



**IBM System Storage N series
Gateway Implementation Guide for IBM SAN
Volume Controller Storage**

Copyright and trademark information

Copyright information

Copyright ©1994 - 2008 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2008 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.

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

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

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

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

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

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

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

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

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:
This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR

CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

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

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

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

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

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

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

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

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

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

Notice of any changes or modifications to the W3C files, including the date changes were made.

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED “AS IS,” AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS

FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

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

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

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

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

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

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

Trademark information

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

Table of Contents

	Prefacexi
Chapter 1	Using SAN Volume Controller with Gateways	1
	Commonly used SAN Volume Controller terms	3
	Overview of SAN Volume Controller with gateways	5
	Rules for SAN Volume Controller with gateways	8
	Zoning requirements	11
	Supported configurations	16
Chapter 2	Configuration Overview	35
Appendix A	Array LUN Sizing for SAN Volume Controller	39
	Index	43


Preface

About this guide

This guide provides information about how to set up your storage array to work with an IBM® System Storage™ N series gateway running Data ONTAP® software, including configuration guidelines and sample configurations. The information in this guide pertains to all supported gateway platforms.

Relationship of this guide to other guides

This guide is intended to be used in conjunction with other information in the gateway library.



The following table describes the relationships between this guide and other gateway documentation.

	Guide name	Information includes...
1	<i>Installation Requirements, Quick Start, and Reference Guide</i>	<ul style="list-style-type: none">◆ General guidelines for creating and making array LUNs available to gateways◆ Quick start installation instructions◆ Reference information

	Guide name	Information includes...
2	<i>Implementation Guides</i>	<ul style="list-style-type: none"> ◆ Vendor-specific details ◆ More detailed configuration examples than are provided in the <i>Installation Requirements, Quick Start, and Reference Guide</i>.
3	<i>Setup, Installation, and Management Guide</i>	Detailed steps for setting up the gateway and installing Data ONTAP software (for installers new to Data ONTAP setup and installation).
4	<i>Planning Guide</i>	Detailed background information, for example, about aggregate and Data ONTAP volume use, array LUN size and layout in aggregates, and checksums

Attention

In this document, *gateway* describes IBM N series models that are used in conjunction with other disk storage systems, which include disk storage systems from IBM, HP®, Hitachi Data Systems®, and Hitachi®, Limited. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system.

In this document, the term *filer* describes IBM N series models that either contain internal disk storage or attach to the disk storage expansion units specifically designed for the IBM N series storage systems.

Audience

This guide is for users who are familiar with operating systems such as UNIX®, Windows® 95, Windows NT®, Windows 2000, and Windows XP and who will be installing gateways. This guide does not discuss basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the gateway.

Terminology

An *active/active configuration* is a pair of gateways configured to serve data for each other if one of the two systems becomes impaired. In gateway documentation, Data ONTAP documentation, and other information resources, active/active configurations are sometimes also referred to as *clusters*.

Supported features IBM N series products 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 at the following Web site:

www.ibm.com/storage/support/nas/

A listing of currently available N series products and features can be found at the following Web site:

www.ibm.com/storage/nas/

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 that you have taken these steps to try to solve the problem yourself:

- ◆ Check all cables to make sure that they are connected properly.
- ◆ Check the power switches to make sure that the system is turned on.
- ◆ Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.

Using the documentation Information about the N series hardware products is available in printed documents and a documentation CD that comes with your system. The same documentation is available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Web sites IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates.

- ◆ For NAS product information, go to the following Web site:
www.ibm.com/storage/nas/
- ◆ For NAS support information, go to the following Web site:

www.ibm.com/storage/support/nas/

- ◆ For AutoSupport information, go to the following Web site:

www.ibm.com/storage/support/nas/

- ◆ For the latest version of publications, go to the following Web site:

www.ibm.com/storage/support/nas/

Accessing online technical support

For online Technical Support for your IBM N series product, visit the following Web site:

www.ibm.com/storage/support/nas/

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following Web site for support telephone numbers:

www.ibm.com/planetwide/

Supported servers and operating systems

IBM N series products attach to many servers and many operating systems. To determine the latest supported attachments, visit the following Web site:

www.ibm.com/storage/support/nas/

Firmware updates

As with all devices, it is recommended that you run the latest level of firmware. For information on firmware updates, visiting the following Web site:

www.ibm.com/storage/support/nas/

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support. See the *Gateway Upgrade Guide* for your version of Data ONTAP for more information on updating firmware.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Attention

Attention contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com or use the Readers' Comments form at the back of this publication. Be sure to include the following:

- ◆ Exact publication title
- ◆ Form number (for example, GC26-1234-02)
- ◆ Page numbers to which you are referring

If the Readers' Comments Form in the back of this manual is missing, you can direct your mail to:

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

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

About this document

This document includes an overview of SAN Volume Controller and provides information to help you plan and implement a deployment with SAN Volume Controller serving storage to gateways. This document is intended to be used in conjunction with other information in the gateway library. In particular, refer to the additional documents in the following table.

For information about...	See...
Data ONTAP releases that support SAN Volume controller, supported switches, supported firmware, capacity, and maximum array LUN count	<i>Gateway Support Matrix</i> at http://www.ibm.com/storage/nas/
How the gateway operates and what you need to plan for a successful deployment with the gateway	<i>Gateway Planning Guide</i>
Rules for the following: <ul style="list-style-type: none">◆ Setting up access to array LUNs◆ Assigning array LUNs to gateways Common Data ONTAP terms that you might encounter in the documentation	<i>Gateway Installation Requirements, Quick Start, and Reference Guide</i>
Maximum and minimum array LUN sizes that the gateways supports for IBM storage arrays	Appendix A, “ Array LUN Sizing for SAN Volume Controller ,” on page 39
How to connect a gateway to a SAN Volume Controller	The gateway guide <i>Connecting Your Gateway</i>

For information about...	See...
How to configure a gateway to interact with a storage array	<p data-bbox="827 239 1231 401"><i>Gateway Installation Requirements, Quick Start, and Reference Guide</i> (quick start instructions for readers familiar with Data ONTAP setup and installation).</p> <p data-bbox="827 430 1231 591"><i>Gateway Setup, Installation, and Management Guide</i> (detailed procedures for readers who are not familiar with Data ONTAP setup and installation)</p>

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Commonly used SAN Volume Controller terms](#)” on page 3
- ◆ “[Overview of SAN Volume Controller with gateways](#)” on page 5
- ◆ “[Rules for SAN Volume Controller with gateways](#)” on page 8
- ◆ “[Zoning requirements](#)” on page 11
- ◆ “[Supported configurations](#)” on page 16

Gateway perception of SAN Volume Controller

Conceptually, a gateway interacts with SAN Volume Controller as if it is just another storage array, although SAN Volume Controller is a different type of device. Many of the rules in the gateways documentation about gateways interaction with storage arrays apply to the SAN Volume Controller also. This chapter provides information specific to SAN Volume Controller deployed with gateway.

Note

SAN Volume Controller defines a storage array as a device that coordinates and controls the operation of one or more disk drives. In gateway documentation, storage array is the term used to describe a device that coordinates and controls the operation of one or more disks (array LUNs) presented to gateways, for example, an IBM DS4xxx or DS8xxx storage array or a Hitachi storage array.

Commonly used SAN Volume Controller terms

cluster	One to four pairs of SAN Volume Controller nodes. The nodes in the SAN Volume Controller cluster are connected to the same SAN and present virtual disks (VDisks) to the hosts, including gateways. The hosts have no knowledge of the underlying physical hardware of the cluster.
extent	Fixed size unit of data that is used to manage the mapping of data between managed disks (MDisks) and VDisks.
host	The computer systems in front of SAN Volume Controller to which SAN Volume Controller presents VDisks. Hosts see VDisks as array LUNs. Gateway, Solaris, and Windows systems are examples of hosts to SAN Volume Controller.
input/output (I/O) group	<p>Defines which SAN Volume Controller nodes provide I/O access to a VDisk. An I/O group contains two SAN Volume Controller nodes.</p> <p>For I/O purposes, SAN Volume Controller nodes within the cluster are grouped into pairs, called I/O groups, with a single pair being responsible for serving I/O on a given VDisk. One node within the I/O group represents the preferred path for I/O to a given VDisk. The other node represents the non-preferred path. As each VDisk is created within an I/O group, this assigned preference alternates between nodes to balance the workload evenly between the two nodes.</p>
managed disk (MDisk)	A LUN that a back-end storage array presents to SAN Volume Controller. A SAN Volume Controller cluster manages MDisks. An MDisk is not visible to host systems on the SAN.
managed disk group (MDG)	Defines which MDisks from the back-end storage arrays make up a specific VDisk. It is recommended that an MDG contain LUNs from only a single storage array.

nodes

The individual servers in a SAN Volume Controller cluster on which the SAN Volume Controller software runs. The nodes are always installed in pairs, with one to four pairs of nodes constituting a cluster. Each node in a pair is configured to back up the other. Each pair of nodes is known as an I/O group.

The preferred node does not have absolute ownership. The partner node in the I/O group accesses the data if the preferred node fails or if the preferred node workload becomes too high.

virtual disk (VDisk)

A SAN Volume Controller logical entity that appears to host systems attached to the SAN as a SCSI disk, logical volume, or array LUN. An individual VDisk is a member of one MDisk group (MDG) and one I/O group.

VDisks are set up in the SAN Volume Controller configuration from the pool of storage (MDisks) that the back-end storage arrays provide to SAN Volume Controller. The pool is split into VDisks that are visible to the host systems that use them.

Overview of SAN Volume Controller with gateways

Advantage of deploying SAN Volume Controller with gateways


A major advantage of deploying SAN Volume Controller with gateways is that you can use Data ONTAP to manage storage provided by storage arrays that a gateway does not support itself.

Overview of SAN Volume Controller

SAN Volume Controller is a virtualization storage appliance from IBM. SAN Volume Controller manages a number of back-end storage arrays, pooling the physical storage from those storage arrays and mapping it to logical disk images that hosts in the SAN can see. The storage arrays behind the SAN Volume Controller are invisible to the hosts in front of it.

SAN Volume Controllers are deployed in pairs for redundancy. Each SAN Volume Controller in the pair is a *node*. One to four pairs of nodes make up a SAN Volume Controller *cluster*. A SAN Volume Controller cluster, therefore, can contain up to eight nodes. The nodes in a SAN Volume Controller cluster see the LUNs presented to them by the back-end storage arrays as a number of disks, referred to as managed disks (*MDisks*).

The following illustration provides an overview of the relationship of SAN Volume Controller to the hosts in front of it and to the back-end storage arrays.



Details for the illustration are as follows.

1 Array LUN presentation by the back-end storage arrays: The storage array administrator configures the storage array to see SAN Volume Controller as a host. The back-end storage arrays present LUNs to SAN Volume Controller.

The WWNs of the hosts in front of SAN Volume Controller are not identified as hosts in the storage array configuration.

2 Managed disks: SAN Volume Controller automatically discovers the LUNs presented by the storage arrays as MDisks. The MDisks from the different storage arrays become part of the SAN Volume Controller storage pool. The SAN Volume Controller administrator creates managed disk groups (MDGs) and assigns MDisks from the storage pool to the MDGs, grouping the MDisks in a way that achieves the organization's performance and reliability requirements.

3 VDisks: The SAN Volume Controller administrator creates virtual disks (VDisks) from the extents in the MDGs and maps the VDisks to hosts.

Mapping a VDisk to a host makes the VDisk accessible to the WWNs identified in the host object configuration and, therefore, to the host itself. Hosts can access only the VDisks that SAN Volume Controller allocates to them.

4 Presentation of storage to hosts: The hosts see each VDisk that SAN Volume Controller presents to them as a LUN. Gateways see the LUNs as storage from which the gateway can build DOT aggregates.

Hosts on the front-end of SAN Volume Controller have no knowledge of the back-end storage arrays.

**Number of SAN
Volume Controllers
that can supply
VDisks**

One SAN Volume Controller cluster or multiple SAN Volume Controller clusters can supply VDisks to one or more gateways.

Rules for SAN Volume Controller with gateways

SAN Volume Controller firmware version

See the gateway *Support Matrix* for information about supported firmware.

Requirements for SAN Volume Controller arrays

If you are planning on using an IBM SAN Volume Controller in conjunction with a gateway, you are strongly encouraged to use the information, tools and services available from IBM to ensure that your SAN Volume Controller implementation (and that of the storage arrays behind it) are properly designed and configured to ensure the maximum performance and availability. The gateways can generate a substantial I/O workload on the SAN Volume Controller and on the storage arrays behind it. Careful planning of the SAN Volume Controller configuration is recommended to ensure a successful implementation.

The information, tools and services available from IBM for SAN Volume Controller include:

- ◆ Formal product publications and redbooks
- ◆ Capacity and performance sizing tools
- ◆ Professional implementation services

Requirements for switches

You can deploy SAN Volume Controller with gateways only in switched environments. SAN Volume Controller ports are connected to switches so that SAN Volume Controller can access the hosts in front of it and the storage arrays behind it. The nodes in a SAN Volume Controller cluster also communicate with each other over the SAN Volume Controller ports.

Requirements for switches are as follows:

- ◆ All switches that SAN Volume Controller uses to access hosts and storage arrays must be on the SAN Volume Controller *Support Matrix*.
See the SAN Volume Controller *Support Matrix* at <http://www-03.ibm.com/servers/storage/support/software/sanvc/installing.html>.
- ◆ The gateway must support the switches between the gateways and SAN Volume Controller.
See the gateway *Interoperability Matrix* at <http://www.ibm.com/storage/nas/>.

- ◆ SAN Volume Controller requires that all switches in the same fabric be from the same vendor.

For information about zoning the switches, see “[Zoning requirements](#)” on page 11.

Identification of the gateway host in SAN Volume Controller

In the SAN Volume Controller configuration, you must assign all FC initiator ports of the gateways in a gateway neighborhood to the same SAN Volume Controller host object; that is, all the gateways that must see the same VDisks (LUNs).

You do not need to specify a specific host type in the SAN Volume Controller configuration for gateways.

Rules for VDisk creation

All types of VDisks—striped, sequential, image, and cache disabled—are supported in deployments with SAN Volume Controller and gateways.

Adhere to the rules in the following table when creating VDisks that you are going to assign to gateways.

Regarding...	Comments
Number of VDisks on the SAN Volume Controller that can be allocated to gateways	<p>The maximum number of array LUNs that a gateway supports depends on the gateway model. The total number of VDisks that you allocate for the gateways in the gateway neighborhood cannot exceed the lowest value of either: (a) the number of array LUNs that the gateway model can support or (b) the number of VDisks that SAN Volume Controller can support per I/O group and across the SAN Volume Controller cluster.</p> <p>See the gateway <i>Support Matrix</i> for information about the maximum number of array LUNs that each gateway model supports. See the SAN Volume Controller documentation for information about the number of VDisks that SAN Volume Controller supports per I/O group and across the SAN Volume Controller cluster.</p>
Maximum VDisk size	<p>The size of the VDisks that you can create on the SAN Volume Controller is limited by the maximum array LUN size that Data ONTAP supports. See Appendix A, “Array LUN Sizing for SAN Volume Controller,” on page 39.</p>

Supported and unsupported SAN Volume Controller VDisk operations

VDisk operations that the gateway supports: A host does not see the MDisk-to-VDisk mapping; it only sees the VDisks that SAN Volume Controller presents to the host. Therefore, you can change the mapping between an MDisk and a VDisk without interrupting gateway access to the VDisk. For example, assume that the VDisk with the logical block address X is on MDisk1. On SAN Volume Controller, you can remap VDisk X to MDisk2 without the gateway being aware that the move occurred.

VDisk operations that the gateway does not support: The gateway does not support any SAN Volume Controller VDisk management features that would sever the relationship between a VDisk and a gateway, for example

- ◆ Unmapping a VDisk from the gateway host object
- ◆ Moving a VDisk from one I/O group to another

Attention

If you move a VDisk from one I/O group to another, the **gateway** panics because the move creates a conflict in the paths over which the **gateway** sees the VDisks. Data loss might occur.

Array LUN assignment to aggregates

You cannot mix LUNs (VDisks) presented by SAN Volume Controller in the same Data ONTAP aggregate as LUNs from other storage arrays.

LUNs (VDisks) from different SAN Volume Controller I/O groups can be mixed in the same aggregate. (This is the equivalent to mixing LUNs from two storage arrays in the same vendor family in the same aggregate.)

Zoning requirements


Why zoning is required

All VDisks are visible on all SAN Volume Controller ports. Therefore, you must use zoning to prevent non gateway hosts from accessing VDisks for the gateways and to prevent gateways from accessing VDisks through more than one port per fabric.

Note

Configuring zoning on a Fibre Channel switch enables you to restrict visibility and connectivity between devices connected to a common Fibre Channel SAN. Switches work by connecting individual ports to individual entities in a virtual point-to-point connection. The zones on the switch work by exclusion. A member of the zone has connectivity to any other member within the zone; a device that is not a member of that zone can only communicate with a member of the zone if both devices are included together in another zone.

The following illustration shows how zoning in a deployment with SAN Volume Controller establishes the SAN Volume Controller as the storage device in the middle between the hosts and the back-end storage arrays.



This illustration shows two host zones and one storage zone.

Host zones: Hosts that use SAN Volume Controller for storage are zoned to SAN Volume Controller. In the illustration, each host is in a separate zone with SAN Volume Controller because they are different hardware platforms. The hosts see only the VDisks that they are allowed to access. To be able to access a VDisk, a host must be in the zone that includes the SAN Volume Controller nodes in the I/O group of which the VDisk is a member and a mapping must exist on SAN Volume Controller between the VDisk and the host object.

Storage zone: The storage zone includes SAN Volume Controller and back-end storage arrays. (In the SAN Volume Controller documentation, this “storage zone” consists of two overlapping zones: the SAN Volume Controller zone, which contains the SAN Volume Controller nodes in the cluster, and the disk zone, which contains the storage arrays.)

- ◆ The back-end storage arrays that provide storage to SAN Volume Controller are zoned to the SAN Volume Controller instead of to the gateways and non gateway hosts in front of SAN Volume Controller. The storage arrays see SAN Volume Controller as their host.
- ◆ The nodes in a SAN Volume Controller cluster are zoned so that they always see each other.

Note

See [“Scenario 2: the gateway and SAN Volume Controller share a back-end storage array”](#) on page 21 for an example of a storage array that provides storage directly to a gateway and to SAN Volume Controller, which also provides storage to the same gateway.

Zoning requirements

This section provides information about zoning requirements in a deployment with a gateway and SAN Volume Controller. Each scenario in [“Supported configurations”](#) on page 16 includes a zoning example.

Attention

When you configure zoning, you are zoning by host or storage array port and not by the entire host or storage array.

Requirements for SAN Volume Controller nodes: All nodes in a SAN Volume Controller cluster must be

- ◆ In the same zone so that they can communicate cluster information among themselves
- ◆ Able to see the same set of ports on each back-end storage array

The back-end storage arrays must be in the same zone as SAN Volume Controller nodes so that the nodes in the cluster can see and address the LUNs presented to them by the storage arrays.

Requirements for back-end storage arrays:

- ◆ Storage arrays that provide storage only through SAN Volume Controller must be zoned to see only SAN Volume Controller as their connected host.
- ◆ If a storage array provides storage to both SAN Volume Controller and directly to a gateway that also uses SAN Volume Controller, you must set up zoning so that the gateway sees both SAN Volume Controller and the back-end storage array through different zones. The zones are the following:
 - ❖ Host and SAN Volume Controller
 - ❖ Host and back-end storage
 - ❖ SAN Volume Controller and back-end storage

This type of zoning is required because SAN Volume Controller requires all of its ports to be zoned with the back-end storage arrays but the host must be zoned with a subset of the SAN Volume Controller ports, typically providing a total of two or four paths from the host to each VDisk. The number of paths depends on the requirements of the multipathing software on the host.

For an example of zoning when a gateway share a back-end storage array, see [“Scenario 2: the gateway and SAN Volume Controller share a back-end storage array”](#) on page 21.

Requirements for hosts for which SAN Volume Controller provides storage:

- ◆ You must set up zoning so that the hosts see SAN Volume Controller as their storage provider, not the back-end storage arrays that provide the storage for SAN Volume Controller.

This zoning ensures that hosts are not able to directly “see” or operate LUNs on the storage arrays that are assigned to the SAN Volume Controller cluster. As a result, conflicts that would result from the SAN Volume Controller and the hosts trying to manage the same storage are avoided.

Note

If a storage array that provides storage to SAN Volume Controller also maps array LUNs directly to a gateway that uses SAN Volume Controller for some storage, you must set up zoning so that the gateway sees both SAN Volume Controller and the back-end storage array through different zones. See [“Scenario 2: the gateway and SAN Volume Controller share a back-end storage array”](#) on page 21.

- ◆ HBAs that are in dissimilar hosts or dissimilar HBAs that are in the same host must be in separate zones. Dissimilar means that the hosts are running different operating systems or that they are different hardware platforms. Therefore, a host zone cannot contain a gateway together with other types of hosts.
- ◆ All VDisks are visible on all ports on the nodes in a SAN Volume Controller cluster. You must use host mapping in SAN Volume Controller to prevent hosts from seeing VDisks that are not allocated to them.

Gateway specific requirements:

- ◆ A specific VDisk is a member of only one I/O group in a SAN Volume Controller cluster. That VDisk is visible on every port of both nodes in the I/O group. You must use zoning to limit the number of paths to a VDisk that a specific gateway FC initiator port can use. Data ONTAP requires that a single gateway FC initiator port access a specific VDisk on only one path. Therefore, for each gateway FC initiator port that is allowed to see a specific VDisk, you must configure a zone that includes only that FC initiator port and one port from only one node in that VDisk's I/O group.
- ◆ Use single initiator zoning for a gateway because the gateway FC initiator ports do not need to and should not see each other. Create a separate zone for each initiator-target combination. This way the gateway FC initiator ports are isolated from each other so that they do not see each other. Each gateway FC initiator port should only see the target SAN Volume Controller port.
- ◆ Each gateway must see both nodes in every I/O group that you want to provide storage to the gateway.
- ◆ From a gateway perspective, each SAN Volume Controller I/O group is conceptually a different storage array.
- ◆ All gateway FC initiator ports attaching to the same I/O group must be in the same host object.
- ◆ A gateway can be mapped to VDisks from more than one I/O group. In this case, the gateway FC initiator ports are zoned to multiple I/O groups. However, in a given gateway redundant port pair, both gateway FC initiator ports in the port pair must be zoned to the same I/O group; they cannot be zoned to different I/O groups.
- ◆ Each gateway can have only one connection to each node in an I/O group. Be sure that you zone for redundancy by including redundant gateway port pairs, redundant switch port pairs, and redundant SAN Volume Controller node port pairs.

Gateway specific path requirements: Each SAN Volume Controller node presents a VDisk to the SAN through multiple paths, typically four, and each VDisk is accessible from the two SAN Volume Controller nodes in an I/O group.

However, gateways expect SAN Volume Controller to provide access to a specific VDisk on *only* two paths: one as the primary path and the other as the alternate path. You must zone the fabric to achieve this.

Supported configurations

Examples of configurations that are supported for storage access

Organizations have different goals for how they want to manage storage. This section shows several options for deploying SAN Volume Controller with gateways and the zoning to support these scenarios. For simplicity, the first few scenarios show a deployment with only one gateway and with only one I/O group on SAN Volume Controller.


- ◆ [“Scenario 1: all storage is provided through SAN Volume Controller”](#) on page 16
- ◆ [“Scenario 2: the gateway and SAN Volume Controller share a back-end storage array”](#) on page 21
- ◆ [“Scenario 3: SAN Volume Controller provides storage only for storage arrays the gateway does not support”](#) on page 25
- ◆ [“Scenario 4: the gateway and a gateway client are both attached to SAN Volume Controller”](#) on page 28
- ◆ [“Scenario 5: active/active configuration with multiple SAN Volume Controller I/O groups”](#) on page 31

Limitation for MetroCluster

MetroCluster is not currently supported with storage provided through SAN Volume Controller. The gateway *Support Matrix* is updated when new support is added.

Scenario 1: all storage is provided through SAN Volume Controller

In this scenario, the SAN Volume Controller pools the storage from all the back-end storage arrays in the environment, which could include storage arrays that the gateway supports. Organizations might use this type of configuration because they want to manage all storage from one storage pool.




When SAN Volume Controller provides all the storage for the hosts, the storage arrays are all configured to see only SAN Volume Controller as their host. SAN Volume Controller is configured see all the hosts, including the gateway, and presents the storage as VDisks to the hosts.

Zoning for this scenario: To discuss zoning for this scenario, assume the following:

- ◆ There are two storage arrays behind SAN Volume Controller, one an IBM DS4300 storage array and the other an EMC storage array. Both provide storage only to SAN Volume Controller.
- ◆ SAN Volume Controller has only one I/O group in the cluster.

The following illustration shows the SAN Volume Controller, gateway, and back-end storage array ports in the zones that are configured on each switch. The same information is summarized in a table following the illustration.



-  Ports in the storage zones
-  Ports in the gateway 1 host zone with SAN Volume Controller
- - - Indicates that the ports at each end of the dashed line are in the same zone
- C1, C2 Represents alternate controllers (clusters) on the storage subsystem

Each switch has the following zones:

- ◆ *Storage zones* (Zones 1 and 2 in this example): The storage zone on each switch includes ports in the SAN Volume Controller I/O group and ports for the EMC and DS4300 storage arrays.
- ◆ *Host zones* (Zones 3 and 4 in this example): Each gateway FC initiator port is in a separate zone. Each host zone includes a port from one node in the SAN Volume Controller I/O group.

Note

Notice that both Zone 1 (a gateway host zone) and Zone 3 (a storage zone) include Node 1, port 1 in I/O group 0. The same port can be included in more than one zone. A member of a zone can connect to any other member of the zone but not to a device outside the zone unless both devices are included together in another zone.

Similarly, Zone 2 and Zone 4 both include Node 2, port 4. However, Gateway 1 port 0c can communicate only to Node 1, port 4 because it is the only SAN Volume Controller port that is explicitly configured in Zone 4.

Redundancy is achieved by setting up zoning as follows:

- ◆ The two gateway FC initiator ports are from different cards (or buses) on the gateway.
- ◆ The gateway sees both nodes in the I/O group. Then, if one node fails, the gateway can access the VDisk through the other node.
- ◆ The connections from the gateway to the SAN Volume Controller nodes are through redundant switch port pairs.
- ◆ The ports on the storage array are from alternate controllers (clusters).

Attention


Gateways expect SAN Volume Controller to provide access to a specific VDisk on *only* two paths. Each gateway can have only one connection to each of the two nodes in an I/O group.

The following table summarizes the zoning for this scenario.

Zone	Switch	The zone includes this Gateway FC initiator port	The zone includes this SAN Volume Controller node and port		The zone includes this storage array, controller, and port	
Zoning pair						
1	1	—	Node 1	port 1	EMC	C1, port 1
			Node 1	port 3	EMC	C2, port 1
			Node 2	port 1	DS4300	C1, port 1
			Node 2	port 3	DS4300	C2, port 1
2	2	—	Node 1	port 2	EMC	C1, port 2
			Node 1	port 4	EMC	C2, port 2
			Node 2	port 2	DS4300	C1, port 2
			Node 2	port 4	DS4300	C2, port 2
Zoning pair						
3	1	0a	Node 1	port 1	—	—
4	2	0c	Node 2	port 4	—	—

Scenario 2: the gateway and SAN Volume Controller share a back-end storage array

The following illustration shows a configuration in which the same back-end storage array presents LUNs to both the SAN Volume Controller (as MDisks, which it virtualizes to hosts) and directly to the gateway.



A gateway can share a back-end storage array with SAN Volume Controller if the following requirements are met:

- ◆ The storage array is supported by both SAN Volume Controller and the gateway.
- ◆ Appropriate LUN partitioning is configured on the storage array.
- ◆ The back-end storage array does not implement “LUN grouping.”

When a storage array implements LUN grouping, all LUNs that are specified to be in a particular LUN group are treated as one unit. Access to the LUNs in that LUN group is restricted to one system. You cannot remap only some of the LUNs in that LUN group to another system. If one LUN is remapped to a different system, all LUNs are automatically remapped to the other system.

Note

LUN grouping is a different function than LUN masking.

This type of configuration is typical in an environment with existing gateways where SAN Volume Controller is being installed for data migration. Some customers prefer to perform data migration in stages to reduce the impact of the data migration on a production network, particularly when the computer systems must be running 24 hours a day 7 days a week. Typically this configuration is transitory and is used just during migration. You can leave this “split controller access” in place after data migration. However, if you do so, you might not be able to take advantage of all the benefits of Data ONTAP or SAN Volume Controller.

Zoning for this scenario: To discuss zoning for this scenario, assume the following:

- ◆ On the DS4300, LUNs 21-50 are allocated to and mapped to the gateway so that the DS4300 can provide some storage to the gateway directly. The gateway is identified as a host in the DS4300 configuration.
- ◆ On the DS4300, LUNs 1-20 are allocated to and mapped to SAN Volume Controller. SAN Volume Controller is identified in the DS4300 configuration as a host.
- ◆ In the SAN Volume Controller configuration, the same gateway is identified as a host.
- ◆ The gateway is a N5200 2864 or N5500 2865 model.


The zoning is set up to provide two paths, and only two paths, to

- ◆ Each gateway VDisk on the SAN Volume Controller
- ◆ Each LUN on the DS4300 that is provided for the gateway to access directly (that is, not through SAN Volume Controller).

Note

From the perspective of the gateway, SAN Volume Controller and the DS4300 (whose storage the gateway accesses directly) are two different types storage arrays. LUNs from these two different storage array types cannot be mixed in the same Data ONTAP aggregate.

The following illustration shows zoning for this scenario. The same information is summarized in a table following the illustration.




The following table summarizes the zoning for this scenario.

Zone	Switch	The zone includes this Gateway FC initiator port	The zone includes this SAN Volume Controller node and port		The zone includes this storage array, controller, and port	
Zoning pair						
1	1	—	Node 1	port 1	DS4300	C1, port 1
			Node 1	port 3	DS4300	C2, port 1
			Node 2	port 1		
			Node 2	port 3		
2	2	—	Node 1	port 2	DS4300	C1, port 2
			Node 1	port 4	DS4300	C2, port 2
			Node 2	port 2		
			Node 2	port 4		
Zoning pair						
3	1	0a	Node 1	port 1	—	
4	2	0c	Node 2	port 4	—	
Zoning pair						
5	1	0a	—		DS4300	C1, port 1
6	2	0c	—		DS4300	C2, port 2

Scenario 3: SAN Volume Controller provides storage only for storage arrays the gateway does not support


In this example, SAN Volume Controller provides storage only from the storage arrays that the gateway does not support. The storage arrays that the gateway supports present LUNs to the gateways directly; they are configured with the gateways as hosts.



Zoning for this scenario: To discuss zoning for this example, assume the following:

- ◆ SAN Volume Controller, which has one I/O group, obtains LUNs from a storage array that the gateway cannot support itself.
- ◆ The gateway obtains LUNs directly from an HDS 9980V subsystem.
- ◆ The gateway is a N5200 2864 or N5500 2865 model.

The following illustration shows zoning for this scenario. The same information is summarized in a table following the illustration.




-  Ports in the storage zones
-  Ports in the gateway 1 host zone with SAN Volume Controller
-  Ports in the zones for gateway assessing the HDS 9980V directly
- - - Indicates that the ports at each end of the dashed line are in the same zone
- C1, C2 Represents alternate controllers (clusters) on the storage subsystem

The following table summarizes the zoning for this scenario.

Zone	Switch	The zone includes this Gateway FC initiator port	The zone includes this SAN Volume Controller node and port		The zone includes this storage array and port	
Zoning pair						
1	1	—	Node 1	port 1	EMC	C1, port 1
			Node 1	port 3	EMC	C2, port 1
			Node 2	port 1		
			Node 2	port 3		
2	2	—	Node 1	port 2	EMC	C1, port 2
			Node 1	port 4	EMC	C2, port 2
			Node 2	port 2		
			Node 2	port 4		
Zoning pair						
3	1	0a	Node 1	port 1	—	
4	2	0c	Node 2	port 4	—	
Zoning pair						
5	1	0a	—		HDS 9980V	C1, port 1
6	2	0c	—		HDS 9980V	C2, port 2

Scenario 4: the gateway and a gateway client are both attached to SAN Volume Controller


The following illustration shows a configuration in which the gateway and a gateway client both attach to SAN Volume Controller. This configuration is possible if you do not map the same VDisk to both the gateway and the gateway client.



Zoning for this scenario: To discuss zoning for this scenario, assume the following:

- ◆ On the SAN Volume Controller, different VDisks are mapped to the gateway and the gateway client.
- ◆ There is only one SAN Volume Controller I/O group.
- ◆ One storage array, a DS4300, provides storage for SAN Volume Controller.
- ◆ The gateway is a N5200 2864 or N5500 2865 model.

The following illustration shows zoning for this scenario. The same information is summarized in a table following the illustration.



-  Ports in the storage zones
-  Ports in the gateway 1 zones with SAN Volume Controller
-  Ports in the gateway client zones with SAN Volume Controller
- - - Indicates that the ports at each end of the dashed line are in the same zone
- C1, C2 Represents alternate controllers (clusters) on the storage subsystems

Unlike the a gateway, the gateway client is not restricted to two paths. Each switch fans out from one port on the client to two ports in the nodes in the SAN Volume Controller I/O group.

The following table summarizes the zoning for this scenario.

Zone	Switch	The zone includes this Gateway FC initiator port		The zone includes this SAN Volume Controller node and port		The zone includes this storage array and port	
Zoning pair							
1	1	Gateway 1	port 0a	I/O group 0	Node 1, port 1	—	
2	2	Gateway 1	port 0c	I/O group 0	Node 2, port 4	—	
Zoning pair							
3	1	Gateway client	port 1	I/O group 0	Node 1, port 1 Node 2, port 1	—	
4	2	Gateway client	port 2	I/O group 0	Node 1, port 4 Node 2, port 4	—	
Zoning pair for the storage zone							
5	1	—		I/O group 0	Node 1, ports 1 and 3 Node 2, ports 1 and 3	DS4300 DS4300	C1, port 1 C2, port 1
6	2	—		I/O group 0	Node 1- ports 2 and 4 Node 2- ports 2 and 4	DS4300 DS4300	C1, port 2 C2, port 2

**Scenario 5:
active/active
configuration with
multiple SAN
Volume Controller
I/O groups**

This scenario includes a gateway active/active configuration and two I/O groups in the SAN Volume Controller cluster.

Zoning for this scenario: Assume the following for this scenario:


- ◆ The gateway active/active configuration accesses all storage through SAN Volume Controller.
- ◆ The goal for this scenario is to distribute the gateway I/O load over the ports. Distributing the I/O load over the ports might be a performance advantage because multiple gateways would not be competing for bandwidth over a single SAN Volume Controller port. The nodes in the gateway active/active configuration are still able to access the VDisks of the partner node if one gateway node fails.
- ◆ The gateway is a N5200 2864-G20 or a N5500 2865-G20 model.

Zoning is set up so that

- ◆ All FC initiator ports on the two gateways are used to connect to SAN Volume Controller for redundancy
Because there are two I/O groups, you can use two FC initiator port pairs on each gateway.
- ◆ Each gateway FC initiator port goes to a single SAN Volume Controller port
- ◆ Each SAN Volume Controller port is servicing only one gatewayport.

Whether this type of zoning is possible for your organization depends on whether you have enough free SAN Volume Controller ports. If you add more gateway active/active configurations to SAN Volume Controller, you might need to share the ports, which reduces the performance advantage. However, the advantage of the I/O distribution over the SAN Volume Controller ports remains.

The following illustration shows zoning for this scenario. The same information is summarized in a table following the illustration.



-  Ports in the storage zones
-  Ports in the gateway 1 host zones with SAN Volume Controller
-  Ports in the gateway 2 zones with SAN Volume Controller
- - - Indicates that the ports at each end of the dashed line are in the same zone
- C1, C2 Represents alternate controllers (clusters) on the storage subsystem

The following table summarizes the zoning for this scenario.

Zone	Switch	The zone includes this Gateway FC initiator port		The zone includes this SAN Volume Controller node and port		The zone includes this storage array and port	
Zoning pair							
1	1	Gateway 1	port 0a	I/O group 0	Node 1, port 1	—	
2	2	Gateway 1	port 0c	I/O group 0	Node 2, port 4	—	
Zoning pair							
3	1	Gateway 1	port 0b	I/O group 1	Node 1, port 1	—	
4	2	Gateway 1	port 0d	I/O group 1	Node 2, port 4	—	
Zoning pair							
5	1	Gateway 2	port 0a	I/O group 0	Node 2, port 1	—	
6	2	Gateway 2	port 0c	I/O group 0	Node 1, port 4	—	
Zoning pair							
7	1	Gateway 2	port 0b	I/O group 1	Node 2, port 1	—	
8	2	Gateway 2	port 0d	I/O group 1	Node 1, port 4	—	
Zoning pair for the storage zone							
9	1	—		I/O group 0	Node 1, ports 1 and 3	DS4300	C1, port 1
				I/O group 1	Node 2, ports 1 and 3		
				I/O group 1	Node 1, ports 1 and 3	DS4300	C2, port 1
					Node 2, ports 1 and 3		

Zone	Switch	The zone includes this Gateway FC initiator port	The zone includes this SAN Volume Controller node and port		The zone includes this storage array and port	
10	2	—	I/O group 0	Node 1-ports 2 and 4 Node 2-ports 2 and 4	DS4300 DS4300	C1, port 2 C2, port 2
			I/O group 1	Node 1-ports 2 and 4 Node 2-ports 2 and 4		

Prerequisites for configuring SAN Volume Controller to provide storage to the gateway

This overview assumes that

- ◆ The LUNs that are to be used by SAN Volume Controller were created on the back-end storage arrays.
- ◆ The storage arrays have presented LUNs to SAN Volume Controller.
- ◆ MDisk groups were created in the SAN Volume Controller configuration.
- ◆ Zoning is configured so that only two paths to a VDisk are available to a gateway.

Configuration overview

The following table provides an overview of what you need to do to configure storage on SAN Volume Controller for gateways. For details about how to configure SAN Volume Controller, see your SAN Volume Controller documentation. For information about how to connect your gateway to SAN Volume Controller, see *Connecting Your Gateway*.

Stage	Process
1	Determine the storage capacity that you need for the gateways and the VDIsks that you want to assign to the gateways.
2	Confirm that the firmware on SAN Volume Controller meets the minimum version level required by the gateways and SAN Volume Controller. See the gateway <i>Support Matrix</i> for version requirements.
3	<p>Create the VDIsks on the SAN Volume Controller that you plan to use for the gateways.</p> <p>Ensure that the VDIsks that you create meet the minimum size required by the gateway and that they do not exceed the maximum LUN size described in Appendix A, “Array LUN Sizing for SAN Volume Controller,” on page 39.</p> <p>Attention</p> <p>If you specify that SAN Volume Controller is to format the VDIsks at the time you are creating the VDIsks on SAN Volume Controller, be sure to wait for this process to complete before you map the VDIsks that SAN Volume Controller is formatting for the gateway.</p>

Stage	Process
4	<p>Gather the following information that you will need to create a host object for the gateways and to complete the VDisk-to-host mapping:</p> <ul style="list-style-type: none"> ◆ The name you want to assign to the host object It is recommended that you assign a name that enables you to easily recognize that the host object is for the gateways. ◆ If necessary, the WWNs of the FC initiator ports of the gateways Typically, SAN Volume Controller automatically finds WWNs for the host systems after the hosts are connected. If you want to configure the gateway host object before connecting your gateway to SAN Volume Controller, see the gateway <i>Integration Guide for IBM Storage</i> for information about obtaining gateway WWNs. ◆ The names of the VDIsks that you want to assign to the gateways.
5	<p>Create a single host object for all gateways so that all gateways in the neighborhood can access the VDIsks that you plan to assign to the gateways.</p> <p>Add the WWNs for all the FC initiator ports of the gateways in the gateway neighborhood.</p>
6	<p>Map the VDIsks for the gateways to the gateway host object (that is, the VDIsks that you want all gateways in the gateway neighborhood to see as LUNs).</p>

When you are ready to set up and configure Data ONTAP

You can begin Data ONTAP setup and configuration any time after assigning VDIsks to the gateways and connecting SAN Volume Controller and the gateway together. The gateway *Software Setup, Installation, and Management Guide* describes how to set up and configure Data ONTAP.

If the gateway is...	Then...
Powered on	To start the setup program, enter <code>bye</code> at the boot prompt (CFE, <code>ok</code> , or <code>LOADER</code> depending on the gateway model).

If the gateway is...	Then...
Not powered on	See the instructions in the gateway <i>Software Setup, Installation, and Management Guide</i> for how to power up the gateway.

Gateway minimum and maximum array LUN sizes

The size of the LUNs that you can create on the storage array is limited by the minimum and maximum LUN sizes that the gateway supports. The Data ONTAP definition of a GB might not match the definition of a GB for your storage array.

The Data ONTAP definition of a gigabyte (GB) is that one GB is equal to $1000 \times 1024 \times 1024$ bytes.

Attention

The minimum array LUN size shown in the following table does not apply to the array LUN 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 *Installation Requirements, Quick Start, and Reference Guide*. 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.

The following table shows the gateway minimum and maximum array LUN sizes.

Data ONTAP release	Minimum array LUN size	Maximum array LUN size
7.2.4	1 GB	1 TB Calculated as: $1000 * 1000 * 1024 * 1024 = 1,048,576,000,000$ bytes
7.2.3	1 GB	750 GB Calculated as: $750 * 1000 * 1024 * 1024 = 786,432,000,000$ bytes
Prior to 7.2.3	1 GB	500 GB Calculated as: $500 * 1000 * 1024 * 1024 = 524,288,000,000$ bytes

Maximum array LUN value usable with SAN Volume Controller

IBM calculates a units of measure for SAN Volume Controller differently than Data ONTAP. The maximum usable values shown in this section are based on the assumption that the units of measurement for SAN Volume Controller are calculated as follows.

Unit	Formula for calculating..
GB	1024 x 1024 x 1024 bytes
MB	1024 x 1024 bytes
KB	1024 bytes

If you plan to use a large-sized LUN that is close to the maximum capacity that Data ONTAP supports, ensure that the size you specify does not exceed the size shown in the “Maximum usable value” column in the following tables.

Note

Storage arrays vary as to how you can specify LUN size (that is, in GB, MB, or 512-byte blocks).

See the gateway *Planning Guide* for guidelines about the implications of different size array LUNs on Data ONTAP storage.

Values for Data ONTAP 7.2.4:

If you are specifying in	Maximum usable value
GB	976 GB
MB	975,000 MB
512-byte blocks	2,047,500,000 512-user blocks

Values for Data ONTAP 7.2.3:

If you are specifying in	Maximum usable value
GB	732 GB
MB	749,000 MB
512-byte block	1,535,500,000 512-user blocks

Values for Data ONTAP 7.2.2 and earlier:

If you are specifying an	Maximum usable value
GB	488.281 GB
MB	500,000 MB
512-user block	1,024,000,000 512-user blocks

Index

A

- active/active configurations
 - defined xii
 - supported deployments with SAN Volume Controller 31
- aggregates, rules for assigning VDisks to 10
- array LUN size supported by the gateway 39

C

- cluster (SAN Volume Controller)
 - defined 3, 4, 5
 - number gateway supports 7
- configurations, SAN Volume Controller 16

E

- extent, defined 3

F

- firmware
 - level supported for SAN Volume Controller 8

G

- gigabyte, defined 39

H

- host object in SAN Volume Controller
 - gateway 9, 36
 - name for 36
- host zones, defined 12
- hosts (SAN Volume Controller)
 - defined 3
 - how storage is presented 7

I

- I/O group
 - defined 3
 - number of nodes in 3

- requirements for gateway zoning 14
- input/output group. *See* I/O group

L

- LUNs (array)
 - access by gateways
 - number of paths supported 15, 19
 - sizing guidelines
 - size 39
 - sizes supported by the gateway 39

M

- managed disk group, defined 3
- managed disk. *See* MDisk
- MDG. *See* managed disk group
- MDisk, defined 3, 5
- MetroCluster, SAN Volume Controller restriction 16

N

- nodes (SAN Volume Controller) 13
 - defined 3, 4, 5
 - zoning requirements 12

P

- paths to LUNs
 - number Data ONTAP supports 15, 19

S

- SAN Volume Controller 1
 - changing mapping, implications 10
 - cluster, defined 3, 4, 5
 - configuration overview 35
 - configurations gateways support 16
 - discovery of array LUNs 6
 - extent, defined 3
 - firmware level 8
 - gateway interaction with 2

- hosts, defined 3
- I/O group, defined 3
- identifying gateways as hosts 9
- LUNs presented to by storage arrays 6
- major advantage with the gateway 5
- managed disk group, defined 3
- managed disk. *See* MDisk
- MDisk, defined 3, 5
- MetroCluster restriction 16
- moving between I/O groups, implications 10
- nodes, defined 3, 4, 5
- number of clusters the gateway supports 7
- overview 5
- redundancy with 5
- relationship to hosts 5
- relationship to storage arrays 5
- storage presentation to hosts 7
- supported configurations 16, 21, 25, 28, 31
- switch requirements 8
- terminology 3
- unmapping, implications 10
- VDisks
 - creating 7, 35
 - defined 4
 - mapping 7
 - maximum with gateway 9
 - operations gateway does not support 10
 - operations gateway supports 10
 - size restriction 9
- zoning
 - host zones defined 12
 - requirements for hosts 13
 - requirements for SAN Volume Controller nodes 12
 - requirements for storage arrays 13
 - requirements for the gateway 14
 - scenario 17, 22
 - storage zones defined 12
 - why required 11
- storage zones, defined 12
- switch zoning. *See* zoning

- switches
 - requirements with SAN Volume Controller 8

T

- terminology, SAN Volume Controller 3

V

VDisks

- changing mapping, implications 10
- creating 7, 35
- defined 4
- mapping 7
- maximum with gateway 9
- moving between I/O groups, implications 10
- operations gateway does not support 10
- operations gateway supports 10
- rules for assigning to aggregates 10
- size restriction 9
- types gateway supports 9
- unmapping, implications 10

- virtual disk. *See* VDisk

Z

zoning

- I/O group
 - requirement for gateway zoning 14
- overview 11
- requirements 11
- SAN Volume Controller
 - requirements 11
 - requirements for hosts 13
 - requirements for SAN Volume Controller nodes 12
 - requirements for storage arrays 13
 - requirements for the gateway 14
 - scenario 17, 22, 25, 28, 31
 - switches 11

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

IBM System Storage N series
Gateway Implementation Guide for IBM SAN Volume Controller Storage

Publication No. GC53-1151-00

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

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

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

Comments:

Thank you for your support.

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

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

Name

Address

Company or Organization

Phone No.

E-mail address



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

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



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



NA 210-03926_A0, Printed in USA

GC53-1151-00

