

IBM System Storage N series



FlexArray Virtualization Implementation Guide for Third-Party Storage

Contents

Preface	7
About this guide	7
Supported features	7
Websites	8
Getting information, help, and service	8
Before you call	8
Using the documentation	9
Hardware service and support	9
Firmware updates	9
How to send your comments	9
How this guide fits into your implementation planning	11
Where to find interoperability and limits information for	
configurations with storage arrays	13
Interoperability information about support for storage arrays	13
Limits information for configurations with storage arrays	14
Use of storage array advanced features	15
Systems running Data ONTAP that can use array LUNs on storage	
arrays	16
General configuration guidelines for all storage arrays	17
Requirement for 8-Gb array port initialization	17
EMC CLARiiON and VNX storage arrays	18
Required parameters for EMC CLARiiON and VNX storage arrays with	
systems running Data ONTAP	18
General configuration guidelines	18
How EMC CLARiiON and VNX storage arrays control access to data	18
Limitation on array LUN numbering with EMC CLARiiON and VNX	
arrays	19
Guidelines for enabling ALUA on EMC CLARiiON and VNX storage	
arrays	19
EMC CLARiiON storage array families	19
EMC VNX storage array families	20
EMC Symmetrix storage arrays	22

Required parameters for EMC Symmetrix storage arrays with systems running Data ONTAP	22
General configuration guidelines	23
Requirements for implementing LUN security on EMC Symmetrix storage arrays	23
Caution about using the VCMDB LUN	24
Caution about using the ACLX LUN	24
Restriction on using gatekeeper LUNs	24
EMC Symmetrix storage array families	24
Fujitsu ETERNUS storage arrays	26
Requirements for configuring Fujitsu ETERNUS storage arrays with systems running Data ONTAP	26
Host Response pattern requirement for Fujitsu storage arrays	26
Required host configuration parameters for Fujitsu DX8xxx storage arrays	26
Required host configuration parameters for Fujitsu DX4xx storage arrays	27
General configuration guidelines	28
Fujitsu ETERNUS storage array families	28
Hitachi storage arrays	30
Required parameters for Hitachi storage arrays with systems running Data ONTAP	30
General configuration guidelines	31
Requirements for creating LUNs in dynamic pools	31
Requirements for formatting array LUNs	31
Requirements for using disks external to the storage array	32
Guidelines for configuring ports on Hitachi storage arrays	32
Hitachi storage array families	33
HP EVA storage arrays	35
Required parameters for HP EVA arrays with systems running Data ONTAP	35
General configuration guidelines	35
Requirements for configuring the HP EVA Storage Management Server	35
Preparing HP EVA storage arrays for NDU	35
ALUA support for HP EVA storage arrays	36
HP EVA storage array families	37
HP XP storage arrays	38

Required parameters for HP XP storage arrays with systems running Data	
ONTAP	38
General configuration guidelines	39
Requirements for creating LUNs in dynamic pools	39
Requirements for formatting array LUNs	39
Requirements for using disks external to the storage array	39
Guidelines for configuring ports on HP XP storage arrays	40
HP XP storage array families	40
IBM DS storage arrays	41
Required settings for IBM DS storage arrays with systems running Data	
ONTAP	41
General configuration guidelines	42
Requirements for volume groups on DS8xxx storage arrays	42
Requirements for direct-attached configurations with all IBM DS arrays	42
IBM DS storage array families	42
IBM XIV storage arrays	44
Requirements for configuring IBM XIV arrays with systems running Data	
ONTAP	44
IBM XIV storage array families	44
3PAR storage arrays	46
Required parameters for 3PAR arrays with systems running Data ONTAP	46
Workaround for the unknown device (L254/SES) problem with 3PAR	
storage arrays	47
General configuration guidelines	47
3PAR storage array families	47
Supported fabric-attached configurations	49
Stand-alone basic configuration	49
Stand-alone with two 2-port array LUN groups	50
Single 4-port array LUN group	50
Two 4-port array LUN groups	51
Eight-port array LUN group configuration	52
Data ONTAP FC initiator port attaching to multiple targets ports	54
Supported direct-attached configurations	56
Two-port direct-attached stand-alone configurations	56
Four-port direct-attached HA pair configuration	57
Terminology comparison between storage array vendors	58

Copyright information 62
Trademark information 63
Index 66

Preface

About this guide

This document applies to IBM N series systems running Data ONTAP. If the term *7-Mode* is used in the document, it refers to Data ONTAP operating in 7-Mode, which has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families. If the terms *Cluster-Mode* or *clustered Data ONTAP* are used in the document, they refer to the Data ONTAP features and functionality designed for clusters, which are different from 7-Mode and prior Data ONTAP 7.1, 7.2, and 7.3 release families.

Starting with Data ONTAP 8.2.1, the capability of using LUNs on a storage array, formerly identified as *gateway functionality*, has a new name—*Data ONTAP FlexArray Virtualization Software*.

The term *gateway* describes IBM N series storage systems prior to Data ONTAP 8.2.1 that have been ordered with the FlexArray Virtualization Software functionality (formerly known as 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.

Although the gateway product name has been retired as of Data ONTAP 8.2.1, Data ONTAP can access LUNs on all the same storage arrays as before, with the same functionality that the gateway product provided.

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. New filer storage systems released with Data ONTAP 8.2.1 and later can support third-party disk storage systems if the appropriate license is installed.

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 (accessed and navigated as described in [Websites](#) on page 8).

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 N series support website (accessed and navigated as described in [Websites](#) on page 8) 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 N series support website (accessed and navigated as described in [Websites](#) on page 8).

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 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, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Note: If you do not see new firmware updates on the 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 email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title

10 | Implementation Guide for Third-Party Storage

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

How this guide fits into your implementation planning

When planning a deployment with systems running Data ONTAP and storage arrays, it is important to review the general requirements for setting up all storage arrays to work with Data ONTAP and also the specific requirements for your vendor and storage array vendor, which are discussed in this guide.

Systems running Data ONTAP can be deployed with a variety of storage arrays. Data ONTAP requirements for working with individual storage arrays can vary by vendor and storage array. This guide provides information about the storage arrays that Data ONTAP supports at the time of publication. Not all models described in this guide are supported in all Data ONTAP releases. To determine which storage array models are supported in a particular Data ONTAP release, see the Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Note: The *Interoperability Matrix* is the final authority on information such as storage arrays, firmware, switches, and features that Data ONTAP supports.

What to read before this guide

When planning for setting up your configuration to work with storage arrays, you should start by reading the *FlexArray Virtualization Installation Requirements and Reference Guide*, which provides the general requirements for setting up Data ONTAP to work with storage arrays. For example, it contains general guidelines about how to configure array LUNs to work with Data ONTAP and how to set up paths to array LUNs. After you understand the requirements that apply to all vendors, you should read this guide, which provides the additional information that is specific to your vendor.

How this guide is organized

The information in this guide is organized in separate chapters for each storage array vendor. Each chapter includes information about the following:

- The requirements for parameters that must be set for each supported storage array from your vendor to be able to work with Data ONTAP
- Additional configuration guidelines, such as restrictions specific to an aspect of your vendor's storage array configuration
- The configurations that are supported for the storage arrays from your vendor

The supported configurations shown for a vendor might not apply to all storage array models from the vendor. Any exceptions for individual models from a vendor are shown in the Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Where to find out about planning for and setting up Data ONTAP features

The FlexArray Virtualization guides provide information about how to set up Data ONTAP to work with storage arrays and how to set up the storage arrays to work with Data ONTAP. You should also refer to the Data ONTAP document library, which provides information about Data ONTAP features. The Data ONTAP guides also contain some procedures for installing and configuring your system to work with array LUNs.

Where to find interoperability and limits information for configurations with storage arrays

When planning your configuration to use systems running Data ONTAP with storage arrays, you need to check sources in addition to product documentation for interoperability details information.

Tools available on the N series support website (accessed and navigated as described in [Websites](#) on page 8) provide, in a central location, specific information about which features, configurations, and storage array models are supported in particular releases.

Interoperability information about support for storage arrays

Not all Data ONTAP releases support the same features, configurations, system models, and storage array models. During your deployment planning, you must check Data ONTAP support information to ensure that your deployment conforms to Data ONTAP hardware and software requirements for all systems in the deployment.

Previously all support information used to set up deployments with storage arrays was included in the Gateway Interoperability Matrix. The information is divided into different tools, as shown in the following table:

For information about...	You should look here...
Data ONTAP working with devices, including the following: - Supported storage arrays and storage array firmware - Supported switches and switch firmware - Whether your storage array supports nondisruptive (live) upgrade of the storage array firmware - Whether a MetroCluster configuration is supported with your storage array	Interoperability Matrix at the N series support website (accessed and navigated as described in Websites on page 8) Note: The <i>Interoperability Matrix</i> shows the Brocade and Cisco switches that are supported. You can find guides for configuring specific vendor switches on the N series support website (accessed and navigated as described in Websites on page 8).

For information about...	You should look here...
Data ONTAP limits for releases and platforms, including the following: • Minimum and maximum array LUN sizes, including the minimum array LUN size for the root volume and spare core array LUNs • Minimum aggregate size for aggregates with array LUNs • Supported block size • Minimum and maximum capacity • Neighborhood limits	<i>IBM System Storage N series Introduction and Planning Guide</i> at the N series support website (accessed and navigated as described in Websites on page 8)
What is supported for specific arrays, including supported configurations	<i>FlexArray Virtualization Implementation Guide for Third-Party Storage</i>

Limits information for configurations with storage arrays

The *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in [Websites](#) on page 8) contains information about limits that you need to consider when planning for a configuration with storage arrays.

The *IBM System Storage N series Introduction and Planning Guide* includes the following limits that do not apply to native disks:

- Minimum and maximum array LUN size that Data ONTAP supports
- Minimum size for the array LUN for the root volume
- Spare core array LUN minimum size
- Limits for RAID groups with array LUNs
- Minimum aggregate size for an aggregate of array LUNs
- Maximum number of array LUNs and disks combined, per platform

The *IBM System Storage N series Introduction and Planning Guide* also contains storage limits that are the same for native disks and array LUNs.

Note: Limits for gateways were previously provided in the *Gateway Limits Reference for Third-Party Storage* and the Gateway Interoperability Matrix.

Use of storage array advanced features

Data ONTAP does not support storage array advanced features unless otherwise specified in the Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Systems running Data ONTAP that can use array LUNs on storage arrays

Gateway systems and new filer platforms released in Data ONTAP 8.2.1 and later can use array LUNs if the proper license is installed. In discussions in the Data ONTAP and FlexArray Virtualization documentation, these systems are collectively referred to as systems running Data ONTAP when it is necessary to make it clear which information applies to them and what information applies to storage arrays.

Note: Starting with Data ONTAP 8.2.1, the capability of using LUNs on a storage array, formerly identified as gateway functionality, has a new name—*Data ONTAP FlexArray Virtualization Software*. Although the gateway product name has been retired as of 8.2.1, Data ONTAP can access LUNs on all the same storage arrays as before, with the same functionality that the gateway product provided.

Systems prior to Data ONTAP 8.2.1 that can use array LUNs

The only systems released prior to Data ONTAP 8.2.1 that can use array LUNs are gateways. A gateway is an open storage controller that virtualizes storage from storage array vendors, native disks, or both into a single heterogeneous storage pool.

Note: Almost all Data ONTAP platforms released prior to Data ONTAP 8.2.1 were released with filer and gateway equivalent models (for example, an N7950T filer and an N7950T gateway). (For a few systems, there were no gateway equivalent models.) Although both types of models could access native disks, only the gateways could attach to storage arrays.

Systems in Data ONTAP 8.2.1 and later that can use array LUNs

Starting with Data ONTAP 8.2.1, the model for how platforms are released and the storage they can use changes. Attaching to storage arrays is no longer limited to gateways.

Starting with Data ONTAP 8.2.1, all new platforms are being released as a single hardware model. This single hardware model has an NS prefix; there are no longer separate gateway and filer models for new platforms. If the V_StorageAttach license package is installed on a new filer model, it can attach to storage arrays. (This is the same license required on a gateway.)

Important: Filers released prior to Data ONTAP 8.2.1 cannot use LUNs on storage arrays, even if they are upgraded to 8.2.1; only the gateway equivalent of a platform can use array LUNs.

General configuration guidelines for all storage arrays

This section includes general configuration guidelines for any of the storage arrays supported by Data ONTAP.

Requirement for 8-Gb array port initialization

In a typical Fibre Channel (FC) fabric, storage array ports initialize as N-Ports (node ports) and their corresponding switch ports initialize as F-Ports (fabric ports). In some cases, these F-ports might go offline and resume operation as L-Ports (loop ports), resulting in fabric errors.

This behavior is caused by topology auto-negotiation issues in the Brocade switch.

To ensure that switch ports always initialize as F-ports, you must run the `portcfggport <portnumber> 1` command on Brocade switch ports that are connected to storage arrays. You must run this command when you start configuring the system to avoid this problem, and you must fix the problem if it occurs.

To promote efficiency on the fabric as a whole, you can also use this workaround on any Brocade ports attached to Data ONTAP initiator ports.

EMC CLARiiON and VNX storage arrays

You must meet specific requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on your storage arrays and deploying only supported configurations.

Required parameters for EMC CLARiiON and VNX storage arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

Host configuration parameters that must be set on the storage array are listed in the following table.

Parameter	Setting
Initiator type	Clariion Open or Clariion VNX
Array Com Path	Enabled
Failover mode	1
Unit serial number	LUN
Host Name	User-supplied host name and port number.
IP address	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.

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

How EMC CLARiiON and VNX storage arrays control access to data

EMC CLARiiON and VNX arrays use Storage Groups to control access to data. A Storage Group is one or more LUNs within the storage array that can be accessed only by the host or hosts that you

associate with the array LUNs. A host cannot access or modify data in any array LUNs that are not part of its Storage Group.

Data ONTAP supports multiple Storage Groups if these rules are followed:

- Switch zoning must define which target ports the FC initiator ports on the system running Data ONTAP use to access each array LUN group.
- LUN masking must be used to restrict host access to array LUNs.
- Storage Groups must define which array LUN groups are presented to each FC initiator port.
- One FC initiator port pair on each system running Data ONTAP is required for each array LUN group.

If array LUN neighborhoods are used, the gateways in the same neighborhood must be in the same Storage Group.

Limitation on array LUN numbering with EMC CLARiiON and VNX arrays

EMC CLARiiON (CX) and VNX storage arrays only support array LUN numbering from 0 through 255. Array LUNs that are numbered beyond that range are not visible to Data ONTAP.

Guidelines for enabling ALUA on EMC CLARiiON and VNX storage arrays

Starting with clustered Data ONTAP 8.2, support has been added for ALUA (Asymmetric Logical Unit Access) with CLARiiON and VNX arrays.

By default, support for ALUA is enabled in clustered Data ONTAP. But to use it, ALUA must be enabled on the storage array.

You should enable ALUA only on new configurations. You should not enable ALUA in an existing configuration.

If you are setting up ALUA on the storage array, you need to ensure that all hosts in a Storage Group are in the same failover mode, that is failover mode 4 (ALUA).

EMC CLARiiON storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

Note: CX storage arrays use asymmetric active-active (ALUA) or active-passive failover mode; they are not active-active storage arrays.

The following list shows CLARiiON storage arrays organized by family:

- Family 1
 - CX300
 - CX500
 - CX700
- Family 2
 - CX3-20
 - CX3-40
 - CX3-80
- Family 3
 - CX4-120
 - CX4-240
 - CX4-480
 - CX4-960

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

EMC VNX storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

Note: VNX storage arrays use asymmetric active-active (ALUA) or active-passive failover mode; they are not active-active storage arrays.

The following list shows the supported storage arrays, which are all in the same family:

- Family 1
 - VNX5100
 - VNX5300
 - VNX5500
 - VNX5700
 - VNX7500

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

EMC Symmetrix storage arrays

You must meet specific requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on your storage arrays and deploying only supported configurations.

Required parameters for EMC Symmetrix storage arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

Required host channel director port configuration parameters

The host channel director port configuration parameters that must be set on the storage array are shown in the following table:

Parameter (names might differ between GUI and CLI)	Setting
Common SN (Common Serial Number or C-bit parameter)	Enable
PP (Point-to-Point parameter)	Enable
SC3 (SCSI-3)	Enable
SPC-2 (SCS2_Protocol_version, SCSI Primary Command 2 parameter, or “Allow inquiry data to be compiled to the standard”)	Enable
UWN (Unique Worldwide Name)	Enable
Volume Set Addressing	Disable

The Volume Set Addressing parameter must be set the same way on all channel director ports to which the LUN is mapped. If the settings are different, Data ONTAP reports this as a LUN ID mismatch in storage errors show output and in an EMS message. See the *FlexArray Virtualization Installation Requirements and Reference Guide* for information about troubleshooting this error.

Required host configuration parameters

The settings for host configuration parameters are shown in the following table:

Parameter	Setting
Host Type Format	Server or Open systems
SCSI3_persist_reserv (Persistent Reservations)	On

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

Requirements for implementing LUN security on EMC Symmetrix storage arrays

You use LUN security to eliminate the possibility of a host writing data to a LUN that is not owned by that host.

To eliminate the possibility of a non-Data ONTAP host overwriting EMC Symmetrix (DMX and VMAX) array LUNs owned by a system running Data ONTAP or vice versa, you must present the Symmetrix logical devices through the host (channel) director ports in one of the following ways:

- Present only the Symmetrix logical devices for Data ONTAP on specific Symmetrix host (channel) director ports that are dedicated to Data ONTAP use.
If ports cannot be dedicated to Data ONTAP, you should confirm that all other hosts using those ports are compatible with Data ONTAP requirements. This is because each host connected to the Symmetrix array has requirements for different port attribute settings. Sharing ports between multiple hosts that are connected to the Symmetrix storage array might result in an impossible-to-implement configuration.
- For DMX storage arrays, use LUN masking to associate all the DMX LUNs with the FC initiator ports on the system running Data ONTAP and with the host (channel) director ports that are dedicated to Data ONTAP use.
- For VMAX storage arrays, create masking views for required mapping and masking by creating port groups, Storage Groups, and initiator groups.
To achieve this, you must first enable the ACLX port attribute on the VMAX storage array ports.

Attention: Do not present the VCMDB LUN 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.

Caution about using the VCMDB LUN

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

If the VCMDB LUN is mapped to a system running Data ONTAP, Data ONTAP periodically logs a message that the VCMDB LUN is less than the minimum size required, and it marks the VCMDB LUN as failed. The system running Data ONTAP continues to function normally after logging this error message, but it cannot use the LUN.

A VCMDB LUN should be unmapped from a system running Data ONTAP.

Caution about using the ACLX LUN

On VMAX arrays, the ACLX (Access Control Logix) LUN is created during initialization if the customer requests that ACLX be used. The ACLX LUN is not a storage LUN, so it should not be mapped to Data ONTAP.

If the ACLX LUN is mapped to a system running Data ONTAP, Data ONTAP logs a message that the ACLX LUN is less than the minimum size required, and it marks the LUN as failed. The system running Data ONTAP continues to function normally after logging this error message, but it cannot use the LUN.

The ACLX LUN should be unmapped from the front-end director ports on the system running Data ONTAP.

Restriction on using gatekeeper LUNs

If a gatekeeper logical device (LUN) is presented, you must not map it to the system running Data ONTAP. Systems running Data ONTAP cannot use gatekeeper LUNs. A gatekeeper LUN is a Symmetrix logical device through which SYMAPI or the ControlCenter agent communicates with the storage array.

EMC Symmetrix storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the supported arrays organized by families:

- Family1
 - VMAX SE
 - VMAX 40k
 - VMAX/VMAX 20k
 - VMAXe/VMAX 10k
- Family 2
 - DMX 3
 - DMX 4

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

Fujitsu ETERNUS storage arrays

You must meet specific requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on your storage arrays and deploying only supported configurations.

Requirements for configuring Fujitsu ETERNUS storage arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

Host Response pattern requirement for Fujitsu storage arrays

For your system running Data ONTAP to work with Fujitsu storage arrays, you need to create a new Host Response pattern and set the parameters required by the system running Data ONTAP.

Fujitsu provides Host Response pattern #00, which is a read-only version of the specification that you edit to match the requirements of your storage array model.

Required host configuration parameters for Fujitsu DX8xxx storage arrays

You must set certain configuration parameters for the storage array to be able to communicate with the system running Data ONTAP.

Host configuration parameters that must be set on the storage array are listed in the following table:

Parameter	Setting
Host Response Name	Name (name of the system running Data ONTAP)
Command Time-out Interval	Standard (25 seconds)
Data Transfer Timer	10 Seconds (Default)
Load Balance Response	Unit Attention
Byte-0 Inquiry Response	No Conversions (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)

Parameter	Setting
Asymmetric / Symmetric Logical Unit Access	Active/Active (Default)
Host Specific Mode	Normal Mode (Default)
LUN Mapping Changes	No Report (Default)
LUN Capacity Expansion	No Report (Default)
Vendor Unique Sense Code	No report (Default)
Sense Code Conversion Pattern	No Conversion (Default)
The following additional parameters must be set if you are using firmware version V10L50 or later:	
iSCSI Discovery Reply	All - Reply All Ports (Default)
TPGS Mode	Enable (Default)
TPG Referrals	Disable (Default)
iSCSI Reservation Range	CA Port (Default)

Required host configuration parameters for Fujitsu DX4xx storage arrays

You must set certain configuration parameters for the storage array to be able to communicate with the system running Data ONTAP.

Host configuration parameters that must be set on the storage array are listed in the following table:

Parameter	Setting
Host Response Name	Name (name of the system running Data ONTAP)
Byte-0 of Inquiry response	No conversion (Default)
Inquiry VPD ID Type	Type 1 + Type 3
Inquiry Standard Data Version	Default (Version 05)
Command Timeout Interval	Default (25 sec)
Load Balance Response	Default (Unit Attention)
Reservation Conflict Response for Test Unit Ready	Default (Normal)
Change Volume Mapping	Default (No Report)
Volume Capacity Expansion	Default (No Report)

Parameter	Setting
Vendor Unique Sense Code	Default (No Report)
Host Specific Mode	Default (Normal)
Asymmetric/Symmetric Logical Unit Access	Active/Active/Preferred_Path (Default)
Sense Data Conversion	Default (no conversion)
The following additional parameters must be set if you are using firmware version V10L50 or later:	
iSCSI Discovery Reply Mode	All - Reply All Ports (Default)
TPGS Mode	Enable (Default)
TPG Referrals	Disable (Default)
iSCSI Reservation Range	CA Port (Default)

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

Fujitsu ETERNUS storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the storage arrays supported in various Data ONTAP 8.x releases organized by family

- DX Generation 1 family
 - DX8000 (DX8700, DX8400)
 - DX400 1st generation (DX410, DX440)
- DX Generation 2 family
 - DX400 S2 (DX410, DX440)

- DX8000 S2 (DX8700)

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

Note: This guide discusses only the Fujitsu storage arrays that are supported in the Data ONTAP 8.x releases. See the *Interoperability Matrix* for information about the ETERNUS family of storage arrays, which are supported in releases earlier than Data ONTAP 8.x. For Data ONTAP 8.0.2 and earlier, the *Gateway Implementation Guide for Fujitsu ETERNUS Storage* provides information about setting up the ETERNUS family models to work with Data ONTAP.

Hitachi storage arrays

You must meet certain requirements when configuring Hitachi storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for Hitachi storage arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

Required system parameter for NSC55, USP, and VSP storage arrays

The required parameter setting is shown in the following table:

Parameter	Setting
Host type	Standard or 00

Required system parameters for AMS and HUS storage arrays

The parameters that must be set are listed in the following table:

Parameter	Setting
Mapping mode	Enabled
Host group security	Enabled
Product ID	DF600F
Mode Setting > Common Setting	Standard Mode
Platform	Not specified
Middleware	Not specified

Required system parameters for HUS VM storage arrays

The parameters that must be set are listed in the following table:

Parameter	Setting
Host mode	Standard or 00

Parameter	Setting
Port security	Enabled
Note: A host group that is separate from the default host groups should be created for each initiator-target port pair.	

The storage array parameters should be set and LUNs should be mapped to the host groups before the switch is zoned to connect the system running Data ONTAP to the Hitachi storage array.

If Data ONTAP FC initiator ports are zoned with Hitachi array target ports before the storage array parameters are set and the LUNs are mapped to the host groups, you might not be able to see any LUNs presented to the Data ONTAP interface.

In this situation, Data ONTAP generates one of the following EMS messages:

```
[fci.device.login.reject:info] The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x9 'Invalid R_CTL Field',
explanation code 0x29 'Insufficient Resources to Support Login'
[fci.device.login.reject:info]: The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x3 'Nx_Port Not Available,
Temporary', explanation code 0x29 'Insufficient Resources to Support Login'
```

Workaround

If zoning was configured prior to setting the parameters and mapping the LUNs to the host groups, you can resolve the problem of LUNs not being visible by performing a path re-initialization to force a rediscovery of the array LUNs that are mapped from the Hitachi storage array.

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

Requirements for creating LUNs in dynamic pools

When creating LUNs in the dynamic pools, do not over-provision the pool capacity.

Requirements for formatting array LUNs

You can perform different types of array LUN formatting. Each type has an impact on the availability of LUNs or the entire storage array.

If you must present a new LUN to the system running Data ONTAP immediately, you should use quick formatting. If you plan to present LUNs to the system running Data ONTAP in the future, quick formatting is recommended but Online formatting is acceptable.

Quick formatting is recommended because both Online and Offline formatting delay disk discovery until after the LUN formatting is complete, as described in the following list:

- Online formatting takes the array LUN that is being formatted completely offline, and the array LUN does not respond to any SCSI commands.
Formatting is performed on one array LUN at a time until all array LUNs are formatted.
- Offline formatting takes the entire storage array offline and the storage array does not respond to any SCSI commands.
The array LUNs are formatted six at a time per controller until the formatting is complete.
The storage array does not respond until all formatting is complete.
- If you are creating array LUNs (for any host) after initial setup of the system running Data ONTAP and you are using Offline formatting, the system panics.
Any array LUNs that are mapped to the system running Data ONTAP, including array LUNs in the root volume, become unavailable until formatting is complete.
- Online and quick formatting of LUNs for other hosts do not impact the system running Data ONTAP.

Requirements for using disks external to the storage array

Data ONTAP supports using the same disks behind the storage arrays that the vendor supports (that is, disks external to the storage array). Data ONTAP assumes that the configuration between the storage array and the external disks behind it is correct.

The following list shows restrictions and best practice recommendations for deploying storage arrays that use external disks:

- Root volumes and spare core LUNs are not supported on external disks.
- Array LUNs from external and internal disks should not be in the same Data ONTAP aggregate.
- Array LUNs from SATA drives and array LUNs from FC drives should not be in the same aggregate.
- Rules for the number of paths to array LUNs on external disks are the same as for array LUNs on disks on the storage array.
In the case of external disks, the paths go from the system running Data ONTAP through the storage array to the external disks. See the *FlexArray Virtualization Installation Requirements and Reference Guide* for information about the number of paths supported for different Data ONTAP releases and modes.
- Only block checksum is supported for array LUNs on external disks.

Guidelines for configuring ports on Hitachi storage arrays

You should follow certain guidelines when configuring ports on Hitachi storage arrays.

The guidelines for configuring ports on Hitachi storage arrays are as follows:

- You should add only one host group per port.
- You should ensure that a host group includes all the FC initiator ports of the gateways in an array LUN neighborhood.

Hitachi storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the storage arrays organized by family:

- Family 1
 - USP-V
 - USP-VM
- Family 2
 - VSP
- Family 3
 - AMS 200
 - AMS 500
 - AMS 1000
- Family 4
 - AMS 2100
 - AMS 2300
 - AMS 2500
- Family 5
 - HUS 110
 - HUS 130
 - HUS 150
- Family 6
 - USP 1100
 - USP 600
 - USP 100
 - NSC55
- Family 7
 - HUS VM

34 | Implementation Guide for Third-Party Storage

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

HP EVA storage arrays

You must meet certain requirements when configuring storage arrays to work with systems running Data ONTAP. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for HP EVA arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

System parameter settings are provided in the following table:

Parameter	Setting
Redundancy	Any RAID level except VRAID0
Preferred path/mode	No preference (default)
Host type	SUN Solaris

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

Requirements for configuring the HP EVA Storage Management Server

Sharing FC initiator ports present on a system running Data ONTAP and initiator ports present on an HP EVA Storage Management Server with common EVA target ports causes compatibility issues.

The compatibility issues result from the different host settings for the FC initiator ports present on a system running Data ONTAP and the EVA Storage Management Server initiator ports.

HP recommends that any HP EVA Storage Management Server that resides in a fabric be configured in a zone that is separate from all operating systems.

Preparing HP EVA storage arrays for NDU

HP EVA storage arrays have a tendency to bring target ports offline during a firmware upgrade. Starting in Data ONTAP 8.2, Data ONTAP provides commands that increase Data ONTAP resiliency so that the system running Data ONTAP is not disrupted if target ports go offline during

the firmware upgrade. These commands can be used only with storage systems running clustered Data ONTAP.

About this task

You must follow the guidelines provided by HP for performing the firmware upgrade, including guidelines about the load recommendation (HP EVA). You can upgrade only to firmware that HP supports. This procedure pertains to the HP EVA firmware upgrade type *online upgrade*.

Note: During the firmware upgrade, both HP EVA controllers are rebooted.

This procedure instructs you to use Data ONTAP commands that increase Data ONTAP resiliency throughout the firmware upgrade. After the firmware upgrade is complete, you use the Data ONTAP commands again to return the ports to their normal operating mode.

Steps

1. On the system running Data ONTAP, set the privilege level of the command session to advanced:
set -privilege advanced
2. Set the `is-upgrade-pending` parameter to `true` for the HP EVA storage array that is undergoing a firmware upgrade:
storage array modify -name *array_name* -is-upgrade-pending true
3. On the storage array, start the firmware upgrade.
4. After the firmware upgrade is complete, set the privilege level to advanced again, if necessary, and then on the system running Data ONTAP, set the `is-upgrade-pending` parameter to `false` to return the storage array ports to normal operation:
storage array modify -name *array_name* -is-upgrade-pending false

If you did not exit advanced mode in Step 2 you do not need to access it again.

The `is-upgrade-pending` parameter contains the following features to ensure proper use of the command:

- If you attempt to set `is-upgrade-pending` to `false` while the upgrade is in progress on the storage array, the command fails and an EMS message is returned.
- If the `is-upgrade-pending` status is not returned to `false` within 60 minutes from the time it is set to `true`, EMS messages are logged hourly until `is-upgrade-pending` is returned to `false`.

ALUA support for HP EVA storage arrays

Starting with Data ONTAP 8.1, support has been added for asymmetric logical unit access (ALUA) with HP EVA arrays.

By default, support for ALUA is enabled in Data ONTAP as well as on all the HP EVA storage arrays.

HP EVA storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the storage arrays organized by family:

- Family 1
 - P6300
 - P6350
 - P6500
 - P6550
- Family 2: HP EVAx100
 - EVA 4100
 - EVA 6100
 - EVA 8100
- Family 3
 - EVA 4000
 - EVA 6000
 - EVA 8000
- Family 4
 - EVA 4400
 - EVA 6400
 - EVA 8400

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

HP XP storage arrays

You must meet certain requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for HP XP storage arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

The required parameter setting is shown in the following table:

Parameter	Setting
Host type	standard or 00

The storage array parameters should be set and LUNs should be mapped to the host groups before the switch is zoned to connect the system running Data ONTAP to an HP XP storage array.

If Data ONTAP FC initiator ports are zoned with HP XP array target ports before the storage array parameters are set and the LUNs are mapped to the host groups, you might not be able to see any LUNs presented to the Data ONTAP interface.

In this situation, Data ONTAP generates one of the following EMS messages:

```
[fci.device.login.reject:info] The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x9 'Invalid R_CTL Field',
explanation code 0x29 'Insufficient Resources to Support Login'
[fci.device.login.reject:info]: The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x3 'Nx_Port Not Available,
Temporary', explanation code 0x29 'Insufficient Resources to Support Login'
```

Workaround

If zoning was configured prior to setting the parameters and mapping the LUNs to the host groups, you can resolve the problem of LUNs not being visible by performing a path re-initialization to force a rediscovery of the array LUNs that are mapped from the HP XP storage array.

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

Requirements for creating LUNs in dynamic pools

When creating LUNs in the dynamic pools, do not over-provision the pool capacity.

Requirements for formatting array LUNs

You can perform different types of array LUN formatting. Each type has an impact on the availability of LUNs or the entire storage array.

If you must present a new LUN to the system running Data ONTAP immediately, you should use quick formatting. If you plan to present LUNs to the system running Data ONTAP in the future, quick formatting is recommended but Online formatting is acceptable.

Quick formatting is recommended because both Online and Offline formatting delay disk discovery until after the LUN formatting is complete, as described in the following list:

- Online formatting takes the array LUN that is being formatted completely offline, and the array LUN does not respond to any SCSI commands.
Formatting is performed on one array LUN at a time until all array LUNs are formatted.
- Offline formatting takes the entire storage array offline and the storage array does not respond to any SCSI commands.
The array LUNs are formatted six at a time per controller until the formatting is complete.
The storage array does not respond until all formatting is complete.
- If you are creating array LUNs (for any host) after initial setup of the system running Data ONTAP and you are using Offline formatting, the system panics.
Any array LUNs that are mapped to the system running Data ONTAP, including array LUNs in the root volume, become unavailable until formatting is complete.
- Online and quick formatting of LUNs for other hosts do not impact the system running Data ONTAP.

Requirements for using disks external to the storage array

Data ONTAP supports using the same disks behind the storage arrays that the vendor supports (that is, disks external to the storage array). Data ONTAP assumes that the configuration between the storage array and the external disks behind it is correct.

The following list shows restrictions and best practice recommendations for deploying storage arrays that use external disks:

- Root volumes and spare core LUNs are not supported on external disks.

- Array LUNs from external and internal disks should not be in the same Data ONTAP aggregate.
- Array LUNs from SATA drives and array LUNs from FC drives should not be in the same aggregate.
- Rules for the number of paths to array LUNs on external disks are the same as for array LUNs on disks on the storage array.

In the case of external disks, the paths go from the system running Data ONTAP through the storage array to the external disks. See the *FlexArray Virtualization Installation Requirements and Reference Guide* for information about the number of paths supported for different Data ONTAP releases and modes.

- Only block checksum is supported for array LUNs on external disks.

Guidelines for configuring ports on HP XP storage arrays

You should follow certain guidelines when configuring ports on HP XP storage arrays.

The guidelines for configuring ports on HP XP storage arrays are as follows:

- You should add only one host group per port.
- You should ensure that a host group includes all the FC initiator ports of the gateways in an array LUN neighborhood.

HP XP storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the storage arrays organized by family:

- Family 1
 - P9500
- Family 2
 - XP24000
 - XP20000
 - XP12000
 - XP10000

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

IBM DS storage arrays

You must meet certain requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required settings for IBM DS storage arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

Required host type for DS8xxx storage arrays

The required host type setting is shown in the following table:

Parameter	Setting
Host type	nseries

Required host type for DS3950, DS4xxx, and DS5xxx storage arrays

The required host type setting is shown in the following table:

Parameter	Setting
Host type	AIX

Requirements for configuring DS8300 9A2 LPAR models

When setting up the DS8300 9A2 LPAR (system logical partition) model to interact with systems running Data ONTAP, you must set up access to each array LUN so that the redundant paths are both accessing the same LPAR.

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

Requirements for volume groups on DS8xxx storage arrays

You must be aware of limitations on the number of volume groups that you can use with a LUN group when configuring DS8xxx arrays.

You must use a single volume group for each LUN group on a DS8xxx storage array to guarantee that the DS8xxx array LUNs are consistently presented to all the FC initiators (that are on the system running Data ONTAP) that access them.

Attention: If array LUNs are not presented consistently, there is a potential for data corruption.

Requirements for direct-attached configurations with all IBM DS arrays

Direct-attached connections to IBM DS storage arrays, supported only for existing deployments, require specific configuration settings on the storage array.

The following list describes the requirements that must be met when configuring direct-attached IBM DS arrays.

- Host Attachment files for the array ports must be created.
- The Host Attachment Port type must be set to FcAL (Fibre Channel Arbitrated Loop).
- The Storage Image I/O port must be set to FcAL .

Note: Direct-attached configurations are not supported for new deployments. This information is provided as a reference for configurations that might already exist in the field.

IBM DS storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows arrays organized by family:

- Family 1
 - DS4800
 - DS4700
 - DS4200
- Family 2
 - DS5100
 - DS5300
- Family 3
 - DS8300
 - DS8100
- Family 4
 - DS5020
 - DS3950
- Family 5
 - DS8800 model 951
 - DS8700 models 941, 94E
 - DS8870

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

IBM XIV storage arrays

You must meet certain requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Requirements for configuring IBM XIV arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

System parameter settings

System parameter settings for IBM XIV Gen2 and Gen3 arrays are provided in the following table:

Parameter	Setting
Type	Default
Host	User-supplied host name
Cluster	User-supplied cluster name

IBM XIV storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the storage arrays organized by family:

- IBM XIV Gen 2
 - 2810-A14
 - 2812-A14

- IBM XIV Gen 3
 - 2810-114
 - 2812-114

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

3PAR storage arrays

You must meet certain requirements when configuring storage arrays to work with systems running Data ONTAP that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for 3PAR arrays with systems running Data ONTAP

Certain parameter settings are required on the storage array for the storage array to work successfully with systems running Data ONTAP.

Host persona settings

For 3PAR arrays with firmware 2.3.1 and later, the required host persona values must be set as shown in the following table:

Connection type	Host persona
Direct connect	10
Fabric attached	10

The host definition, with the required host persona, should be created on a 3PAR storage array before the switch is zoned to connect the system running Data ONTAP to a 3PAR storage array. If zoning is configured before the host definition, Data ONTAP discovers a 3PAR LUN as LUN 254 and product type SES, which Data ONTAP cannot use. The result is that newly mapped 3PAR LUNs cannot be assigned to system running Data ONTAP and some 3PAR LUNs that might have already been added to an aggregate cannot be used because Data ONTAP sees a LUN 254 instead of a 3PAR LUN. If this problem occurs, Data ONTAP generates the following EMS message:

```
mlm.array.unknown:warning - Array LUN XXXXXXXX.126L254 [S/N 'XXXXXXX'
vendor '3PARdata' product 'SES'] is unknown and is not supported in this
version of Data ONTAP.
```

If zoning was configured prior to the host definition, you can resolve the problem by doing path re-initialization to force a rediscovery of the array LUNs that are mapped from the 3PAR storage array. See [Workaround for the unknown device \(L254/SES\) problem with 3PAR storage arrays](#) on page 47 for the workaround for this problem.

Port persona settings

For 3PAR arrays with firmware 2.2.x, the required port persona values must be set as shown in the following table:

Connection type	Port persona
Direct connect	18
Fabric attached	19

Workaround for the unknown device (L254/SES) problem with 3PAR storage arrays

The host definition, with the host persona specified in the [Required parameters for 3PAR arrays with Data ONTAP systems](#) on page 46, should be created on a 3PAR storage array before the switch is zoned to connect the system running Data ONTAP to a 3PAR storage array. If zoning is configured before the host definition, Data ONTAP discovers a 3PAR array LUN as LUN 254 and product type SES, which Data ONTAP cannot use.

When Data ONTAP discovers 3PAR array LUNs as LUN 254 and product type SES, newly mapped 3PAR LUNs cannot be assigned to systems running Data ONTAP and some 3PAR LUNs that might have already been added to an aggregate cannot be used because Data ONTAP sees a LUN 254. Data ONTAP generates EMS messages for the LUN it discovered as LUN 254 and product type SES, reporting some LUNs as unknown devices.

If zoning was configured before the host definition, you can resolve the LUN 254/SES unknown device problem by forcing a rediscovery of the array LUNs that are mapped from the 3PAR storage array. The procedure you use depends on whether aggregates with 3PAR array LUNs have already been configured on the system running Data ONTAP. (In the following procedures, assume that Node-1 and Node-2 are the names of the Data ONTAP nodes.)

General configuration guidelines

Information in this section includes cautions, restrictions, special requirements, and information about array characteristics that you might need to consider when configuring or using your storage array.

3PAR storage array families

Data ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

The following list shows the storage arrays organized by family:

- Family 1
 - E200
- Family 2
 - S400
 - S800
- Family 3
 - T400
 - T800
- Family 4
 - F200
 - F400
- Family 5
 - P10000 - V400
 - P10000 - V800
 - 3PAR StoreServ 7200
 - 3PAR StoreServ 7400

The storage arrays and families shown are current at the time of publication. The Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8) is the final authority for which storage arrays are supported for each vendor.

Supported fabric-attached configurations

The fabric-attached configurations included in this section are supported for all the storage array vendors unless mentioned otherwise.


For additional information about configuration with specific storage arrays, see the Interoperability Matrix at the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Stand-alone basic configuration

The stand-alone basic configuration of a system running Data ONTAP that uses array LUNs is a simple, fabric-attached configuration with a single FC initiator port pair accessing a single LUN group.

This configuration is supported for use with all storage arrays listed in the *Interoperability Matrix* as supported for the release and mode of Data ONTAP running on your system.

The following illustration shows this configuration:




Stand-alone with two 2-port array LUN groups

In a stand-alone system running Data ONTAP which is in a fabric-attached simple configuration, each FC initiator port pair present on the system running Data ONTAP accesses a separate array LUN group.

This configuration is supported for use with all storage arrays listed in the *Interoperability Matrix* as supported for the release and mode of Data ONTAP running on your system.

The following illustration shows this configuration:




Single 4-port array LUN group

This configuration contains a single, 4-port LUN group with each target port accessed by a single Data ONTAP FC initiator port from the HA pair. Zoning ensures that there are only two paths to a specific array LUN from each system running Data ONTAP.

This configuration is supported for use with all storage arrays listed in the *Interoperability Matrix* as supported for the release and mode of Data ONTAP running on your system.

The following illustration shows this configuration:




Two 4-port array LUN groups

In this configuration, each Data ONTAP FC initiator port pair accesses a separate array LUN group. The zoning is a single Data ONTAP FC initiator to a single array target port.

This configuration is supported for use with all storage arrays listed in the *Interoperability Matrix* as supported for the release and mode of Data ONTAP running on your system.

The following illustration shows a block diagram of this configuration:



Eight-port array LUN group configuration

Starting in Data ONTAP 8.2, an eight-port LUN group configuration is supported on clustered gateways and, starting in Data ONTAP 8.2.1, also on systems running Data ONTAP that can use array LUNs.

This configuration is supported for use with all storage arrays listed in the *Interoperability Matrix* as supported for the release and mode of Data ONTAP running on your system.


There are two ways of deploying this configuration: crossed and uncrossed back-end connections.

Crossed back-end connections

In this configuration with the back-end connections crossed, the FC connections from the same storage array controller go to both fabric switches (redundant).

In this illustration of crossed back-end connections, note how the nodes are attached to the switches and to the storage array. Vs1 uses switch 1 when attaching to the storage array Controller 1 port 1A and Controller 2 port 2C, and uses switch 2 when attaching to storage array Controller 2 ports 2A and

Controller 1 port 1C. This optimizes the use of switch ports and array ports, which reduces the impact of a switch or storage array controller failure.




Note: The following problem can occur with Active-Passive storage arrays in crossed back-end connections: If one of the fabric switches is taken offline or goes down and a path failure from both systems running Data ONTAP occurs, the systems running Data ONTAP panic even though the alternate path from each system is still online. For example, if switch 2 is down and the "0b" paths on both vs1 and vs2 go down, vs1 and vs2 panic even though the "0a" paths are still online.

Back-end connections are not crossed

In this configuration in which the back-end connections are not crossed, the FC connections from the same storage array controller go to only one fabric switch.

The following illustration shows this configuration when the back-end connections are not crossed.



Data ONTAP FC initiator port attaching to multiple targets ports

Starting in Data ONTAP 8.2, a Data ONTAP FC initiator port can be attached to multiple target ports on separate storage arrays within the same family. The *FlexArray Virtualization Installation Requirements and Reference Guide* contains the rules for setting up this configuration. This configuration is not supported with MetroCluster configurations.


This configuration is supported for use with all storage arrays listed in the *Interoperability Matrix* as supported for the release and mode of Data ONTAP running on your system.

The following illustration shows a shared Data ONTAP FC initiator port configuration with an HA pair. However, sharing an FC initiator port with multiple target ports is supported with both stand-alone systems and HA pairs.

For details about sharing an FC initiator port with multiple target ports and zoning, see the information in the *FlexArray Virtualization Installation Requirements and Reference Guide*.

One Data ONTAP FC initiator port connecting to target ports on *separate* storage arrays

The following example shows an HA pair in which a single Data ONTAP FC initiator port connects to multiple target ports on different storage arrays:




Supported direct-attached configurations

Direct-attached configurations are not supported for new deployments. For older deployments, direct-attached configurations are limited to some storage arrays and some Data ONTAP releases. See the *Interoperability Matrix* for information about configurations supported for direct attachment.

Two-port direct-attached stand-alone configurations

Direct-attached configurations are not supported for new deployments. For older deployments, direct-attached configurations are limited to some storage arrays and some Data ONTAP releases. See the *Interoperability Matrix* for information about configurations supported for direct attachment.


The following illustration is shown as a reference for configurations that might already exist in the field. Two Data ONTAP FC initiator ports provide access to LUNs on the storage array.



Four-port direct-attached HA pair configuration

Direct-attached configurations are not supported for new deployments. For older deployments, direct-attached configurations are limited to some storage arrays and some Data ONTAP releases. See the *Interoperability Matrix* for information about configurations supported for direct attachment.

The following illustration is shown as a reference for configurations that might already exist in the field. The illustration shows a direct-attached HA pair with four ports providing access to array LUNs on the storage array.



Terminology comparison between storage array vendors

Different storage array vendors occasionally use different terms to describe similar concepts. Conversely, the meaning of the same term might differ between array vendors.

The following table provides a mapping between some common vendor terms:

Term	Vendor	Definition
host group	Hitachi, Sun	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 system running Data ONTAP 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	
cluster	IBM XIV	
host affinity group	Fujitsu ETERNUS4000, ETERNUS6000, ETERNUS8000, ETERNUS DX8000, ETERNUS DX400	
host definition	3PAR	
host	3PAR, HP EVA, HP XP, Hitachi	
—	IBM ESS	No concept of <i>host group</i> . You must create a host in the ESS user interface for each FC initiator port that you plan to connect to the storage array and map each host to a port.
parity group	IBM DS8xxx, IBM ESS, Hitachi, HP XP, Sun	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	

Term	Vendor	Definition
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 <i>chunklets</i> . (A chunklet is a 256-MB block of contiguous space on a physical disk.)
cluster	Data ONTAP	In clustered Data ONTAP 8.x, a grouping of nodes that enables multiple nodes to pool their resources into a large virtual server and to distribute work across the cluster.
	Hitachi, HP XP, Sun	A hardware component on the storage arrays that contains the ports to which hosts attach.
	IBM XIV	An entity that groups multiple hosts together and assigns the same mapping to all the hosts.
controller	Data ONTAP	The component of a storage system 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 EVA, HP XP, IBM	Hardware on the storage array on which the target ports are located.

Term	Vendor	Definition
interface module	IBM XIV	A hardware component on the storage arrays that contains the ports to which hosts attach.
node	3-PAR	
FEBE Board	EMC Symmetrix	
Storage processor (SP)	EMC CLARiiON	
Controller Module	Fujitsu ETERNUS	
LUN	Many storage arrays	A grouping of one or more disks or disk partitions into one span of disk storage space. In the Data ONTAP documentation, this is referred to as <i>array LUN</i> .
LDEV	Hitachi, HP XP	
LUN	Data ONTAP	The system running Data ONTAP can virtualize the storage attached to it and serve the storage up as LUNs to external applications and clients (for example, through iSCSI and FC). Clients are unaware of where a front-end LUN is stored.
LUN, virtual disk	HP EVA	A virtual disk (called a <i>Vdisk</i> 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	Data ONTAP documentation, Data ONTAP storage management tools	The Data ONTAP documentation uses the term <i>array LUN</i> to distinguish LUNs on the storage arrays from front-end LUNs (Data ONTAP LUNs).

Term	Vendor	Definition
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, IBM XIV	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 (FC), and Internet SCSI (iSCSI).
	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.

Copyright and trademark information

Copyright ©1994 - 2014 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2014 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, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak,

RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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

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

NetApp, Inc. 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.

Index

- 3PAR storage arrays
 - configuration requirements [46](#)
 - families (storage array) [47](#)
 - troubleshooting unknown device problems [47](#)
- 8-port array LUN group configuration
 - support for [52](#)

A

- ACLX LUN
 - caution about using [24](#)
- ALUA support for HP EVA [36](#)
- ALUA support with EMC CX and VNX
 - guidelines for enabling [19](#)
- AMS storage arrays
 - required system parameters [30](#)
- array LUNs
 - systems running Data ONTAP that can use [16](#)
- array port initialization
 - requirements for [17](#)

C

- channel director ports
 - parameter settings [22](#)
- configuration requirements
 - 3PAR [46](#)
 - EMC Symmetrix [22](#)
 - Fujitsu ETERNUS [26](#), [27](#)
 - Hitachi [30](#)
 - HP EVA [35](#)
 - HP XP [38](#)
 - IBM DS [41](#)
 - IBM XIV [44](#)
- configurations supported with storage arrays
 - 8-port array LUN group [52](#)
 - FC initiator port attaching to multiple targets ports [54](#)
 - four-port direct-attached HA pair [57](#)
 - single 4-port array LUN group [50](#)
 - stand-alone basic configuration [49](#)
 - stand-alone with two 2-port array LUN groups [50](#)
 - two 4-port array LUN groups [51](#)
 - two-port direct-attached stand-alone [56](#)

D

- direct-attached configurations
 - requirements for IBM DS arrays [42](#)
- DS8xxx storage arrays
 - volume group number requirements [42](#)
- DX4xx storage arrays
 - host parameter settings [27](#)
- DX8xxx storage arrays
 - host parameter settings [26](#)

E

- EMC CLARiiON storage arrays
 - ALUA support [19](#)
 - array LUN numbering [19](#)
 - families (storage array) [19](#)
 - required parameters with Data ONTAP systems [18](#)
 - Storage Group use [18](#)
- EMC Symmetrix storage arrays
 - caution about using the ACLX LUN [24](#)
 - caution about using the VCMDB LUN [24](#)
 - channel director port parameter settings [22](#)
 - configuration requirements [22](#)
 - gatekeeper LUN use restriction [24](#)
 - host configuration parameters [22](#)
- EMC VNX storage arrays
 - ALUA support [19](#)
 - array LUN numbering [19](#)
 - families (storage array) [20](#)
 - required parameters with Data ONTAP systems [18](#)
 - Storage Group use [18](#)

F

- families (storage array)
 - 3PAR [47](#)
 - EMC CLARiiON [19](#)
 - EMC Symmetrix [24](#)
 - EMC VNX [20](#)
 - Fujitsu ETERNUS [28](#)
 - Hitachi [33](#)
 - HP EVA [37](#)
 - HP XP [40](#)
 - IBM DS [42](#)
 - IBM XIV [44](#)

67 | Implementation Guide for Third-Party Storage

- FC initiator ports
 - sharing target ports
 - supported configuration [54](#)
- formatting array LUNs
 - Hitachi [31, 39](#)
 - HP XP [31, 39](#)
- four-port direct-attached HA pair configuration
 - support for [57](#)
- Fujitsu ETERNUS storage arrays
 - configuration requirements [26, 27](#)
 - DX8xxx parameter settings [26](#)
 - families (storage array) [28](#)
 - Host Response pattern [26](#)

G

- gatekeeper LUNs
 - restrictions for using [24](#)
- gateway functionality
 - name change to FlexArray Virtualization [16](#)
- Gateway Interoperability Matrix
 - See Interoperability Matrix
- gateways
 - See LUNs (array)

H

- Hitachi storage arrays
 - configuration requirements [30](#)
 - Data ONTAP requirements for external disks [32, 39](#)
 - families (storage array) [33](#)
 - formatting array LUNs [31, 39](#)
 - port configuration guidelines [32](#)
- host configuration parameters
 - EMC Symmetrix [22](#)
- host parameter settings
 - DX4xx storage arrays [27](#)
- host persona [46](#)
- Host Response pattern [26](#)
- HP EVA storage arrays
 - ALUA support [36](#)
 - configuration requirements [35](#)
 - families (storage array) [37](#)
 - preparing for NDU [35](#)
 - zoning requirements [35](#)
- HP XP storage arrays
 - configuration requirements [38](#)
 - Data ONTAP requirements for external disks [32, 39](#)
 - families (storage array) [40](#)
 - formatting array LUNs [31, 39](#)

- port configuration guidelines [40](#)
- HUS storage arrays
 - required system parameters [30](#)

I

- IBM DS storage arrays
 - configuration requirements [41](#)
 - families (storage array) [42](#)
 - host type setting [41](#)
 - requirements for direct-attached connections [42](#)
- IBM XIV storage arrays
 - configuration requirements [44](#)
 - families (storage array) [44](#)
 - parameter settings [44](#)
- implementation planning
 - information needed for [11](#)
- initiator ports
 - sharing with multiple target ports
 - supported configuration [54](#)
- Interoperability Matrix
 - support information in [13](#)

L

- limits for storage arrays
 - where to find information needed for planning [14](#)
- LUN limits
 - where to find information needed for planning [14](#)
- LUN security
 - EMC Symmetrix storage arrays
 - LUN security methods [23](#)
- LUNs (array)
 - systems running Data ONTAP that can use [16](#)
- LUNs, array
 - See LUNs (array)

N

- NDUs
 - preparing HP EVA storage arrays for [35](#)
- NSC55 storage arrays
 - required system parameters [30](#)

P

- parameters
 - required for EMC CLARiiON and VNX storage arrays with Data ONTAP systems [18](#)

port persona settings [46](#)
 portcfggport command
 use for array port initialization [17](#)

ports
 configuration guidelines for Hitachi [32](#)
 configuration guidelines for HP XP [40](#)
 sharing initiator port with target ports
 supported configuration [54](#)

R

required parameters
 for EMC CLARiON storage arrays with Data
 ONTAP systems [18](#)
 for EMC VNX storage arrays with Data ONTAP
 systems [18](#)

S

shared FC initiator target ports
 supported configuration [54](#)
 single 4-port array LUN group configuration
 support for [50](#)
 stand-alone basic configuration
 support for [49](#)
 stand-alone with two 2-port array LUN groups
 configuration
 support for [50](#)
 storage array families
 3PAR [47](#)
 EMC CLARiON [19](#)
 EMC Symmetrix [24](#)
 EMC VNX [20](#)
 Fujitsu ETERNUS [28](#)
 Hitachi [33](#)
 HP EVA [37](#)
 HP XP [40](#)
 IBM DS [42](#)
 IBM XIV [44](#)
 storage arrays
 configuration guidelines for [17](#)
 preparing for NDUs on HP EVA [35](#)
 supported, where to find information about [13](#)
 Storage Groups
 ALUA support with [19](#)
 guidelines for enabling ALUA with [19](#)

 how used for access control [18](#)
 support for ALUA in HP EVA [36](#)
 supported configurations with storage arrays
 8-port array LUN group [52](#)
 FC initiator port attaching to multiple targets ports
 [54](#)
 four-port direct-attached HA pair [57](#)
 single 4-port array LUN group [50](#)
 stand-alone basic configuration [49](#)
 stand-alone with two 2-port array LUN groups [50](#)
 two 4-port array LUN groups [51](#)
 two-port direct-attached stand-alone [56](#)
 switches
 supported with back-end storage arrays
 finding information about [13](#)

T

terminology
 storage array vendor [58](#)
 two 4-port array LUN groups configuration
 support for [51](#)
 two-port direct-attached stand-alone configuration
 support for [56](#)

U

USP storage arrays
 required system parameters [30](#)

V

VCMDB LUNs
 caution about using [24](#)
 restricting visibility of [23](#)
 volume groups
 requirements for number of [42](#)
 VSP storage arrays
 required system parameters [30](#)

Z

zoning
 compatibility issues with HP EVA storage arrays [35](#)
 shared FC initiator port supported configuration [54](#)



NA 210-06500_A0, Printed in USA

GA32-1050-04

