Zoning setup: The following table shows how zoning is set up for this scenario. You need to zone the switches so that there are only two paths to each array LUN, one unique path from each gateway FC initiator port through each switch. If there are multiple connections between a gateway and the switch, put each connection into a separate zone.

Zone	Gateway port	Storage array controller and port
Switch 1		
z1	Gateway 1, port 0a	Controller 1 port A
z2	Gateway 1, port 0a	Controller 1 port C
z3	Gateway 2, port 0a	Controller 1 port A

Zone	Gateway port	Storage array controller and port
z4	Gateway 2, port 0a	Controller 1 port C
Switch 2		
z5	Gateway 1, port 0c	Controller 2 port A
z6	Gateway 1, port 0c	Controller 2 port C
z7	Gateway 2, port 0c	Controller 2 port A
z8	Gateway 2, port 0c	Controller 2 port C


Array LUN security setup: The following table shows how to set up array LUN security for this scenario.

These storage array ports...	Should have array LUN security set up for these gateway FC initiator ports...
Controller 1 port A Controller 1 port B	gateway 1: FC initiator port 0a gateway 2: FC initiator port 0a
Controller 2 port A Controller 2 port B	gateway 1: FC initiator port 0c gateway 2: FC initiator port 0c

Invalid path setup examples

Too many paths to an array LUN

The following configuration is invalid because there are too many paths to the array LUN. Data ONTAP requires and supports only two paths to an array LUN. In the following example, the path setup is invalid because the same array LUN would be accessed over four paths instead of only two paths, as required by Data ONTAP.



The following table shows how the zoning is set up for this *invalid* scenario.


Gateway system FC initiator port	Switch	Over connection number	Storage array port
0a	1	5	Controller 1 Port A
0a	1	4	Controller 2 Port A
0c	2	5	Controller 1 Port B
0c	2	4	Controller 2 Port B

As the table shows, connection 3 from 0a through Switch 1 is incorrectly zoned to both connection 4 and connection 5. Likewise, connection 3 from 0c through Switch 2 is incorrectly zoned to both connection 4 and connection 5. The result is that array LUN 1 is seen over more than two paths, which Data ONTAP does not allow.

Requirement for this configuration to be correct: For this configuration to be correct, FC initiator port 0a must see *either* Controller 1 port A or Controller 2 Port A on the storage array, but not both. Likewise, FC initiator port 0c must see *either* port Controller 1 Port B or Controller 2 Port B on the storage array, but not both.

Alternate paths are not configured

The following configuration is invalid because it does not provide alternate paths from each FC initiator port on the gateways to each LUN on the storage array. Both FC initiator ports from the same gateway are connected to the storage array through the same switch.



Assume that the following zoning is in place in this invalid example:

- ◆ For gateway 1:
 - ❖ 0a is zoned to see Controller 1 Port A
 - ❖ 0c is zoned to see Controller 1 Port C
- ◆ For gateway 2:
 - ❖ 0a is zoned to see Controller 2 Port A
 - ❖ 0c is zoned to see Controller 2 Port C

The problems that result from not having alternate paths are as follows:

- ◆ Each switch becomes an SPOF.

- ◆ Gateway 1's FC initiator port 0a can access LUNs in LUN group 1 on the storage array's Controller 1 port A, but it cannot access the LUNs in LUN group 2.
- ◆ Gateway 2's FC initiator port 0a can access LUNs in LUN group 1 through the storage array's Controller 1 port A, but it cannot access LUNs in LUN group 2.

To make this a valid configuration, the following changes must be made:

- ◆ Gateway 1's FC initiator port 0c must be connected to Switch 2.
- ◆ Gateway 2's FC initiator port 0a must be connected to Switch 1.
- ◆ Appropriate zoning must be configured. If you are using multiple ports on a storage array that supports configuring a specific set of LUNs on a selected set of ports, ensure that a given FC initiator port sees all array LUNs presented on the fabric.

How automatic port selection and monitoring works

How Data ONTAP automatic monitoring of port usage works

Data ONTAP periodically checks the percentage of I/O on each of the gateway FC initiator ports to ensure that it is distributing the I/O load efficiently over the system's initiator ports. If Data ONTAP determines that the loading over the ports is inefficient, it redistributes the servicing of array LUNs over the FC initiator ports. This redistribution can change the path that Data ONTAP uses to access some array LUNs.

Note

Data ONTAP redistribution of loading over the ports is according to how busy an initiator port is. It is not load balancing in the sense of balancing the *number* of array LUNs over initiator ports. One initiator could have a lot more array LUNs than another initiator, but that initiator with more LUNs could have less I/O than an initiator with fewer LUNs.

How Data ONTAP works with active-passive and active-active storage arrays

With *active-passive* storage arrays, the storage array assigns a primary path for a given array LUN. Data ONTAP always uses the primary path supplied by the active-passive storage array unless the primary path fails. If a failure occurs, the gateway fails over to the secondary path but resumes using the primary path again when the primary path is restored.

With *active-active* storage arrays, the storage array does not assign a preferred or default controller to a LUN; the LUN can be accessed through any controller without a performance impact. Data ONTAP automatically chooses one of the storage array ports that you specified for a given LUN as the primary path and the other port as the secondary (alternate) path. Data ONTAP uses its own algorithm to balance traffic to the array LUNs over its FC initiator ports. You cannot manually change the primary and secondary paths to an array LUN.

Note

I/O to a given array LUN is always through one path at any given time, whether the storage array supplies the preferred path to the LUN or Data ONTAP determines the preferred path to the LUN.

How Data ONTAP selects the path to an array LUN at bootup

When the gateway boots up, it queries the storage array to determine whether the storage array has assigned a primary path to the LUN. If the storage array sends back the primary path to the LUN, as is the case with active-passive storage arrays, Data ONTAP selects that path to the LUN. Otherwise, Data ONTAP selects the path on which the LUN initially becomes available as the primary path.

How Data ONTAP selects the path to an array LUN after bootup

After bootup, the path that Data ONTAP chooses automatically changes in the following circumstances:

- ◆ When a failure occurs in the primary path to the array LUN
If a failure occurs in the primary path to the array LUN, the gateway automatically switches to the other path (which then becomes the primary path). For active-passive storage arrays, the gateway begins to service the array LUN again over its primary path when the primary path becomes available.

- ◆ When some FC initiator ports are heavily loaded and others are lightly loaded

Data ONTAP monitors its FC initiator ports for the load. After the first load transaction threshold is reached, Data ONTAP reevaluates the usage of each FC initiator port. If it determines that an FC initiator port is heavily loaded and another FC initiator port is lightly loaded, it moves servicing of some of the array LUNs that are on the heavily loaded FC initiator port to the lightly loaded FC initiator port.

For example, assume that 100 array LUNs on the storage array are allocated to the gateways. In the gateway configuration, the gateway nodes in the HA pair have been configured to “own” those array LUNs as follows:

- ❖ LUNs 0 through 49 are assigned to Gateway 1.
- ❖ LUNs 50 through 99 are assigned to Gateway 2.

Gateway 1 sees LUNs 0 through 49 (which it owns) and LUNs 50 through 99 (which its partner owns). (Each node in the HA pair must be able to see the same LUNs.)

Assume that Gateway 1 FC initiator port 0a is very busy servicing LUNs 0 through 49. LUNs 50 through 99 are on Gateway 1 FC initiator port 0c, but there is no activity to them. The gateway determines that 0a is busier than 0c and then changes the path to some of the LUNs from 0a to 0c in an attempt to balance the load over the Gateway 1 FC initiator ports.

Viewing the path used at a given time

Because Data ONTAP automatically handles load balancing over the paths to the array LUN, you do not need to monitor the path to a LUN that is being used. However, if you are interested in seeing which path to a LUN is being used at a *given time*, you can use the `storage show disk -p` command. Data ONTAP changes the path as necessary for load balancing so the primary and secondary paths for a given array LUN can change at different times when you issue the command. See [“Using the storage show disk -p command to check paths”](#) on page 43.

What happens when a link failure occurs


**Data ONTAP
response when a
link failure occurs**

Data ONTAP monitors the link’s usage periodically. The following table shows what occurs in the case of a failure in a fabric-attached configuration.

If a failure occurs in the link between the...	Then...
Gateway and the switch	Data ONTAP receives notification immediately and sends traffic to the other path immediately.
Switch and the storage array	Data ONTAP is not immediately aware that there is a link failure because the link is still established between the gateway and the switch. Data ONTAP becomes aware that there is a failure when the I/O times out. Data ONTAP retries three times to send the traffic on the original path, then it fails over the traffic to the other path.

Failure scenario 1

The following illustration shows what occurs when a link in the primary path to a LUN in LUN group 1 fails.




In this scenario, the primary path to LUN group 1 from vs1 is through vs1's FC initiator port 0a, to Switch 1, and then to the storage array's Controller 1 port 1A. If a failure occurs in the link between vs1 FC initiator port 0a and Switch 1 when vs1 tries to access LUN group 1, Data ONTAP automatically switches to the alternate path through vs1's FC initiator port 0c. Gateway vs1 can then access LUN group 1 through Switch 2, and then through storage array Controller 2 port 2A.

Until the link failure is fixed, there is only one interface to the storage. When connectivity is restored, Data ONTAP redistributes the array LUNs over the paths.

Failure scenario 2

The following illustration shows what Data ONTAP does when the primary path to a LUN fails in a configuration for an HA pair with two 2-port array LUN groups with fan-in.



Failover in this example with two LUN groups works the same way as in the first failure scenario (assuming the primary path from vs1 to a LUN in LUN group 1 is through vs1's FC initiator port 0a). If a failure occurs in the link between vs1's FC initiator port 0a and Switch 1. Data ONTAP automatically switches to the alternate path through vs1's FC initiator port 0c, which enables vs1 to access the LUN in LUN group 1 through Switch 2 and then through the storage array's Controller 2 port A.

Until the link failure is fixed, there is only one interface to the storage. When connectivity is restored, Data ONTAP redistributes the array LUNs over the paths.

How paths are reflected in array LUN names

Format of array LUN names

The array LUN name is a path-based name that includes the devices in the path between the gateway and the storage array. By looking at the array LUN name as it is displayed in Data ONTAP output, you can identify the following:

- ◆ Devices in the path between the storage system and the storage array
- ◆ Ports used
- ◆ The LUN identifier that the storage array presents externally for mapping to hosts

The format of the array LUN name differs for direct-attached and fabric-attached configurations, as shown in the following table.

Configuration	Array LUN name format	Component descriptions
Direct-attached	For Data ONTAP 7.x releases and earlier, and Data ONTAP 8.0 7-Mode: <i>adapter.idlun-id</i> For Data ONTAP 8.0 Cluster-Mode: <i>node-name.adapter.idlun-id</i>	<i>adapter</i> is the adapter on the gateway <i>id</i> is the channel adapter port on the storage array. <i>lun-id</i> is the array LUN number that the storage array presents to hosts. <i>node-name</i> is the name of the Cluster-Mode node. In Cluster-Mode, the node name is prepended to the LUN name so that the path-based name will be unique within the cluster. Example : 0a.0L0 or, for Data ONTAP 8.0 Cluster-Mode, node1.0a.0L0

Configuration	Array LUN name format	Component descriptions
Fabric-attached	For Data ONTAP 7.x releases and earlier, and Data ONTAP 8.0 7-Mode: <i>switch-name:port.idlun-id</i> For Data ONTAP 8.0 Cluster-Mode: <i>node-name.switch-name:port.idlun-id</i>	<i>switch-name</i> is the name of the switch. <i>port</i> is the switch port that is connected to the target port (the end point). <i>id</i> is the device ID. <i>lun-id</i> is the array LUN number that the storage array presents to hosts. <i>node-name</i> is the name of the Cluster-Mode node. In Cluster-Mode, the node name is prepended to the LUN name so that the path-based name will be unique within the cluster. Example: mcdat3:6.127L0 or, for Data ONTAP 8.0 Cluster-Mode, node1.mcdat3:6.127L0

The array LUN names consist of a path component and the SCSI LUN ID on that path. For example, in the array LUN name example for the fabric-attached configuration, mcdat3:6.127 is the path component and L0 is the SCSI LUN ID. On a gateway, there are two names for each LUN because there are two paths to each LUN—for example, mcdat3:6.127L0 and brocade15:6.127L0.

How the array LUN name changes

Keep the following in mind when you are looking at Data ONTAP displays that show array LUNs:

- ◆ Array LUN names are relative to the gateway from which the array LUN is viewed. Therefore, the name of an array LUN might be different from each gateway in an HA pair or cluster because the path to the LUN is different.
- ◆ On each gateway there are multiple valid names for a single array LUN, one per path.

The name for a given LUN that is displayed in Data ONTAP can change depending on which path is active at a given time. For example, if the primary path becomes unavailable and Data ONTAP switches to the alternate path, the LUN name that is displayed changes.
- ◆ Each node in an HA pair typically accesses a given LUN through a different storage array port to limit contention on a single port.

Note

It is possible for different gateways to show the same name for different array LUNs. For example, this could occur if both switches have the same name.

When to check the path setup to array LUNs

Data ONTAP expects and requires that a storage array provide access to a specific array LUN on only two of the storage array ports; that is, through only two redundant paths. A given array LUN is accessed through only one port at a time.

You need to ensure that two paths to an array LUN exist between the gateways and the storage array when performing the following activities:

- ◆ During initial installation
- ◆ Before performing fabric maintenance, for example:
 - ❖ Before, during, and after an infrastructure upgrade
 - ❖ Before taking a switch out of service for maintenance

You want to ensure that the paths were configured as redundant paths before you remove a switch between the gateways and the storage array so that access to the array LUNs is not interrupted.
- ❖ Before maintaining hardware on a storage array, for example, maintaining the hardware component on which host adapters and ports are located (the name of this component varies on different storage array models.)

Commands for checking paths to array LUNs

You can check access to array LUNs from each gateway in the following ways:

- ◆ With the Data ONTAP `storage show disk -p` command
- ◆ With the Data ONTAP `sysconfig -v` command
- ◆ For systems running Data ONTAP 8.0 or later, with the `storage array show-config` command.

Using the storage show disk -p command to check paths

You can use the Data ONTAP `storage show disk -p` command to check that each array LUN is visible through both paths.

Use the `all` variable (`storage show disk -p all`) to see the adapters. (The adapter column shows the FC initiator port over which the LUN is accessed.) Looking at the adapters enables you to easily see when all paths are on a single adapter. For example, both paths through the gateway's `0c` port indicates that the back-end zoning is redundantly crossed, which is not supported.

Note

When you use the `all` variable, you cannot see unassigned array LUNs.

Data ONTAP changes the path as necessary for load balancing, so the primary and secondary paths for a given array LUN can change at different times when you issue the command.

Example output showing two paths

The following example shows two paths to each array LUN.

```
vse-v825mc1-1> storage show disk -p all
```

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	ADAPTER
vss_b3250_1:2.126L2	-	vss_b3250_2:2.126L2	-	-	-	0a
vss_b3250_1:2.126L5	-	vss_b3250_2:2.126L5	-	-	-	0a
vss_b3250_1:2.126L7	-	vss_b3250_2:2.126L7	-	-	-	0a
vss_b3250_1:2.126L9	-	vss_b3250_2:2.126L9	-	-	-	0a
vss_b3250_1:2.126L11	-	vss_b3250_2:2.126L11	-	-	-	0a
vss_b3250_1:2.126L13	-	vss_b3250_2:2.126L13	-	-	-	0a
vss_b3250_1:2.126L15	-	vss_b3250_2:2.126L15	-	-	-	0a
vss_b3250_1:2.126L17	-	vss_b3250_2:2.126L17	-	-	-	0a
vss_b3250_1:2.126L19	-	vss_b3250_2:2.126L19	-	-	-	0a
vss_b3250_1:5.126L254	-	vss_b3250_2:5.126L254	-	-	-	0a
vss_b3250_1:5.126L255	-	vss_b3250_2:5.126L255	-	-	-	0a
vss_b3250_2:2.126L1	-	vss_b3250_1:2.126L1	-	-	-	0c
vss_b3250_2:2.126L3	-	vss_b3250_1:2.126L3	-	-	-	0c
vss_b3250_2:2.126L6	-	vss_b3250_1:2.126L6	-	-	-	0c
vss_b3250_2:2.126L8	-	vss_b3250_1:2.126L8	-	-	-	0c
vss_b3250_2:2.126L10	-	vss_b3250_1:2.126L10	-	-	-	0c
vss_b3250_2:2.126L12	-	vss_b3250_1:2.126L12	-	-	-	0c
vss_b3250_2:2.126L14	-	vss_b3250_1:2.126L14	-	-	-	0c
vss_b3250_2:2.126L16	-	vss_b3250_1:2.126L16	-	-	-	0c
vss_b3250_2:2.126L18	-	vss_b3250_1:2.126L18	-	-	-	0c
vss_b3250_2:2.126L20	-	vss_b3250_1:2.126L20	-	-	-	0c

**What the output
looks like if there is
only one path**

If there is only one path to the LUN, the Secondary (path) column is empty.

Using the `sysconfig -v` command to check paths

You can use the Data ONTAP `sysconfig -v` command on each gateway to make sure that all the gateway LUNs on the storage array are accessible from each gateway. This command lists the name of each array LUN that the gateway can access (for example, `vnbr4100s31:5.126L4`). Compare the list of array LUNs from each gateway to ensure that each gateway can access all its array LUNs.

Using the storage array show-config command to check paths

For systems running Data ONTAP 8.0 or later, use the storage array show-config command to list the array LUNs that the storage system can access and to check that each array LUN is visible through both paths.

Example output showing two paths

The following example shows output from a gateway connected to two storage arrays. Each LUN group is comprised of LUNs that share the same two paths. Groups 0 and 1 contain a total of 8 LUNs on the HP_HSV210_1 array and Group 2 contains 50 LUNs on the HP_HSV200_1 array.

```
> storage array show-config
LUN Group      Array Name      Array Target Ports      Switch Port  Initiator
Group 0 (4 LUNS) HP_HSV210_1    50:00:1f:e1:50:0a:86:6d vnmc4300s35:11 0b
                50:00:1f:e1:50:0a:86:69 vnbr4100s31:15 0c
Group 1 (4 LUNS) HP_HSV210_1    50:00:1f:e1:50:0a:86:68 vnbr4100s31:1 0a
                50:00:1f:e1:50:0a:86:6c vnmc4300s35:6 0d
Group 2 (50 LUNS) HP_HSV200_1    50:00:1f:e1:50:0d:14:6d vnbr4100s31:5 0a
                50:00:1f:e1:50:0d:14:68 vnmc4300s35:3 0d
```

What the output looks like if there are not two paths

Array LUNs that are not configured with two paths are shown as one or more LUNs with a single path, similar to the following example. The incorrectly configured LUNs are the 20 LUNs not belonging to a group and showing only a single path.

LUN Group	Array Name	Array Target Ports	Switch Port	Initiator
Group 2 (50 LUNS)	HP_HSV200_1	50:00:1f:e1:50:0d:14:68	vnmc4300s35:3	0d
		50:00:1f:e1:50:0d:14:6d	vnbr4100s31:5	0a
	HP_HSV200_1	50:00:1f:e1:50:0d:14:69	vnmc4300s35:2	0e

How to get details about individual LUNs

If you see array LUNs in the output from the storage array show-config command with only one path, you need to use the storage array show-luns command to show information about each array LUN on the storage array. This information allows you to determine which array LUNs are members of groups and which are incorrectly configured. The output from the storage array show-luns HP_HSV200_1 produces output similar to the following (the output is abbreviated):

Name	WWPNs
vnmc4300s35:3.127L1	50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnmc4300s35:3.127L2	50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnmc4300s35:3.127L3	50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
.	
.	
.	
vnbr4100s31:5.126L49	50:00:1f:e1:50:0d:14:6d, 50:00:1f:e1:50:0d:14:68
vnmc4300s35:3.127L50	50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnmc4300s35:3.127L51	50:00:1f:e1:50:0d:14:69,
vnmc4300s35:3.127L52	50:00:1f:e1:50:0d:14:69,
.	
.	
.	
vnbr4100s31:5.126L53	50:00:1f:e1:50:0d:14:69,
vnbr4100s31:5.126L70	50:00:1f:e1:50:0d:14:69,

LUNs 1 - 50 make up Group 2, the array LUNs configured with two ports as shown with the `storage array show-config` command. LUNs 51 through 70 are the 20 LUNs that contain a single path connected only to Port 50:00:1f:e1:50:0d:14:69 of the storage array.

A single path indicates a configuration problem. The second path can have a problem at the gateway port, the switch, or the storage array port.

About this chapter

This chapter provides a best practice recommendation for how to set up zoning in fabric-attached configurations with gateways and provides information about use of specific switches with gateways.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“General requirements for fabric setup”](#) on page 50
- ◆ [“Switch zoning”](#) on page 51
- ◆ [“Examples of zoning a gateway with storage arrays”](#) on page 52
- ◆ [“Guidelines for Brocade switches”](#) on page 56
- ◆ [“Guidelines for Cisco fabrics”](#) on page 57
- ◆ [“Guidelines for McDATA switches and directors”](#) on page 58
- ◆ [“Guidelines for switches with a MetroCluster”](#) on page 59

General requirements for fabric setup

Where to find information about supported switches

See the gateway *Interoperability Matrix* to determine which switches that gateways support.

Requirement for redundant switches

Gateways connect to the storage array through a redundant Fibre Channel network. Two Fibre Channel networks are required so that fabric ports or switches can be taken offline for upgrades and replacements without impacting the gateways.

Number of switch ports needed

When determining the number of ports that you need, you must consider how performance is impacted by the number of ports you use.

Switch zoning

About switch zoning

Configuring zoning on a Fibre Channel switch enables you to define paths between connected nodes, restricting visibility and connectivity between devices connected to a common Fibre Channel SAN.

A zone grants or restricts access only to a given *port* on a storage array. More than one host can be zoned to the same storage array port. With zoning, you overlay a security map on the network that dictates which hosts can see and have access to specific targets, thereby reducing the risk of data loss. If, for example, one host were to gain access to an array LUN being used for storage by another host, the data in the array LUN could become corrupted. If a host cannot see an array LUN, it cannot access the LUN to corrupt it.

A single shared device might need to exist in multiple zones at the same time. This typically occurs when multiple hosts require access to the same target device. In this case, you can usually configure each host to reside in a separate zone with the target device; the target device resides in the zone for each host. This scheme enables each of the hosts that needs to access the common target device to access it, while preventing the hosts themselves from communicating with each other.

You can configure zoning either by specifying worldwide names (WWNs) or ports.

Zoning recommendation for the gateway

It is recommended that you use single-initiator zoning, which limits each zone to a single gateway FC initiator port. The benefits of creating a separate zone for each gateway's FC initiator port and each non-gateway host are as follows:


- ◆ You limit the number of ports over which a specific array LUN can be accessed. You can prevent gateways from accessing a given array LUN for gateways over more than two storage array ports.
- ◆ Single-initiator zoning improves discovery and boot time because the gateway FC initiators do not attempt to discover each other.



Check the gateway *Implementation Guide* for your vendor for any additional details regarding zoning with your vendor's storage arrays.

Examples of zoning a gateway with storage arrays

Zoning example 1 The following illustration shows zoning before the addition of an HA pair.



Zoning on the switches is set up as follows.

Zone	Includes...
z3 and z4	The Linux® server accessing the Vendor 3 storage array and the Vendor 2, Series B storage array (for high availability)
z5	The IBM® server accessing the Vendor 1 storage array
z8 and z9	The Linux server accessing the Vendor 2, Series A storage array and the Vendor 2, Series B storage array

Zoning example 2— addition of an HA pair



The following illustration shows the same deployment as in Example 1, but with a gateway HA pair. For simplicity, the illustration shows only the connections and zones for the gateways and the Linux server. The zones for the Linux server and the gateways overlap at the switch.

z6 and z7, for the gateway HA pair, provide access to the following storage arrays:

- ◆ Vendor 2, Series A
- ◆ Vendor 2, Series B

z8 and z9, for the Linux server, provide access to the following storage array:

- ◆ Vendor 2, Series A
- ◆ Vendor 2, Series B



Overlapping zoning on the switch is set up as shown in the following table.

Zone	Includes...
z6	The two Vendor 2 storage array connections that go to the left switch and the two gateway connections to the same switch
z7	The two Vendor 2 storage array connections that go to the right switch and the two gateway connections to the same switch
z8	The two Vendor 2 storage array connections that go to the left switch and the Linux server connections to the same switch
z9	The two Vendor 2 storage array connections that go to the right switch and the Linux server connections to the same switch

Keep the following in mind:

- ◆ When sharing ports across heterogeneous systems, you must be careful that you do not expose array LUNs from the gateway to other systems, and the reverse. Therefore, when you create host groups on the Vendor 2 storage arrays, you must use array LUN security or array LUN masking to ensure that only the array LUNs that are for gateway storage are visible to the gateways.
- ◆ On the switch, you assign ports on the gateway and ports on the storage array to the same zone. This enables the gateways to see the LUNs on the storage arrays.
- ◆ When assigning array LUN ownership on the gateways in the HA pair, assign LUNs as follows:
 - ❖ To Gateway A: LUNs 0, 4, 8, 9 on Vendor 2, Series A
 - ❖ To Gateway B: LUNs 8, 9, 15 on Vendor 2, Series B

Guidelines for Brocade switches

Firmware upgrade issue on Brocade 3200/3800 switches

On Brocade 3200/3800 switches, when you upgrade the firmware from 3.1.1x to 3.2.0, the numbering of the ports changes. You can correct this problem by reconfiguring the fabric parameter of the switches; you need to set the value of “Switch PID Format: (0..2) [0]” to 1. Making this correction resolves the zoning issue if port-based zoning is used.

Guidelines for Cisco fabrics

Requirement with multiple ISLs on a Cisco fabric

When you use multiple ISLs on a Cisco fabric, make sure to change the vsan setting to Source/Destination load balancing for the VI connection. By default, Cisco uses an advanced source/dest/exchange method, which works with gateway connections to storage arrays but not with the IB cluster interconnect in HA pairs.

Guidelines for McDATA switches and directors

Requirement for OSMS

The gateway uses the name server's Open Systems Management Software (OSMS) to recognize array LUNs on the storage array. By default, McDATA switches are not configured to use the name server's OSMS. You must enable McDATA switches to use OSMS. Otherwise, the gateway bypasses the name services or the gateway cannot see array LUNs. The result is that the gateway cannot boot.

Note

OSMS 6.0.2 is available from your switch vendor at no extra charge.

Port setting requirement

You must set the port setting on McDATA switches to "F" for gateway connectivity.

Recommendation if you have automated scripts

In McDATA switch firmware version 9.01, a login banner is displayed before the user name prompt. If you are running scripts in your environment, it is recommended that you disable the login banner. The reason is that this login banner has been known to cause automated scripts to fail.

To disable the login banner, enter the following command at the CLI prompt:

```
config switch banner setCLIState disable
```

Requirements to use QPM blades with a McDATA 6140 switch

If you want to use QPM blades with a McDATA 6140 switch

- ◆ You need to upgrade the switch firmware to version 9.00 or later. You can then install the QPM blades in any available slot.
- ◆ If you want 2 Gb/s guaranteed bandwidth, no special port settings are required.
- ◆ If you want 4Gb/s guaranteed bandwidth, it is recommended that you configure the switch port as "F_port" and "4 Sustained." With this configuration, the port adjacent to the configured port is disabled.

Guidelines for switches with a MetroCluster

Switch option for a fabric-attached MetroCluster

Customers with MetroCluster configurations that use Brocade 200E switches and 2 GB cards should set the following switch option:

```
fabric.ops.mode.fcpProbeDisable:1
```

If you set this option and the FC-VI port hangs (becomes stuck in the down state) during a fabric event, you can bring the port up by removing the cable and re-inserting it into the port. If you do not set this option, recovering from a hung FC-VI port requires a reboot of the gateway.

Requirements for gateway fabric-attached MetroCluster configurations

Fabric-attached configurations require in-order delivery of FC-VI traffic. See the gateway *MetroCluster Guide* for information about when you need to configure switches, hardware, or both to support in-order delivery of FC-VI traffic.

About this chapter

This chapter provides information to help you determine the number and size of array LUNs that you need for your aggregates. The general guidelines in this chapter apply to provisioning array LUNs, regardless of storage array.

Note

If your gateway configuration does not use third-party storage, skip this chapter.



For any additional details specific to setting up array LUNs for the gateway on your storage array model, check the gateway *Implementation Guide* for your vendor.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Optimizing storage array RAID groups for LUNs for Data ONTAP”](#) on page 62
- ◆ [“Types of LUNs you can map to gateways”](#) on page 64
- ◆ [“Guidelines for determining the number and size of array LUNs”](#) on page 65
- ◆ [“Ensuring that spare core LUNs are available for core dumps”](#) on page 70
- ◆ [“How the checksum type affects usable space in an array LUN”](#) on page 71

Optimizing storage array RAID groups for LUNs for Data ONTAP

RAID group defined

On a storage array, a RAID group is the arrangement of disks that together form the defined RAID level. Each RAID group supports only one RAID type. The number of disks that you select for a RAID group determines the RAID type that a particular RAID group supports.

Note

Different vendors use different terms for *RAID group*. Equivalent terms among vendors include parity group, disk group, RAID set, array, and Parity RAID group. Data ONTAP RAID groups are similar in concept to RAID groups on a storage array.

Sharing array LUNs in a RAID group

You can map different array LUNs within a storage array RAID group to gateways and non gateway hosts (that is, share the RAID group). However, a single array LUN cannot be shared between gateways and non gateway hosts.

Recommendation for sizing array LUNs in the same RAID group


All the array LUNs you create for the same RAID group on the storage array should be the same size. When they are the same size, they divide equally into the capacity of the RAID group, thereby minimizing wasted disk space and enabling WAFL® to stripe across the disks efficiently.

Minimum number of array LUNs for smaller disks

If you are using smaller disks, such as 72-GB disks, the minimum number of array LUNs you need to create for each RAID group is one array LUN if the RAID group size is within the maximum array LUN size limit for your storage array. (See the information about maximum array LUN sizes for your storage array in the gateway *Implementation Guide* for your vendor.)

The following illustration shows one array LUN created for the RAID group when the disks are 72 GB.

7+1 ratio of data to parity




72-GB disks

Minimum number of array LUNs for larger disks

If the RAID group you create gives you a RAID group size greater than the maximum array LUN size that the gateway supports, you must create more than one array LUN for the RAID group. For example, for 150-GB disks in a 7 + 1 RAID group, the minimum number is two LUNs. For 300-GB disks in a 7 + 1 RAID group, the minimum number is four LUNs. See [“Minimum and maximum array LUN sizes supported by Data ONTAP”](#) on page 67 for a definition of a Data ONTAP GB.

The following illustration shows four array LUNs created for the 7 + 1 RAID group when the disks are 300 GB.

7+1 ratio of data to parity



300-GB disks

Types of LUNs you can map to gateways

Requirement for storage array LUNs

Map only storage array LUNs to the gateways. A gateway cannot use non-storage LUNs.

Attention

Some storage arrays have a non-storage “command” LUN. Do not map a “command” type LUN to a gateway.

Constraints for mapping LUN 0 to the gateway

Follow these guidelines about mapping LUN 0 to gateways.

If you have not yet mapped LUNs to gateways: Do not map LUN 0 to gateways even if LUN 0 is a storage LUN.

If you already mapped a LUN 0 storage LUN to gateways: You do not need to remove LUN 0. However, be sure to follow these rules:

- ◆ Before you change the LUN number from LUN 0 to another number, be sure that you shut down the gateway if it is up and running.
- ◆ Be sure to follow the gateway requirements for the minimum and maximum array LUN size.

Note

If an array LUN does not meet the Data ONTAP size requirements, Data ONTAP issues an error message. Do not use the array LUN identified in the error message. Remove the mapping of that array LUN to the gateway so that it cannot be used.

See the gateway *Implementation Guide* for your vendor for the minimum and maximum array LUN size requirements for your storage array.

Guidelines for determining the number and size of array LUNs

Considerations for the number and size of array LUNs you need

The following table summarizes the factors to consider when determining the number and size of array LUNs you need.

Factor	Consideration
Minimum number of array LUNs required	See “Minimum number of array LUNs per gateway” on page 66.
How you want to lay out LUNs in your aggregates.	How you plan to use aggregates for your third-party storage impacts the number and size of array LUNs that you need. See “Implications of LUN size and number on Data ONTAP RAID groups” on page 67.
Minimum and maximum LUN size allowed by Data ONTAP	See “Minimum and maximum array LUN sizes supported by Data ONTAP” on page 67 and the gateway <i>Implementation Guide</i> for your vendor.
Usable space in an array LUN	The amount of usable space in an array LUN depends on which checksum type you specify for a LUN when you assign the LUN to a gateway. Be sure to consider the impact of the checksum type when determining the number and size of array LUNs that you need. See “How the checksum type affects usable space in an array LUN” on page 71.
Relationship between the size of the LUNs and the number you need	The smaller the array LUNs, the more LUNs that you need for the storage that you want. See “Best practice recommendations for array LUN size” on page 68.

Factor	Consideration
The maximum assigned device limit supported per platform for a stand-alone system or HA pair	For every gateway platform, there is a hard-coded maximum limit for the number of devices that can be assigned to it. This maximum limit is a combination of the number of disks and array LUNs, with the LUNs from all storage arrays that are providing storage counted toward the total. This limit is the same for a stand-alone systems and an HA pair (because a node taking over for its partner must be able to accommodate both its assigned storage and its partners storage without exceeding the platform limit). Note The platform maximum device assigned limit for a gateway has no relationship to the number of storage array ports that are used or to the number of array LUNs that a particular storage array can map to a gateway.
The maximum number of array LUNs supported per neighborhood	If your configuration is limited to one stand-alone system or one HA pair, <i>neighborhoods</i> are not applicable. For gateways only, you can group up to six systems into a logical neighborhood. Using a neighborhood simplifies storage management. If you use neighborhoods, you must determine how many LUNs storage array administrators can present to the systems in the same neighborhood. See “Using Neighborhoods for Storage Management” on page 83 for more details about neighborhoods and how to determine the maximum number of LUNs you can assign to the systems in a neighborhood.

Minimum number of array LUNs per gateway

If your gateway is not using native disk shelves on which the root volume is installed, the minimum number of array LUNs that a gateway requires is as follows:

- ◆ A stand-alone gateway must own at least one array LUN.
- ◆ An HA pair must own at least two array LUNs.

Implications of LUN size and number on Data ONTAP RAID groups

Part of planning for aggregates is to plan the size and number of Data ONTAP RAID groups you need for those aggregates, and the size and number of array LUNs for the Data ONTAP RAID groups. Setting up Data ONTAP RAID groups for array LUNs includes coordination with the storage array administrator.

Planning for Data ONTAP RAID groups involves the following:

1. Planning the size of the aggregate that best meets your data needs.
2. Planning the number and size of RAID groups that you need for the size of the aggregate.
RAID groups in the same aggregate should be the same size with the same number of array LUNs in each RAID group. Use the default RAID group size if possible.
3. Planning the size of the array LUNs that you need in your Data ONTAP RAID groups.
 - ❖ To avoid a performance penalty, all array LUNs in a particular Data ONTAP RAID group should be the same size.
 - ❖ The array LUNs should be the same size in all RAID groups in the same aggregate.
4. Communicating with the storage array administrator to create the number of array LUNs of the size you need for the aggregate.
The array LUNs should be optimized for performance, according to the instructions in the storage array vendor documentation.

For more recommendations about setting up Data ONTAP RAID groups for use with third-party storage, including minimum and maximum RAID group size, see the *Data ONTAP Storage Management Guide*.

Minimum and maximum array LUN sizes supported by Data ONTAP

The size of the array LUNs that you can create on the storage array is limited by the minimum and maximum array LUN sizes that the gateway supports. The Data ONTAP definition of a gigabyte (GB) is as follows:

One GB is equal to 1000 x 1024 x 1024 bytes.

The Data ONTAP definition of a GB might not match the definition of a GB for your storage array. You must determine the minimum and maximum array LUN sizes for your storage array that are the equivalent to the minimum and maximum array LUN sizes that the gateway supports. The minimum array LUN for your vendor might be less than how Data ONTAP calculates the minimum array LUN size.



To ensure that you do not exceed the array LUN size that Data ONTAP supports for the gateway, see the gateway *Implementation Guide* for your vendor for the minimum and maximum size LUNs that you can create on your storage array according to how your vendor calculates GBs. Different vendors use different formulas for calculating units of measurement.

Note

The minimum and maximum array LUN sizes, according to Data ONTAP units of measurement, are in the gateway *Interoperability Matrix*.

Minimum array LUN size for the root volume

It is strongly recommended that you do not set the size of a root volume below the minimum root volume size shown in the gateway *Interoperability Matrix*. The reason is that you want to ensure that there is sufficient space in the root volume for system files, log files, and core files. If a system problem occurs, you need to provide these files to technical support.

Note

Both the minimum array LUN size for the root volume and the minimum array LUN size for other volumes are shown in the gateway *Interoperability Matrix*. The minimum array LUN size for a non-root volume is considerably smaller than for the root volume, so be sure that you look at the information about the minimum array LUN size for the root volume.

Best practice recommendations for array LUN size

Ideally, creating one big array LUN from a given storage array RAID group is recommended. However, this might not always be possible. If you cannot create one big array LUN from the RAID group, create array LUNs that take advantage of as many disk spindles as possible. Best practice, for performance reasons, is to create fewer and larger array LUNs rather than more LUNs that are smaller.

For example, assume that a RAID group has 40 array LUNs of 10 GB each, for a total raw space of 400 GB. The gateway sees each of the 40 array LUNs as a separate disk. Therefore, the gateway must manage 40 disks.

However, if you have fewer, larger array LUNs—for example, one 400 GB, two 200 GB, or four 100 GB array LUNs rather than 40 array LUNs of 10 GB each—the gateway has to manage fewer disks. Each array LUN incurs some WAFL overhead, but the fewer array LUNs that Data ONTAP sees, the better the overall performance.

What happens if the LUN does not meet the array LUN size requirements?

Be sure to follow the Data ONTAP requirements for the minimum and maximum array LUN size. If an array LUN does not meet the Data ONTAP size requirements, Data ONTAP issues an error message identifying the LUN. The LUN is not available for use by Data ONTAP, and Data ONTAP fails the LUN if you attempt to assign it to a gateway with the `disk assign` command. Remove the mapping of that array LUN to the gateway so that it cannot be used.

Ensuring that spare core LUNs are available for core dumps

What a core dump file is

A core dump file contains the contents of memory and NVRAM. When a hardware or software failure causes a gateway to crash, the gateway typically creates a core file that technical support can use to troubleshoot the problem.

When a spare core LUN is required

A spare core array LUN is required in the following circumstances:

- ◆ If the `cf.takeover.on_panic` option is enabled (set to `on`). The `cf.takeover.on_panic` option is enabled by default in the following circumstances:
 - ❖ When FCP is licensed
 - ❖ When iSCSI is licensed

What the `cf.takeover.on_panic` option does

The `cf.takeover.on_panic` option controls whether a node in an HA pair immediately takes over for a panicked partner. When the `cf.takeover.on_panic` option is set to `on`, the time between the initial failure and when service is fully restored is shorter because an HA pair can fail over in less time than it takes for a gateway to recover from a panic. However, the subsequent giveback causes another brief outage. If service disruption is an issue (for example, if you have CIFS clients), you might not want to incur the initial takeover.

If the `cf.takeover.on_panic` option is enabled, you must have a spare core array LUN that meets the minimum required spare core LUN size. Otherwise, no core dump file is produced on failure.

Array LUN space reserved for core dumps

One percent of the space on each array LUN is reserved for core dumps. A minimum system capacity of 272 GB results in 2.5 GB of core dump space (accounting for loss of space due to checksum formatting). Typically, this 1 percent allocation is enough for core dumps.

Where to find the minimum spare core LUN size

See the gateway *Interoperability Matrix* for the minimum spare core LUN size for each gateway.

How the checksum type affects usable space in an array LUN

Checksum types for array LUNs

When an array LUN from the storage array is mapped to be used by a gateway, the gateway treats it as a “raw,” unformatted disk (array LUN). When you assign an array LUN to a gateway, you need to tell Data ONTAP how to format the “raw” array LUN by specifying its *checksum type*.

Two checksum types are available for assignment to array LUNs—block checksum type (BCS) and zoned checksum type (ZCS). When considering the size and number of array LUNs you need for Data ONTAP, you should consider the advantages and disadvantages of the two checksum types.

Note

A *checksum* is a form of redundancy check, a simple measure for protecting the integrity of data through error detection. It is used mainly in data storage and networking protocols. It adds up the basic components of a message, typically the bytes, and stores the resulting value. Later, the authentic checksum can verify that the message was not corrupted by doing the same operation on the data, and checking the sum.

Checksum type for native disks

Only BCS type is available for native disks.

Checksum type of aggregates

Each aggregate has a checksum type associated with it. For aggregates comprised of array LUNs, the checksum type of the aggregate is determined by the checksum type of first array LUN you assign to the aggregate. Therefore, when planning how you are going to use checksums for your array LUNs, you need to decide both of the following:

- ◆ The best checksum type to assign to each array LUN
- ◆ The best checksum type for each aggregate

Required checksum type when using the deduplication feature

You must assign the BCS type to all array LUNs that will be added to FlexVol volumes on which deduplication will be run. ZCS type is not supported with deduplication.

Characteristics of checksum types

The characteristics of BCS and ZCS types are as follows.

Checksum type	Description
Block	Data ONTAP uses block checksums by default because they provide better performance in certain types of workloads. Space requirement: With block checksums, a checksum entry is appended to each WAFL block. In the case of gateways, the entry uses up a sector for every WAFL block—the sector following the WAFL block. Data ONTAP reserves 12.5 percent of the space in the array LUN for checksum. Type of aggregate for block type LUNs: An array LUN of type “block” is expected to be used with block checksum type aggregates.
Zoned	Zoned checksums provide an alternative mechanism for ensuring data integrity while also providing better storage capacity utilization than is provided for block checksum types. Note_____ At certain workloads, zoned checksum type array LUNs have a performance impact; random-read intensive workloads are affected the most. Contact your Sales Engineer for details. _____ Space requirement: With zoned checksums, entries for 63 WAFL blocks are stored in a 4K block following each set of 63 WAFL blocks. The 1/64 space for the zoned checksum is taken from the 10% WAFL reserve. Therefore, the net effect is that no additional storage is lost above the core dump and WAFL reserve. Type of aggregate for zoned type LUNs: An array LUN of type “zoned” is expected to be used with zoned checksum type aggregates.

To determine the best checksum type to assign, consider which of the following is the most important to you:

- ◆ Performance

If your primary concern is performance, block is the appropriate choice.
- ◆ Space availability in the array LUN

If having virtually all the space in the array LUN available is the most important to you, zoned is the appropriate choice. See [“Elements that reduce usable space in an array LUN”](#) on page 73.

Elements that reduce usable space in an array LUN

The usable space of an array LUN is reduced by the elements shown in the following table.

If you are using block checksums	If you are using zoned checksums
◆ 10%–WAFL reserve ◆ 1%– Core dump ◆ 20%–Volume-level Snapshot copy (default setting; optional) ◆ 5%–Aggregate-level Snapshot copy (default setting; optional) ◆ 12.5%–Checksum	◆ 10%–WAFL reserve ◆ 1%– Core dump ◆ 20%–Volume-level Snapshot copy (default setting; optional) ◆ 5%–Aggregate-level Snapshot copy (default setting; optional)

Note
Space for aggregate-level and volume-level Snapshot reserves are optional and changeable through Data ONTAP. See the Data ONTAP *Data Protection Online Backup Guide* for more information about Snapshot copies.

Examples of determining array LUN size when using checksums

You can determine array LUN size in the following respects:

- ◆ Usable capacity (the actual capacity available for storage)
- ◆ Desired capacity (the exact array LUN size desired)

Calculating the usable capacity: To determine the usable capacity for storage in a array LUN, after taking into account the Snapshot reserve, use the following formula, where Y is the usable capacity for storage and N is the total capacity of the array LUN.

Note
The Data ONTAP definition of GB might not match your storage array’s definition of GB. In the following formulas and examples, make the appropriate conversions for your storage array’s definition of GB. See “[Minimum and maximum array LUN sizes supported by Data ONTAP](#)” on page 67 for the Data ONTAP definition of GB.

Array LUN type	Formula
block	$Y = N \times \{0.875 \times 0.9 \times 0.99 \times \text{Snapshot reserve}\}$
zoned	$Y = N \times \{0.9 \times 0.99 \times \text{Snapshot reserve}\}$

Example: Total capacity of array LUN is 4 GB with a volume Snapshot reserve set at default:

Array LUN type	Formula
block	$4 \times \{0.875 \times 0.9 \times 0.99 \times 0.8\} = 2.5 \text{ GB}$
zoned	$4 \times \{0.9 \times 0.99 \times 0.8\} = 2.85 \text{ GB}$

Calculating desired capacity: To determine the array LUN capacity needed to obtain the desired capacity, while taking a Snapshot reserve into account, use the following formula, where *Y* is the exact desired array LUN size.

Array LUN type	Formula
block	$Y \text{ divided by } \{0.875 \times 0.9 \times 0.99 \times \text{Snapshot reserve}\} = \text{Actual capacity needed}$
zoned	$Y \text{ divided by } \{0.9 \times 0.99 \times \text{Snapshot reserve}\} = \text{Actual capacity needed}$

Examples: Desired array LUN capacity is 10 GB of usable capacity.

- ◆ Calculating actual array LUN size without accounting for Snapshot reserve.

Array LUN type	Formula
block	$10 \text{ GB divided by } \{0.875 \times 0.9 \times 0.99\} = 12.8 \text{ GB}$
zoned	$10 \text{ GB divided by } \{0.9 \times 0.99\} = 11.22 \text{ GB}$

- ◆ Volume Snapshot reserve left on default setting (20 percent).

Array LUN type	Formula
block	$10 \text{ GB divided by } \{0.875 \times 0.9 \times 0.99 \times 0.8\} = 16 \text{ GB}$
zoned	$10 \text{ GB divided by } \{0.9 \times 0.99 \times 0.8\} = 14 \text{ GB}$

How you change the checksum type of an array LUN

You can change the checksum type for an array LUN, for example, because your remaining array LUNs are all block checksum type and you need them to be zoned. You can change the checksum type by removing ownership of the spare

array LUN (with the `disk remove ownership` command) and using the `disk assign` command again to reassign LUN ownership with the other checksum type. See the *Data ONTAP Storage Management Guide* for details about how to change the checksum type of a LUN.

Choosing the best checksum type for an aggregate

The rules for checksum type as it applies to aggregates are as follows:

- ◆ The checksum type of the aggregate is determined by the checksum type of the first array LUN that you use to create the aggregate.
- ◆ The checksum type applies to an entire aggregate (that is, to all FlexVol volumes in the aggregate or to an entire traditional volume).
- ◆ When you create the root volume during initial setup, if an array LUN with the block checksum type is available, the system automatically selects the block checksum array LUN. If multiple block checksum array LUNs are available, the system picks the smallest of them.
- ◆ You cannot convert an aggregate from one checksum type to the other.
- ◆ For spare array LUNs, you can remove ownership of a LUN from a gateway, change the checksum type, reassign ownership of the LUN to the gateway, and add it to the aggregate.

If you want to create an aggregate that is either of the types shown in the following table, the following additional rules apply.

Type...	Then...
Block checksum	◆ To use block checksums when you create a new aggregate for FlexVol volumes, the same number of or more block checksum array LUNs must be available as the number of array LUNs you specified when you created the aggregate. ◆ The aggregate can contain only block checksum type array LUNs. If you add any array LUNs of zoned checksum type to the aggregate, the aggregate becomes a zoned checksum type. ◆ If you are specifying a number for the array LUNs to be included (rather than specific LUN IDs), ensure that there are enough block checksum array LUNs available to match the number you specified. If you specify a number of array LUNs that is greater than the number of available block array LUNs, and zoned array LUNs are available, the checksum type for the aggregate becomes zoned. The block checksum array LUNs incur an additional 12.5 percent space usage because they are part of a zoned checksum aggregate.
Zoned checksum	◆ When you create the aggregate, first specify an array LUN that is zoned. Thereafter, you can add array LUNs that are either block or zoned checksum types. ◆ If you add a block type array LUN to an aggregate that is a zoned checksum type, you lose 12.5 percent of the space in the array LUN. It is better to convert the array LUN to a zoned checksum type so that you do not lose the 12.5 percent. Attention _____ Take care when adding a block checksum array LUN to an aggregate of zoned checksum type. If you add a block checksum array LUN in error, you must destroy the aggregate to remove it. You can use the <code>aggr status -c</code> command to check the checksum type. _____

Example of how the aggregate checksum type is determined

In this example, look at the effect of entering array LUNs in different ways. Assume that when you enter the `aggr status -s` command to see the list of available array LUNs, you see a list of the following array LUNs (the checksum type you see is what you assigned earlier with the `disk assign` command).

Block array LUNs	Zoned array LUNs
L1	L6
L2	L7
L3	L8
L4	L9
L5	L10

When you enter the `disk-list` parameter when creating the aggregate, you could specify either of the following:

- ◆ The number of array LUNs you want to be in the aggregate
- ◆ The LUN names of the array LUNs you want to be in the aggregate

If you were to specify 5 as the number of array LUNs, Data ONTAP would take block array LUNs L1 through L5 and create a block checksum type aggregate. The reason is that, by default, Data ONTAP first attempts to take block checksum array LUNs if no zoned checksum array LUNs are specified.

If you were to specify 6 as the number of array LUNs, Data ONTAP would first take zoned array LUN L6, then it would take block array LUNs L1 through L5. Although block checksum array LUNs can be added to zoned checksum aggregates, zoned checksum array LUNs cannot be added to block checksum aggregates. Therefore, the result is a zoned checksum aggregate. The block checksum array LUNs incur an additional 12.5 percent space usage because they are part of a zoned checksum aggregate.

If you were to enter the array LUN names, the same rules apply. If you specify the array LUN names for only the block checksum array LUNs, the aggregate is a block checksum type. If you specify the array LUN names for L4, L5, L6, and L7, for example, the aggregate is a zoned checksum type (because you cannot mix block and zoned checksum array LUNs in a block checksum aggregate).

Changing checksum type after an aggregate is created

If you want to change the checksum type of an aggregate after you create it, you must destroy the aggregate and re-create the aggregate of the desired checksum type.

About this chapter

This chapter explains what you need to configure on a gateway to enable it to use third-party storage, native disks, or both. It also describes what you need to plan for before you start configuring your system. The *Data ONTAP Storage Management Guide* provides more details about the Data ONTAP disk ownership functionality and aggregates, which are discussed in this chapter. Understanding how ownership works enables you to maximize storage redundancy and manage your spares effectively.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Requirement for a gateway to use array LUNs”](#) on page 80
- ◆ [“Requirement for a gateway to be able to use a native disk”](#) on page 82

Requirement for a gateway to use array LUNs

To make array LUNs accessible (available) for gateways, you must create the LUNs on the storage array and map the LUNs to ports on the storage array to which the gateway is attached, either directly or through a switch.

Prerequisites for using an array LUN for storage

Although you make a LUN on the storage array accessible to Data ONTAP, you must complete both of the following configuration tasks in Data ONTAP before a gateway can *use* a particular array LUN for storage:

1. You must assign one gateway to be the *owner* of the array LUN.
2. You must add the array LUN to an *aggregate*.

What ownership means

When you assign an array LUN to a gateway, Data ONTAP writes data to the array LUN to identify the assigned system as the *owner* of the LUN. When you assign an array LUN to a gateway, it becomes a *spare LUN* owned by that gateway and it is no longer available to any other gateway. Only one gateway can be the owner.

Ownership of the array LUN is from the perspective of Data ONTAP, not from the perspective of the storage array. From the point of view of Data ONTAP, this logical relationship between an array LUN and a gateway is referred to as disk ownership because Data ONTAP considers an array LUN to be a virtual disk.

How the array LUN becomes available for storage

For the owning system to use the array LUN for storage, you must add it to an aggregate. The array LUN remains a spare LUN owned by that gateway until you add it to an aggregate. After the array LUN is added to an aggregate, Data ONTAP ensures that only the owner of the array LUN can write data to and read data from the LUN.

Note

An *aggregate* is a Data ONTAP “physical” storage unit that organizes a collection of LUNs from the storage array.

In an HA pair, both nodes in the pair must be able to see the same storage. But only one node is the owner of the disk or array LUN. The partner node takes over read/write access to a disk or array LUN in case of a failure of the owning node. But, the original owning node takes over again after the problem that caused unavailability of the node is fixed.

Considerations when planning for disk ownership

If you are deploying multiple gateways, you must determine which gateways will “own” which array LUNs. Consider the following when planning array LUN assignment:

- ◆ The maximum assigned device limit supported by your platform
The gateway *Interoperability Matrix* shows the maximum assigned device limit that is supported for different platforms. This is a hard-coded limit. If your system uses both array LUNs and disks, this maximum limit is the maximum of disks and array LUNs *combined*. You must account for both types of storage when determining how many array LUNs and disks you can assign to a system.
- ◆ Amount of load that you expect to be generated by different applications used in your environment
Some types of applications are likely to generate a lot of requests whereas other applications (for example, archival applications) generate fewer requests. You might want to consider weighing ownership assignments based on expected load from specific applications.

Array LUN assignment changes

You can change which gateway owns which array LUNs, if needed. If the LUN has already been added to an aggregate, you must destroy the aggregate to which the LUN is assigned before you can change the ownership of the LUN.

Load balancing among systems that see the same array LUNs

If your configuration includes more than one stand-alone system or more than one HA pair, you can use the functionality provided by a logical *neighborhood* for load balancing among up to six gateways. Neighborhoods are applicable only for gateways, and only with the storage provided by the third-party storage arrays. Load balancing among disks is limited to the nodes in an HA pair. See Chapter 7, “[Using Neighborhoods for Storage Management](#),” on page 83.

Requirement for a gateway to be able to use a native disk

Although you always have to assign ownership of array LUNs to a gateway, this is not the case for any native disks attached to your gateway. By default, Data ONTAP automatically assigns ownership to native disks attached to your system. When disk autoassignment runs, Data ONTAP looks for any unassigned disks and assigns them to the same system as all other disks on their loop or stack.

As is the case for array LUNs, a disk must be assigned to an aggregate before it can be used for storage.

If your gateway uses both third-party storage and disks, be sure to consider how many disks and array LUNs combined can be assigned to your system without exceeding the supported maximum assigned device limit for your system.

See the gateway *Implementation Guide for Native Disk Shelves* and the Data ONTAP *Storage Management Guide* for information about disk ownership for storage on native disk shelves connected to a gateway.

About this chapter

This chapter describes what a neighborhood is, how you can use neighborhoods, and what you need to plan for to establish a neighborhood.

Note

Neighborhoods do not pertain to configurations with only one stand-alone system or only one HA pair and no other systems. Skip this chapter if your gateway configuration is a single stand-alone system, a single HA pair, or you do not use third-party storage.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Overview of a neighborhood”](#) on page 84
- ◆ [“Neighborhood requirements”](#) on page 87
- ◆ [“How a neighborhood is established”](#) on page 89
- ◆ [“Limits to consider when planning for a neighborhood”](#) on page 91
- ◆ [“Factors that impact the neighborhood maximum LUN limit”](#) on page 93
 - ❖ [“Factor 1—Mixing gateway platforms in a neighborhood”](#) on page 94
 - ❖ [“Factor 2—Storage arrays specific restrictions”](#) on page 95
 - ❖ [“Factor 3—Using both native disks and array LUNs”](#) on page 96

Overview of a neighborhood

What a neighborhood is

A neighborhood is a logical entity that enables gateways without an HA pair relationship to see the same *array LUNs*.

Ordinarily systems that are not part of an HA pair have no relationship with each other; they cannot see the same array LUNs and you cannot load balance between them. However, with a neighborhood, you can easily reassign ownership of array LUNs from one neighborhood member to another through Data ONTAP—among up to *six* systems that see the same array LUNs. You can also transparently load balance data service among the gateways in a neighborhood by using SnapMover to move vFiler units among neighborhood members. The physical storage always remains on the storage array.

Neighborhood functionality is limited to third-party storage. Although neighborhood members see the same array LUNs, the systems outside of an HA pair cannot see each other's disks.

The neighborhood relationship does not provide any failover between neighborhood members if a member becomes unavailable. Failover of services is a function of the relationship between two nodes in an HA pair, and can occur only between the two nodes in an HA pair.

When a neighborhood is not applicable

Neighborhoods do not apply to the following:


- ◆ Configurations with only one stand-alone system or one HA pair
Neighborhoods have no relevance to a gateway configuration that includes only one stand-alone system or only one HA pair and no other systems. The functionality provided by an HA pair enables you to load balance between the two nodes in the pair and change ownership among the neighborhood members.
- ◆ Native disks on gateways
The neighborhood functionality applies only to array LUNs. However, the number of disks attached to a gateway in the neighborhood must be factored in when you calculate the number of array LUNs that the systems in your neighborhood can support. See [“Factor 3—Using both native disks and array LUNs”](#) on page 96 for more information about calculating neighborhood LUN limits.
- ◆ Storage arrays for which neighborhoods are not supported

Neighborhoods are not supported with all storage arrays. If a neighborhood is not supported with a storage array, the configuration is limited to either one stand-alone system or one HA pair. The gateway *Interoperability Matrix* and the gateway *Implementation Guides* provide information about whether neighborhoods are supported with specific vendors and storage arrays.

◆ MetroCluster configurations

Example neighborhood

The following example shows a gateway neighborhood that includes three gateways—an HA pair (vs1 and vs2) and a stand-alone system (vs3). Two storage arrays (Storage 1 and Storage 2) provide LUNs to the three gateways that form the gateway neighborhood.



In the example, the output from the `disk show -v` command on all three gateways shows that all three gateways can see the same LUNs from Storage 1 and Storage 2. All gateways in the neighborhood must be able to see the same LUNs (although in Data ONTAP only one system is assigned as the owner of the LUN).

You calculate the number of array LUNs that would be presented to the neighborhood by adding the 50 LUNs that Storage 1 mapped to these systems to the 100 LUNs Storage 2 mapped to these systems. The total number of LUNs presented to the neighborhood from the two storage arrays is 150. This total LUN

number cannot exceed the neighborhood maximum LUN limit specified in the gateway *Interoperability Matrix* for the platform type and adjusted, as necessary, for other factors. (See “[Factors that impact the neighborhood maximum LUN limit](#)” on page 93.)

Number of neighborhoods you can have

There is no limit to the number of gateway neighborhoods that you can deploy within a given SAN. If you use multiple neighborhoods, ensure that only the gateways in one neighborhood can see a particular set of array LUNs.

To illustrate use of multiple neighborhoods, assume that Company ABC has two data centers. One data center has four gateways and one storage array, Hitachi A. The other data center has six gateways and one storage array, Hitachi B. The company wants to manage the storage at each data center separately, using SnapMover to balance the traffic among its gateways in each data center.

The company establishes two different neighborhoods, one for each data center. Each neighborhood includes only the gateways at the one data center. The gateways in one neighborhood cannot see the array LUNs on the storage array in the neighborhood at the other data center.

Neighborhood requirements

Rules for a neighborhood

The following table summarizes the requirements for neighborhoods and what is supported.

Factor	Rules
Visibility of storage	All gateways in the same neighborhood must be able to see the <i>same</i> array LUNs.
Maximum number of gateways in a neighborhood	Six
Configurations supported	You can include any combination of stand-alone systems and HA pairs in a neighborhood. Note _____ Neighborhoods are not supported with MetroClusters.
Platforms in a neighborhood	You can mix platform types in a neighborhood.
Neighborhood maximum LUN limit	You cannot exceed the maximum LUN limit for <i>your</i> neighborhood. Although there is base neighborhood maximum LUN limit by platform type in the gateway <i>Interoperability Matrix</i> , other factors might reduce that limit for your neighborhood. See “ Factors that impact the neighborhood maximum LUN limit ” on page 93 for more information.

Factor	Rules
Storage arrays supported	Data ONTAP does not support neighborhoods for all storage arrays. See the gateway <i>Implementation Guide</i> for your vendor and the gateway <i>Interoperability Matrix</i> to ensure that neighborhoods are supported for your storage array.
Data ONTAP releases supported	See the gateway <i>Implementation Guide</i> for your vendor and the gateway <i>Interoperability Matrix</i> to ensure that neighborhoods are supported for your Data ONTAP release.
Protocols	You should use the NFS protocol so that changing ownership of array LUNs among neighborhood members is transparent.
Additional software	To move vFiler units among neighborhood members quickly, without copying the data, use MultiStore and SnapMover.

Possible configurations if neighborhoods are not supported

If neighborhoods are not supported for your storage array or Data ONTAP release, your configuration must be either a single stand-alone system or a single HA pair.

How a neighborhood is established

A neighborhood exists only when a gateway can see array LUNs that belong to another gateway and that system is not its partner in an HA pair. To establish a neighborhood, you must configure the storage arrays and switches to enable all gateways that you want in the same neighborhood to see the same array LUNs.

Data ONTAP configuration to support neighborhoods

No explicit configuration is required on a gateway to support neighborhoods. The underlying functionality that enables the gateways to operate as a neighborhood is the Data ONTAP disk ownership feature (assigning array LUNs to a specific gateway).

Note

After you determine the maximum number of LUNs that the storage arrays can present to your neighborhood, be sure to communicate that information to the storage array administrators.

Storage array configuration to support neighborhoods

The storage array administrator must configure one or more storage arrays to present the same LUNs to the gateways that you want to be in the neighborhood. The method by which a storage array administrator creates and presents LUNs to hosts varies on different storage arrays. A typical method is that the storage array administrator specifies the FC initiator ports of a number of gateways to be in the same “host group” on the storage array. This host group configuration enables all the systems to see the same LUNs.

Note

The term “host group” is used on some storage arrays to describe an entity in the storage partition topology that defines a logical collection of host computers that require shared access to one or more logical drives. Different storage arrays use different terms to describe this configuration entity.

The storage array administrator must also set the storage array access controls so that all the gateways can see the same array LUNs.

Switch configuration to support neighborhoods

If your configuration is fabric attached, you must zone switch ports that connect to your gateway's FC initiator ports so that all gateways in the neighborhood can see the same array LUNs.

Note

It is recommended that you use single-initiator zoning, which limits each zone to a single gateway FC initiator port.

Limits to consider when planning for a neighborhood

When determining the number of LUNs that storage arrays can present to the gateway in your neighborhood, you need to consider the following two Data ONTAP limits:

- ◆ Maximum total number of devices that can be assigned to the platforms that you want to be neighborhood members
- ◆ Neighborhood maximum LUN limit

Platform maximum assigned device limit

For every platform, Data ONTAP hard codes the maximum number of devices (disks and array LUNs combined) that a stand-alone system or HA pair can support. You can think about this limit as the *platform maximum assigned device limit* because the combination of disks and LUNs that you *assign* to a gateway (through the Data ONTAP disk ownership feature) cannot exceed this limit.

The platform maximum assigned device limit does not limit the number of disks and array LUNs that the systems in a gateway neighborhood can *see* (the *visible* limit); it limits only the number of disks and LUNs that you can assign to the platform. The visible limit is determined by the neighborhood maximum LUN limit.

Note

The platform maximum assigned device limit applies whether a gateway is in a neighborhood or not. The platform maximum assigned device limit is the same for a stand-alone system and HA pair because each node in the pair must be able to handle its storage and its partner's storage if the partner becomes unavailable.

Neighborhood maximum LUN limit

The neighborhood maximum LUN limit has the following two components:

- ◆ It is the maximum *visible* limit for the neighborhood.

This limit is the maximum number of the *same* array LUNs that gateways in a neighborhood are allowed to *see*. All members of the neighborhood see all the same array LUNs.

Individual gateways in the neighborhood that have disks attached cannot see more array LUNs and disks combined than the maximum visible limit. See [“Factor 3—Using both native disks and array LUNs”](#) on page 96 for more information.
- ◆ It is the *maximum assigned LUN limit* for all the systems in the neighborhood combined.

The platform maximum assigned device limit, not the neighborhood maximum LUN limit, dictates the maximum number of disks and LUNs that can be assigned to a stand-alone system or HA pair. However, in the context of the neighborhood, you might not be able to assign the maximum number of devices that the platform can support because the combined total of assigned devices must not exceed the neighborhood limit.

For example, assume that you have two stand-alone systems in the neighborhood, and the maximum assigned limit for the platform is 600 devices. If the neighborhood LUN limit was 1,000 array LUNs, you could not assign 600 array LUNs to each system because the total assigned LUNs for the two systems would be 1,200, which is 200 LUNs more than the 1,000 neighborhood maximum LUN limit.

The gateway *Interoperability Matrix* provides a neighborhood maximum LUN limit for each gateway platform type. You can use the neighborhood maximum LUN limit shown if your neighborhood members are using only array LUNs, and no other factors would reduce the neighborhood maximum LUN limit for your neighborhood. There are a number of factors that can reduce the neighborhood maximum LUN limit. Therefore, the limit in the *Interoperability Matrix* might not be the actual limit for your neighborhood. See “[Factors that impact the neighborhood maximum LUN limit](#)” on page 93.

What the neighborhood maximum LUN limit means for storage arrays

If you have one storage array, that storage array cannot present to the neighborhood members more LUNs than the neighborhood maximum LUN limit. If multiple storage arrays are presenting LUNs to the neighborhood members, the combination of LUNs from all those storage arrays cannot exceed the neighborhood maximum LUN limit.

Factors that impact the neighborhood maximum LUN limit

What impacts the neighborhood maximum LUN limit

Neighborhood members can never *see* more array LUNs than the neighborhood maximum LUN limit that is shown in the *gateway Interoperability Matrix* for the platform type. However, certain factors can reduce the maximum LUN limit from the number shown in the *Interoperability Matrix*. These factors are:

- ◆ The systems in the neighborhood are mixed platform types.
- ◆ There are neighborhood restrictions related to the storage arrays presenting LUNs to the neighborhood systems.
- ◆ Disks are connected to the gateways in the neighborhood.

Attention

The lowest limit based on any factor becomes the maximum LUN limit for your neighborhood.

Factor 1—Mixing gateway platforms in a neighborhood

If the systems in your neighborhood are different platform types, the platform with the lowest maximum LUN limit establishes the neighborhood maximum LUN limit for the neighborhood. In the gateway *Interoperability Matrix*, look for the limit for the lowest platform in your neighborhood.

Factor 2—Storage arrays specific restrictions

For some storage array models and vendors, there is a maximum limit on the number of LUNs that are allowed per storage array or per “host group” on a storage array.

Assume that your configuration includes two of the same model storage arrays and the maximum host group limit for that storage array model is 256 LUNs. Each storage array has one host group with 256 LUNs. You would factor the host group limit into your neighborhood LUN limit calculations as follows:

- ◆ There are two host groups, each with the maximum number of LUNs allowed for the host group. You multiply 256 (LUNs) X 2 (host groups). The maximum number of LUNs that would be presented to the neighborhood is 512.
- ◆ If 512 LUNs is the same as or lower than the maximum LUN limit shown in the gateway *Interoperability Matrix* for the platforms in the neighborhood, presenting 512 LUNs to the neighborhood systems is acceptable.



Check the gateway *Implementation Guide* for your vendor to determine whether there are maximum LUN limits per host group for your vendor or storage array model.

Factor 3—Using both native disks and array LUNs

Effect on platform maximum assigned limit

You can never exceed the platform maximum assigned limit when *assigning* disks and array LUNs to your system. The number of assigned disks and the number of assigned array LUNs combined counts toward the platform maximum assigned limit.

Effect on neighborhood maximum assigned limit

The functionality provided by a neighborhood—load balancing and changing ownership among up to *six* systems—does not apply to native disks. However, when disks are attached to a system in the neighborhood that uses both disks and array LUNs, you must count the disks toward the visible limit. When disks are connected to neighborhood members, you calculate the visible limit differently than the way you calculate the visible limit when neighborhood members use only array LUNs as follows.

When members of the neighborhood use only array LUNs: Every neighborhood member sees the same LUNs, up to the neighborhood maximum LUN limit in the gateway *Interoperability Matrix* (adjusted according to other factors).

When disks are attached to the systems in the neighborhood: The neighborhood maximum LUN limit shown in the *Interoperability Matrix* is really the *visible device limit*. For any given stand-alone system or HA pair in the neighborhood, you combine the number of array LUNs that it can see and the disks attached to it (which it also sees). The total cannot exceed the neighborhood maximum LUN limit shown in the *Interoperability Matrix* for the lowest platform type in the neighborhood.

When disks are attached to the neighborhood systems you calculate the visible device limit from the perspective of each stand-alone system or HA pair node to which disks are connected. You combine the number of disks and array LUNs that the system sees and check to make sure that number does not exceed the neighborhood maximum LUN limit in the *Interoperability Matrix*.

The reason you apply the limit from the perspective of an individual stand-alone system or HA pair is because not all neighborhood members can see all the disks attached neighborhood systems. A stand-alone system is the only system that can

see the disks that are connected to it; the nodes in an HA pair are the only systems that can see the disks connected to the nodes in the pair. (In contrast, all members of the neighborhood see the all the same array LUNs.)

Calculation of visible limit for both array LUNs and disks

Assume that in this neighborhood example there are two HA pairs and one stand-alone system, with disks connected to each system. The following table shows the maximum used in this example for each limit.

Limit type	Maximum allowed
Neighborhood maximum LUN limit (<i>visible device limit</i>)	1200
Platform maximum assigned limit (disks and array LUNs combined) for the platform type	500

The following table shows calculations to determine whether the planned *visible* limit of array LUNs and disks this neighborhood is within the supported limits. For simplicity, assume there are no restrictions related to the storage arrays that present LUNs to the neighborhood members. (This table does not show the number of disks and LUNs assigned to each system.)

System	Number of array LUNs seen	Number of disks connected to it	Total disks and array LUNs	Within visible limits?
HA pair 1:				
vs-X	1,000	200	1,200	Yes
vs-Y	1,000	150	1,150	Yes
HA pair 2:				
vs-A	1,000	200	1,200	Yes
vs-B	1,000	350	1,350	No. (The maximum visible device limit is 1200)
Stand-alone system:				

System	Number of array LUNs seen	Number of disks connected to it	Total disks and array LUNs	Within visible limits?
vs-standalone	1,000	100	1,100	Yes

All neighborhood members see all the same array LUNs. Therefore, for every system in this example, the 1,000 LUNs is counted toward the maximum visible limit of 1,200. When disks and array LUNs are both used for storage, think of this as the device limit.

Disks are counted on a system-by-system basis toward the visible device limit. Therefore, the disks are counted toward the limit for the system to which the disks are connected, and not toward the visible device limit of any other systems.

For four systems in this example neighborhood, the visible device limit is the same as or less than the visible device limit. However, if 1,000 array LUNs were presented to the neighborhood, you would not be able to connect 350 disks to system vs-B in HA pair 2 because the total of array LUNs and disks seen by system vs-B would be 1,350, which is more than the visible device limit of 1,200.

No more than 500 array LUNs and disks combined can be *assigned* to any one system in this example (so that the *platform maximum assigned limit* is not exceeded). However, you cannot assign 500 devices to each of the systems in the neighborhood because the total assigned devices would be 2,500. The combined assigned devices for the neighborhood cannot exceed the neighborhood LUN limit, so the total assigned devices must be 1,200 or less.

About this chapter

This chapter describes which storage you can and cannot mix in the same aggregate, including the rules about mixing array LUNs from different storage arrays in the same aggregate.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Rules about mixing storage in aggregates”](#) on page 100
- ◆ [“Rules for aggregates if multiple storage arrays are deployed”](#) on page 102

Rules about mixing storage in aggregates

The following table summarizes the rules about adding array LUNs, disks, or both to Data ONTAP aggregates.

Note

There are restrictions on the types of array LUNs that you can mix in the same aggregate, which you must observe when you add array LUNs to an aggregate. Data ONTAP does not *prevent* you from mixing different types of array LUNs in an aggregate.

If...	The rule is...
The array LUNs are from different storage array vendors or they are different storage array types from the same vendor	A stand-alone gateway or a gateway HA pair can connect to multiple storage arrays from the same vendor or from different vendors. However, you cannot assign array LUNs from different vendors or from different model families of the same vendor to the same Data ONTAP aggregate. See “Rules for aggregates if multiple storage arrays are deployed” on page 102 for more information.
The array LUNs are from different storage array drive types	Array LUNs from a storage array with Fibre Channel disks and array LUNs from a storage array with SATA disks cannot be assigned to the same aggregate, even if the model is in the same family. You must treat the disks as if they are from different families.
Your configuration includes a Hitachi storage array or an HP XP storage array with external drives	Be sure to check the gateway <i>Implementation Guide for Hitachi Storage</i> or <i>Implementation Guide for HP XP Storage</i> (as appropriate) for additional details.
Your system is using both array LUNs and native disks	You cannot assign array LUNs and disks to the same aggregate.
The native disks your system uses are different disks types or different speeds	Do not assign disks of different drive types or speeds to the same aggregate.

Note

Setting up and managing native disks with gateways and IBM N series storage systems is virtually the same. See the gateway *Implementation Guide* for gateway-specific requirements for using native disks. See the Data ONTAP *Storage Management Guide* for detailed information about disks on all storage systems that run Data ONTAP.

Rules for aggregates if multiple storage arrays are deployed

The rules about assigning array LUNs to aggregates differ depending on whether the storage arrays in your configuration are from the same vendor and whether they are they are from the same storage array model family. The basic rule is that you cannot mix array LUNs from different storage array families in the same aggregate.

What a family is

Storage arrays in the same *family* share the same performance and failover characteristics. For example, members of the same family all perform active-active failover or they all perform active-passive failover. Storage arrays with 4 GB HBAs are not considered to be in the same family as storage arrays with 2 GB HBAs.

Each *gateway Implementation Guide* contains information about which storage array models are in the same family.

When to contact a sales engineer

If your configuration is large, contact a sales engineer for help with planning for deploying multiple storage arrays behind the same gateway. Factors such as which storage arrays you are mixing behind a gateway and the load that is being handled (for example, for archive or database processing) might make a difference in your deployment.


Rules when the storage arrays are from the same family

When the storage arrays behind the gateway are of the same family (and are from the same vendor), you can do the following:


- ◆ Assign all LUNs from multiple storage arrays of the same family to the same aggregate (Example 1).
- ◆ Create multiple aggregates but assign the LUNs from only one storage array to a specific aggregate (Example 2).
- ◆ Distribute and mix LUNs from various storage arrays over multiple aggregates (Example 3).

The following illustrations show how you can assign array LUNs to aggregates when all the vendor's storage arrays are from the same family. For simplicity, the illustrations show only two storage arrays; your deployment can include more storage arrays.


Example 1: One aggregate with all array LUNs from the same storage array family assigned it.



Example 2: Multiple aggregates, each of which contains array LUNs from a single storage array even though the storage arrays are from the same family.



Example 3: Multiple aggregates, with array LUNs from the storage arrays in the same family mixed in the aggregates.



Example 3 is not supported if one of the storage arrays has Fibre Channel drives and the other storage array has SATA drives.

Rules when the storage arrays are from different family or different vendors

When the storage arrays from the same vendor are from a different family or the storage arrays are from different vendors, the rules are as follows:


- ◆ You cannot mix array LUNs from a different storage array family or from a different storage array vendor in the same aggregate.
- ◆ The aggregate with the root volume can be associated with any of the family types.
- ◆ If when you configure an aggregate you accidentally mix array LUNs from a different storage array family or from a different storage array vendor, you must destroy the aggregate and re-create it.

Note


Be sure that when you create your aggregate you explicitly specify the IDs of the array LUNs for the aggregate. Do not use the parameters for specifying number and size of array LUNs (disks) to be picked up because the system might automatically pick up LUNs from the different family or from a different vendor's storage array.

The following illustrations show examples of how you can assign array LUNs to aggregates when the storage arrays are from a different family or from a different vendor.

Example 1: LUNs from the two storage arrays cannot be mixed in the same aggregate because the storage arrays are a different family. If the storage arrays are from different vendors, you cannot mix LUNS from the different storage arrays in the same aggregate.



Example 2: LUNs from the same storage array family can be mixed in an aggregate but LUNs from a different storage array family (or different vendor) cannot be mixed.



In Example 2, one storage array is from Family A and two storage arrays are from Family B. The LUNs from the Family A storage array cannot be assigned to the same aggregate as the LUNs from a Family B storage array because they are in a different family. However, LUN 1 of storage array 3 can be assigned to aggregate 2, which also contains LUNs from storage array 2, because the two storage arrays are in the same family.

About this chapter

This chapter enables you to identify which FC initiator ports on your gateway are redundant and provides guidelines for use of the ports on your gateway.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Identifying the ports on your gateway”](#) on page 108
- ◆ [“Guidelines for gateway port usage”](#) on page 112

Identifying the ports on your gateway

How FC initiator ports are labeled

All FC initiator ports on gateways are identified by a number and a letter, as shown in the following table.

When the ports are on...	The number identifying the port...
Motherboard (not in expansion slots)	Starts with the number 0, and the letter that follows 0 distinguishes the port, for example, 0a, 0b, 0c, and 0d.
Cards in expansion slots	Refers to the slot in the chassis and the letter distinguishes the different ports. For example, if you install a two-port card into slot 3, the ports on that card are identified as 3A and 3B. Note_____The FC initiator ports on the gateway are labeled 1 and 2. However, the software refers to them as A and B. You see these labels in the user interface and system messages displayed on the console.

What a port pair is

A port pair consists of two different FC initiator ports that are used to access the same set of array LUNs or disks. To ensure redundancy, select port pairs when attaching FC initiator ports to array LUNs or native disks.

Location of ports, by gateway

The following table shows the location of FC initiator ports, by gateway.

Note_____The terminology used for the hardware component on which host adapters and ports are located varies on different storage array models. The following table uses cluster (controller) to refer to that hardware component. Ensure that the two ports on the storage array that you select to access a given array LUN are from different components that could represent a SPOF.

Gateway	FC initiator port location on the gateway...	To create a redundant port pair...
N5xxx and N6xxx gateways	Four onboard FC initiator ports and ports on FC cards in expansion slots. The chip sets for the onboard ports are as follows: ◆ Chip set 1: 0a and 0b ◆ Chip set 2: 0c and 0d	Redundant onboard port pairs are: ◆ 0a, 0c ◆ 0a, 0d ◆ 0b, 0c ◆ 0b, 0d For example, when connecting to a storage array, you would: ◆ Connect one cable from port 0a on the gateway to cluster (controller) 1 port 1 on the storage array. ◆ Connect another cable from port 0c on the gateway 1 to cluster (controller) 2 port 1 on the storage array.
N6200 series gateways	Two onboard FC initiator ports and ports on FC cards in expansion slots. Note_____ The onboard FC ports (0c, 0d) are not on independent busses. They do not provide a redundant port pair. Some port failures could cause the system to panic. To create a redundant port pair use a supported FC card in an available expansion slot. _____	Only the FC card in an expansion slot provides a redundant port pair. Use only two initiator ports on a single HBA. The ports should be on separate busses. For example, use ports 1a and 1c. In a MetroCluster configuration, you must use an add-on HBA. Be sure to use redundant port pairs for the local and remote storage. For example, use ports 2a and 2c for the local storage and 2b and 2d for the remote storage.

Gateway	FC initiator port location on the gateway...	To create a redundant port pair...
N7600, N7700, N7800, N7900 systems, and N7x50T series systems	Eight onboard FC initiator ports and ports on FC cards in expansion slots A N7600, N7700, N7800, N7900 and N7x50T series system has two buses. Each chip set on the two buses has two ports, as follows: ◆ Chip set 1: Ports 0a and 0b ◆ Chip set 2: Ports 0c and 0d ◆ Chip set 3: Ports 0e and 0f ◆ Chip set 4: Ports 0g and 0h Every four onboard ports share one PCI bus (for example, 0a, 0b, 0c, and 0d are on one bus)	For optimal redundancy, best practice is to select a port from each bus to create the redundant port pair. For example, when connecting to a storage array, you would: ◆ Connect one cable from port 0a on the gateway to cluster (controller) 1 port 1 on the storage array. ◆ Connect another cable from port 0e on the gateway 1 to cluster (controller) 2 port 1 on the storage array. If you must form a port pair by using ports from the same bus, be sure that each port in the pair is from a different chip set on that bus. If a bus fails, both ports also fail.

Note

See the *Hardware and Service Guide* for your product for information about which slots can be used for which cards.


Connector type required

Gateway FC initiator ports require LC connectors. However, the ports on your storage array might require a connector type different from an LC connector. If so, consult your storage array documentation for details.

How to handle unused ports on the gateway

Handle unused ports as follows:

- ◆ For N7000 series gateways, use the optional `fcadmin config -t unconfig adapter_name` command to set any unused embedded ports offline.
- ◆ For other platforms, use a loopback plug for any unused Fibre Channel ports (see the following figure).



Guidelines for gateway port usage

Using redundant port pairs

Use redundant port pairs when connecting your gateway to array LUNs and native disks. See “[Location of ports, by gateway](#)” on page 108 for guidelines about which ports on your system are redundant.

Initiator port setting requirement for HBAs

All gateway HBAs that are used to access disks or array LUNs must be set to initiator ports. The factory might configure some HBAs as target ports. Therefore, when you set up your system you need to check the HBA settings and, if necessary, change a target HBA setting to an initiator. See “[Step 2: Install the gateway](#)” on page 141 for information about how to check the setting of the HBAs and change them if necessary.

Number of ports required to connect to a LUN group on a storage array

The gateway requires two paths to each array LUN. Use two redundant gateway FC initiator ports to access each LUN group on the storage array.

Restriction for sharing an initiator port for multiple storage arrays

You cannot access multiple storage arrays through the same gateway FC initiator port.

Restrictions for sharing initiator ports for different storage and devices

Use a separate gateway initiator port to connect each of the following:

- ◆ Disk shelves
- ◆ Array LUNs
- ◆ Tape devices

If your gateway does not have enough internal ports for your needs, you must order an additional HBA.

Number of ports required to connect to a disk shelf

You need to use the following number of ports to connect the gateway to a disk shelf:

For a stand-alone system:

- ◆ If you are using one loop, use one FC initiator port
- ◆ If you are connecting two loops, use two FC initiator ports, one for each loop.

For an HA pair:

- ◆ If you are using Multipath Storage, use two FC initiator ports for each loop
- ◆ If you are not using Multipath Storage, use one FC initiator port for each controller for each loop in the configuration

General port usage guidelines for systems accessing array LUNs and disks

Consider redundancy and load balancing when selecting which initiator ports to use to connect to array LUNs and disk shelves. You should use redundant port pairs to access storage. For example, on N5200, N5300, N5500, N5600 gateways, the redundant port pairs are as follows:

- ❖ 0a and 0c
- ❖ 0b and 0d

If you are using 0a and 0c as the redundant port pair to access array LUNs, 0b and 0d is the redundant port pair to access the disk shelf.

See “[Location of ports, by gateway](#)” on page 108 for information about which ports are redundant on specific gateways.

Where to obtain more information

You can obtain additional information about port usage from the following sources.

For more information about...	See...
Expansion slots and ports	◆ The appropriate hardware and service guide
For HA pairs, ports for connections to disk shelves	◆ <i>Data ONTAP Active/Active Configuration Guide</i> (for 7.3.x releases) or <i>High Availability Guide</i> (for 8.0 and later releases)

For more information about...	See...
Additional best practice recommendations for port usage	www.ibm.com/storage/support/nseries/

About this chapter

This chapter provides instructions for how to connect your stand-alone gateway or HA pair to a back-end storage array.

Note

This chapter does not include cabling guidelines for connecting MetroCluster configurations. See the Gateway *MetroCluster Guide* for connection information.

Topics in this chapter

This document contains the following topics:

- ◆ [“When to connect your gateway to third-party storage”](#) on page 116
- ◆ [“Planning a port-to-port connectivity scheme”](#) on page 118
- ◆ [“Connection locations on a gateway”](#) on page 119
- ◆ [“Connecting a stand-alone gateway”](#) on page 125
- ◆ [“Connecting a gateway HA pair”](#) on page 128
- ◆ [“Connecting a MetroCluster”](#) on page 133
- ◆ [“Settings for Connecting to an ASCII Terminal Console”](#) on page 134

Terminology

Cluster (controller) is used in this chapter as a generic term to describe the hardware component on which host adapters and ports on the storage array are located. The terminology for this component varies among vendors and even among array models from the same vendor. To provide redundant paths to an array LUN for the gateway, ensure that the two ports that you select to access the array LUN are from different components that could each represent a single point of failure.

Determining which gateways are supported in which Data ONTAP releases

Not all gateways described in this guide are supported by all Data ONTAP releases. See the gateway *Interoperability Matrix* to determine which gateways are supported in a particular Data ONTAP release.

Note

The gateway *Interoperability Matrix* is the final authority about which gateways are supported in particular Data ONTAP releases.

When to connect your gateway to third-party storage

The sequence for setting up a third-party storage array and connecting it to your gateway differs depending on how your ordered your gateway.

If...	Then...
Your gateway WAS ordered with native disk shelves	When you order your system with disk shelves, the factory installs the root volume, licenses, and the Data ONTAP software for you. You use the setup program on your gateway to complete basic setup (identify the gateway on the network) before you start configuring the gateway to use third-party storage. If desired, you can wait until after basic gateway setup to prepare third-party storage for Data ONTAP use and to connect your system to third-party storage. If you are waiting to set up your gateway with third-party storage, use the following information, in the order shown: 1. See the basic instructions about setting up a gateway with native disk shelves, in the gateway <i>Software Setup, Installation, and Management Guide</i> (for Data ONTAP 7.3.x) or the Data ONTAP <i>setup guides</i> (for Data ONTAP 8.0 or later). 2. Use the instructions in this chapter to connect your system to the storage array. 3. See the Data ONTAP <i>Storage Management Guide</i> for information about how to assign array LUNs to your gateway.

If...	Then...
If your gateway <i>WAS NOT</i> ordered with native disk shelves	You must prepare the storage for Data ONTAP on the storage array and connect the storage array to the gateway before you start to set up the gateway. The storage array must present LUNs to your gateway before you can do the following: ◆ Assign array LUNs to your gateway in Data ONTAP configuration ◆ Use the setup program to identify your system to Data ONTAP ◆ Install licenses and Data ONTAP software on your system. Use the following information, in the order shown: 1. Use the instructions in this chapter to connect your system to the storage array, 2. See the gateway <i>Software Setup, Installation, and Management Guide</i> (for Data ONTAP 7.3.x) or the <i>Data ONTAP Setup Guide</i> (for Data ONTAP 8.0 or later) for instructions about how to set up a gateway with third-party storage.

Planning a port-to-port connectivity scheme

Resources to help you plan a connectivity scheme

When planning your configuration, you must devise a port-to-port connectivity scheme. The following resources will be particularly helpful when developing your port-to-port connectivity scheme.

For information about...	See...
Gateway FC initiator port usage guidelines	“Guidelines for Gateway FC Initiator Port Usage” on page 107
The configurations supported for your storage array and any special connectivity requirements for your storage array	<i>Gateway Implementation Guide</i> for your vendor <i>Gateway Interoperability Matrix</i>
Configurations supported for gateway MetroClusters	<i>Gateway MetroCluster Guide</i>

Requirement for redundant connections

- You must set up your connections correctly to ensure redundancy, as follows:
- ◆ Each connection in a redundant pair must be attached to a different FC initiator port on your gateway.
 - ◆ Use redundant ports on the Fibre Channel switches.
 - ◆ The two ports on the storage array that you select to access a given LUN should from different components that could represent a single point of failure (SPOF)—for example, from alternate controllers, clusters, or enclosures. (The terminology used for the hardware component on which host adapters and ports are located varies on different storage array models.)

Connection locations on a gateway

Connections on N7000 series gateways


The following illustration shows the locations of the onboard ports, PCI slots, and AC connections on all configurations of the N7000 series gateways.

- ◆ Stand-alone gateway—NVRAM6 adapter (slot 2)
- ◆ HA pair—NVRAM6 adapter (slot 1)
- ◆ Stretch MetroCluster — NVRAM6 adapter (slot 1)
- ◆ Fabric-attached MetroCluster—NVRAM6 adapter (slot 2) and VI-MC cluster adapter (slot 1)

Note

For more details about slots used, see the *Hardware and Service Guide* for your product.

The following illustration shows a N7000 series gateway.



Connections on N7x50T series gateways

The PCI slots on all configurations of the N7x50T series gateways are as follows:

- ◆ Stand-alone gateway—NVRAM6 adapter (slot 2)
- ◆ HA pair—NVRAM6 adapter (slot 1)
- ◆ Stretch MetroCluster—NVRAM6 adapter (slot 1)

- ◆ Fabric-attached MetroCluster—NVRAM6 adapter (slot 2) and VI-MC cluster adapter (slot 1)

Note

For more details about slots used, see the *Hardware and Service Guide* for your product.

Connections on N5000 series gateways


The following illustrations show the locations of the onboard ports, PCI slots, and AC connections on all configurations of the N7000 series gateways.

Note

For more details about slots used, see the *Hardware and Service Guide* for your product.


N5200 and N5500 series gateways using NVRAM5: The following illustration shows the locations of the onboard ports, PCI slots, and AC connections on all configurations of the N5200 and N5500 gateways.

- ◆ Stand-alone gateway—NVRAM5 adapter (slot 1)
- ◆ HA pair—NVRAM5 adapter (slot 3)
- ◆ Stretch MetroCluster configuration—NVRAM5 adapter (slot 3)
- ◆ Fabric-attached MetroCluster—NVRAM6 adapter (slot 1) and VI-MC cluster adapter (slot 2)



N5600 and N5300 gateways using NVRAM6: The following illustration shows the locations of the onboard ports, PCI slots, and AC connections on all configurations of the N5600 and N5300 gateways.

- ◆ Stand-alone gateway—NVRAM6 adapter (slot 1)
- ◆ HA pair—NVRAM6 adapter (slot 3)
- ◆ Stretch MetroCluster — NVRAM6 adapter (slot 3)
- ◆ Fabric-attached MetroCluster—NVRAM6 adapter (slot 1) and VI-MC cluster adapter (slot 2)




Connections on N6000 series gateways

N6000 series gateways are available as single controller systems and dual controller systems.

Unlike other gateways, N6000 series gateways do not have a PCI-X or PCIe card to provide NVRAM. Instead, the NVRAM—known as NVRAM7—is integrated into the motherboard of each controller.

In addition to preserving unwritten data, NVRAM7 provides InfiniBand (IB) paths in dual-controller systems for traffic between the two controllers in an HA pair. The IB connection is made over the system backplane, completely inside the chassis; there are no external IB connections.


The following illustration shows the location of the onboard FC ports and other connections on each controller module in a N6000 series gateway.



Connections on N6200 series gateways

N6200 series gateways are available as single controller systems and dual controller systems. All controllers in this family can be configured as HA pairs.

The following illustration shows a single-chassis HA pair. The locations of the onboard ports and expansion slots on the systems are visible.



Initiator ports can be configured using the on-board Fibre Channel ports 0c and 0d.

Attention

The on-board Fibre Channel ports (0c, 0d) are not on independent busses and do not provide redundant port pairs. Some port failures could cause the system to panic. To configure redundant port pairs, use a Fibre Channel HBA in an available PCI expansion slot.

Use only two initiator ports on a single HBA. The ports should be on separate busses. For example, use ports 1a and 1c.

In a MetroCluster configuration, you must use an add-on HBA. Be sure to use redundant port pairs for the local and remote storage. For example, use ports 2a and 2c for the local storage and 2b and 2d for the remote storage.

N6210 configuration restrictions: The following restrictions apply to N6210 configurations.

- ◆ N6210 configurations only support connection to a single third-party storage array.
- ◆ N6210 configurations only support a single LUN group.

Connecting a stand-alone gateway

Connecting a stand-alone system

To connect a stand-alone gateway to a storage array, either directly or through a switch, complete the following steps.

1. Identify onboard ports and expansion adapter ports for your gateway, using the information in “[Connection locations on a gateway](#)” on page 119 for reference.

Note

“[How FC initiator ports are labeled](#)” on page 108 describes how FC initiator ports are identified on different gateways.

2. Locate the ports on the storage array that you want to use to connect to the gateway, either directly or through a switch.

3. For a direct-attached stand-alone configuration

Connect the gateway to the storage array, using redundant FC initiator ports, as described in “[Location of ports, by gateway](#)” on page 108.

- a. Connect one cable from one FC initiator port on the gateway to cluster (controller) 1 port 1 on the storage array.
- b. Connect a cable from a redundant FC initiator port on the gateway 1 to cluster (controller) 2 port 1 on the storage array.

Attention

If your storage array supports fewer array LUNs per host group per port than the number of LUNs that the gateways will be using, you need to add additional cables between the gateway and the storage array. See the gateway *Implementation Guide* for your vendor for information about port-to-port connectivity requirements for your storage array type.

4. For a fabric-attached stand-alone configuration

Using redundant FC initiator ports on the gateway, as described in “[How FC initiator ports are labeled](#)” on page 108, connect the gateway to the switches as follows:

- a. Connect one cable from the one FC initiator port on the gateway to Switch 1.

- b. Connect another cable from the redundant port on the gateway to Switch 2.

Connect the switches to the storage array as follows:

- a. Connect Switch 1 to the storage array cluster (controller) 1, port 1.
- b. Connect Switch 2 to the storage array cluster (controller) 2, port 1.

Attention

The gateway supports only two paths to an array LUN. Ensure that the two ports on the storage array that you select to access a given LUN are from different components that could represent a single point of failure (SPOF), for example, from alternate controllers, clusters, or enclosures. You want to ensure that you do not lose all access to a LUN if the one component fails.

- 5. (Optional) Connect the gateway to a tape backup device through a separate FC initiator port or SCSI tape adapter.
- 6. Connect a console cable to the console port on the gateway. Use the RJ-45 to DB-9 adapter that is included with your system. Connect the console cable to the adapter.

If you are connecting to an ASCII terminal, see “[Settings for Connecting to an ASCII Terminal Console](#)” on page 134 for information about settings.

- 7. Install the cable management tray by pinching the arms of the tray and fitting the holes in the arms through the motherboard tray pins. Then push the cables into the cable holders, thread the adapter cables through the top rows of the cable holders, and thread the port cables through the lower cable holders.
- 8. Connect the gateway to an Ethernet network by plugging the network cable into the networking port. If you are connecting more than one network cable to the network, connect to the ports sequentially. Use the cable management tray to direct all the cabling from your system.
- 9. (Optional) Connect Remote LAN Module (RLM) from the back of the gateway to the network using an Ethernet cable.

Attention

The network switch port for the RLM connection must negotiate down to 10/100 or autonegotiate.

- 10. Verify that the storage array is configured and connected properly, and that it is powered on.

Attention

Your configured and connected storage array must be powered on before you power on your gateway. See your storage array documentation for how to power on the storage array. The gateway expects these units to be ready for input/output when it powers on and performs its reset and self-test.

11. If your deployment includes switches, make sure that all switch IDs are set, then turn on each switch 10 minutes apart from one another.
12. If applicable, turn on any tape backup devices.
13. For each power supply on the gateway
 - a. Ensure that the power switch is in the Off (O) position.
 - b. Connect the socket end of the power cord to the power plug on the power supply.
 - c. Secure the power cord with the retaining adjustable clip on the power supply.
 - d. Plug the other end of the power cord into a grounded electrical outlet.

Attention

To obtain power supply redundancy, you must connect the second power supply to a separate AC circuit.

14. Start a communications program.

You must use some form of communications program to be able to perform initial network setup and gateway configuration. You can start a communications program through RLM or through the console after connecting to the serial port.
15. Turn the power switch on the gateway to the On (I) position.

Result: The system verifies the hardware and loads the operating system.
16. If the storage array does not automatically discover gateway WWNs after you connect the gateway to the storage array, see Appendix C, “[Obtaining WWNs Manually](#),” on page 153.
17. Continue with the appropriate setup of your gateway and Data ONTAP, as described in “[When to connect your gateway to third-party storage](#)” on page 116.

Connecting a gateway HA pair

About HA pairs

Nodes in an HA pair have the following characteristics:

- ◆ They are connected to each other through a cluster interconnect consisting of adapters and cable through which they do the following tasks:
 - ❖ Continually check whether the other node is functioning
 - ❖ Mirror log data for each other's NVRAM
 - ❖ Synchronize each other's time
- ◆ The cluster adapter supports dual cluster interconnect cables.
 - ❖ If one cable fails, the heartbeat and NVRAM data are automatically sent over the second cable with no delay or interruption.
 - ❖ If both cables fail, the failover capability is disabled but both nodes continue to serve data to their respective applications and users.

Connecting an HA pair

To connect an HA pair to a storage array, either directly or through a switch, complete the following steps.

1. Identify onboard ports and expansion adapter ports for your gateway, using the information in [“Connection locations on a gateway”](#) on page 119 for reference.

Note

[“How FC initiator ports are labeled”](#) on page 108 describes how FC initiator ports are identified on different gateways.

2. See the *Hardware and Service Guide* for your product to ensure that your cluster interconnect adapter is in the correct slot for your system in an HA pair.

Note

For systems that use NVRAM5 or NVRAM6, the NVRAM card functions as the cluster interconnect adapter.

3. Plug one end of the optical cable into one of the local node's cluster adapter ports, then plug the other end into the partner node's corresponding cluster adapter port.

Attention

You must not cross-cable the cluster interconnect adapter. Cable the local node ports only to the identical ports on the partner node.

4. Repeat [Step 3](#) for the two remaining ports on the cluster adapters.
5. Locate the ports on the storage array that you want to use to connect to the gateway to the storage array, either directly or through a switch.
6. For a direct-attached HA pair:

Connect the gateway to the storage array, using redundant FC initiator ports, as described in “[How FC initiator ports are labeled](#)” on page 108.

Connecting Gateway 1:

- a. Connect one cable from one FC initiator port on gateway 1 to cluster (controller) 1, port 1, on the storage array.
- b. Connect a cable from a redundant FC initiator port on gateway 1 to cluster (controller) 2, port 1, on the storage array.

Connecting Gateway 2:

- a. Connect one cable from one FC initiator port on gateway 2 to cluster (controller) 1, port 1, on the storage array.
- b. Connect a cable from a redundant FC initiator port on gateway 2 to cluster (controller) 2, port 1, on the storage array.

Attention

If your storage array supports fewer LUNs per host group per port than the number of LUNs that the gateways will be using, you need to add additional cables between the gateway and the storage array, as required. See the gateway Implementation Guide for your vendor for information about port-to-port connectivity requirements for your storage array type.

7. For a fabric-attached HA pair:

Connect the gateway to the storage array, using redundant FC initiator ports, as described in “[How FC initiator ports are labeled](#)” on page 108.

Connecting Gateway 1 to the switches:

- a. Connect one cable from one FC initiator port on gateway 1 to Switch 1.
- b. Connect another cable from a redundant FC initiator port on gateway 1 to Switch 2.

Connecting Gateway 2 to the switches:

- a. Connect one cable from one FC initiator port on gateway 2 to Switch 1.
- b. Connect another cable from a redundant FC initiator port on gateway 2 to Switch 2.

Connecting the switches to the storage array:

- a. Connect Switch 1 to the storage array cluster (controller) 1, port 1.
- b. Connect Switch 2 to the storage array cluster (controller) 2, port 1.

Attention

Between the switch and the storage array, the gateway only two paths to an array LUN. Ensure that the two ports on the storage array that you select to access a given LUN are from different components that could represent a single point of failure (SPOF), for example, from alternate controllers, clusters, or enclosures. You want to ensure that you do not lose all access to a LUN if the one component fails.

8. (Optional) Connect the gateway to a tape backup device through a separate FC initiator port or SCSI tape adapter.
9. Connect a console cable to the console port on the gateway. Use the RJ-45 to DB-9 adapter that is included with your system. Connect the console cable to the adapter.

If you are connecting to an ASCII terminal, see “[Settings for Connecting to an ASCII Terminal Console](#)” on page 134 for information about settings.

10. Install the cable management tray by pinching the arms of the tray and fitting the holes in the arms through the motherboard tray pins. Then push the cables into the cable holders, thread the adapter cables through the top rows of the cable holders, and thread the port cables through the lower cable holders.
11. Connect the gateway to an Ethernet network by plugging the network cable into the networking port. If you are connecting more than one network cable to the network, connect to the ports sequentially. Use the cable management tray to direct all the cabling from your system.
12. (Optional) Connect Remote LAN Module (RLM) from the back of the gateway to the network using an Ethernet cable.

Attention

The network switch port for the RLM connection must negotiate down to 10/100 or autonegotiate.

13. Verify that the storage array is configured and connected properly, and that it is powered on.

Attention

Your configured and connected storage array must be powered on before you power on your gateway. See your storage array documentation for how to power on the storage array. The gateway expects these units to be ready for input/output when it powers on and performs its reset and self-test.

14. If your deployment includes switches, make sure that all switch IDs are set, then turn on each switch 10 minutes apart from one another.
15. If applicable, turn on any tape backup devices.
16. For each power supply on the gateway
 - a. Ensure that the power switch is in the Off (O) position
 - b. Connect the socket end of the power cord to the power plug on the power supply.
 - c. Secure the power cord with the retaining adjustable clip on the power supply.
 - d. Plug the other end of the power cord into a grounded electrical outlet.

Attention

To obtain power supply redundancy, you must connect the second power supply to a separate AC circuit.

17. Start a communications program.

You must use some form of communications program to be able to perform initial network setup and gateway configuration. You can start a communications program through RLM or through the console after connecting to the serial port.

18. Turn the power switch on the gateway to the On (I) position.

Result: The system verifies the hardware and loads the operating system.

19. If the storage array does not automatically discover gateway WWNs after you connect the gateway to the storage array, see Appendix C, “[Obtaining WWNs Manually](#),” on page 153.
20. Continue with the appropriate setup of your gateway and Data ONTAP, as described in “[When to connect your gateway to third-party storage](#)” on page 116.

Connecting a MetroCluster

For information about how to connect a MetroCluster to third-party storage, see the gateway *MetroCluster Guide*. MetroCluster configuration with gateways using native-disk shelves are not supported.

Settings for Connecting to an ASCII Terminal Console

About the ASCII terminal console

The ASCII terminal console enables you to monitor the boot process and helps you configure your gateway after it boots. You can attach an ASCII terminal console through the serial port on the back of your gateway if you want to do local system administration.

ASCII terminal console cable wiring

The ASCII terminal console is connected to your gateway with a DB-9 serial adapter, attached to an RJ-45 converter cable. The DB-9 adapter connects into the DB-9 serial port on the back of your gateway.

The following table lists how the DB-9 serial cable is wired. Input indicates data flow from the ASCII terminal to your gateway and output indicates data flow from your gateway to the ASCII terminal.

Pin number	Signal	Data flow direction	Description
1	DCD	Input	Data carrier detect
2	SIN	Input	Serial input
3	SOUT	Output	Serial output
4	DTR	Output	Data terminal ready
5	GND	N/A	Signal ground
6	DSR	Input	Data set ready
7	RTS	Output	Request to send
8	CTS	Input	Clear to send
9	RI	Input	Ring indicator

Communications parameters

The following table shows the communications parameters for connecting an ASCII terminal console to a gateway. You need to set the following communications parameters to the same values for both your gateway and the ASCII terminal.

Parameter	Setting
Baud	9600
Data bit	8
Parity	None
Stop bits	1
Flow control	None

Note

See your terminal documentation for information about changing the ASCII console terminal settings.

Limiting the Target Queue Depth

What the target queue depth is

The target queue depth limits the number of commands (read and write requests) that the storage array port must handle from gateways and non gateway hosts. When multiple initiators are accessing a target port on the storage array, you do not want the outstanding commands in the queue buffer, from all initiators together, to exceed what the storage array can handle. Otherwise, the performance of gateways and non gateway hosts can suffer. Storage arrays differ in the number of commands that they can handle in the queue buffer.

Non gateway hosts also provide a means for limiting the target queue length. See the documentation for the host for information about limiting target queue length.

Note

Target queue depth might also be referred to as target queue length, Q-Depth or Max Throttle.

When you need to configure the target queue depth

The default value for the Data ONTAP target queue depth setting is 256. If this value is not appropriate for your organization, you need to change the value. On gateways you use the `disk.target_port.cmd_queue_depth` option to set the number of commands that the gateway can send to the target port on the storage array.

Guidelines for determining the appropriate target queue depth

You need to consider the impact of all the initiators on the storage array port when you are planning the configuration for a specific gateway or a specific non gateway host. If your deployment includes more than one initiator, you need to configure each initiator so that the total number of commands by all initiators does not exceed the maximum that the storage array port can handle.

Some guidelines for determining the appropriate target queue length are as follows:

- ◆ 0 (zero) means no limit on the outstanding commands. Therefore, do not configure a value of 0.
- ◆ You could divide 256 by the number of gateways and non gateway hosts that will be acting as initiators on the target port on the storage array, configuring each gateway and non gateway host with the value that resulted.
- ◆ You could consider the volume of commands that specific initiators would be likely to send to the target port. You could then configure higher values

for initiators likely to send a greater number of requests and a lower value for initiators likely to send a lesser number of requests.

- ◆ Configure non gateway hosts according to the guidelines provided for those hosts.

Setting the target queue depth

To set the `disk.target_port.cmd_queue_depth` option, complete the following step.

Step	Action
1	Enter the following at the Data ONTAP command line: <pre>options disk.target_port.cmd_queue_depth value</pre> Do not use a value of 0 (zero) because it means no limit on the outstanding commands.

About this appendix The instructions in this appendix provide quick start instructions for setting up gateway to work with a storage array when the gateway is using third-party storage only. These instructions are for installing a simple two-port configuration with any gateway, any switch, and any storage array, regardless of vendor.

Where to find detailed instructions

If you are using only third-party storage with your system and you need detailed setup and installation instructions, see the following guides instead of this appendix:

- ◆ For Data ONTAP 7.3.x releases: *Gateway Setup Installation, and Management Guide*
- ◆ For Data ONTAP 8.0 releases: Data ONTAP setup guides

If your system is using both third-party storage and native disk shelves, see the following guides instead of this appendix:

For Data ONTAP 7.3.x releases:

- ◆ *Gateway Setup Installation, and Management Guide*
- ◆ *Gateway Implementation Guide for Native Disk Shelves*

For Data ONTAP 8.0 releases:


- ◆ *Data ONTAP Setup Guide*
- ◆ *Gateway Implementation Guide for Native Disk Shelves*
- ◆ *Data ONTAP Storage Management Guide*
- ◆ *Clustered Storage Management and Data Protection Guide*

Note

The setup and installation procedures differ depending on whether you order your system with native disk shelves or not. If you order your system with disk shelves, the factory configures the root volume and installs licenses and Data ONTAP software for you.

Example configuration for this quick start

You can use the following illustration of a basic configuration as a reference while you are using the quick start.



This example includes a Hitachi USP storage array and an N5000 series gateway HA pair. However the quick start applies to setting up your configuration with any storage array and any gateway.

USP-GP1 in this example represents a group of array LUNs allocated for the gateway. CLS1 and CLS2 represent clusters on the USP storage array. These are the hardware components on which the ports are located. Different vendors and different storage array models use different terminology to represent these hardware components (for example, *cluster* or *controller* for Hitachi, HP XP, and IBM; *Storage processor* or *SP* for EMC CLARiON; and *Controller Module* for Fujitsu ETERNUS).

To ensure availability, use a redundant port pair on each node in the HA pair (that is, a port from each bus). Then, if one path from a gateway node fails, the other path from the node is used; gateway controller takeover does not occur.

Installation quick start

Step 1: Pre-installation tasks

Before beginning installation, complete the following tasks:

1. See the gateway *Interoperability Matrix* at <http://www.ibm.com/storage/nas/> to ensure compliance with supported storage array models, firmware, switches, and Data ONTAP versions.
2. Ask the storage administrator to create at least four LUNs on the storage array for the gateway.

Two of the array LUNs will be used as root volumes, one on each gateway in the HA pair. The other two array LUNs will be used as the spare LUNs for core dumps, one for each gateway. For information about LUN sizes, see the following:

- ❖ “[Guidelines for determining the number and size of array LUNs](#)” on page 65
 - ❖ gateway *Interoperability Matrix* for information about the minimum size spare core LUN needed for each gateway
3. Ask the storage array administrator to configure any parameters on the storage array that are required to work with the gateway.

See Appendix D, “[Configuration Requirements by Vendor](#),” on page 155 to determine whether the gateway requires any specific settings to work with your storage array (for example, a specific host type or specific parameter settings). Also consult the gateway *Implementation Guide* for your vendor.
 4. See the gateway *Interoperability Matrix* to check the minimum root volume size for your gateway.
 5. Obtain appropriate Data ONTAP software version zip file.

Step 2: Install the gateway

To install each gateway, complete the following tasks:

1. Power on the gateway and interrupt the boot process by pressing Ctrl-C when you see the following message on the console:

```
Starting Press CTRL-C for special boot menu
```

2. Select “Maintenance mode boot” on the menu.

Do not proceed any further with gateway installation and setup at this time.

3. Check the settings of the gateway HBAs to ensure that they are configured as initiators.

Note

The factory might configure some gateway HBAs as target ports. On N5200, N5300, N5500, N5600 gateways, for example, typical gateway installations use 0a and 0c as a port pair or 0b and 0d as a port pair. If your installation uses port pairs in this manner, you must change the ports so that they are initiators. Most gateways do not have targets.

- a. Enter the following command to see which ports are configured as target ports:

```
fcadmin config
```

For each target port that you want to change to an initiator port, enter the following command:

```
fcadmin config -t initiator port#
```

4. Install the Fibre Channel cables to connect the gateway to the switches and to connect the switches to the storage array.

For cabling details, see Chapter 10, “[Connecting Your System to a Storage Array](#),” on page 115.

Step 3: Set up the switches

Switch configuration is typically done by the storage or SAN administrator.

Setting up zoning: The switches need to be zoned so that the gateways and the storage arrays can see each other. Use single-initiator zoning so that the gateway FC initiator ports do not see each other. Some customers use hard zoning and other customers use soft zoning.

If you are configuring hard zoning, complete the following step:

1. Create the zones using the switch zoning commands.

If you are configuring soft zoning, complete the following steps:

1. Obtain the gateway Worldwide Port Names (WWPNs).

To eliminate the possibility of errors when typing the WWPNs into the switch configuration, it is recommended that you let the SAN switches automatically discover the WWPNs. The following must occur before automatic discovery is successful:

- ❖ The devices in the configuration must be cabled together.
- ❖ The storage array must be powered on.

- ❖ The gateways must be in Maintenance Mode.

If you need to obtain WWPNs manually, see Appendix C, “[Obtaining WWPns Manually](#),” on page 153.

2. Log on to the storage array and obtain the WWPns of the FC adapters of the storage array.
3. Use the Fibre Channel switch commands to zone each switch so that the storage array and the gateways see each other’s WWPns.

Example: In the configuration example at the beginning of this chapter, the zones are as follows if you are using soft zoning.

Switch	Zones
SW1	VS1:0a, CLS1:1A VS2:0a, CLS1:1A
SW2	VS1:0c, CLS2:2A VS2:0c, CLS2:2A

If your configuration includes McDATA switches: If your configuration includes McDATA switches, be sure that Open Systems Management Software (OSMS) is enabled on *both* switches. If you do not configure OSMS, the gateway bypasses the name services. The result is that the gateway cannot see array LUNs and, therefore, cannot boot. See Chapter 4, “[Guidelines for a Fabric-Attached Configuration](#),” on page 49 for more details about OSMS.

Step 4: Set up LUN security

Have the storage array administrator do the following:

1. Set up LUN security on the storage array.

The concept of LUN security is similar to zoning except that this step is performed on the storage array. LUN security keeps different servers from using each others storage on the SAN.

Note

LUN security might also be referred to as *LUN masking*.

2. Create host groups for the gateway, or the equivalent.

The term *host group* is used on some storage arrays to describe a configuration parameter that enables you to specify host access to specific ports on the storage array (see Appendix F, “[Terminology Comparison Between Vendors](#),” on page 175). Different storage arrays use different

terms to describe this configuration parameter. Each storage array vendor has its own process for creating a host group or the equivalent.

Step 5: Assign the array LUN to the gateway and create the root volume

For each gateway, you assign an array LUN to it and create the root volume for it so that you can install Data ONTAP. At this point it is easiest if you assign only one array LUN to each gateway. You can add additional array LUNs after you have installed Data ONTAP on each of the gateways and you have tested paths and cluster failover.

Best practice: For Data ONTAP 8.0, the root volume must be a FlexVol volume. For Data ONTAP 7.x releases, it is recommended that the root volume be a FlexVol volume. It is recommended that you assign only one array LUN to the aggregate with the root volume. The FlexVol volume can then use all the space for the root volume.

To assign the LUN and create the root volume, complete the following steps.

1. Return to the gateway console. It should still be in Maintenance Mode.
2. Use the `disk show -v` command to confirm that you can see the array LUNs created by the storage administrator.

Attention

You must see array LUNs at this point to be able to continue. If you do not see the array LUNs, reboot the gateway into Maintenance Mode. If you still do not see the array LUNs, double-check that the array LUNs exist, that the host groups were created correctly, and that zoning is correct.

3. Use the `disk assign` command to assign the first array LUN to the gateway. For example, on gateway 1, enter `disk assign name.L1`. On gateway 2, enter `disk assign name.L2`.

Best practice recommendation: Use the block checksum type, which is the default.

Note

Ignore any messages that say “unable to obtain owner name” and tell you to rerun the command.

4. Enter `disk show -v` and confirm that the system ID of the gateway is shown as the owner of the array LUN (which indicates that the array LUN was assigned to the gateway).
5. Enter `halt` to exit Maintenance Mode.
6. At the boot environment prompt, enter `bye`.

7. Press Ctrl-C to interrupt the boot process.
8. Create a root volume by choosing option 4a for a FlexVol volume.
9. Enter `y` to install a new file system.
10. Enter `y` at the message “This will erase all data on the disks, are you sure?”
Setup begins. Complete the setup by entering a hostname, IP address, and other information that is requested.
11. Continue to “[Step 6: Install Data ONTAP software and required licenses](#)” on page 145.

Step 6: Install Data ONTAP software and required licenses

Perform Data ONTAP software installation using either of the following methods:

- ◆ Map the gateway C\$ share to your laptop as a CIFS share.
- ◆ If an HTTP server is available, use the `software install http://ipaddr/file.zip` command to install the software.

The following steps describe the CIFS share method:

1. Install the CIFS and gateway licenses. You can install other licenses now or later.
2. Run CIFS setup in workgroup mode.
3. Map the gateway C\$ to your laptop.
4. Make an `/etc/software` directory (or enter `software list` to create the `/etc/software` directory).
5. Copy the Data ONTAP executable to the `/etc/software` directory.
6. Enter `software install` to run the software executable `<release_setup_e.exe>`.
7. Download the Data ONTAP software.
8. Reboot the gateway.
9. If your gateway is a node in an HA pair, repeat the setup and Data ONTAP software installation steps on the partner node before starting the testing described in “[Step 7: Test your setup](#)” on page 145.

Step 7: Test your setup

Before putting a gateway into a production environment, it is imperative that you ensure that everything works.

Checking your system: To check your system to ensure that the configuration is as expected, the following commands are available.

Command	This command enables you to do this...
<code>disk show -v</code>	Check whether all array LUNs are visible.
<code>sysconfig -v</code>	Check which gateway FC initiator ports, switch ports, and array LUNs are used.
<code>storage array show-config</code>	(In Data ONTAP 8.0 or later) Check which FC initiator ports, switch ports, and array LUNs are used.
<code>sysconfig -r</code>	Look at the aggregate configuration and ensure that spare array LUNs are available.

Verifying that there are two paths to storage: You want to ensure that there are two paths to each array LUN so that the gateway can continue to work when running on a single path, for example, if a switch or an array port is taken offline for upgrades.

1. To verify that there are two paths to each array LUN, enter the following command:

❖ For systems running Data ONTAP 7.x, enter:

```
storage show disk -p -all
```

Note

You can see the adapters only when you use the `all` variable. When you use the `all` variable, you cannot see unassigned LUNs.

❖ For systems running Data ONTAP 8.0 or later, enter:

```
storage array show-config
```

2. Check whether two paths to the array LUNs are shown.

If you do not see two paths to the array LUNs, check zoning, “host group” configuration, and cabling. (See Chapter 3, “[Guidelines for Setting Up Access to Array LUNs](#),” on page 17 for example output for each command that you can use to confirm that there are two paths to array LUNs.)

3. Look at the adapters shown to see whether all paths are on a single adapter.

If you see both paths through only one gateway FC initiator port (the gateway’s 0c port, for example) this is an indication that the back-end zoning is redundantly crossed. This is not a supported configuration.

Attention

Do not continue with testing until you see two paths.

Testing normal cluster failover:

1. On the partner node (vs2 in the example in this chapter), enter the following commands to perform a cf takeover: A takeover of vs1 should occur.

```
vs2 > cf status
vs2> cf takeover (vs1 goes down; vs2 takes over)
vs1/vs2> df (vs2 looks at vs1)
vs1/vs2> partner (to switch to the partner)
vs2> sysconfig
vs2 (takeover)> cf status
vs2 (takeover)> cf giveback (to return to normal operation)
```

2. Repeat the same commands on the local gateway (vs1 in the example configuration).

Testing path failover and cluster failover: When you use this procedure, you are testing path failure when the first port is taken offline; array LUN access is from a single gateway FC initiator port at that time.

1. On the local gateway (vs1 in the example configuration), enter the following commands:

```
fcadmin offline 0a
storage show disk -p (You should see only one path) or storage array
show-config
fcadmin offline 0c (Cluster takeover should occur)
```

2. On the partner node (vs2 in the example configuration), enter the following commands:

```
cf takeover
cf giveback
```

3. After both the local node (vs1) and the partner node (vs2) are back online, go to the partner node (vs2 in the example) and repeat the procedure.

Step 8: Additional setup

After you complete Data ONTAP installation, you might want to do the following:

1. Assign additional array LUNs to the gateways as required. (Use the `disk assign` command).

After the basic gateway setup is complete, you can have the storage administrator create additional array LUNs for the gateway, as needed. (See “[Minimum array LUN size for the root volume](#)” on page 68.)

Note

Before assigning the new array LUNs to the gateway, use the `disk show -v` command to verify that the new array LUNs are visible to the gateway.

2. Create Data ONTAP aggregates and volumes as desired.
3. Set up additional Data ONTAP features on your gateway, for example, features for backup and recovery.

See the Data ONTAP manuals for more information.

About this appendix This appendix provides instructions for how to use netboot with gateways to complete a basic set up a gateway that is using third-party storage only. It includes instructions for creating the root volume, assigning array LUNs to the gateway, and obtaining and installing Data ONTAP software.

When to use this procedure You can use netboot as an alternative way to boot a gateway from a Data ONTAP software image that is stored on a TFTP server. You might need to use netboot in the following circumstances:

- ◆ You need to boot an alternative kernel because the new CompactFlash (CF) card does not have a boot image, or the CF card is damaged.
- ◆ You need to boot an alternative kernel because the boot image does not support a particular hardware configuration that is to be installed.

When not to use this procedure Do not use this procedure in the following circumstances:

- ◆ If you ordered your gateway with disk shelves
The procedure in this appendix does not cover basic setup of a gateway ordered with disk shelves. When a gateway is ordered with disk shelves, the factory installs the root volume, licenses, and Data ONTAP software for you. (This procedure covers those steps.)
- ◆ If your system is running Data ONTAP 8.0
In Data ONTAP 8.0, netboot is not a supported functionality except for CF card replacement. In Data ONTAP 8.0, the boot menu includes an option called “Install new software first.” If the Data ONTAP 8.0 software on the boot device does not include support for the storage array you want to use, this option can be used to obtain a version of the software that supports your storage array and install it on your system. Contact Technical Support for assistance using this option or if you need to boot your system from a Data ONTAP version stored on a remote server.
- ◆ If you want to use an HTTP server instead of a TFTP server, or if you want to use DHCP to obtain the gateway IP address.
In these circumstances, see the netboot procedures in the *Data ONTAP System Administration Guide*. The netboot procedure is exactly the same for gateway and storage systems that run Data ONTAP software.

Preparing for using netboot

In this netboot procedure, you will use a Trivial File Transfer Protocol (TFTP) to move the netboot image from your laptop to the gateway and you will obtain a static IP address for the gateway (rather than obtaining the IP address through DHCP).

When you use netboot, you can select the Data ONTAP version that you want. You can access maintenance mode or perform a normal boot, just as you can when you boot from a CF card.

To use the netboot option to start your gateway, you must have the following:

- ◆ A TFTP server available on your network.
- ◆ A Data ONTAP boot image on a server or your TFTP server
- ◆ Networking configured with the gateway IP address.

Note

Network interfaces configured in the firmware environment are not persistent across system halts and reboots. You must configure the network interface each time you use the netboot option.

Setting up the gateway

To enable netboot, you must configure networking with the gateway static IP address and place the boot image on a configured TFTP server. (Your laptop computer can be your TFTP server.) When netbooting, you will see a TFTP attach session on the TFTP server during the TFTP connection.

Complete the following steps.

1. Download the netboot image and the system files for the Data ONTAP release you need.

You can copy the boot image from the system boot directory, `/etc/boot/ibm-mips`, or download it from the IBM web site at <http://www.ibm.com/storage/nas>.

The release number is included in the netboot file name and system file name. For example, for the Data ONTAP 7.0.5 release, the netboot file name is “705_netboot.e” and the system file name is “705_setup_e.exe.”

2. Download to your TFTP server the latest version of the TFTP zip file (`tftpd32.###.zip`) from <http://tftpd32.jounin.net/>.
3. Unzip the TFTP file, for example, to `/tmp/tftp`.
4. Launch the `tftpd32.exe` application.
5. Copy the netboot image to the directory where you unzipped TFTP (for example, to `/tmp/tftp/`).

6. Set up the hardware for netboot (this step assumes you are using a laptop).
 - a. Connect a laptop to use for console access and Ethernet access to the gateway.
 - b. Connect the laptop's GigE port to the gateway e0a port.
 - c. Assign an IP address to the laptop, for example 192.168.1.x.
 - d. Power on the gateway and interrupt the boot process so that you stop at the boot environment prompt.
 - e. Run the following commands on the gateway to manually assign the gateway IP address and start netboot.

```
CFE> ifconfig e0a -addr=gateway_IP_address
```

```
CFE> ping TFTP_server_IP_address
```

```
CFE> netboot TFTP_server_IP_address:netboot_executable
```

Example: CFE> netboot 192.168.1.x:705_netboot.e

During the netboot session, you can interrupt the boot process and stop at the boot options (1-5) menu or continue with normal boot.

7. Assign the array LUN for the root volume to this gateway.
 - a. Boot into maintenance mode.
 - b. Use the `disk assign` command to assign the array LUN for the root volume to this gateway.
 - c. Enter the following commands to confirm that array LUN was assigned properly.

```
disk show -v  
halt  
bye
```

8. Install the version of the Data ONTAP software that you need.

Note

This step describes how to install the software using the CIFS share method.

- a. Repeat Step 6d to manually assign the gateway IP address and start netboot again.
- b. Interrupt the boot process by pressing Ctrl-C when you see the message “Press Ctrl-C for special boot menu” on the console, then select 4a on the boot options menu to install a new file system.

- c.** Install the CIFS and gateway licenses. You can install other licenses now or later.
- d.** Run CIFS setup in workgroup mode.
- e.** Map the gateway C\$ to your laptop.
- f.** Make an /etc/software directory (or enter `software list` to create the /etc/software directory).
- g.** Copy the Data ONTAP executable to the /etc/software directory.
- h.** Enter `software install` to run the software executable `<release_setup_e.exe>`.
- i.** Download the Data ONTAP software.
- j.** Reboot the gateway.
- k.** If your gateway is a node in an HA pair, follow the same setup and Data ONTAP installation steps on the partner node

When to use this procedure

If the gateway is not connected to the SAN switch, use this procedure to obtain the World Wide Port Names (WWPNs) of the gateway FC initiator ports that will be used to connect the gateway to the switch.

Having the switch automatically discover WWPNs is the preferred method of obtaining WWPNs because you can avoid potential errors resulting from typing the WWPNs into the switch configuration.

Obtaining the WWPNs from your gateway

To manually obtain the gateway WWPNs, complete the following steps:

1. Connect the gateway console connection to a laptop computer.
2. Power on your gateway.

Interrupt the boot process by pressing Ctrl-C when you see the following message on the console:

Starting Press CTRL-C for floppy boot menu

3. Select the Maintenance mode option on the boot options menu.
4. Enter the following command to list the WWPNs of the gateway FC initiator ports:

```
storage show adapter
```

To list a specific adapter WWPN, add the adapter name, for example,

```
storage show adapter 0a
```

5. Record the WWPNs that will be used and leave the gateway in maintenance mode.

About this appendix This appendix summarizes the configuration requirements for each vendor for setting up storage array models to work with gateways.

Topics in this appendix

This appendix contains the following topics:

- ◆ [“Gateway requirements for EMC CLARiiON storage arrays”](#) on page 156
- ◆ [“Gateway requirements for EMC Symmetrix storage arrays”](#) on page 158
- ◆ [“Gateway requirements for Fujitsu ETERNUS storage arrays”](#) on page 162
- ◆ [“Gateway requirements for Hitachi storage arrays”](#) on page 166
- ◆ [“Gateway requirements for HP storage arrays”](#) on page 168
- ◆ [“Gateway requirements for IBM storage arrays”](#) on page 169
- ◆ [“Gateway requirements for 3PAR storage arrays”](#) on page 172

Restriction for storage array replication and copying products

Do not use storage array replications and copying products with LUNs assigned to gateways. If you use the storage array copy and replication products, the source data is not altered. However, because this operation is being done outside of Data ONTAP, the data on the target is not synchronized in Data ONTAP. Therefore, the gateway cannot use the data on the target. Additionally, there might be a performance impact on the source.

Gateway requirements for EMC CLARiiON storage arrays



See the gateway *Implementation Guide for EMC CLARiiON Storage* for additional details regarding configurations with CX storage arrays.

CX configuration requirements to work with a gateway

When you manually register the gateway FC initiator port names as hosts, you set the parameters as shown in the following table.

Setting	Required value
Initiator Type	CLARiiON Open
Array Com Path	Enabled
Failover mode	1
Unit serial number	LUN
Host Name	If this section is available in the dialog box, specify New Host . To simplify management, it is recommended that you enter a host name and port number under in the host name fields, for example, gateway N7600 2866-G10-2b-0a .
IP address	Enter a unique fake IP address. Be sure that you have not entered this IP address anywhere else in the storage array configuration and that it is not an IP address that is present on the network.

Note

The Access Logix feature installed on a shared CX storage array enables you to set up data access and create Storage Groups. You need to manually register the gateway FC initiator ports as hosts in the array before you assign them to the gateway Storage Group.

See the gateway *Interoperability Matrix* at www.ibm.com/storage/support/nseries/ for any updates to information about requirements for CX storage arrays.

Rules for use of Storage Groups with the gateway

If you are using Storage Groups on your CX storage array, you must adhere to the following rules:

- ◆ For Data ONTAP releases 7.3 and later, the gateway supports more than one Storage Group per CX storage array.
- ◆ For releases earlier than Data ONTAP 7.3, the gateway supports only one Storage Group (per CX storage array).
- ◆ All the gateways in the same neighborhood must be in the same Storage Group (because all the gateways in the neighborhood must be able to see the same array LUNs).
- ◆ The maximum number of LUNs per Storage Group is 256.

Gateway requirements for EMC Symmetrix storage arrays



See the gateway *Implementation Guide for EMC Symmetrix Storage* for additional details regarding configurations with DMX storage arrays.

You must set both parameters related to SCSI-3

You must set two different parameters related to SCSI-3—one at the port level and one at the logical device (LUN) level. The two SCSI-3 flags that you need to set are as follows:

- ◆ SC3 Flag or SCSI-3 Flag. You must set this flag on each host (channel) director port to which the gateway connects.
- ◆ SCSI3_persist_reserv or PER Flag. You must set this flag on each logical device that the gateway will use.

Attention

Both SCSI-3 parameters are required for the gateway to work.

The SCSI-3 flag enables the connecting port to understand the SCSI-3 “dialect” and semantics. In addition to the SCSI-3 flag being enabled, each LUN must have the ability to remember and enforce persistent reservations requested in the SCSI-3 “dialect”—the SCSI3_persist_reserv flag setting for each logical device enables each LUN to enforce the persistent reservations requests.

Parameters to set on a host (channel) director port basis

Configure the following settings for each host (channel) director port to which the gateway connects.

Note

Parameter names might be different depending on whether they appear in the GUI or CLI.

Parameter	Value to set	Comments
Common SN (Common Serial Number flag, also known as the C-bit Flag)	Enable this parameter so that each array LUN is presented across the various Fibre Channel director ports.	This parameter assigns a common serial number across the logical devices that are exposed through multiple target ports. Note —— Most multipathing software uses the LUN serial number to determine whether a LUN has multiple paths.
PP Flag (Point-to-Point flag)	Enable this flag on the DMX FA ports if the gateways and DMX storage array are connected through switches.	Do not enable the PP Flag if the gateways and the DMX storage array are directly connected.
SC3 Flag (SCSI-3)	Enable this flag.	Describes the command sets for block-oriented direct-access devices—for example, disk drives. The gateway supports SCSI-3 command blocks.
SPC-2 Flag (SCS2_Protocol_version, SCSI Primary Command 2 flag, or “Allow inquiry data to be compiled to the standard”)	Enable this flag.	Contains the second-generation definition of the basic commands for all SCSI devices. SPC-2 is used in conjunction with several standards for the specific device type, including SCSI-3 Block Commands (SBC-3).
UWN Flag (Unique World Wide Name flag)	Enable this flag for each target port so that each FA port that communicates with the gateway has a Unique World Wide Name.	—

Parameters to set for each logical device for the gateway

When you create a logical device that you intend to present to gateways, you must set the values of the parameters in the following table as shown.

Parameter	Value to set	Comments
Host Type Format	Default format type for Open Systems, which is “Server” in the EMC ControlCenter.	Each logical device in a DMX storage array is formatted using a specific format type.
SCSI3_persist_reserv (Persistent Reservations flag or PER Flag)	Set this flag to On.	This flag, commonly known as SCSI-3 Persistent Reservations, is set on a DMX storage array to support hosts that require SCSI-3 persistent reservations. The gateway uses the “Write Exclusive Registrants Only” persistent reservation type.

Requirement for LUN access control

To eliminate the possibility of a non gateway overwriting DMX array LUNs owned by a gateway, or the reverse, you must present the DMX logical devices through the host (channel) director ports in one of the following ways:

Method 1: Present only the DMX logical devices for the gateway on specific DMX host (channel) director ports that are dedicated to gateway use, and prevent the logical devices for the gateway from being mapped through other host (channel) director ports.

Method 2: Use the LUN masking capability to associate all DMX logical devices to all host (channel) director ports, but allow only the gateway FC initiator ports to see the LUNs for the gateway.

If you are setting up a configuration in which you are using multiple array LUN groups for the gateway, you will set up a separate “host group” for each group of LUNs for the gateway.

Attention

If you use Method 2, do not present the (Volume Configuration Management DataBase (VCMDB) to all hosts by default. Configure the global setting to restrict visibility to the VCMDB unless it has been specifically made visible to a particular host.

The results of Method 2 are the same as Method 1. But the benefit is that masking is at the logical level (initiator port WWN) rather than at the physical or switch zone level.

**Recommendation
regarding port
sharing**

It is strongly recommended that you do not share DMX array ports between gateway and non gateway clients. The reason is that there is great potential for incompatible port requirements when DMX ports are shared between gateway and non gateway clients. Each host connecting to the DMX storage array has different requirements for port attribute settings, which can result in a configuration that is impossible to implement.

If the ports on the DMX storage array cannot be dedicated to the gateway, confirm that the port requirements for all other hosts using that port are compatible with gateway port requirements.

**Caution about
using the VCMDB
LUN**

For Volume Configuration Management DataBase (VCMDB) to be enabled, the VCMDB LUN must exist. The VCMDB LUN is a “command” type LUN, not a storage LUN. VCMDB is typically mapped to LUN 0. However, the VCMDB LUN can be an array LUN other than LUN 0.

If you map the VCMDB LUN to the gateway, the gateway periodically logs a message that the VCMDB LUN is less than the minimum size required and it marks the VCMDB LUN as failed. The gateway will continue to function normally after logging this error message.

**Caution about
using a gatekeeper
LUN**

If a gatekeeper logical device (LUN) is presented, do not map it to the gateway. A gateway cannot use a gatekeeper LUN. A gatekeeper LUN is smaller than gateway can handle. You cannot assign the gatekeeper LUN to the gateway with the `disk assign` command.

Note

A gatekeeper is a DMX logical device, accessible by a host, through which SYMAPI or the ControlCenter agent communicates with the DMX storage array.

Gateway requirements for Fujitsu ETERNUS storage arrays



See the gateway *Implementation Guide for Fujitsu ETERNUS Storage* for additional details regarding configurations with Fujitsu ETERNUS storage arrays.

Configuration requirements for the ETERNUS3000

You need to change some of the parameters for the Host Response function to ensure that the gateway recognizes the ETERNUS3000. You configure the Host Response function through the Host Response Settings screen.

To create the new Host Response Pattern for the gateway, make a copy of the read-only Host Response Pattern#00 that Fujitsu supplies and change the settings, as necessary, to match the settings in the following table. The other parameters should remain as the default values.

Parameter	Value
Host Response Name	<i>name</i> It is recommended that you enter a name that makes it easy to identify the gateway, for example, gateway.
Response status to host when overloaded	Queue Full
Response to inquiry command (Peripheral Qualifier/Peripheral Type)	Default
Inquiry Standard Data NACA Function	Disable
Inquiry Standard Data Version Field	03 (Default)
Inquiry Command Page 83	Type 01 & 03 Note _____ This might be set to Type 01 by default. Be sure to change it to Type 01 & 03. _____

Parameter	Value
Reservation Conflict Response to Test Unit Ready Commands	Enable
Target Port Group Access Support	Disable (default)
Response Sense at Firmware Hot Switching:	Disable

Configuration requirements for the ETERNUS6000

You need to change some of the parameters for the Host Response function to ensure that the gateway recognizes the ETERNUS6000. You configure the Host Response function through the Set Host Response screen.

The following table shows the required settings for the gateway to operate with the ETERNUS6000. The other parameters should remain as the default values.

Parameter	Value
Host Response Name	<i>name</i> It is recommended that you enter a name that makes it easy to identify the gateway, for example, gateway.
response status in overload	<i>Value is supplied by the management software.</i>
byte0 of Inquiry response	default (no conversion)
Inquiry Standard Data Version	Version 03 (default)
response data type for Inquiry PageCode 0x83	type 1 + type3
Reservation Conflict response for Test Unit Ready	conflict response (default)
Ordered Support Mode	Off (default)
host specific mode	normal moded (default)
pattern of sense code conversion	default (no conversion)

Required Host Response parameters for the gateway

For the gateway to operate with the ETERNUS4000 and ETERNUS8000, you need to create a new Host Response Pattern and set the parameters as required by the gateway. Fujitsu provides Host Response Pattern#00, which is read only.

Attention

If you do not change the Host Response setting for the gateway, the ETERNUS4000 or ETERNUS8000 assigns the default pattern to the gateway.

The Host Response Profile differs depending on the storage array firmware version. Select the following table for the Host Response Profile that matches your storage array firmware.

Host Response Profile parameter settings specification 1:

Parameter	Value
Host Response Name	<i>name</i> It is recommended that you enter a name that makes it easy to identify the gateway, for example, gateway.
Response status to host when overloaded	Unit Attention
Response to inquiry command (Peripheral Qualifier/Peripheral Type)	Default
Inquiry Standard Data NACA Function	Disable
Inquiry Standard Data Version Field	03 (Default)
Inquiry Command Page 83	Type 01 & 03 (Default)
Reservation Conflict Response for Test Unit Ready	Normal
Asymmetric / Symmetric Logical Unit Access	Active/Active (Default)

Host Response Profile parameter settings specification 2:

Parameter	Value
Host Response Name	<i>name</i> It is recommended that you enter a name that makes it easy to identify the gateway, for example, gateway.
Command Time-out Interval	Standard (25 seconds)
Load Balance Response	Unit Attention
Byte-0 of Inquiry Response	No Conversion (Default)
Inquiry VPD ID Type	Type1 + Type3 (Default)
Inquiry Standard Data Version	Version 05 (Default)
Reservation Conflict Response for Test Unit Ready	Normal Response (Default)
Host Specific Mode	Normal Mode (Default)
Asymmetric / Symmetric Logical Unit Access	Active/Active (Default)
LUN Mapping Changes	No Report (Default)
LUN Capacity Expansion	No Report (Default)
Vendor Unique Sense Code	No Report
Sense Code Conversion Pattern	No Conversion (Default)

See the gateway *Interoperability Matrix* for any updates to information about requirements for Fujitsu storage arrays.

Gateway requirements for Hitachi storage arrays



See the gateway *Implementation Guide for Hitachi Storage* for additional details regarding configurations with Hitachi storage arrays.

Required parameters

The following table lists the required parameters for a Hitachi storage array to work with a gateway.

Platform	Required system parameter	Required Host Group option
AMS200, AMS500, AMS1000, AMS2500, AMS2300	—	—
USP, NSC55	—	—
9910, 9960	System mode 254 on	—
9970V, 9980V	System mode 254 on	—
9570V	—	No_RSV_ConfMode Enabled
9580V, 9585V	—	No_RSV_ConfMode Enabled

See the gateway *Interoperability Matrix* at www.ibm.com/storage/support/nseries/ for any updates to information about system parameters.

Additional settings

The following table shows additional settings for your storage array to work with a gateway.

Models	Settings
AMS1000, AMS500, and AMS200, AMS2500, AMS2300	◆ Enable Mapping mode ◆ Enable Host group security ◆ Change the default Product ID in the AMS configuration from DF600F to DF700
9570V, 9580V, 9585V	◆ (Recommended) Enable LUN mapping ◆ Use data sharing mode

Models	Settings
99xx	Use quick formatting

Gateway requirements for HP storage arrays



See the gateway *Implementation Guide for HP XP Storage* or the gateway *Implementation Guide for HP EVA Storage* for additional details regarding configurations with HP storage arrays.

Requirements for HP XP storage arrays

The following table shows the system parameter settings that are required for an HP XP storage array to work with a gateway.

Platform	Configuration setting
XP24000, XP20000, XP12000, XP10000	—
XP1024, XP128, XP512, XP48	Set System mode 254 on

See the gateway *Interoperability Matrix* at www.ibm.com/storage/support/nseries/ for any updates to information about system parameters.

See the gateway *Implementation Guide for HP XP Storage* for information about requirements for using external disks with HP XP storage arrays.

Requirements for HP EVA storage arrays

The following table shows the system parameter settings that are required for an HP EVA storage array to work with a gateway.

Parameter	Value
Redundancy	Vraid5
Preferred path/mode	Default
Host Type	SUN Solaris

Gateway requirements for IBM storage arrays



See the gateway *Implementation Guide for IBM Storage* for additional details regarding configurations with IBM storage arrays.

ESS series

Use the following information to plan for configuring an ESS storage array to work with a gateway.

For...	Value...
Number of LUNs on the ESS storage array that can be allocated to gateways	The maximum number of LUNs that you can allocate for gateways depends on the gateways that you are deploying, not on what the ESS models can support. ESS storage arrays support more LUNs than the high-end gateways support. See the gateway <i>Interoperability Matrix</i> at http://www.ibm.com/storage/nas/ for information about the maximum number of LUNs that each gateway supports.
Host type for LUNs mapped to the gateway.	RS/6000

Note

ESS logical volume IDs are in hexadecimal format. Data ONTAP LUN names are in decimal format. To correlate ESS logical volumes to Data ONTAP LUNs, you must convert the hexadecimal numbers to decimal.

DS8xxx series

Use the following information to plan for configuring a DS8xxx storage array to work with a gateway.

For...	Value...
Number of LUNs on the DS8xxx storage array that can be allocated to gateways	256 LUNs per host group, regardless of the number of LUNs that the gateway or DS8xxx supports. This is a gateway limitation with DS8xxx storage arrays.
Host type for LUNs mapped to the gateway	Sun-Solaris Note ————— When a host type of Sun-Solaris is used, the DS8xxx logical volume IDs are in hexadecimal format. Data ONTAP LUN names are in decimal format. To correlate DS8xxx logical volumes to Data ONTAP LUNs, you must convert the hexadecimal numbers to decimal. —————

Recommended number of volume groups: It is recommended that you use a single volume group *volume group* on a DS8xxx storage array.

Attention —————
You can use two volume groups if you ensure that the LUN numbers between the volume groups do not overlap. If you add FC initiator ports and LUNs to the volume groups after your initial (working) volume group configuration, be very careful that the LUN numbers do not overlap. If the LUN numbers overlap, the gateway panic with a misconfiguration error and both nodes in the HA pair fail. The gateway panic on reboot until the misconfigured LUNs are fixed or removed.

Note—————
A DS8xxx volume group is used to control the hosts that can access LUNs (LUN masking) by associating host attachments or port groups with the LUNs that they are allowed to access. You assign the desired FC initiator ports and the LUNs to be accessed to the same volume group.

LUN access with the DS8300 9A2 LPAR (system logical partitions):

When setting up the DS8300 9A2 LPAR (system logical partition) model to interact with gateways, ensure that you set up access to a given LUN so that the redundant paths are both accessing the same LPAR.

DS4xxx/DS5xxx series

Use the following information to plan for configuring all DS4xxx/DS5xxx series storage array models to work with a gateway.

For...	Value...
Number of LUNs on the DS4xxx/DS5xxx that can be allocated to gateways	256 LUNs per host group, regardless of the number of LUNs that the gateway supports. This is a gateway limitation with DS4xxx/DS5xxx storage arrays.
Host type for LUNs mapped to the gateway.	AIX

Limitation in number of paths between a gateway and a switch:

There cannot be more than two paths between a gateway and a switch when the gateways are deployed with DS4xxx/DS5xxx storage arrays. There can be only two paths between the switch and a LUN on the storage array. This is true whether your configuration includes a stand-alone gateway, a gateway HA pair, or both.

Gateway requirements for 3PAR storage arrays



See the gateway *Implementation Guide for 3PAR Storage* for additional details regarding configurations with 3PAR storage arrays.

Required system parameters for 3PAR storage arrays

For a 3PAR storage array to work with a gateway, you must change the port persona to one of the Data ONTAP personas shown in the following table. The 3PAR port persona number represents a profile of the host system to which the target port is attached. Changing the port persona to a Data ONTAP persona changes the 3PAR port-persona pair from initiator to gateway persona target ports. You must set the port persona before you assign array LUNs to the gateway.

Configuration	Port persona
Direct-connect	18
Fabric-attached	19

Attention

After you change the 3PAR InServ® Storage Server control port-persona pair to **18** or **19** (the gateways persona target port), you cannot assign any other third-party initiators to that port. You can only assign additional gateway initiators to that port.

See the gateway *Interoperability Matrix* for any updates to information about requirements for 3PAR storage arrays.

About this appendix This appendix is a reference to equivalent models between storage array vendors.

**Equivalent Hitachi,
HP XPxxx, and Sun
models**

Hitachi	HP XP	Sun
USP-V	XP24000, XP20000	—
USP	XP12000	StoreEdge 9990
NSC	XP10000	StoreEdge 9985
9980	XP1024	—
9970	XP128	—
9960	XP512	StoreEdge 9980
9910	XP48	StoreEdge 9980
95xx	—	StoreEdge 9970

Terminology Comparison Between Vendors



Different vendors that the gateway supports sometimes use different terms to describe the same things. This appendix provides a mapping between some common vendor terms.

Term	Vendor	Definition
host group	Hitachi	A configuration entity that enables you to specify host access to ports on the storage array. You identify the FC initiator port WWNs for the gateway that you want to access the LUNs; the process differs according to vendor and sometimes differs for different storage array models of the same vendor.
	IBM DS4xxx/DS5xxx	
	EMC DMX	
	HP XP	
volume group	IBM DS8xxx	
Storage Group	EMC CX	
zone	Fujitsu ETERNUS3000	
host affinity group	Fujitsu ETERNUS4000, ETERNUS6000, and ETERNUS8000	
host definition	3PAR	
host	3PAR HP EVA	
—	IBM ESS	No concept of “host group.” You must create a host in the ESS user interface for each gateway FC initiator port that you plan to connect to the storage array and map each host to a port.

Term	Vendor	Definition
parity group	IBM DS8xxx, IBM ESS, Hitachi, HP XP	The arrangement of disks in the back-end that together form the defined RAID level.
RAID group	Data ONTAP, EMC CX, Fujitsu ETERNUS	
array, RAID set	IBM DS4xxx/DS5xxx	
Parity RAID, Parity RAID group	EMC DMX	A DMX feature that provides parity data protection on the disk device level using physical parity volumes.
disk group	HP EVA	A set of physical disks that form storage pools from which you can create virtual disks.
parity set, RAID set	3PAR	A group of parity-protected chunklets. (A chunklet is a 256-MB block of contiguous space on a physical disk.)
cluster	Data ONTAP	In Data ONTAP 6.x and some 7.x releases, the term formerly used for an active/active configuration or HA pair.
	Hitachi, HP XP	A hardware component on the storage arrays that contains the ports to which hosts attach.
controller	Data ONTAP	The component of a gateway that runs the Data ONTAP® operating system and interacts with back-end storage arrays. Controllers are also sometimes called <i>heads</i> , or <i>CPU modules</i> .
	Hitachi, HP XP HP EVA	A hardware component on the storage arrays that contains the ports to which hosts attach.

Term	Vendor	Definition
LUN	Many storage arrays	A grouping of one or more disks or disk partitions into one span of disk storage space.
	Data ONTAP	The gateway can virtualize the storage attached to it and serve the storage up as LUNs to applications and customers outside the gateway (for example, through iSCSI and FCP). Clients are unaware of where a <i>Data ONTAP-served LUN</i> is stored.
LUN, virtual disk	HP EVA	A virtual disk (called a Vdisk in the user interface) is a simulated disk drive created in a disk group. You can assign a combination of characteristics to a virtual disk, such as a name, redundancy level, and size. Presenting a virtual disk offers its storage to a host.
array LUN	Gateway documentation	The gateway documentation uses the term array LUN to distinguish LUNs on the storage arrays from Data ONTAP LUNs.
vLUN	3PAR	(volume-LUN) A pairing between a virtual volume and a logical unit number (LUN). For a host to see a virtual volume, the volume must be exported as a LUN by creating VLUNs on the storage array.
volume	IBM	Equivalent to what other storage array vendors call a LUN.
	Data ONTAP	A logical entity that holds user data that is accessible through one or more of the access protocols supported by Data ONTAP, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Fibre Channel Protocol (FCP), and Internet SCSI (iSCSI). The gateway treats an IBM volume as a disk.
	EMC DMX	A general term referring to a storage device. A physical volume corresponds to a single disk device.
virtual volume	3PAR	A virtual storage unit created by mapping data from one or more logical disks.

Glossary

active/active configuration

Term used in Data ONTAP 7.x releases for two storage systems (nodes) running Data ONTAP that are connected to each other either directly or through switches. This is the same as what is called an HA pair. See *HA pair*.

active-active storage arrays

With *active-active* storage arrays, the storage array does not assign a preferred or default controller to an array LUN; the array LUN can be accessed through any controller without a performance impact. A gateway automatically chooses one of the storage array ports that you specified for a given array LUN as the primary path and the other port as the secondary (alternate) path. A gateway uses its own algorithm to balance traffic to the array LUNs over its FC initiator ports. You cannot manually change the primary and secondary paths to an array LUN.

I/O to a given array LUN is always through one path at any given time, whether the storage array supplies the primary path to the array LUN or the gateway determines the primary path to the array LUN.

active/passive configurations

In a Data ONTAP HA pair, the passive node has only a root volume, while the active node has all of the remaining storage and services all data requests during normal operation. The passive node responds to data requests only if it has taken over the active node.

active-passive storage arrays

With *active-passive* storage arrays, the storage array assigns a primary path for a given array LUN. The gateway always uses the primary path supplied by the active-passive storage array, unless the primary path fails. If a failure occurs, the gateway fails over to the secondary path but resumes using the primary path when the primary path is restored.

aggregate

In Data ONTAP, a user-defined grouping of physical storage resources. Aggregates provide storage to the volume or volumes that they contain.

array LUN

(Logical unit) The term used in the gateway documentation to refer to a grouping of one or more disks or disk partitions on the storage array into one span of disk storage space. Some vendors use the term LUN.

checksum

A form of redundancy check, a simple measure for protecting the integrity of data through error detection. It is used mainly in data storage and networking protocols. It adds up the basic components of a message, typically the bytes, and stores the resulting value. Later, the authentic checksum can verify that the message was not corrupted by doing the same operation on the data, and checking the sum.

cluster**Data ONTAP definition:**

Term used in Data ONTAP 6.x releases and some early Data ONTAP 7.x releases for two storage systems (nodes) running Data ONTAP that are connected to each other either directly or through switches. This term was changed to active/active configuration and then to HA pair. See *HA pair*.

Each node in the cluster can see the same volumes as any other node in the cluster, assuming that a single virtual server exists across the entire cluster. The total file-system namespace, which comprises all of the volumes and their resultant paths, is global across the cluster.

Some storage array vendors' definition:

The hardware component on which host adapters and ports are located. Some vendors refer to this component as a *controller*.

controller**Data ONTAP definition:**

The component that runs the Data ONTAP operating system and interacts with back-end storage arrays. Controllers are also sometimes called *gateways*, *heads*, or *CPU modules*.

Some storage array vendors' definition:

The hardware component on which host adapters and ports are located. Some vendors refer to this component as a *cluster*.

core dump	A core dump is a valuable tool for technical support to use while troubleshooting why a system running Data ONTAP crashed. When a system running Data ONTAP crashes, a core dump file is saved to the root volume if adequate space is available on the root volume.
Data ONTAP-served LUN	See <i>front-end LUN</i> .
data set	A logically related group of files and directories organized in a single subtree, for example, all the files related to an Oracle database instance.
disk group	The arrangement of disks in the back-end that together form the defined RAID level. See also Appendix F, “ Terminology Comparison Between Vendors ,” on page 175.
disk ownership	The software-based scheme that Data ONTAP uses to control which gateway can access a specific LUN that a storage array has made available. Only one gateway can access a particular LUN.
disks, disk shelves	The terms <i>disk</i> and <i>disk shelf</i> in this document refer to native storage connected to the gateway. These terms do not refer to disks or disk shelves on a third-party storage array.
family	Storage arrays in the same <i>family</i> share the same performance and failover characteristics. For example, members of the same family all perform active-active failover or they all perform active-passive failover. Storage arrays with 4 GB HBAs are not considered to be in the same family as storage arrays with 2 GB HBAs. When you set up a Data ONTAP aggregate, you cannot assign array LUNs from different storage array families or different vendors to the same aggregate.
flexible volume	See <i>FlexVol volume</i> .

FlexVol volume	A volume that is loosely coupled with its containing aggregate, which enables you to manage the volume separately from the aggregate and gives you a lot more flexibility in managing the size of the volume.
front-end LUN	Data ONTAP can virtualize the storage attached to it and serve the storage up as LUNs to applications and customers outside the gateway (for example, through iSCSI and FCP). In the gateway documentation, these front-end LUNs are referred to as <i>LUNs</i> or <i>Data ONTAP-served LUNs</i> to distinguish them from array LUNs on the back-end storage arrays. Clients are unaware of where a front-end LUN is stored.
gateway neighborhood	A set of one or more gateways that all see the same LUNs on the storage array. The neighborhood feature enables you to define up to six gateways as a logical group to facilitate storage management. For details about neighborhoods, see Chapter 7, “Using Neighborhoods for Storage Management,” on page 83.
giveback	The process that the HA pair goes through when returning to normal mode.
HA pair	Two storage systems (nodes) running Data ONTAP that are connected to each other either directly or through switches. The nodes in an HA pair are connected to each other through a cluster adapter or an NVRAM adapter, which allows one node to serve data to the array LUNs or disks of its failed partner node. Each node continually monitors its partner. Both gateways in an HA pair must be able to see the same LUNs from the storage arrays. However, only one node has read and write access to the LUN because it is the configured owner. A Data ONTAP HA pair was previously referred to as an <i>active/active configuration</i> or a <i>cluster</i>.
hard zone	A hard zone is created using the names of the switch ports. See also, soft zone.
host group	A term used on some storage arrays to describe an entity in the storage partition topology that defines a logical collection of host computers that require shared access to one or more logical drives. Different storage arrays use different terms to describe this configuration entity.

Some storage arrays allow you to define multiple “host groups” for hosts, which enables you to use two port pairs on a gateway to access two different sets of array LUNs on a storage array. The gateway neighborhood supports multiple host groups on the same storage array. See also Appendix F, “[Terminology Comparison Between Vendors](#),” on page 175.

igroup

Initiator group. A collection of unique iSCSI node names of initiators (hosts) in an IP network that are given access to “front-end LUNs” when they are mapped to those LUNs. (Array LUNs on a storage array that provide storage for gateways can be considered “back-end LUNs.”)

LUN

Data ONTAP definition:

Data ONTAP can virtualize the storage attached to it and serve the storage up as LUNs to applications and customers outside the gateway (for example, through iSCSI and FCP).

Term as defined by many storage array vendors:

(Logical unit) A grouping of one or more disks or disk partitions into one span of disk storage space. A LUN on a storage array looks like an individual disk to the gateway. A gateway or non gateway host reads data from or writes data to a LUN on a storage array.

A LUN on a storage array has a RAID type and properties that define it.

Some vendors use other terms than LUN to describe this storage space. For example, some IBM storage arrays describe this storage space as a *volume*. The gateway documentation uses the term *array LUN* to distinguish the storage on storage arrays from LUNs that Data ONTAP serves.

See also Appendix F, “[Terminology Comparison Between Vendors](#),” on page 175.

LUN group

A set of logical devices on a storage array that a gateway accesses over the same paths. Each gateway must have exactly two paths to a LUN group. The storage array administrator configures a set of logical devices as a group in order to define which host WWPNs can access them.

LUN access control

See *LUN security*.

LUN masking	Controls which hosts that are zoned to the same port can access specific LUNs on the storage array port. You allow a specific host to access specific array LUNs. You prevent other hosts from accessing those same array LUNs by hiding (or masking) those same array LUNs from the other hosts.
LUN ownership	See <i>disk ownership</i> .
LUN security	A control that prevents one host from accessing an array LUN that has been allocated to another host. When LUN security is applied, the array LUN is exposed only to the host to which it has been allocated. Therefore, the array LUN is not accessible to other hosts.
MetroCluster	MetroCluster is a special type of HA pair that provides the capability to force a takeover when an entire node is destroyed or unavailable. Not all platforms, switches, storage subsystems, or Data ONTAP versions are supported in MetroCluster configurations. Before configuring a MetroCluster configuration, see the gateway Interoperability Matrix and the gateway <i>MetroCluster Guide</i> to make sure your configuration is valid and supported.
native disks, native disk shelves	Disks and disk shelves that are sold as local storage for systems that run Data ONTAP software.
neighborhood	A set of one or more gateways that all see the same LUNs on the storage array. The neighborhood feature enables you to define up to six gateways as a logical group to facilitate storage management. For details about neighborhoods, see Chapter 7, “Using Neighborhoods for Storage Management,” on page 83
node	Data ONTAP definition: In a gateway HA pair, either the local or partner gateway.
panic	A serious error causing the gateway to halt. It is similar to a system crash in a Windows environment.

parity group

The arrangement of disks in the back-end that together form the defined RAID level. See also Appendix F, “[Terminology Comparison Between Vendors](#),” on page 175.

parity RAID group

The arrangement of disks in the back-end that together form the defined RAID level. See also Appendix F, “[Terminology Comparison Between Vendors](#),” on page 175.

persistent reservations

Persistent reservations are a SCSI-3 feature for restricting access to storage media. The logical drive reservation is a feature of the cluster software, which allows one or more host ports to reserve a logical drive, thereby preventing other host ports to access the same logical drive. The actual reservation of the logical drive is handled by the host application.

The benefit of the persistent reservations feature is that it allows the storage to integrate with cluster solutions that use shared logical drives for increased availability, scalability, and performance.

Gateways use the “Write Exclusive Registrants Only” persistent reservation type. This type of reservation is well-suited to SAN multipath, multihost environments because it enables cooperating hosts to concurrently access the array LUN while excluding access to other hosts. Therefore, it is ideally suited for clusters of nodes (gateway nodes or a cluster of nodes for non gateway hosts).

With the “Write Exclusive Registrants Only” persistent reservation type, the nodes register keys and form a shared reservation (one key per initiator or HBA and the keys do not need to be unique). Access to the array LUN is granted only to the hosts (initiators) that registered with the LUN. With gateways, the reservation owner is the gateways that owns the array LUN. For gateways, the registration key contains the information that Data ONTAP uses to determine which node is holding the reservation. This information includes the NVRAM serial number.

plex

A collection of one or more RAID groups that together provide the storage for one or more Data ONTAP volumes. If you are not using SyncMirror, you do not need to plan for and configure plexes; Data ONTAP creates one plex automatically. If you plan to use SyncMirror, there are specific configuration requirements for the gateway related to setting up plexes for SyncMirror.

pool

For data to be mirrored, one plex of an aggregate must reside on each storage array or be physically isolated on the same storage array. The scheme that Data ONTAP uses to specify the location of the plexes is through the assignment of SyncMirror *pools*. In Data ONTAP, pools are used to split a set of array LUNs into two groupings so that the two groupings do not share common points of failure.

RAID

(Redundant array of independent disks) RAID is a way of grouping individual physical drives together to form one bigger drive called a RAID group. RAID can make many smaller disks appear as one large disk to a server. The RAID group represents all the smaller physical drives as one logical disk to your server. The logical disk is called an array LUN, or *logical unit number*. Using RAID has two main advantages—better performance and higher availability.

RAID group

Data ONTAP RAID group:

One or more Data ONTAP RAID (Redundant Array of Independent Disks) groups within an aggregate. From the perspective of gateways, Data ONTAP uses RAID groups to determine where to allocate data to the array LUNs.

Storage array RAID group:

The arrangement of disks in the back-end that together form the defined RAID level. Each RAID group supports only one RAID type. The number of disks that you select for a RAID group determines the RAID types that a particular RAID group supports. See also Appendix F, “[Terminology Comparison Between Vendors](#),” on page 175.

Different vendors use different terms for RAID group. Equivalent terms among vendors include parity group, disk group, and Parity RAID group.

root volume

A special volume on each gateway, which can be a traditional volume or a FlexVol volume. The root volume contains system files and configuration information. It is required for the system to be able to boot and to function properly. The root volume can also contain data if you want it to. Core dump files, which are important for troubleshooting, are written to the root volume if there is enough space.

Snapshot copy	A frozen, read-only image of a traditional volume, a FlexVol volume, or an aggregate that reflects the state of the file system at the time the Snapshot copy was created. Snapshot copies are your first line of defense for backing up and restoring data.
soft zone	A soft zone is created using a list of WWNs of the HBAs on the servers or storage resources connected to the SAN. See also, hard zone.
SyncMirror	A software product that creates two physically separated copies of an aggregate. Customers who choose to use SyncMirror typically have two storage arrays and cannot tolerate any down time. You must purchase specific licenses to use SyncMirror
takeover	The process by which one gateway in an HA pair undergoes a system failure and cannot reboot, and the partner gateway takes over the failed gateway functions and serves data.
traditional volume	A volume that is contained in a single, dedicated aggregate. You manage the aggregate and traditional volume together.
units of measurement for storage	Data ONTAP defines units of measurement for storage as follows: Kilobyte (KB) = 1,024 bytes Megabyte (MB) = 1,024 x 1,024 bytes Gigabyte (GB) = 1,000 bytes x MB Terabyte (TB) = 1,000 bytes x GB Storage array vendors might use different units of measurement. See the gateway <i>Implementation Guide</i> for your vendor for information about how to calculate the units of measurement for your vendor that are equivalent to the Data ONTAP limits. All limits used in the gateway document library use the Data ONTAP-defined units of measurement.

vFiler unit	A virtual gateway that you create using MultiStore, which enables you to partition the storage and network resources of a single gateway so that it appears as multiple gateways on the network.
volumes	Data ONTAP definition: A Data ONTAP volume is a logical entity that holds user data that is accessible through one or more of the access protocols supported by Data ONTAP, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Fibre Channel Protocol (FCP), and Internet SCSI (iSCSI). The gateway treats an IBM volume as a disk. IBM definition: IBM uses the term <i>volume</i> to describe the area on the storage array that is available for a gateway or non gateway host to read data from or write data to. The gateway documentation uses the term array LUN to describe this area. See also Appendix F, “Terminology Comparison Between Vendors,” on page 175.
World Wide Name (WWN)	A unique identifier assigned to a Fibre Channel device.
World Wide Port Name (WWPN)	A unique identifier assigned to a port on a Fibre Channel device. Ports on some storage arrays, for example, Hitachi storage arrays, have a unique WWPN.
volume group	On an IBM DS8xxx storage array, a volume group is used to control the hosts that can access array LUNs (LUN masking) by associating host attachments or port groups with the array LUNs that they are allowed to access. You assign the desired FC initiator ports and the array LUNs to be accessed to the same volume group. See also Appendix F, “ Terminology Comparison Between Vendors ,” on page 175.
zoning	Configuring zoning on a Fibre Channel switch enables you to restrict visibility and connectivity between devices connected to a common Fibre Channel SAN. With zoning, you overlay a security map on the network that dictates which hosts can see and have access to specific targets, thereby reducing the risk of data loss.

If, for example, one host were to gain access to an array LUN being used for storage by another host, the data in the array LUN could become corrupted. If a host cannot see an array LUN, it cannot access the LUN to corrupt it.

Index

Numerics

- 3PAR storage arrays
 - required settings 172
- 95xxV storage arrays
 - No RSV Conf mode requirement 166
 - required settings 166
- 99xxV storage arrays
 - required settings 166
 - System mode 254 requirement 166

A

- accessible LUNs, defined 80
- access to array LUNs
 - controlling through LUN masking 18
 - controlling through zoning 51
 - number of paths allowed 20, 42
 - planning for 17
 - restricting to one gateway 81
 - setting up 20, 42
- active/active configurations *See* HA pair
- aggregates
 - checksum type of LUNs, impact of 75
 - defined 80
 - LUNs added to
 - example of checksum type 77
 - rules for checksum type 75
 - planning RAID groups for 67
 - relationship of LUN checksum type 71
 - rules for assigning array LUNs to 100
- AMS storage arrays
 - required settings 166
- ASCII terminal console
 - cable wiring 134
 - cabling to your system 135
 - connecting to 134
- assigning storage array ports 20
- available LUNs, defined 80

B

- block checksum type LUNs

- aggregate to be used in 72
- layout of 72
- performance considerations 72
- space reserved in 72

- Brocade switches, guidelines for 56

C

- capacity for gateway 14
- C-bit Flag 159
- cf.takeover.on_panic option
 - minimum required size of spare LUN 70
 - purpose of 70
 - requirement for spare LUN 70
- characteristics of HA pairs 128
- checksum type of a LUN
 - block 71
 - changing 74
 - deduplication requirement 71
 - zoned 71
- checksum, defined 71, 180
- Cisco switches, guidelines for 57
- clusters on storage array
 - redundancy requirement 20, 108
 - specifying ports on 118, 126, 130
- Common Serial Number flag 159
- communication parameters for ASCII terminal console 135
- connecting
 - a stand-alone gateway 125
 - an HA pair 128
 - planning for 118
 - redundancy requirement 118
 - your system to a storage array 115
- connections on gateways
 - N5000 119
 - N6000 122
 - N7000 119
 - N7x050T 119
- connections to storage arrays
 - direct-attached 6
 - fabric attached 6

- methods 6
- MetroCluster 7
- connections, redundancy requirement 118
- connector type 110
- console
 - cabling 135
 - wiring 134
- controllers on storage array
 - defined 180
 - redundancy requirement 20, 108
 - specifying ports on 118, 126, 130
- copying products, use with gateways 155
- core dumps
 - described 70
 - file
 - description 70
 - how it is used 181
 - in root volume 186
 - space requirement in LUNs 70
- CX storage arrays
 - required parameters 156
 - Storage Groups, rules fo 157

D

- Data ONTAP RAID groups. *See* RAID groups (Data ONTAP)
- Data ONTAP-served LUN, defined 2, 181
- data set, defined 181
- deduplication, checksum requirement 71
- desired space in a LUN, calculating 73
- direct-attached configurations
 - advantages 6
 - disadvantages 6
- disk group
 - defined 62
 - minimum number of LUNs in 62, 63
 - relationship to LUN 62
 - sharing LUNs in 62
- disk ownership
 - defined 80
 - planning for 81
 - with array LUNs 80
 - with native disks 82
- disk.target_port.cmd_queue_depth option 137

- disks on gateways
 - Data ONTAP releases that support 1, 3
- DMX storage arrays
 - gatekeeper LUN caution 161
 - LUN access control 160
 - port sharing 161
 - required parameters 158
 - VCMDB LUN 161
 - defined 161
 - results of mapping it 161
- DS4xxx storage arrays
 - maximum number of LUNs supported 171
 - rules for gateways 171
- DS5xxx storage arrays
 - maximum number of LUNs supported 171
 - rules for gateways 171
- DS8xxx storage arrays
 - DS8300 9A2 LPAR, path setup with 171
 - logical volume IDs 170
 - maximum number of LUNs supported 170
 - rules for gateways 170
- dual paths to array LUNs 33

E

- EMC CX storage arrays
 - required settings 156
 - Storage Group rules 157
- EMC DMX storage arrays
 - gatekeeper caution 161
 - LUN access control 160
 - port sharing 161
 - required parameters 158
 - VCMDB LUN 161
 - defined 161
 - results of mapping it 161
- enclosures on storage arrays
 - redundancy requirement 20, 108
- ESS storage arrays
 - logical volume IDs 169
 - maximum LUNs supported 169
 - required settings 169
 - rules for gateways 169
- ETERNUS3000, required settings 162
- ETERNUS6000, required settings 163

F

- fabric-attached configurations
 - advantages 6
 - supported switches 50
- fabric-attached MetroCluster
 - defined 7
 - in-order-delivery FC-VI traffic with 59
- family, defined 8, 102, 181
- FC initiator ports
 - by gateway 109
 - how identified 108
 - using multiple pairs 21
- FC-VI traffic, MetroCluster requirement 59
- Fibre Channel disks, assigning LUNs from 100
- FlexVol volumes, defined 182
- front-end LUN, defined 2
- Fujitsu storage arrays
 - ETERNUS3000 required settings 162
 - ETERNUS6000 required settings 163

G

- gatekeeper LUN for DMX
 - caution 161
 - defined 161
- gateway Interoperability Matrix
 - information in 1
 - where to find it 1
- gateway neighborhood. *See* neighborhood
- gateway ports
 - guidelines for determining usage 113
 - required number to disk shelf 113
 - required number to storage array 112
 - required setting 112
 - restrictions for sharing 112
- gateways
 - connecting to an ASCII terminal 135
 - connection locations
 - N5000 gateway series 119, 120
 - N6000 122
 - N7000 119
 - N7x50T 119
 - maximum capacity 14
 - minimum capacity 14
 - multiple storage arrays connecting to 100

- obtaining WWNs 153
- gigabyte, defined 67, 187
- giveback, defined 182

H

- HA pairs, defined 182
- Hitachi storage arrays
 - required settings 166
- host group
 - defined 89, 143, 182
 - equivalent terms between vendors 175
- host type
 - DS4xxx storage arrays 171
 - DS5xxx storage arrays 171
 - DS8xxx storage arrays 170
 - ESS storage arrays 169
 - HP EVA storage arrays 168
- Host Type Format 160
- HP EVA storage arrays
 - required settings 168
- HP XP storage arrays
 - required settings 168

I

- IBM storage arrays
 - required settings 169
- in-order delivery of FC-VI traffic 59
- installing
 - back-end implementation 9
 - front-end implementation 9
- ISLs
 - requirements with Cisco fabric 57
 - requirements with MetroCluster 59

K

- kilobyte, measurement defined 187

L

- LC connectors 110
- link failure
 - Data ONTAP response to 36
 - scenario, impact on pathing 36, 37

- load balancing in a neighborhood 84
 - logical unit number, defined 2
 - LUN 0
 - constraints for using 64
 - do not map to gateway 161
 - LUN group
 - defined 21, 183
 - LUN masking
 - controlling access to array LUNs 18
 - defined 18, 184
 - LUN numbers, consistency when mapping 21
 - LUN ownership
 - defined 80, 181
 - planning for 81
 - LUN security
 - defined 18, 184
 - methods 18
 - requirement for 18
 - LUNs
 - See *LUNs (array)*, *LUNs (front-end LUNs)*, or *LUNs (Data ONTAP-served LUNs)*
 - LUNs (array)
 - accessible, defined 80
 - assigning to aggregates 100
 - Hitachi storage array consideration 100
 - HP XP storage array consideration 100
 - available, defined 80
 - checksum type
 - block 72
 - changing 74
 - comparing types 71
 - described 71
 - example of adding to aggregates 77
 - relationship to aggregates 71
 - rules for adding to aggregates 75
 - zoned 72
 - constraints for LUN 0 64
 - core dump space on 70
 - Data ONTAP calculation of GB 65
 - defined 2, 177, 180, 183
 - desired space in, calculating 73
 - gatekeeper LUN for DMX 161
 - how available for storage 80
 - minimum number allowed 66
 - minimum number in disk group 62, 63
 - minimum number in parity group 62, 63
 - minimum number in RAID group 62
 - name format 39
 - number considerations 65
 - ownership changes 81
 - paths to
 - Data ONTAP response to link failure 36
 - errors when setting up 23
 - how Data ONTAP chooses 34
 - how to check access 42
 - invalid setup examples 29, 30
 - viewing path being used 35
 - planning access to 17
 - relationship to disk groups 62
 - relationship to parity groups 62
 - requirement for gateways to use 80
 - requirement for redundant paths 20
 - restriction in hosts sharing 62
 - rules for setting up paths to 42
 - sharing in disk group 62
 - sharing in parity group 62
 - sizing 62
 - best practice 68
 - calculating desired space 73
 - calculating usable space 73
 - error when size not met 69
 - guidelines 65, 67
 - root volume 68
 - space reserved for core dumps 70
 - spare core LUNs
 - allocating 141
 - when required 70
 - type you can map to the gateway 64
 - usable space in
 - calculating 73
 - factors that impact 73
 - VCMDB, caution about using 161
 - LUNs (Data ONTAP-served), defined 2
 - LUNs (front-end), defined 2
- ## M
- mapping array LUNs
 - constraints for LUN 0 64
 - LUN number requirement 21

- LUN types to map to the gateway 64
- maximum capacity of gateways 14
- McDATA switches
 - guidelines for use with 58
 - OSMS requirement 58
 - port setting requirement 58
 - QPM blades with 6140 58
 - recommendation with firmware 9.01 58
- megabyte, measurement defined 187
- MetroCluster
 - fabric-attached, defined 7
 - FC-VI traffic requirement 59
 - neighborhood, not supported with 85, 87
 - restriction for native disks 7
 - stretch, defined 7
 - switches with 59
 - types supported 7
- minimum capacity of gateways 14
- multiple port pairs, use of 21

N

- N5200 gateways, connecting 120
- N5300 systems, connecting 121
- N5500 gateways, connecting 120
- N5600 gateways, connecting 121
- N6000 gateways, connecting 122
- N7000 gateways, connecting 119
- N7x050T gateways, connecting 119
- naming, format for array LUNs 39
- native disk shelves, defined 2, 184
- native disks
 - defined 2, 184
 - RAID protection for 5
 - restriction for MetroCluster 7
- native storage, defined 2, 181
- neighborhood
 - access control considerations 89
 - calculating limits for, example 97
 - configurations supported 87
 - defined 84, 182, 184
 - determining maximum LUNs for 93
 - example 85
 - host group configuration for 89
 - how established 89

- impact of native disks on systems in 96
- limits 91, 92
- maximum LUN limit 87
- maximum systems in 87
- mixed platforms in 94
- number supported 86
- planning for 91
- requirements to establish 87
- SnapMover, moving vFiler units with 84
- storage array configuration for 89
- storage array specific restrictions 95
- types of limits compared 92
- what impacts neighborhood limit 93
- when neighborhoods not supported 88
- when not applicable 84
- zoning considerations 90
- neighborhood maximum LUN limit 87, 97
 - calculating 97
 - defined 91
 - relationship to platform limit 92
 - what reduces limit 93
- No RSV Conf mode requirement 166
- node, defined 184
- NSC55 storage arrays
 - required settings 166
- NVRAM
 - N5200 gateways 120
 - N5300 gateways 121
 - N5500 gateways 120
 - N5600 gateways 121
 - N6000 gateways 122
 - N7000 gateways 119

O

- OSMS, requirement for 58
- ownership
 - defined 80
 - planning for 81
 - with native disks 82

P

- panic, defined 184
- parity group
 - defined 62

- minimim number of LUNs in 62, 63
 - relationship to LUN 62
 - sharing LUNs in 62
 - sizing LUNs 62
- Parity RAID group defined 62
- paths to array LUNs
 - checking 43, 46
 - commands for 42
 - with storage array show-config 46
 - errors when setting up 23
 - how Data ONTAP chooses 34
 - how Data ONTAP chooses, failure response 36
 - how shown in array LUN names 39
 - invalid setup examples 29, 30
 - number allowed 20, 42
 - redundant 20, 42
 - rules for 42
 - troubleshooting 29, 30
 - viewing path used 35
- Persistent Reservations flag 160
- persistent reservations, defined 185
- planning
 - access to storage resources 11
 - LUN sizing 61
 - major tasks 11
 - RAID groups for aggregates 67
- platform maximum assigned device limit
 - defined 91
 - relationship to neighborhood limit 92
- plexes
 - described 185, 187
 - relationship to pools 186
- Point-to-Point flag 159
- pools (SyncMirror)
 - defined 186
 - relationship to plexes 186
- port monitoring, Data ONTAP
 - difference between array types 33
 - how it works 33
- port selection, Data ONTAP
 - after bootup 34
 - at bootup 34
 - how path is selected after bootup 33
 - how path is selected at bootup 33
- ports (storage array)

- connecting to redundant 20
- ports on gateway
 - guidelines for determining usage 113
 - how identified 108
 - redundant 109
 - required number to disk shelf 113
 - required number to storage array 112
 - required setting 112
 - restrictions for sharing 112
 - unused, how to handle 110
- PP Flag 159
- provisioning array LUNs, guidelines 61

Q

- QPM blades, requirements to use 58
- queue depth
 - command 137
 - defined 137
 - limit 137
 - setting 138
- quick start 139

R

- RAID groups (Data ONTAP)
 - defined 186
 - LUN number implications 67
 - LUN size implications 67
- RAID groups (storage arrays)
 - defined 5, 62, 176, 186
 - minimim number of LUNs in 62
- RAID protection
 - for array LUNs 5
 - for native disks 5
- RAID0
 - data protection, none 5
 - how Data ONTAP uses 5
- RAID4 protection for native disks 5
- RAID-DP protection for native disks 5
- redundant configuration
 - required number of paths 20
 - requirement for 20
- redundant paths to array LUNs
 - requirement 20, 42
 - why necessary 20

- redundant port pairs on gateway 109
- replication products, use with gateways 155
- root volume

- core dump files in 186
 - defined 186
 - sizing LUN for 68

S

- SATA disks, assigning array LUNs from 100
- SC3 Flag 159
- SCSI3_persist_reserv 160
- single-initiator zoning, benefit of 51
- size of array LUNs, gateway rules for 65
- SnapMover
 - load balancing in a neighborhood 84
- Snapshot copies, defined 187
- spare core LUN (for core dumps)
 - configuring array LUN for 141
 - relationship to cf.takeover.on_panic option 70
 - when required 70
- SPC-2 Flag 159
- storage
 - how presented to clients 4
 - virtualization 3
- storage array show-config command
 - checking paths to LUNs 42, 46
 - example output 46
- storage array show-luns command
 - checking paths to LUNs 46
- storage arrays
 - connecting to redundant ports 20
 - multiple behind gateway 7
 - different series 104
 - same series 102
- Storage Groups
 - maximum LUNs in 157
 - number supported 157
- storage show disk -p command
 - output example 43
 - use for checking paths to LUNs 42, 43
 - when used 35
- storage, planning access to 17
- stretch MetroCluster
 - defined 7

- switch zoning. *See* zoning
- switches

- Brocade, guidelines for 56
 - Cisco, guidelines for 57
 - McDATA
 - guidelines for 58
 - OSMS requirement 58
 - need for redundancy with 20
 - requirements in a gateway configuration 49
 - with MetroCluster 59
 - WWNs for zoning 51
 - zoning
 - defined 188
 - example 52, 53
 - ports for 51
 - purpose of 51
 - recommendation 51

- SyncMirror
 - plexes, described 187
 - pools with 186
 - why use 187
- sysconfig -v command
 - checking paths to LUNs 42, 45
- System mode 254 requirement 166

T

- takeover, defined 187
- target queue depth
 - defined 137
 - limit 137
 - setting 138
- target queue length 137
- terabyte, measurement defined 187
- traditional volumes, defined 187
- troubleshooting
 - pathing problems 29, 30

U

- Unique World Wide Name 159
- units of measurement, Data ONTAP 187
- usable space in a LUN
 - calculating 73
 - factors that impact 73
- USP storage arrays, required settings 166

V

VCMDB LUN

- defined 161

- results of mapping it 161

vFiler units, defined 188

visible device limit 96

volumes (Data ONTAP)

- defined 177, 188

- FlexVol, defined 182

- root

 - defined 186

 - minimum LUN size for 68

- traditional, defined 187

volumes (IBM), defined 188

W

WWPNs, obtaining manually 153

Z

zoned checksum type LUNs

- aggregate to be used in 72

- compared to block checksum 72

- layout of 72

- performance considerations 72

zoning

- defined 51

- devices sharing zones 51

- examples 52, 53

- overview 51, 188

- purpose of 51

- recommendation 51

- recommendation for 51

- single-initiator zoning, benefit of 51

- specifying ports 51

- specifying WWNs 51



NA 210-05286_A0, Printed in USA

GC53-1148-05

