



NetApp®
Go further, faster

NetApp University

Accelerated NCDA Boot Camp Data ONTAP 8.0 7-Mode

Student Guide





NETAPP UNIVERSITY

Accelerated NCDA Boot Camp Data ONTAP 8.0 7-Mode

Student Guide

Course Number: STRSW-ILT-ANCDA-D87M
Catalog Number: STRSW-ILT-ANCDA-D87M-SG
Content Version: 1.0

ATTENTION

The information contained in this guide is intended for training use only. This guide contains information and activities that, while beneficial for the purposes of training in a closed, non-production environment, can result in downtime or other severe consequences and therefore are not intended as a reference guide. This guide is not a technical reference and should not, under any circumstances, be used in production environments. To obtain reference materials, please refer to the NetApp product documentation located at <http://now.netapp.com/> for product information.

COPYRIGHT

© 2010 NetApp, Inc. All rights reserved. Printed in the U.S.A. Specifications subject to change without notice.

No part of this book 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.

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 or materials described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product or materials 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. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND

NetApp Documentation is protected by Copyright and is provided to U.S. Government Agencies with LIMITED RIGHTS as defined at FAR 52.227-14(a). Use, duplication, or disclosure by the U.S. Government is subject to the restrictions as set forth therein. In the event of use by a DOD agency, the Government's rights in Documentation are governed by the restrictions in the Technical Data Commercial Items clause at DFARS 252.227-7015 and the Commercial Computer Software and Commercial Computer Software Documentation clause at DFARS 252.227-7202.

TRADEMARK INFORMATION

NetApp, the NetApp logo, Go Further, Faster, Data ONTAP, Appliance Watch, ASUP, AutoSupport, Bolt Design, Center-to-Edge, ComplianceClock, ComplianceJournal, ContentDirector, Cryptainer, Data Motion, DataFabric, DataFort, Decru, Decru DataFort, Evolution of Storage, Exec-Vault, FAServer, FilerView, FlexCache, FlexClone, FlexShare, FlexVol, FPolicy, Get Successful, gFiler, LockVault, Manage ONTAP, MultiStore, NearStore, NetApp Availability Assurance, NetApp IT As A Service, NetApp ProTech Expert, NetCache, NOW, NOW (NetApp on the Web), ONTAPI, Raid-DP, Replicator-X, SANscreen, SecureAdmin, SecureShare, Shadow Tape, Simulate ONTAP, SmartClone, SnapCache, SnapCopy, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, Snapshot, SnapStore, SnapSuite, SnapValidator, SnapVault, Spinnaker Networks, Spinnaker Networks logo, SpinCluster, SpinFlex, SpinFS, SpinHA, SpinMove, SpinServer, SpinStor, StoreVault, SyncMirror, Tech OnTap, Topio, vFiler, VFM, VFM (Virtual File Manager), WAFL, and Web Filer are either trademarks, registered trademarks, or service marks of NetApp, Inc. in the United States and/or other countries.

Not all common law marks used by NetApp are listed on this page. Failure of a common law mark to appear on this page does not mean that NetApp does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Apple and QuickTime are either trademarks or registered trademarks of Apple Computer, Inc. in the United States and/or other countries.

Microsoft and Windows Media are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, RealVideo, RealMedia, RealProxy, and SureStream are either trademarks or registered trademarks of RealNetworks, Inc. in the United States and/or other countries.

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

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

TABLE OF CONTENTS

WELCOME	1
MODULE 1: NCDA OVERVIEW	1-1
MODULE 2: NFS OVERVIEW	2-1
MODULE 3: NFS SETUP	3-1
MODULE 4: EXPORTS AND MOUNTS	4-1
MODULE 5: CIFS OVERVIEW	5-1
MODULE 6: CIFS WORKGROUPS	6-1
MODULE 7: CIFS SHARES AND SESSIONS	7-1
MODULE 8: CIFS ACCESS CONTROL	8-1
MODULE 9: CIFS DOMAINS	9-1
MODULE 10: NAS MULTIPROTOCOL	10-1
MODULE 11: NAS TROUBLESHOOTING	11-1
MODULE 12: SAN OVERVIEW	12-1
MODULE 13: FC CONNECTIVITY	13-1
MODULE 14: ISCSI CONNECTIVITY	14-1
MODULE 15: LUN ACCESS	15-1
MODULE 16: AVAILABILITY OVERVIEW	16-1
MODULE 17: SNAPSHOT COPIES	17-1
MODULE 18: SNAPRESTORE	18-1
MODULE 19: SNAPVAULT	19-1
MODULE 20: OPEN SYSTEMS SNAPVAULT	20-1
MODULE 21: HIGH AVAILABILITY	21-1
MODULE 22: METROCLUSTER	22-1
MODULE 23: SNAPMIRROR	23-1
MODULE 24: PERFORMANCE	24-1
APPENDIX A: PROTECTION MANAGER OVERVIEW	A-1



Go further, faster®

Accelerated NCDA Boot Camp Data ONTAP 8.0 7-Mode

Part Number:
STRSW-ILT-ANCDA-D87M



ACCELERATED NCDA BOOT CAMP DATA ONTAP 8.0 7-MODE



Logistics and Safety

Logistics

- Introductions
- Schedule (start time, breaks, lunch, close)
- Telephones and messages
- Food and drinks
- Restrooms

Safety

- Alarm signal
- Evacuation route
- Assembly area
- Electrical safety

© 2010 NetApp, Inc. All rights reserved.

LOGISTICS AND SAFETY



Course Objectives

By the end of this course, you should be able to:

- Configure a storage system in an NFS environment
- Set up and administer a storage system for CIFS functionality
- Discuss configuring a storage system for a SAN Fibre Channel environment
- Set up and administer a storage system in a SAN iSCSI environment
- Explain and implement backup and recovery methods available in Data ONTAP®
- Describe and implement business continuance methods available in Data ONTAP

© 2010 NetApp, Inc. All rights reserved.

COURSE OBJECTIVES



Course Agenda: Day 1

- Day 1
 - Welcome
 - Module 1: NCDA Overview
 - Module 2: NFS Overview
 - Module 3: NFS Setup
 - Module 4: Exports and Mounts
 - Module 5: CIFS Overview
 - Module 6: CIFS Workgroups

© 2010 NetApp, Inc. All rights reserved.

COURSE AGENDA: DAY 1



Course Agenda: Day 2

- Day 2
 - Module 7: CIFS Shares and Sessions
 - Module 8: CIFS Access Control
 - Module 9: CIFS Domains
 - Module 10: NAS Multiprotocol
 - Module 11: NAS Troubleshooting

© 2010 NetApp, Inc. All rights reserved.

COURSE AGENDA: DAY 2



Course Agenda: Day 3

- Day 3
 - Module 12: SAN Overview
 - Module 13: FC Connectivity
 - Module 14: iSCSI Connectivity
 - Module 15: LUN Access

© 2010 NetApp, Inc. All rights reserved.

COURSE AGENDA: DAY 3



Course Agenda: Day 4

■ Day 4

- Module 16: Availability Overview
- Module 17: Snapshot™ Copies
- Module 18: SnapRestore®
- Module 19: SnapVault®
- Module 20: Open Systems SnapVault

© 2010 NetApp, Inc. All rights reserved.

COURSE AGENDA: DAY 4



Course Agenda: Day 5

- Day 5
 - Module 21: High Availability
 - Module 22: MetroCluster
 - Module 23: SnapMirror®
 - Module 24: Performance

© 2010 NetApp, Inc. All rights reserved.

COURSE AGENDA: DAY 5



NetApp University Information Sources

- NOW® (NetApp on the Web)
 - <http://now.netapp.com>

- NetApp University
 - <http://www.netapp.com/us/services/university/>

- NetApp University Support
 - <http://netappsupport.custhelp.com>

© 2010 NetApp, Inc. All rights reserved.

NETAPP UNIVERSITY INFORMATION SOURCES



Font Styles

Convention	Type of Information
<i>Italic Font</i>	Book titles. Words or characters that require special attention. Variable names or placeholders for information that must be supplied, for example: An ifstat command looks like this: <code>ifstat -z -a <interface></code> The name of the interface for which you want to view statistics is <i>interface</i> .
Monospaced font	Command names, daemon names, and option names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters that are typed, for example: Enter the following command: <code>options httpd.enable on</code> <code>license add <code1> <code2></code>

© 2010 NetApp, Inc. All rights reserved.

FONT STYLES



Go further, faster®

NCDA Overview

Module 1
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



NCDA OVERVIEW



Module Objectives

By the end of this module, you should be able to:

- Explain the NCDA certification
- Review key concepts from the Data ONTAP® 8.0 7-Mode Administration course

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



NCDA Certification

© 2010 NetApp, Inc. All rights reserved.

NCDA CERTIFICATION



NCDA Certification

- NetApp® certification is proof that you have the skills necessary to manage and deploy NetApp technologies
- NetApp Certified Data Management Administrators (NCDAs) must prove they have in-depth knowledge to administrate NetApp technologies
- Additional certifications are available



© 2010 NetApp, Inc. All rights reserved.

NCDA CERTIFICATION

As a NetApp Certified Data Management Administrator, you will have proven skills in performing in-depth support, administrative functions, and performance management for CIFS, NFS, and FC for SCSI or iSCSI for TCP/IP protocols on NetApp storage systems running the Data ONTAP® operating system in NFS and Windows® (CIFS) multiprotocol environments. You will also be able to implement active-active controller configuration and SyncMirror® software to ensure continuous data availability and rapid recovery of data in the event of a disaster, and use the SnapMirror®, SnapRestore®, and SnapVault® products to manage and protect mission-critical data.



Benefit of Certification

- With certification, storage administrators receive:
 - Recognition of industry achievement
 - Proof of skills needed to manage and deploy NetApp technologies



© 2010 NetApp, Inc. All rights reserved.

BENEFIT OF CERTIFICATION



Recommended Courses

- To prepare for the NCDA certification, NetApp strongly recommends:
 - The instructor-led *Data ONTAP 8.0 7-Mode Administration* course
- Additional recommended courses:
 - *CIFS Administration*
 - *NFS Administration*
 - *SAN Administration*
 - *NetApp Protection Software Administration*
 - *High Availability* (Web-based training only)

© 2010 NetApp, Inc. All rights reserved.

RECOMMENDED COURSES



Data ONTAP 8.0 7-Mode Administration Course

© 2010 NetApp, Inc. All rights reserved.

DATA ONTAP 8.0 7-MODE ADMINISTRATION COURSE



Data ONTAP 8.0 7-Mode Administration

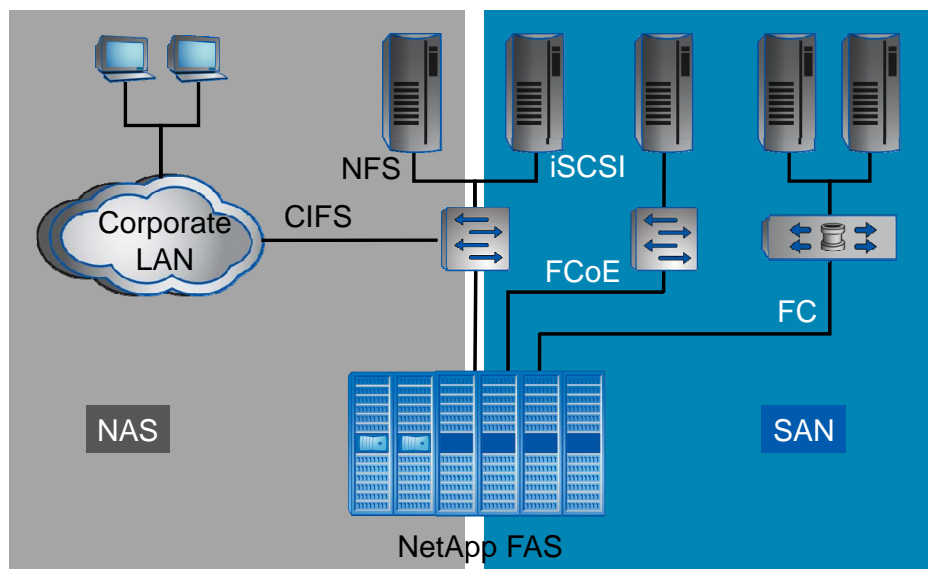
- The following topics are covered in the *Data ONTAP 8.0 7-Mode Administration* course:
 - Introduces NAS and SAN technologies
 - Distinguishes between modes within Data ONTAP 8.0
 - Identifies and discusses the benefits of the NetApp storage architecture
 - Describes role-based access controls
 - Steps to administer a NetApp storage system
- The next slides discuss some (but not all) of the important topics covered in this course

© 2010 NetApp, Inc. All rights reserved.

DATA ONTAP 8.0 7-MODE ADMINISTRATION



NAS and SAN Topology



© 2010 NetApp, Inc. All rights reserved.

NAS AND SAN TOPOLOGY

SAN is a block-based storage system that makes data available over the network using FC, FCoE, and iSCSI protocols.

NAS is a file-based storage system that makes data available over the network using NFS and CIFS protocols.

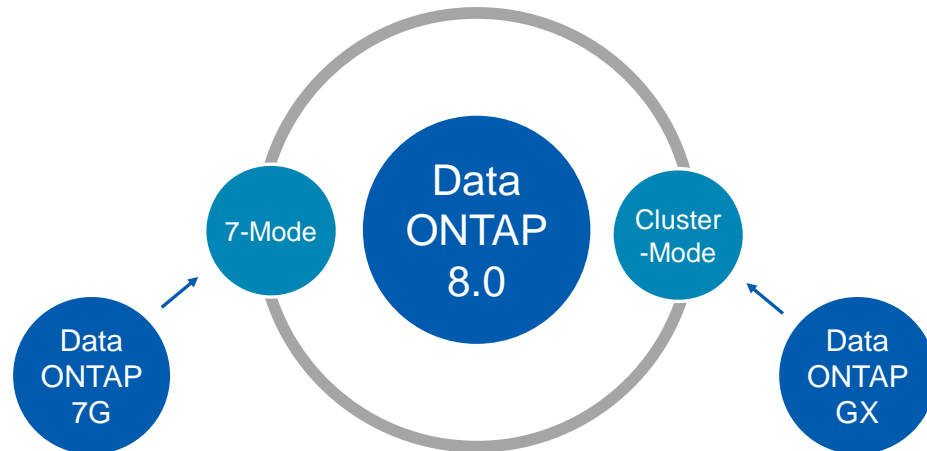
The NetApp SAN and unified storage architecture provides an outstanding level of investment protection and flexibility. The fabric-attached storage (FAS) system at the bottom of the graphic implies one “box.”

However, the actual storage environment includes small and large FAS systems, and NetApp VTL systems.



Data ONTAP 8.0 Review

- Data ONTAP 8.0 comes in two modes:
 - 7-Mode
 - Cluster-Mode



© 2010 NetApp, Inc. All rights reserved.

DATA ONTAP 8.0 REVIEW

Achieve new levels of scalability and storage flexibility, resulting in lower TCO, while providing maximized business agility and 24x7 business continuity.

Accelerate your move to a service-oriented architecture with Data ONTAP 8.0, which enables service levels across a diverse set of applications and extends data center virtualization. Data ONTAP 8.0 provides a single unified, scalable platform to address your NAS, SAN, multi-tier, multi-protocol, and multi-tenant virtualized environments.




Data ONTAP 8.0

■ 7-Mode

- Designed to be a simple transition from Data ONTAP 7G
- Scale-up technology allows aggregates to be up to 100 TB (higher in the future)
- Simple configuration for NAS or SAN


7-Mode



■ Cluster-Mode

- Designed to be a simple transition from Data ONTAP GX
- Scale-out technology allows a pool of storage controllers to manage the storage cluster
- Single NAS shared namespace across the cluster

Cluster-Mode



© 2010 NetApp, Inc. All rights reserved.

DATA ONTAP 8.0

NetApp storage solutions help you manage data in your enterprise environment with a scalable and flexible operating system we call Data ONTAP 8.0 7-Mode. Data ONTAP 8.0 7-Mode provides:

- More efficient storage
- High availability
- Business continuance
- Reduced storage management complexity

Deploy Data ONTAP 8.0 Cluster-Mode for high performance and high capacity. NetApp Data ONTAP 8.0 Cluster-Mode helps you achieve results and get to market faster by providing the massive throughput and scalability you need to meet the demanding requirements of your high-performance computing and digital media content applications. Achieve high levels of performance, manageability, and reliability for your large Linux®, UNIX®, or Microsoft® Windows clusters with Data ONTAP 8.0 Cluster-Mode. The Data ONTAP 8.0 Cluster-Mode operating system includes:


- Multi-node scaling using a global namespace
- NetApp FlexVol® storage virtualization
- Clustered file system
- Snapshot™ replication and mirroring



Storage Architecture

■ Storage architecture

- Aggregate
- Plex
- RAID group
- Disk



```
system> sysconfig -r
...
RAID group /aggr1/plex0/rg0 (normal)
RAID Disk Device HA SHELF BAY CHAN Pool...
-----
parity 0a.24 0a 1 8 FC:A 0...
data 0a.25 0a 1 9 FC:A 0...
...
```

© 2010 NetApp, Inc. All rights reserved.

STORAGE ARCHITECTURE

Data ONTAP 8.0 7-Mode storage architecture is as following:

- Aggregate - provide storage to a volume or volumes that they contain. Each aggregate contains its own plex(es), RAID configuration and a set of assigned physical disks.
- Plex - contain RAID groups and are associated with an aggregate. Normally, an aggregate will have only one plex. Mirrored aggregates using SyncMirror will have two plexes (plex0 and plex1) with plex1 containing a mirror of the plex0's data.
- RAID group - contains physical disks and are associated with a plex. RAID groups will either be RAID4 or RAID-DP configurations.
- Disks - are either parity, double-parity, or data disks.



Creating an Aggregate Using the CLI

- To create a 64-bit aggregate:

```
system> aggr create aggrname -B 64 24
```

 - Creates a 64-bit aggregate called *aggrname* with 24 disks
 - By default, this aggregate uses RAID-DP®
 - 24 disks must be available (spares) for the command to succeed
- To create a 32-bit aggregate:

```
system> aggr create aggrname -B 32 24
```

or

```
system> aggr create aggrname 24
```

© 2010 NetApp, Inc. All rights reserved.

CREATING AN AGGREGATE USING THE CLI

For more information about 64-bit aggregates, please see the Technical Report 3786 found at www.netapp.com/us/library/technical-reports/tr-3786.html.



NetApp System Manager: Aggregate

Select **Create** to create a new aggregate

Select **Aggregates** to administrate aggregates

Name	Disks	Status	Available Space	Used Space	Committed (%)	Total Space
aggr0	3	online	1.29 GB	27.08 GB	95.44	28.37 GB

Name	RAID Type	RAID group
0a.17	parity	rg0
0a.16	parity	rg0
0a.18	data	rg0

© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: AGGREGATE



Create Aggregate Wizard

Welcome to the Create Aggregate Wizard

The Aggregate Wizard steps you through the process of creating a new aggregate. You will be asked for information about the aggregate name, RAID type, disk selection mode and type, disk size etc.

Until you click on the Finish button in summary page, no permanent changes are made to the storage system. At any point in time, you may exit the Aggregate Wizard by pressing the Cancel button, and no changes are made to the storage system.

To continue, click Next.

< Back Next >

Create Aggregate Wizard

Name And RAID Details

Enter aggregate name, RAID type and disk selection parameters.

Aggregate name:

RAID type:

- ☒ Dual parity (recommended)
- ☐ RAID 4

[Tell me more about RAID types](#)

Disk selection:

- ☒ Allow system to select disks automatically based on the required aggregate size
- ☐ Manually select disks

Disk type:

Maximum size:

- ☒ Allow aggregate size to be greater than 16 TB (64-bit aggregate)

[Tell me more about aggregate size](#)

< Back Next > Cancel

Check for a 64-bit aggregate
or leave it blank for a
32-bit aggregate

© 2010 NetApp, Inc. All rights reserved.

CREATE AGGREGATE WIZARD



Create Aggregate Wizard (Cont.)

Create Aggregate Wizard

Aggregate Size
Choose the usable size.

You have 25 spare disks to create an aggregate of size between 29.88 GB and 597.66 GB. Choose the size of the aggregate.

Minimum size: 29.88 GB Maximum size: 597.66 GB

No of disks: 3 Usable size: 29.88 GB

< Back Next >

Create Aggregate Wizard

Aggregate Summary
Review the summary before creating your aggregate.

The following tasks will be performed when you start the process:

Aggregate name: aggr1

RAID type: RAID DP
Block Type: 64 bit aggregate
Disk type: FCAL or SAS
Disk count: 3

Usable size: 29.88 GB
Total size: 33.2 GB

Disk details:
0a:45 (33.47 GB, FCAL, 15000 RPM)
0a:22 (33.47 GB, FCAL, 15000 RPM)
0a:34 (33.47 GB, FCAL, 15000 RPM)

To start creating the aggregate, click Next.

< Back Next > Cancel

© 2010 NetApp, Inc. All rights reserved.

CREATE AGGREGATE WIZARD (CONT.)



Create Aggregate Wizard (Cont.)

NetApp System Manager

Name	Disks	Status	Available Space	Used Space	Committed (%)	Total Space
aggr0	3	online	1.3 GB	27.08 GB	95.43	28.37 GB
aggr1	3	online	29.86 GB	96 KB	0	29.86 GB

Create Aggregate Wizard

Completing the Create Aggregate Wizard

- Creating aggregate 'aggr1'...Successful
- Zeroing disks...Successful
- Setting snapshot reserve to '0%'...Successful
- Setting auto snapshot schedule to '0'...Successful

Aggregate 'aggr1' has been added to the system.
You have successfully completed the aggregate creation process.

To close this wizard, click Finish.

aggr1
online
RAID DP
Aggregate
No
96
31,142
Checksums: block
64-bit aggregate: Yes

Disks:

Name	RAID Type	RAID group
0a:45	data	rg0
0a:22	parity	rg0
0a:34	data	rg0

Details | Space Breakout


© 2010 NetApp, Inc. All rights reserved.

CREATE AGGREGATE WIZARD (CONT.)



Flexible Volumes

- Flexible volumes manage the logical layer independent of the physical layer
- Multiple flexible volumes can exist within a single aggregate



© 2010 NetApp, Inc. All rights reserved.

FLEXIBLE VOLUMES

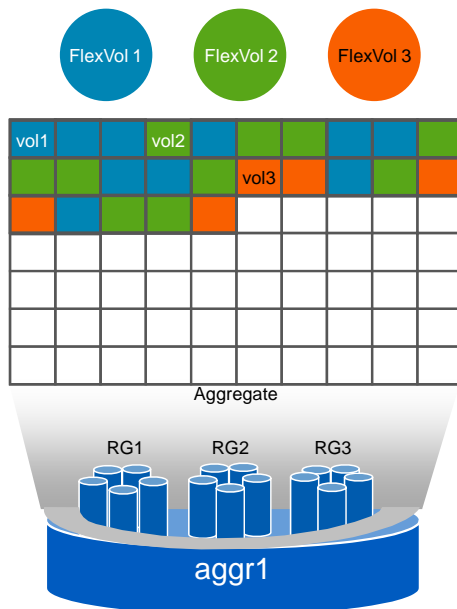
A flexible volume (also called a FlexVol volume) is a volume that is loosely coupled to its container aggregate. Because the volume is managed separately from the aggregate, you can create small FlexVol volumes (20 MB or larger), and then increase or decrease the size of the FlexVol volumes in increments as small as 4 KB.

Advantages of flexible volumes:

- You can create flexible volumes almost instantaneously. These volumes:
 - Can be as small as 20 MB
 - Are limited to aggregate capacity (if guaranteed)
 - Can be as large as the volume capacity supported for your storage system (not guaranteed)
- You can increase and decrease a flexible volume while online, allowing you to:
 - Resize without disruption
 - Size in any increment (as small as 4 KB)
 - Size quickly



Aggregates and FlexVol Volumes



- Create an aggregate
 - RAID groups are created as result
- Create FlexVol 1
 - Only metadata space is used
 - There is no pre-allocation of disk blocks to a specific volume
- Create FlexVol 2
 - WAFL® allocates aggregate space as data is written
- Populate volumes

© 2010 NetApp, Inc. All rights reserved.

AGGREGATES AND FLEXVOL VOLUMES


A FlexVol guarantee is an option of a flexible volume which determine when space is allocated out of the containing aggregate space for a volume or the files within the volume. There are three possible guarantees:

- Volume - is the default option and indicates that space is allocated or “taken away” from the aggregate when the volume is created.
- File - indicates that space is allocated or “taken away” from the aggregate when certain “space-reserved” files (such as a space-reserved LUN) is created.
- None - indicates that space is not allocated or “taken away” from the aggregate until it is used by the file. This is also referred to as “thin provisioning” a FlexVol.



Role-Based Access Control

- Role-based Access Control (RBAC)
 - Mechanism for managing a set of capabilities that an administrator can perform on a storage system
- Steps to implement:
 - Create a role with specific capabilities
 - Create a group with one or more assigned roles
 - Create user(s) assigned to one or more groups



© 2010 NetApp, Inc. All rights reserved.

ROLE-BASED ACCESS CONTROL

Role-based access control (RBAC) specifies how users and administrators can use a particular computing environment.


Most organizations have multiple system administrators, some of whom require more privileges than others. By selectively granting or revoking privileges for each user, you can customize the degree of access that each administrator has to the system.

RBAC allows you to define sets of capabilities that apply to one or more users. Users are assigned to groups based on their job functions, and each group is granted a set of roles to perform those functions.



Interface Groups

- Previously called virtual interfaces (vifs)
- Interface groups allow trunking of one or more Ethernet interfaces
 - IEEE 802.3ad link aggregation
- Types:
 - Single-mode
 - Multi-mode
- Command:
`system> ifgrp...`



© 2010 NetApp, Inc. All rights reserved.

INTERFACE GROUPS

Virtual interfaces (vifs) were renamed in Data ONTAP 8.0 7-Mode to eliminate any confusion with the term vif, which was used in Data ONTAP GX and Data ONTAP 8.0 Cluster-Mode.



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Explain the NCDA certification
- Review key concepts from the Data ONTAP 8.0 7-Mode Administration course

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 1: NCDA Overview
Estimated Time: 15 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- How is Data ONTAP 7G and Data ONTAP GX related in Data ONTAP 8.0?
- What are the two storage topologies supported by Data ONTAP?
- How is SAN different than NAS?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

NFS Overview

Module 2
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



NFS OVERVIEW



Module Objectives

By the end of this module, you should be able to:

- Define Network File System (NFS)
- Differentiate between NFS protocol versions
- Recognize the difference between stateless and stateful protocols
- Describe how the storage system acts as an NFS file server
- List the requirements of NFS

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Protocol Overview

© 2010 NetApp, Inc. All rights reserved.



Network File System

- NFS allows networked computers to access shared files
- Platforms that support NFS
 - Solaris™
 - Linux®
 - HP-UX®
- NFS allows network systems (clients) to access shared files and directories that are stored and administered centrally from a storage system

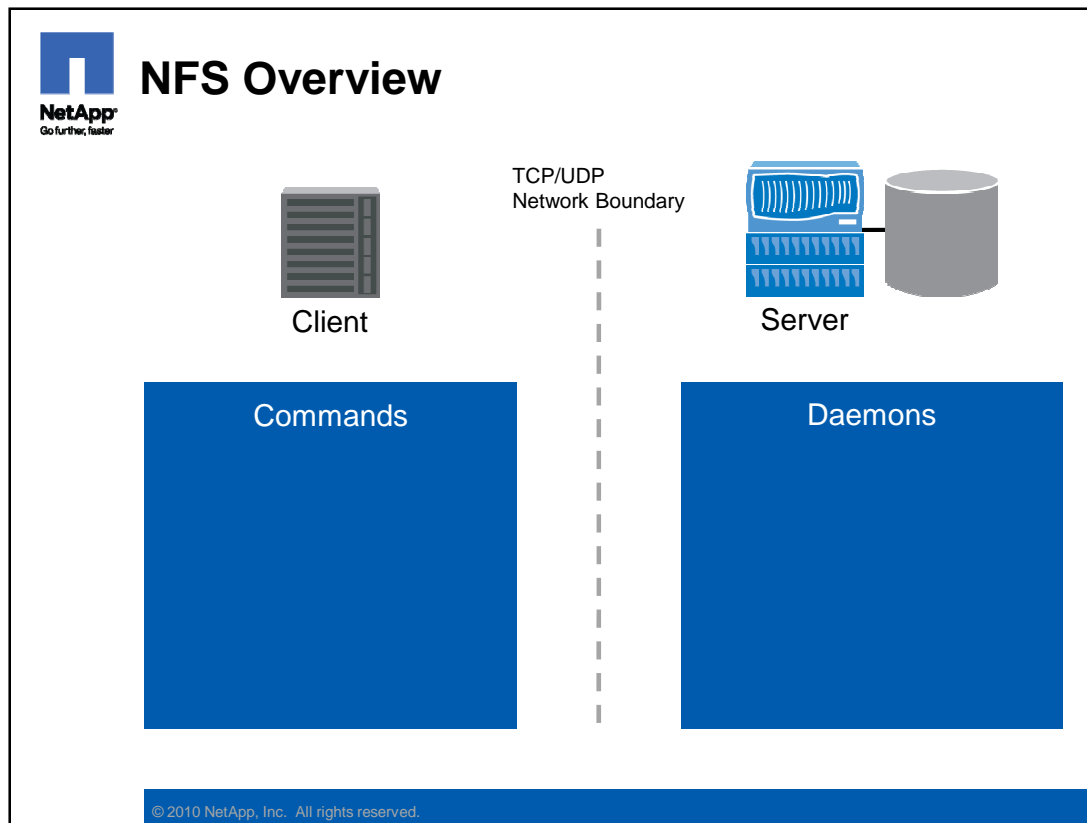
© 2010 NetApp, Inc. All rights reserved.

NETWORK FILE SYSTEM

The Network File System (NFS) is a distributed file system, developed by Sun Microsystems, Inc. in the 1980s, to address the need of sharing resources in a distributed networking environment. Networked computers are able to share files across networks without being in the same physical location as the server.

An NFS server has one or more directories that are mounted by NFS clients; to the NFS clients, the remote directories look like local directories or folders.

A NetApp® storage system in a NAS implementation can act as the NFS server. NetApp storage systems support NFS: v2, v3, and v4 to allow clients running different UNIX® or Linux operating systems to share files using the version of NFS supported in their environment. At this time, most clients are running NFS v3.



NFS OVERVIEW

Client-Server Architecture

The theory of client-server architecture is based on the concept that one computer has the resources that are required by another computer. These resources can be made available to systems that need them through NFS. The system with the resource is called the server and the system that requires the resources is called the client. Examples of resources are mail, database, and files. The client and the server communicate with each other through established protocols.

A distributed network (client-server network) might contain multiple servers and multiple clients, or multiple clients and a single server. The configuration of the network depends on the resource requirement of the environment.


The benefits of client-server architecture include cost reduction due to hardware and space requirements. The local workstations do not need as much disk space because commonly used data can be stored on the server. Other benefits include centralized support (backups, maintenance, and so on) performed on the server.

NFS is a widely used protocol for sharing files across networks. It is designed to be stateless to allow for easy recovery in the event of server failure. In the diagram above, the server in the network is a NetApp storage system, and the client could be one of many versions of a UNIX or Linux operating system.

As a file server, the storage system provides services that include mount daemon, Network Lock Manager (nlm_main), Network File System daemon (nfsd), Status Monitor, quota daemon, and portmap or rpcbind. Each of these services is important for a successful operation of an NFS process. For example, a client cannot mount a resource if mountd is not running on the server. Similarly, if rpcbind is not running on the server, NFS communication cannot be established between the client and the server.



NFS Overview: Mount



NFS OVERVIEW: MOUNT

In the above figure, the NetApp storage system is configured as the NFS server. The NFS client first mounts the required file system using the standard UNIX mount command. The mount command on the client host will first send a query to the portmap / rpcbind daemon asking which port number the mountd daemon is listening to. The portmap daemon will respond with the port number being used by the mountd daemon or a message indicating that the mount service is not registered. First, the client will make a remote procedure call to the portmap or rpcbind daemon running on the server.

REMOTE PROCEDURE CALL

Remote Procedure Call is a client-server programming environment that vendors use for developing platform-independent applications. Remote Procedure Call allows applications (programs) to communicate with each other just like network nodes communicate with each other using TCP or User Datagram Protocol (UDP).

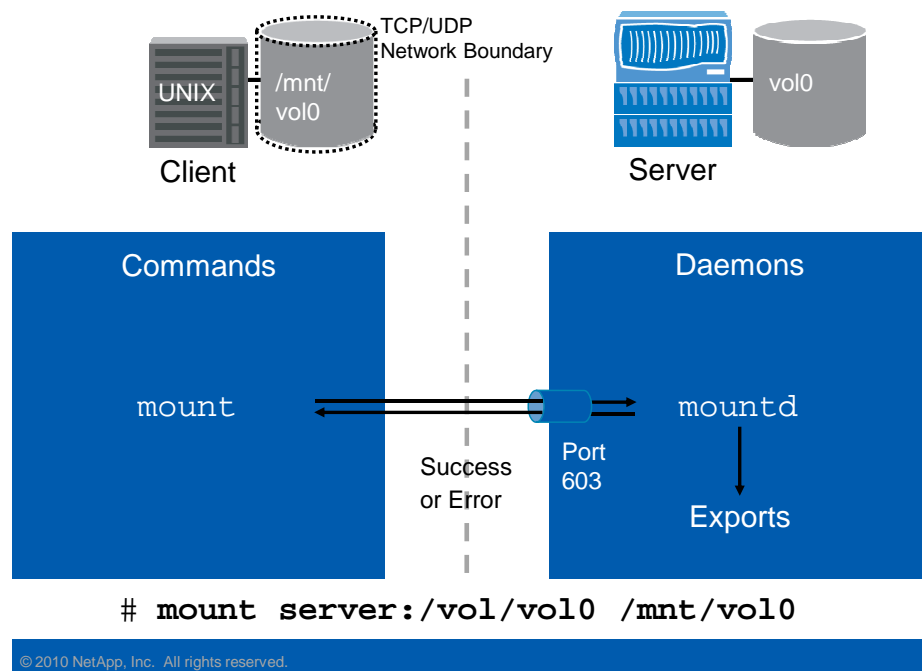
PORTMAP

A portmap, sometimes known as rpcbind, is a Remote Procedure Call service that allows clients and servers to communicate with each other using inter-process communication methods. The rpcbind/portmap daemon is used to translate Remote Procedure Call program numbers into UDP/TCP port numbers. This allows the other daemons (mountd, nfsd, and so on) to listen to ports that are not "well known." Just like network nodes communicate with each other using IP addresses, the portmapper service allows a Remote Procedure Call service (process, program) to communicate with other services using assigned port addresses.

Portmap allows these Remote Procedure Call services to use assigned ports as long as they are registered with the portmapper with program number, version, and transport protocol. The portmapper program is usually registered on port 111 of the TCP and UDP transport protocols. Usually, NFS servers (nfsd) default to port 2049.



NFS Overview: Mount (Cont.)



NFS OVERVIEW: MOUNT (CONT.)

The following ports are found on the storage system with NFS enabled:

- UDP 602 NFS mount daemon (mountd)
- TCP 603 NFS mount daemon (mountd)
- UDP 604 NFS status daemon (statd, statmon)
- TCP 605 NFS status daemon (statd, statmon)
- UDP 606 NFS lock manager (lockd, nlockmgr)
- TCP 607 NFS lock manager (lockd, nlockmgr)
- UDP 608 NFS quota daemon (quotad, rquotad)

The client will then issue a call to the mount daemon (mountd) on the server. Mountd will verify access to the resource, and then record the results in the access cache. Either a successful result or an error is returned. If the mount command was successful, the resource will now be accessed at the mountpoint as shown in this diagram.

NOTE: If the mountpoint has any local files, these files will not be visible or accessible while the file system is mounted.

Other possible mechanisms for mounting resources are:

- Using Solaris as an example, by updating the /etc/fstab file for persistent mounting of the file system across reboots. Other Unix or Linux-based systems will have similar mechanisms.
- Automounters.


NOTE: The mechanism for mounting resources are operating-system dependent.

AUTOMOUNTER

Automounter is an NFS program that mounts the file system on demand and unmounts the files if they are not accessed within a few minutes.



NFS Overview: NFS Call




© 2010 NetApp, Inc. All rights reserved.

NFS OVERVIEW: NFS CALL



NFS Overview: NFS Call (Cont.)



© 2010 NetApp, Inc. All rights reserved.

NFS OVERVIEW: NFS CALL (CONT.)

The client can now issue file system commands (such as `ls`) within the mountpoint. A remote procedure call will be sent to the NFS daemon (`nfsd`) on the server to process the call and return the results.



Three Different Versions of NFS

Version 2	Version 3	Version 4
Based on RFC 1094	Based on RFC 1813	Based on RFC 3530
Uses RPC protocol based on RFC 1057	Uses RPC protocol based on RFC 1057	Uses compound RPC protocol based on RFC 1831
Supports 32-bit file size	Supports 32-bit to 64-bit file size	Supports 32-bit to 64-bit file size
Stateless	Stateless	Stateful, no dependency on NFS v2 or v3

© 2010 NetApp, Inc. All rights reserved.

THREE DIFFERENT VERSIONS OF NFS

NFS V2

The features and functions of NFS v2 were defined in RFC 1094, the Remote Procedure Call is based on RFC 1057, and an External Data Representation (XDR) is based on RFC 1014. The maximum file size is 4 GB.

NFS V3

This version of NFS was developed to minimize the limitations in v2, especially the file-size limitations. NFS v3 is based on RFC 1813, the same Remote Procedure Call as in v2.

NFS V4

NFS v4 is a distributed file system based on RFC 3530, RFC 1831, and supports the same file size as NFS v3. It is designed to use the Internet, support traditional file-access methods, and integrate support for file locking and mount protocol. The XDR is based on RFC 1832. NFS v4 makes provision for end-to-end security, and Kerberos V5 is one of the supported methods.

Data ONTAP® 6.4 and later provides complete v4 server and client (Linux, Solaris™, Hummingbird NFS Maestro Client for ® NT) support.

STATEFUL OR STATELESS

If a protocol is stateless, it means that it does not require that the server maintain any session state between messages; instead, all session states are maintained by the client. With a stateless protocol, each request from client to server must contain all of the information necessary to understand the request and cannot take advantage of any stored content on the server.



Requirements

© 2010 NetApp, Inc. All rights reserved.

REQUIREMENTS



Requirements for NFS

- NFS servers must provide:
 - Resource List
 - Allows clients to discover resources
 - Identification
 - Identifies who is communicating with the storage system
 - Authorization
 - Allows properly identified clients to perform only certain actions

© 2010 NetApp, Inc. All rights reserved.

REQUIREMENTS FOR NFS

A resource list is a group of storage objects such as directories or files available for clients. Identification is the ability to associate IP addresses with host/client names. Authorization assigns permission to identified clients.



Resource List Through Exports

- Exports define what resources are available to which clients
 - Held in memory and used by mountd
- The storage system provides two types of exports:
 - Persistent: defined in */etc/exports*, persistent across reboots
 - Temporary: defined through command, located in memory only

Mount
command → mountd

Path	Rule
/vol/test	ro,root=unix1
/vol/vol1	rw,root=unix1

Storage System

© 2010 NetApp, Inc. All rights reserved.

RESOURCE LIST THROUGH EXPORTS

The export list resides in memory and is used by the mountd process to respond to mount requests. Contents of the memory list are established at start of the NFS service using the persisted cache (*/etc/exports*) and then can be dynamically controlled by way of commands.

EXPORTS

Exports are directories that can be exported to NFS clients.

```
/vol/test      -rw,root=unix1 /vol/vol1      -rw,root=unix1
```

RESOURCES

Resources are destinations to which resources are exported. Examples include:

- **Client** – Typically the UNIX/Linux host system connected to the storage system. The exports can be defined with either the IP address of the client or the host name if the name can be properly resolved.
- **Netgroup** – A netgroup is a network-wide group of machines granted identical access to certain network resources for security and organizational reasons.
- **Subnet** – A subnet is a physical grouping of connected network devices. Nodes on a subnet tend to be located in close physical proximity to each other on a LAN.
- **DNS Subdomain** – A subdomain is a domain that is part of a larger domain. A DNS hierarchy consists of the root-level domain at the top, underneath which are the top-level domains, followed by second-level domains, and finally the subdomains.



Identification

- Identify client hosts (targets) through:
 - IP only
 - Host name resolution
 - Name-to-IP resolution required
 - Local `/etc/hosts` file
 - Network Information Service (NIS)
 - DNS
 - Netgroup resolution
 - `/etc/netgroup` file
 - NIS
 - Lightweight Directory Access Protocol (LDAP)
 - IP subnet
 - DNS subdomains
- Lookup order defined in `/etc/nsswitch.conf` file

© 2010 NetApp, Inc. All rights reserved.

IDENTIFICATION

Network Information Service (NIS): Provides a simple network lookup service consisting of databases and processes. It was formerly known as Sun™ Yellow Pages (YP). The functionality of the two remains the same. Its purpose is to provide information that has to be known throughout the network, to all machines on the network. Information likely to be distributed by NIS is:

- Login names/passwords/home directories (`/etc/passwd`)
- Group information (`/etc/group`)
- Host names and IP numbers (`/etc/hosts`)

With Data ONTAP 7.1 and later, the storage system is capable of becoming an NIS slave. Like Domain Name System, NIS enables you to centrally maintain host information. NIS provides two methods for storage system host-name resolution:

- Using a makefile master on the NIS server, which creates an `/etc/hosts` file and copies it to the storage system's default volume for local host-name lookup
- Using host map, maintained as a database on the NIS server, which the storage system queries in a host lookup request across the network

NIS also enables you to maintain user information. For more information, see the *Data ONTAP 8.0 Network Management Guide*.

Domain Name System (DNS): Domain Name System (or Service or Server), is an Internet service that translates domain names into IP addresses. Because domain names are alphabetic, they're easier to remember. The Internet however, is based on IP addresses. Every time a domain name is used, a DNS service must translate the name into the corresponding IP address. For example, the domain name `www.example.com` might translate to `198.105.232.4`.

DNS enables you to maintain host information centrally. As a result, you do not have to update the `/etc/hosts` file every time you add a new host to the network. If you have several storage systems on your network, maintaining host information centrally saves you from updating the `/etc/hosts` file on each storage system every time you add or delete a host.

A conventional storage system policy for efficient host-name resolution is to do both of the following:

- Maintain a short /etc/hosts file containing local interfaces
- Enable DNS with DNS caching

Netgroup: A netgroup is a local file or NIS entity that associates a group of hosts with a group name. These netgroups are configured on the master NIS server and processed into netgroup maps, which are then propagated to the slave NIS servers. The two netgroup maps of interest are keyed using different fields. The first map “netgroup” is keyed by the netgroup name and has a primary value of the netgroup name followed by a list of hosts and other netgroups (that is, netgroups can be hierarchical) that belong to the netgroup. The second map “netgroup.byhost” is keyed by the host name and has a primary value of the host name (with a potential wildcard domain) followed by a comma-separated list of all the netgroups to which that host belongs.

Lightweight Directory Access Protocol (LDAP): As the name suggests, it is a lightweight protocol for accessing directory services, specifically X.500-based directory services.

The LDAP information model is based on entries. An entry is a collection of attributes that has a globally unique Distinguished Name (DN). Each of the entry’s attributes has a type and one or more values. The types are typically mnemonic strings, like “cn” for common name or “mail” for e-mail address. The syntax of the values depends on the attribute type. For example, a cn attribute might contain the value “Grace Adler.” A mail attribute might contain the value Grace.Adler@netapp.com.

NOTE: The lookup order for the above services is defined in /etc/nsswitch.conf file. This file is addressed in more detail in the subsequent sections.



Authorization

- For NFS v2 and v3, Client hosts (targets) are given access permissions during the mount request in the export definitions
 - Example:
 - Read only
 - Read and write (default if nothing specified)
- Requests for access are honored based on directory and/or file-level permissions

© 2010 NetApp, Inc. All rights reserved.

AUTHORIZATION

Examples of permission available and their usage:

To see the contents of a directory (ls command) you need read access

To create a file, you need read and write access.

To back up a Filesystem, you need read access but NOT write access. In other words, you need root read access because to perform a backup, you need to copy every file of every user. (root=client_LINUX,ro).



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Define Network File System (NFS)
- Differentiate between NFS protocol versions
- Recognize the difference between stateless and stateful protocols
- Describe how the storage system acts as an NFS file server
- List the requirements of NFS

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 2: NFS Overview
Estimated Time: 15 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- NFS is based on client-server architecture. True or false?
- List the three versions of NFS.
- What does stateful mean?
- NFS v3 is a stateful protocol. True or false?
- NFS v4 is a stateful protocol. True or false?
- What is a netgroup?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

NFS Setup

Module 3
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



NFS SETUP



Module Objectives

By the end of this module, you should be able to:

- Configure Network File System (NFS) on a NetApp® storage system
- Add Network Information Server (NIS) to manage users, groups, and name-to-IP resolution
- Administer a storage system to perform Domain Name System (DNS) lookups
- Configure a storage system to access a Lightweight Directory Access Protocol (LDAP) server to centrally manage users and groups
- Set up PC-NFS and WebNFS environments to extend the reach of NFS

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Environments

© 2010 NetApp, Inc. All rights reserved.

ENVIRONMENTS



NFS Environments

- NFS can be configured in many types of environments, for example:
 - NFS with local identification
 - NFS with NIS
 - NFS with DNS and LDAP
 - NFS with PC-NFS
 - NFS with WebNFS
- Your environment will be based upon your requirements

© 2010 NetApp, Inc. All rights reserved.

NFS ENVIRONMENTS



NFS

© 2010 NetApp, Inc. All rights reserved.

NFS



NFS Configuration

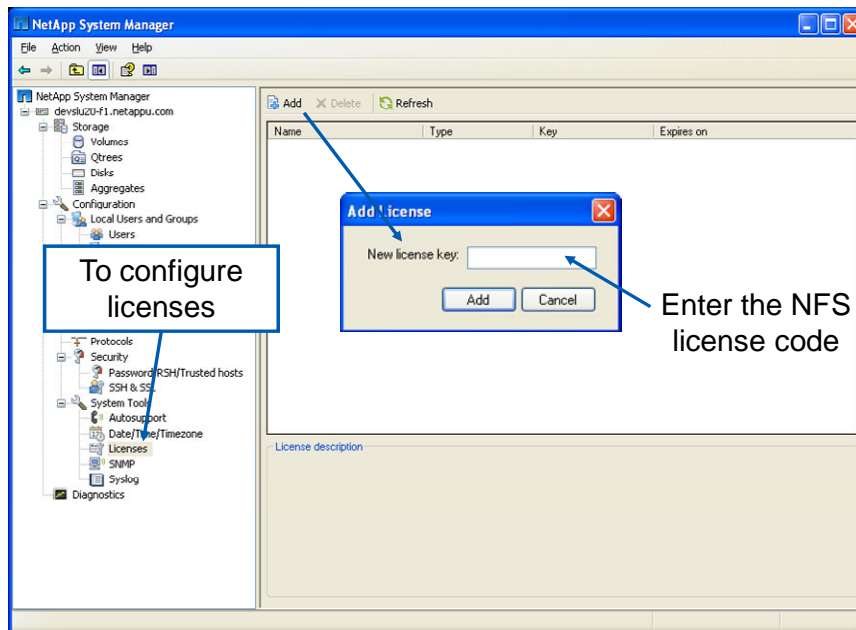
- Setting up NFS on the clients and server (storage system) involves:
 - Step 1: Licensing NFS on the storage system
 - Use `license add` or NetApp System Manager
 - Starts daemons (`mountd` and `nfsd`) that handle NFS remote procedure call protocol
 - Step 2: Configuring NFS
(discussed in the remainder of this module)
 - Step 3: Exporting file systems from the storage system (discussed in the next module)
 - Step 4: Mounting file systems on clients
(discussed in the next module)

© 2010 NetApp, Inc. All rights reserved.

NFS CONFIGURATION



System Manager: NFS Setup



© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: NFS SETUP



System Manager: NFS Setup (Cont.)

NetApp System Manager

devsluzo-f1.netappu.com

Storage

- Volumes
- Shared Folders
- Exports
- Qtrees
- Disks
- Aggregates
- Configuration
- Local Users and Groups
- Users
- Groups
- Network
- DNS
- Interfaces
- Network Files
- NIS
- Protocols
- Security
- Password(RSH)/Trusted hosts
- SSH & SSL
- System Tools
- Autosupport
- Date/Time/Timezone
- Licenses**
- SNMP
- Syslog
- Diagnostics

Add X Delete Refresh

Name	Type	Key	Expires on
UNIX: Exports(NFS)	Evaluation	TPQIZBN	7/17/2009 10:10:43 AM

The newly added license code

Exports Added

[License description](#)

The storage system requires a software license to enable NFS services. The license is installed on the storage system at the factory per your order, so the initial setup of your storage system does not involve entering license codes.

© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: NFS SETUP (CONT.)



NFS Versions

- As stated in Module 2, the storage system can support NFS v2, v3, or v4
- But what version do you choose?
- Version 2 is the default and can't be disabled
- Version 3 is a common choice because:
 - It's backward compatible with v2
 - It supports 64-bit file size
 - It has asynchronous writes, which eliminates the synchronous write system blockages of v2

© 2010 NetApp, Inc. All rights reserved.

NFS VERSIONS

NFS v3 has clear advantages:

- The client and server implementations of NFS v3 provide backward compatibility with NFS v2 by supporting both NFS v2 and NFS v3.
- The 64-bit extensions in NFS v3 support both 32-bit and 64-bit clients and servers.
- NFS v3 asynchronous writes eliminate the synchronous write system blockages in NFS v2.

In NFS v2, all operations that modify the file system must be committed to stable storage before the remote procedure call can be acknowledged. Most servers do not have battery-backed memory; the stable storage requirement means that all written data must be on the disk before the servers can reply to the remote procedure call. For a growing file, an update may require up to three synchronous disk writes: one for the inode to update its size, one for the indirect block to add a new data pointer, and one for the new data itself. Each synchronous write takes several milliseconds; this delay severely restricts the write throughput for any given client file. Version 3 of the NFS protocol eliminates some of the synchronous writes by adding a new asynchronous write remote procedure call request. When such a request is received by the server, it is permitted to acknowledge the remote procedure call without writing the new data to stable storage.

Typically, a client will do a series of asynchronous write requests followed by a commit remote procedure call request when it reaches the end of the file or it runs out of buffer space to store the file. The commit remote procedure call request causes the server to write any unwritten parts of the file to stable storage before acknowledging the commit remote procedure call. The server benefits by having to write the inode and indirect blocks for the file only once per batch of asynchronous writes, instead of on every write remote procedure call request. The client benefits from having higher throughput for file writes. The client does have the added overhead of having to save copies of all asynchronously written buffers until a commit remote procedure call is done, because the server may crash before having written one or more of the asynchronous buffers to stable storage. When the client sends the commit remote procedure call, the acknowledgment to that remote procedure call tells which of the asynchronous blocks were written to stable storage. If any of the asynchronous writes done by the client are missing, the client knows that the server has crashed during the

asynchronous-writing period, and resends the unacknowledged blocks. After all the asynchronously written blocks have been acknowledged, they can be dropped from the client cache.

For more information, see *NFS Version 3 Design and Implementation*, which can be found at http://media.netapp.com/documents/NFSv3_Rev_3.pdf.



NFS Versions (Cont.)

- Unlike previous versions, NFS v4:
 - Integrates file locking
 - Provides strong security
 - Enables better support for access over the Internet
- NFS v4 does it all; it:
 - Eliminates the need for the Network Lock Manager Protocol and the Mount Protocol in v2 and v3
 - Implements mandatory file locking
 - Uses well-defined ports that easily transit through firewalls
 - Groups several remote procedure calls to increase performance

© 2010 NetApp, Inc. All rights reserved.

NFS VERSIONS (CONT.)

NFS v4 introduces a major structural change to the protocol compared to prior versions and the elimination of ancillary protocols. In NFS v2 and v3, the Mount protocol was used to obtain the initial file handle, while file locking was supported by way of the Network Lock Manager protocol. NFS v4 is a single protocol that uses a well-defined port, which, coupled with the use of TCP, allows NFS to easily transit firewalls to enable support for the Internet. As in WebNFS, the use of initialized file handles obviates the need for a separate Mount protocol. Locking has been fully integrated into the protocol—which was also required to enable mandatory locking. The lease-based locking support adds significant state (and concomitant error recovery complexity) to the NFS v4 protocol.

Another structural difference between NFS v4 and its predecessors is the introduction of a COMPOUND remote procedure call procedure that allows the client to group traditional file operations into a single request to send to the server. In NFS v2 and v3, all actions were remote procedure call procedures. NFS v4 is no longer a "simple" remote-procedure-call-based distributed application. In NFS v4, work is accomplished through operations. An operation is a file system action that forms part of a COMPOUND procedure. NFS v4 operations correspond functionally to remote procedure call procedures in former versions of NFS. The server in turn groups the operation replies into a single response. Error handling is simple on the server—evaluation proceeds until the first error or last operation whereupon the server returns a reply for all evaluated operations.

See *The NFS Version 4 Protocol* at <http://www.netapp.com/library/tr/3085.pdf> for more information.

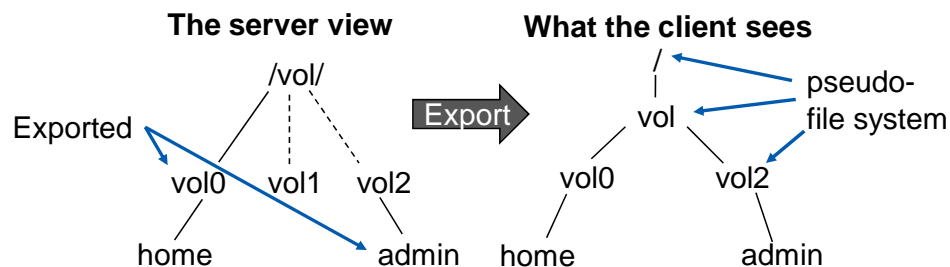


NFS Versions (Cont.)

■ NFS v4

- Requires Kerberos v5 implementation (*discussed in Module 4*)
- Requires TCP
- Creates pseudo-file systems
 - Separate volumes may be exported with a single common root

```
system> mount -t nfs4 system:/ /mnt/system
```



© 2010 NetApp, Inc. All rights reserved.

NFS VERSIONS (CONT.)

On most UNIX® systems, when a system provides files to share, or “export,” the exports are relative to root or /. Because Data ONTAP® provides the ability to create multiple volumes, there was a need to distinguish between the volumes and therefore the /vol/volumename convention was created. In NFS v3, if you do not use the complete path (/vol/) and you mount /, the mount path is assumed to be /vol/vol0 if vol0 is the root. This default was created to maintain compatibility with most UNIX systems. In NFS v4, if you mount /, you will mount the root of a pseudo file system.

To better see this, consider the following example: The storage system has three volumes: vol0, vol1, and vol2 with the following exports: /vol/vol0 and /vol/vol2/admin. In NFS v4, the server provides a single view of the exported file system to the client as shown in this slide regardless of which volume the resource originated from.

When a server chooses to export a disconnected portion of its namespace (that is, vol0 and vol2/admin), the server creates a pseudo-file system to bridge the unexported portions of the namespace allowing a client to reach the export points from the single common root (that is, /). This pseudo-file system is a structure containing only directories that allows a client to browse the hierarchy of exported file systems. The client can notice the underlying volume transitions on the server by observing that the fsID changes. The client's view of the pseudo-file system is limited to those paths that lead to exported file systems. Because /vol/vol1 is not exported in this example, it does not appear to the client during browsing operations as shown in the client's view in this slide. But a pseudo vol2 does appear because it is part of the export /vol/vol2/admin.

To explore this further, we will look at an actual example:

```
system> exportfs /vol/vol0 -sec=sys:krb5, rw /vol/vol2/admin -sec=sys:krb5, rw
```

Only vol0 and vol2/admin are exported. When a client attempts to mount / on the storage system with the command `mount -t nfs4 system:/ /mnt/system` using NFS v4, the storage system will create a pseudo file system. This pseudo-file system will consist of the following:

- /vol – a pseudo-file system

- /vol/vol0 – a real file system
- /vol/vol2 – a pseudo-file system
(note it is a pseudo-file system because vol2 was not explicitly exported)
- /vol/vol2/admin – a real file system

To explore the pseud-file system, navigate to the mount path:

```
# cd /mnt/system
# lsvol0    vol2
# cd vol2
# lsadmin
```

Notice that moving from one file system (/vol) to another (/vol/vol2) is seamless. Back on the storage system, we can see that this is actually changing volumes by referring to the underlying fsID.

```
system*> showfh4 -v /vol/vol (really /vol): exp.fileid=0x00042b
exp.snapgen=0x000001 flags=0x05 snapid=0x0 unused=0x0 fileid=0x00042b
gen=0x000001 fsid=0x000002 handle_type=0x02

system*> showfh4 -v /vol/vol0/vol/vol2 (really /vol/vol2): exp.fileid=0x000040
exp.snapgen=0x1da07d flags=0x00 snapid=0x0 unused=0x0 fileid=0x000040
gen=0x1da07d fsid=0xe802e1 handle_type=0x02

system*> showfh4 -v /vol/vol2/admin/vol/vol2/admin (really /vol/vol2/admin):
exp.fileid=0x000064 exp.snapgen=0x355073 flags=0x00 snapid=0x0 unused=0x0
fileid=0x000064 gen=0x2355073 fsid=0xc1504d handle_type=0x00
```



Configuring NFS Versions

NFS Version	Option	Default Value
v2	<code>nfs.v2.df_2gb_lim</code>	off
v3	<code>nfs.v3.enable</code>	on
v4	<code>nfs.v4.enable</code>	off
	<code>nfs.v4.id.domain</code>	off
	<code>nfs.v4.acl.enable</code>	off
	<code>nfs.v4.read_delegation</code>	off
	<code>nfs.v4.write_delegation</code>	off
	<code>nfs.v4.setattr_acl_preserve</code>	off

© 2010 NetApp, Inc. All rights reserved.

CONFIGURING NFS VERSIONS

By default, NFS v2 and v3 are enabled. You can use the options command to enable versions that are not currently enabled.



Network Configuration

- `nfs.tcp.enable` `on`
 - When enabled, NFS uses TCP as transport
- `nfs.udp.enable` `on`
 - When enabled, NFS uses User Datagram Protocol (UDP) as transport
- `nfs.udp.xfersize` `32768`
 - Data transfer size in bytes
- `nfs.ifc.xmt.high` `16`
 - High-limit transmit-flow control value
- `nfs.ifc.xmt.low` `8`
 - Low-limit transmit-flow control value

© 2010 NetApp, Inc. All rights reserved.

NETWORK CONFIGURATION

`nfs.tcp.enable`

- When enabled, NFS uses TCP as transport.

`nfs.udp.enable`

- When enabled, NFS uses User Datagram Protocol (UDP) as transport.

`nfs.udp.xfersize`

- This is the UDP data transfer size. The default is 32,768.

`nfs.ifc.xmt.high`

- The limit of outstanding requests for which NFS will go into flow control. The default is 16 and the maximum is 64. This is a persistent option.

`nfs.ifc.xmt.low`

- The limit of outstanding requests for which NFS comes out of flow control. The default value for this option is set to 8. Its minimum value is 0. This is a persistent option.

Please see the *Data ONTAP 8.0 Network Management Guide* for details about each configuration option.



Other NFS Configurations

- In an NFS environment with a local identification model:
 - Hostname-to-IP resolution must be configured
 - Netgroups will likely also be configured
 - The resolution mechanism must be verified

© 2010 NetApp, Inc. All rights reserved.

OTHER NFS CONFIGURATIONS



/etc/hosts file

- List of IP addresses followed by the hostname
- There are three types of entries:
 - Local hosts loopback device
 - Local hostname
 - Remote hosts

```
system> rdfile /etc/hosts
#Auto-generated by setup Tue May 8 1:01:01
127.0.0.1      localhost
10.61.77.156   system    system-e0a
#0.0.0.0       system    system-e0b
#0.0.0.0       system    system-e0c
#0.0.0.0       system    system-e0d
10.61.77.122   lux-client
```

© 2010 NetApp, Inc. All rights reserved.

/ETC/HOSTS FILE

Each entry in the /etc/hosts file lists an IP address followed by the hostname and any aliases for that host. The hosts file has three types of entries, containing information about the:

- Local hosts loopback device, which ensures that data packets sent from a machine to itself are not sent on to the network
- Local hostname
- Remote hosts



Using Netgroups

- The following `/etc/netgroup` snippet contains three netgroups:

Case-sensitive file

```
trustedhosts (host1,,) (host2,,)
untrustedhosts (host3,,) (host4,,) (host5,,)
allhosts trustedhosts untrustedhosts
```

- In an exports entry, you can specify the `trustedhosts`, `untrustedhosts`, or `allhosts` netgroup as the export target with the `rw`, `ro`, and `root` options (*discussed in the next module*)

© 2010 NetApp, Inc. All rights reserved.

USING NETGROUPS

`/etc/netgroup` defines network-wide groups used for access permission checking during remote mount request processing.

Each line defines a group and has the format:

- `groupname member-list`

Each element in member-list is either another group name or a triple of the form:

- `(hostname, username, domainname)`

Network groups can also be stored in a network information services, such as LDAP, NIS, or NIS+ (in NIS compatibility mode only).



/etc/nsswitch.conf file

■ /etc/nsswitch.conf file

- Determines the order in which identification systems are queried

```
system> rdfile /etc/nsswitch.conf
```

```
#Auto-generated by setup Fri Jun 30 07:35:27
```

```
hosts:    files    nis    dns
passwd:   files    nis    ldap
netgroup: files    nis    ldap
group:    files    nis    ldap
shadow:   files    nis
```

 /etc/hosts

© 2010 NetApp, Inc. All rights reserved.

/ETC/NSSWITCH.CONF FILE



NIS

© 2010 NetApp, Inc. All rights reserved.

NIS



NFS Versus NFS and NIS

- Architecture options:
 - NFS alone server-client model is used in small networks for easier localized maintenance, but is not scalable
 - NFS server with NIS server-client model is used in large distributed networks, provides centralized maintenance, and is scalable

© 2010 NetApp, Inc. All rights reserved.

NFS VERSUS NFS AND NIS

Both NFS and NIS are client-server applications, which means that they sit at the top layer of the protocol stack and use External Data Representation (XDR) and remote procedure call services.

In addition to NFS servers, NIS servers are typically used in large distributed networks.

A major problem in running a distributed computing environment is maintaining separate copies of common configuration files, such as the passwd, group, and hosts files. Ideally, the network should be consistent in its configuration so that users do not have to worry about where they have accounts or if they will be able to find a new machine on the network. Preserving consistency, however, means that every change to one of these common files must be propagated to every host on the network, which is difficult and not scalable. The NIS addresses these problems. It is a distributed system that replaces copies of commonly replicated configuration files with a centralized management facility. Machines that are using NIS retrieve information from one centralized database that maintains updates and propagates changes to the rest of the network. Files that are generally the same on all hosts in a network, such as /etc/passwd and /etc/hosts, reside on the NIS database.

Typically, NIS is a lookup service that NFS and mount depend on. It performs host lookup from the export maps, reverse lookups, and so on.



NIS Configuration

- NIS is used to resolve:
 - User
 - Hostname-to-IP resolution
 - Netgroup
- To configure NIS:
 - `nis.enable` on
 - Default is off
 - `nis.domainname` `domain_name`
 - `nis.servers` `server_name_or_ip`, `server...`
- Use `nis info` to display configuration information
- Other commands: `ypcat`, `ypgroup`, `ypmatch`, `ypwhich`

NOTE: The storage system will only work with an NIS+ server if NIS+ server is set to NIS-compatibility mode

© 2010 NetApp, Inc. All rights reserved.

NIS CONFIGURATION



NIS Slave Mode

- NIS slave
 - Downloads NIS maps from master servers every 45 minutes to the storage system
 - Handles all NIS lookups from local NIS slave maps
- To configure Data ONTAP:
 - `nis.slave.enable` on
 - Default is off

© 2010 NetApp, Inc. All rights reserved.

NIS SLAVE MODE

Data ONTAP 7.1 or later can be configured as an NIS slave. The slave can be turned on or off using the following option:

```
options nis.slave.enable on | off
```

After the maps are downloaded by the slave, all NIS requests are serviced using the downloaded maps. There will not be any NIS requests going to the NIS servers. If the slave is disabled the storage system will revert back to the client behavior. The slave has two major functions:

- Download the maps from the NIS master: The NIS slave checks every 45 minutes with the master server for updates. If there are updates, these updates are downloaded.
- Service YPPUSH requests. All other NIS/YP requests are denied. If the storage system is configured as a slave on the NIS master, when the maps on the master are updated, the administrator has an option to notify all the slaves.

The downloaded maps are stored under `/etc/yp/<NIS_domain_name>/`. There needs to be sufficient space on the root volume of the storage system to download maps for the slave to function correctly. The amount of space needed depends on the size of the maps. It takes almost the same size as maps on the NIS server. The maps are stored in a database file and you can verify the data in each of the map database files using a `"db_dump185 -p <map_name>"`.

NOTE: The NIS slave mode is for storage system use only, not to serve NIS requests to other NIS clients.

Slave mode caches a copy of the information that would otherwise be on the NIS master server. This allows for better lookup performance.



DNS and LDAP

© 2010 NetApp, Inc. All rights reserved.

DNS AND LDAP



DNS

- DNS performs hostname-to-IP resolution
- To configure:
 - `dns.enable on`
 - `dns.domainname domainname`
 - `dns.cache.enable on`
 - Disabling cache – clears cache
 - `dns flush` command
 - Clears cache without disabling the cache
 - `dns info` command
 - Displays configuration information
 - Modify `/etc/resolv.conf` (*discussed next*)

© 2010 NetApp, Inc. All rights reserved.

DNS

The DNS is the name service provided by the Internet for TCP/IP networks. It was developed so that workstations on the network could be identified with common names instead of Internet addresses. DNS performs naming between hosts within your local administrative domain and across domain boundaries. The collection of networked workstations that use DNS are referred to as the DNS namespace. The DNS namespace can be divided into a hierarchy of domains. DNS uses Secure Sockets Layer (SSL) to authenticate users before they can change definitions.

DNS is the name resolution system used for wide-area networks such as the Internet. DNS in UNIX uses a resolver configuration file (the `/etc/resolv.conf` file). This file lists the domain and name servers available on the local network, which the system can use to resolve name queries of remote machines. The resolver uses the domain list when translating names that are not fully qualified. It queries the name server when attempting to look up a name.



/etc/resolv.conf File

- Has two types of entries:
 - Search (domain) entry lists the names of up to six local DNS domains
 - Nameserver entries
- The resolver queries the DNS servers in the order in which they are listed

```
# cat /etc/resolv.conf
nameserver 215.243.23.25
nameserver 10.61.77.193
```

© 2010 NetApp, Inc. All rights reserved.

/ETC/RESOLV.CONF FILE

For DNS, the /etc/resolv.conf file has two types of entries, with each entry structured as a keyword followed by a value. The file's search entry lists the names of up to six local DNS domains to search. These domains are arranged from specific to general, so subdomains are listed before the parent domains. Search domain names are appended to partially qualified hostnames when a lookup is performed.

Instead of the search keyword, some older systems feature the domain keyword, which specifies the local domain name only.

There are nameserver entries, each indicated by the nameserver keyword listed after the search entry in the /etc/resolv.conf file. These specify the DNS nameservers and their IP addresses.

When attempting to look up a name, the resolver queries the name servers in the order in which they are listed in the file. Therefore, the name server closest to the host should be listed first in this file to ensure faster name resolution times. If a request times out, the system queries the next server listed. If no name server responds, the system starts again with the first name server listed.



LDAP

- LDAP centrally maintains users and groups
- To configure LDAP on a storage system:
 - `ldap.servers.preferred servername,...`
 - `ldap.servers servername, servername...`
 - `ldap.port port_number`
 - `ldap.ssl.enable`
 - If enabled, provide key with `keymgr install root` command
 - Configure `/etc/nsswitch.conf` to use LDAP
 - `netgroup: ldap files nis`

© 2010 NetApp, Inc. All rights reserved.

LDAP

Data ONTAP chooses an LDAP server based on your LDAP server option settings.

See the *Data ONTAP 8.0 Network Administration Guide* for more information.

NOTE: For more information, please see Technical Report 3458 for UNIX authorization using Microsoft® Active Directory LDAP server and Technical Report 3464 for UNIX-based LDAP Servers.



LDAP (Cont.)

- To configure LDAP (Cont.)
 - `ldap.base name`
 - Example name:
“`c=ntap,c=us`”
 - `ldap.base.passwd name`
 - Example name:
“`ou=People,dc=domain,dc=com`”
 - `ldap.base.group name`
 - Example name:
“`ou=Groups,dc=domain,dc=com`”

© 2010 NetApp, Inc. All rights reserved.

LDAP (CONT.)

The LDAP base is the distinguished name of the LDAP tree in which user information is stored. All lookup requests sent to the LDAP server will be limited to the search base and scope specified by the `ldap.base` option value, unless further restricted by a more specific base and scope lookup value, such as `ldap.base.passwd` or `ldap.base.group`.



PC-NFS and WebNFS

© 2010 NetApp, Inc. All rights reserved.

PC-NFS AND WEBNFS



PC-NFS

- PC-NFS allows non-UNIX clients to mount file system paths
- To configure PC-NFS:
 - `pcnfs.enable` on
 - Default is `off`
 - Create local users through `/etc/passwd` file or `/etc/passwd` and `/etc/shadow` files
 - Create local groups through `/etc/groups`
 - To determine the default umask (permissions) setting when PC-NFS client creates files
 - `pcnfsd.umask umask_number`

© 2010 NetApp, Inc. All rights reserved.

PC-NFS

Unlike NFS users, PC-NFS users cannot execute the UNIX `umask` command to set the file mode creation mask (umask), which determines the default file permissions. However, Data ONTAP defines a umask for all PC-NFS users.

The permissions for each file are defined by three octal values, which apply to owner (sometimes called user), group, and other (sometimes called world). When a PC-NFS client creates a new file, Data ONTAP subtracts the umask, which is a three-digit octal number from 666. The results are the file permissions for the new file.

Digit in the umask	Description
0	Read and write permission
2	Write permission
4	Read-only permission
6	No permission



WebNFS

- WebNFS extends NFS to the Internet
- Access files through URLs such as
`nfs://computer.site.com/filedirectory/file`
- To configure WebNFS:
 - `nfs.webnfs.enable` `off` Default
↙
 - Change the default and enable WebNFS
 - `nfs.webnfs.rootdir` `XXX`
 - Root directory for WebNFS
 - `nfs.webnfs.rootdir.set` `off`
 - Enables or disables root directory for WebNFS

© 2010 NetApp, Inc. All rights reserved.

WEBNFS

`nfs.webnfs.enable`

- Enables WebNFS. Default is **off**.

`nfs.webnfs.rootdir`

- Specifies the WebNFS rootdir. Once the rootdir is set, WebNFS clients can issue lookups relative to the rootdir. The default value for this option is `XXX`.

`nfs.webnfs.rootdir.set`

- This option needs to be enabled for the rootdir setting to take effect. Disabling this option disables the existing rootdir setting.



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Configure NFS on a NetApp storage system
- Add NIS to manage users, groups, and name-to-IP resolution
- Administer a storage system to perform DNS lookups
- Configure a storage system to access an LDAP server to centrally manage users and groups
- Set up PC-NFS and WebNFS environments to extend the reach of NFS

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 3: NFS Setup
Estimated Time: 20 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- What command would you use to view configured NFS characteristics?
- Can you use NFS v4 over UDP?
- NFS v4 doesn't require Kerberos. True or false?
- Can Data ONTAP be an NIS slave?
- Will Data ONTAP use an NIS+ environment?
- LDAP can be configured to use SSL. True or false?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

Exports and Mounts

Module 4
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



EXPORTS AND MOUNTS



Module Objectives

By the end of this module, you should be able to:

- Identify exportable resources
- Export and unexport resources to clients, subnets, and netgroups
- Administer storage system with the `exportfs` command
- Create mountpoints and mount exported resources on a client
- Monitor the usage of exported resources

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Exports Using the Command-Line Interface

© 2010 NetApp, Inc. All rights reserved.

EXPORTS USING THE COMMAND-LINE INTERFACE



Exporting Resources

- To export resources with NFS:
 - Modify or update `/etc/exports` file
 - Syntax: `path options`
 - Example: `/vol/volx -rw=host1:host2`
 - Load file with `exportfs -a` or `reboot` or `nfs off` and `nfs on`
 - Run the `exportfs` command on the command-line interface
 - Syntax: `exportfs options path`
 - Example: `exportfs -o rw=host1:host2 /vol/volx`

NOTE: The export must be in memory to be accessible

© 2010 NetApp, Inc. All rights reserved.

EXPORTING RESOURCES

The `/etc/exports` file has the following limitations:

- Contains up to 10,240 entries with no limit to size of entry
- Entry can span multiple lines
- Entry consists of path name and options



Resources

- An export provides resources to targets
- Exportable resources:
 - Volume
 - Directory
 - Qtree
 - File
- Exported resources will have default access security (covered in detail later)
 - Example: read-write (rw) or read-only (ro)

© 2010 NetApp, Inc. All rights reserved.

RESOURCES

Resources are made available to client hosts (targets) through exports. The following storage objects may be exported: volumes, directories, qtrees, and individual files.



Targets

- Target examples from */etc/exports*:
 - Host – use the hostname or IP address
 - `/vol/vol0/home -rw=venus`
 - Netgroup – use the group name
 - `/vol/vol0/home -rw=mygroup`
 - Subnet – specify the subnet address
 - `/vol/vol0/home -rw="192.168.0.0 255.255.255.0"`
 - `/vol/vol0/home -rw=192.168.0.0/24`
 - DNS domain
 - `/vol/vol0/home -rw=".eng.netapp.com"`

© 2010 NetApp, Inc. All rights reserved.

TARGETS

Examples of destinations to which resources are exported include:

Host – Typically the UNIX®/Linux® host system connected to the storage system.

Netgroup – A netgroup is a network-wide group of machines granted identical access to certain network resources for security and organizational reasons. Use an @ sign to indicate that a name is a group if you have host the same name as a netgroup.

Subnet – A subnet is a physical grouping of connected network devices. Nodes on a subnet tend to be located in close physical proximity to each other on a LAN.

DNS Subdomain – A subdomain is a domain that is part of a larger domain. A DNS hierarchy consists of the root-level domain at the top, underneath which are the top-level domains, followed by second-level domains, and finally the subdomains.



Rules for Exporting Resources

- Specify complete path name; must begin with `/vol` prefix
 - Example: `/vol/vol0/home`
- Cannot export `/vol`, which is not a path name to a file, directory, or volume
 - Export each volume separately
- When you export a resource to multiple targets, separate the target names with a colon (`:`)
Example: `/vol/vol0/home -rw=venus:mars`

© 2010 NetApp, Inc. All rights reserved.

RULES FOR EXPORTING RESOURCES

If the storage system has multiple volumes, export each volume separately; it is not recommended to export volumes by specifying `/vol` as the exported directory. Although this method works in a UNIX environment, it doesn't work in a Data ONTAP® environment.

Example: For a storage system with four volumes – `vol0`, `users`, `builds`, and `testnfs`, make sure all four volumes are exported separately.

- `/vol/vol0 -root=toaster,rw=venus:mars`
- `/vol/users -root=toaster,ro=venus:mars`
- `/vol/builds -root=toaster,rw=venus:mars`
- `/vol/testnfs -root=toaster,rw=venus:mars`

You cannot export `/vol` by itself, for example: `/vol -root=toaster, -rw=mars:venus` is not allowed.



Rules for Exporting Resources (Cont.)

- Storage system must resolve hostnames using DNS, NIS, or */etc/hosts* per order in */etc/nsswitch.conf*
- Exporting ancestors and descendants is allowed
 - `/vol/vol1/home -rw=admingroup`
 - `/vol/vol1/home/jim -rw=jim_host`
- Storage system determines permissions by matching the longest prefix

© 2010 NetApp, Inc. All rights reserved.

RULES FOR EXPORTING RESOURCES (CONT.)

To export directories to hosts, the storage system must be able to resolve hostnames into IP addresses. Resolution can be based on DNS, NIS, or */etc/hosts* file entries.

EXPORTING PARENTS AND CHILDREN (ANCESTORS AND DESCENDANTS)

The storage system permits directories that have exported ancestors to be exported. In many implementations of the UNIX operating system, a directory cannot be exported if an exported parent (ancestor) is in the same file system.

The following example shows exports allowed by the storage system:

```
/vol/vol0 -rw=adminhost,root=adminhost/vol/vol0/home -rw=blender:mixer
```

In this example, the volume “vol0” and directory “home” are explicitly specified in the */etc/exports* file, with root access granted to adminhost for “vol0” and access granted to both “blend and mixer” for the “home” directory.

DETERMINING PERMISSIONS BY STORAGE SYSTEM MATCHING LONGEST PREFIX

In the preceding example, a client mounting `/vol/vol0/home/user1` gets permissions for `/vol/vol0/home` because `/vol/vol0/home` is the longest matching prefix. A client mounting `/vol/vol0` gets `-rw=adminhost` and `-root=adminhost` permissions.

For example, the following shows a */etc/exports* file that creates a security breach by enabling any host to mount the `/vol/vol0` directory while restricting specific hosts from mounting the `/vol/vol0/home` directory. In this example, any host can gain access to the `/vol/vol0/home` directory by mounting the `/vol/vol0` directory:

```
/vol/vol0/vol/vol0/home -rw=bashful:dopey:sleepy
```



Access Restrictions

- Access restrictions specify what operations a target can perform on a resource
 - Whether access is read-write (`rw`) or read-only (`ro`)
 - What is the effective user ID (UID) of a client root user
 - What is the actual path of the mountpoint
 - What is the effective UID of all root users
 - Whether files can be created with the `SETUID` or `SETGID` bit
 - Whether UNIX Auth_SYS security is used or Kerberos

© 2010 NetApp, Inc. All rights reserved.

ACCESS RESTRICTIONS

When you export a resource, you can specify the access restrictions that govern how the resource can be mounted. If you export a resource without specifying access restrictions, it can be mounted read-write by all hosts.

You specify access restrictions using export options in the `/etc/exports` file or with command-line options for the `exportfs` command. (These options are identical using either method.) Access restriction options determine:

- Which hosts can mount the resource
- Whether the resource can be mounted read-write (`rw`) or read-only (`ro`)
- Whether the root user on the client can access the resource and its user ID (UID)
- Whether files can be created with the `setuid` or `setgid` bit

When multiple restrictions are applied to the same resource, the most restrictive option always takes precedence.



Access Options

- Default is read-write (`rw`) and UNIX `AUTH_SYS` (`sys`)
- To specify something different:
 - `ro` option provides read-only access to all hosts
 - `ro=` option provides read-only access to specified hosts
 - `rw=` option provides read-write access to specified hosts
 - `root=` option specifies that root on the target has root permissions for this resource when it is mounted from the storage system; use this to specify root access to specific hosts or for all hosts
- If a different option is set, everyone else gets nothing

© 2010 NetApp, Inc. All rights reserved.

ACCESS OPTIONS

THE ROOT OPTION

This option determines the UID for the root user on the client who sends the NFS requests after mounting the resource. The following list describes the effect of including or excluding a host from the root option:

- If you specify a host with the root option, the root user on that host keeps the root UID (0) when accessing the resource.
- The `rw` option gives read-write access to specified hosts; if no host is specified, all hosts have read-write access.
- The `ro` option gives read-only access to specified hosts; if no host is specified, all hosts have read-only access.
- If no root option is specified for the resource, or if the host trying to access the resource is not specified with the root option, access by root on the client is either denied or subject to modification by the anon option.



How the Access Rules Work

Example:

If the `/etc/exports` file contains:

```
/vol/vol172 -ro=host1:host3,rw=host2,root=host2
```

- Only host1, host2, and host3 can access /vol/vol172
- Read-write access is granted to host2
- Read-only access is granted to both host1 and host3
- Root access is granted to host2

NOTE: Certain invalid combinations when entered for options in exports file return errors on Data ONTAP when loaded

© 2010 NetApp, Inc. All rights reserved.

HOW THE ACCESS RULES WORK

Access control in NFS exists at the server as well as the client. The server uses permissions to limit read or write access to exported resources by clients. In addition, the server can limit access by using normal UNIX “rwx” controls.

The `/etc/exports` file contains the permissions assigned to exported resources. By default, if neither the `ro` nor `rw` permissions are specified, read-write is the default permission. If `-rw=` is present, read-only is not the default for all other hosts not listed for the resource.

Also, when specifying permission levels with `ro` and `rw` options, note that the following are invalid combinations for `/etc/exports`:

- You cannot specify the `ro` option in addition to the `ro=` option (that is, you cannot specify everyone `ro`, and then also specify individual hosts with `ro`).
- You cannot specify the `rw` option in addition to the `rw=` option (that is, you cannot specify everyone `rw`, and then also specify individual hosts with `rw`).
- You cannot exclude an NFS client identifier from the `ro=` or `rw=` option and include the same NFS client identifier in the other option.
- The order of precedence of the options is:
 - The `ro` option takes precedence over the `rw` option.
 - The `ro=` option takes precedence over the `rw` option.
 - The `rw=` option takes precedence over the `ro` option.
 - The `ro=` option takes precedence over the `rw=` option.
- A hostname or IP address in the `ro=` or `rw` option takes precedence over a netgroup, subnet, or domain in the other option.
- Hostnames and IP addresses take precedence from left to right within an option.

For details of other permissions and options, see the online manual page for exports.



Actual Path Option

- Helpful with migrations
- Example:
 - A previous server had data at old:/old/server/info
 - Hundreds of NFS clients mounted this point
 - Data then migrated to FAS1:/vol/vol0/data
 - Old server is taken offline
 - FAS1 is given old server hostname/IP or alias created
 - FAS1 then uses actual option for data
 - `exportfs -o actual=/vol/vol0/data /old/server/info`
 - NFS clients' mountpoints work as normal

NOTE: Not supported in NFS v4

© 2010 NetApp, Inc. All rights reserved.

ACTUAL PATH OPTION

Beginning in Data ONTAP 6.4, there is an "actual" option that allows NFS clients to mount a storage system volume using a different pathname than the storage system's pathname. The -actual parameter allows for a server side aliasing of path names in the event that the storage has been moved and clients have not yet been updated to the new path.

An example of how you can export the path using the -actual option:

```
new_system> exportfs -o actual=/vol/vfilers/vf19,rw=pets,root=workers  
/vol/vol9/vf19
```

In the example above, you are actually exporting /vol/vfilers/vf19 on 'new_system' under the name or alias of /vol/vol9/vf19 which may have existed in 'old_system' or could have been simply a bogus name on the same storage system. However, if it existed in 'old_system', the NFS clients are probably still pointing to old_system:/vol/vol9/vf19, so you need to set an alias that allows the clients to be sent to the new storage system and new path: /vol/vfilers/vf19.

Here is how you determine the actual path to storage using the -s option:

```
new_system> exportfs -s /vol/vol9/vf19/vol/vfilers/vf19  
new_system> exportfs -s /vol/vol9/vf19/src/vol/vfilers/vf19/src
```

When you enter the exportfs command or you enter a line in the /etc/exports file, you can specify a symbolic virtual path (which may or may not exist on the storage system, but it needs to start with the /vol/) to actually export another path (the one specified in the -actual option).



Automatic Exports

- If `nfs.export.auto-update` is enabled,
 - `/etc/exports` file is automatically updated when:
 - Volume is created: `vol create volnfs 3g`
 - Volume is renamed: `vol rename volnfs vol2nfs`
 - Volume is destroyed: `vol destroy vol2nfs`
 - **NOTE:** `admin.hosts` (a hidden option) determines the defaults mount permissions
 - If set, then the auto-export feature will only grant RW to machines defined in `admin.hosts` and deny all other hosts
 - If not set, then auto-export feature will grant RW to all hosts
- When providing an adminhost during storage system setup; if it is in a different DNS domain, use a fully qualified domain name (FQDN)
- Qtrees exported by:
 - Using `exportfs -o` command
 - Modifying and loading the `/etc/exports` file

© 2010 NetApp, Inc. All rights reserved.

AUTOMATIC EXPORTS

options `nfs.export.auto-update`

The default for this option is enabled. This option controls whether automatic updates are performed on the `/etc/exports` file. If it is not set, then the commands `vol create`, `vol delete`, and `vol rename` will not automatically rewrite the file.

When security of access to the data is paramount, this value should be off.

USING FULLY QUALIFIED DOMAIN NAME (FQDN)

When providing an adminhost that is in a different domain name during setup, use a fully qualified domain name.



Common `exportfs` Switches

- Displays all current exports in memory
 - `exportfs`
- Adds exports to the `/etc/exports` file and in memory
 - `exportfs -p [options] path`
- Saves existing exports in memory to exports file
 - `exportfs -w pathname`
- Unexports an export and removes it from `/etc/exports`
 - `exportfs -z [path]`
- Reloads only exports from `/etc/exports` files
 - `exportfs -r`
- Unexports a specific export
 - `exportfs -u [path]`
- Unexports all exports with verbose output
 - `exportfs -uav`
- Verifies the actual path to which a volume is exported
 - `exportfs -s pathname`
- Displays export options per file system path
 - `exportfs -q pathname`

© 2010 NetApp, Inc. All rights reserved.

COMMON EXPORTFS SWITCHES

The `exportfs` command can be used to export directories and files, determine the current exports, check the access cache for an export, determine the actual storage path of an export, and revert the export file to a rule version prior to Data ONTAP 6.5.

If no *pathname* is provided, `exportfs` lists all currently exported directories and files. If *pathname* is provided, `exportfs` makes the specified file or directory available or unavailable for mounting by NFS clients.

In UNIX, it is illegal to export a directory that has an exported ancestor in the same file system. Data ONTAP does not have this restriction. For example, you can export the `/vol/vol0` directory and the `/vol/vol0/home` directory. In determining permissions, the storage system uses the longest matching prefix.



Sample Output of Exports

```
system> exportfs
/vol/test          -sec=sys,rw,root=10.254.134.38,nosuid
/vol/flex1/a       -sec=sys,rw
/vol/flex1/ngA     -sec=sys,rw=trusted_hosts
/vol/north         -sec=sys,rw,root=10.254.134.38,nosuid
/vol/flex1/net     -sec=sys,rw=10.254.134.0/24
/vol/vol0/home     -sec=sys,rw,root=10.254.134.38,nosuid
/vol/vol0         -sec=sys,ro,rw=10.254.134.38,root=10.254.134.38,nosuid
/vol/flex1         -sec=sys,rw,root=10.254.134.38,nosuid
/vol/flex1/mktg    -sec=sys,ro,nosuid
/vol/flex1/unix    -sec=sys,rw=10.254.134.38,root=10.254.134.10,nosuid

system> rdfile /etc/exports
...
/vol/flex1         -sec=sys,rw,root=10.254.134.38,nosuid
/vol/vol0         -sec=sys,ro,rw=10.254.134.38,root=10.254.134.38,nosuid
/vol/test         -sec=sys,rw,root=10.254.134.38,nosuid
/vol/north        -sec=sys,rw,root=10.254.134.38,nosuid
/vol/vol0/home    -sec=sys,rw,root=10.254.134.38,nosuid
```

© 2010 NetApp, Inc. All rights reserved.

SAMPLE OUTPUT OF EXPORTS

To verify which resources have been exported, use the `exportfs` command without any flag, and all exported resources will be displayed. Next, use the `rdfile` command to show which resources are in the `/etc/exports` file.

In this example, an administrator has loaded the exports from the `/etc/exports` file but has also added additional exports by way of the `exportfs` command.



Picking the Correct Export

1. Grant root access to /vol/vol0 to adminhost.

2. Grant read-write access to /vol/vol0/home to host1 and host2.

3. Grant read-write access to /vol/vol1 to host1 and read-only access to host3.

/vol/vol1 -rw=host1:host3

/vol/vol0 -rw=adminhost,root=adminhost 1

/vol/vol0/home -rw=host1:host2 2

/vol/vol3 -ro=adminhost,root=adminhost

/vol/vol1 -rw=host1,ro=host3 3

/vol/vol0/home -rw=host1,ro=host2

© 2010 NetApp, Inc. All rights reserved.

PICKING THE CORRECT EXPORT



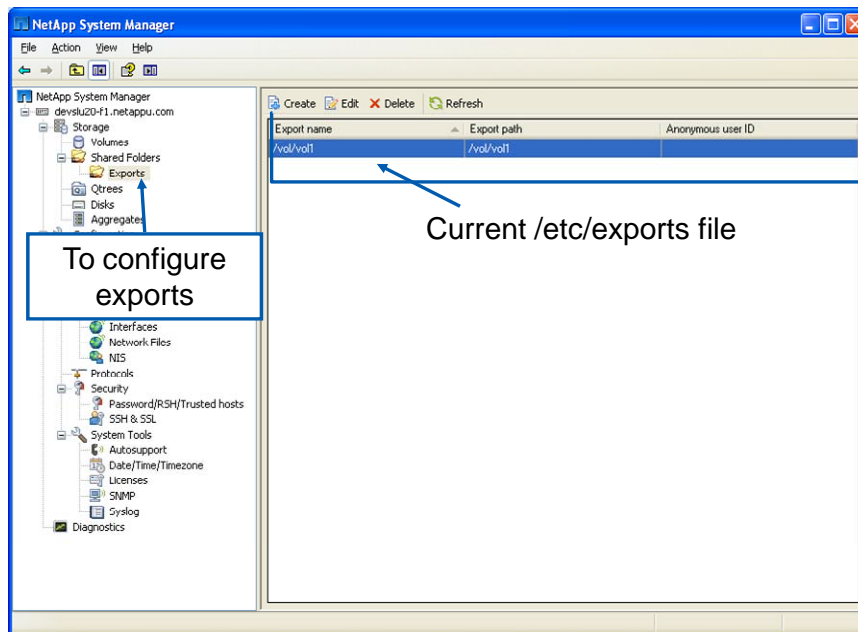
Exports Using NetApp System Manager

© 2010 NetApp, Inc. All rights reserved.

EXPORTS USING NETAPP SYSTEM MANAGER



NetApp System Manager: Exporting

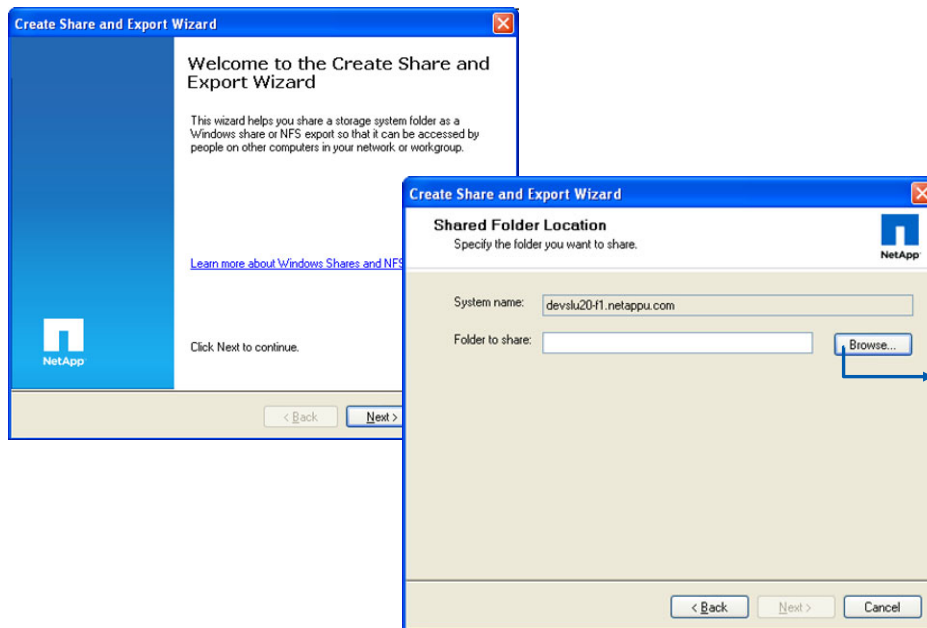


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING



NetApp System Manager: Exporting (Cont.)

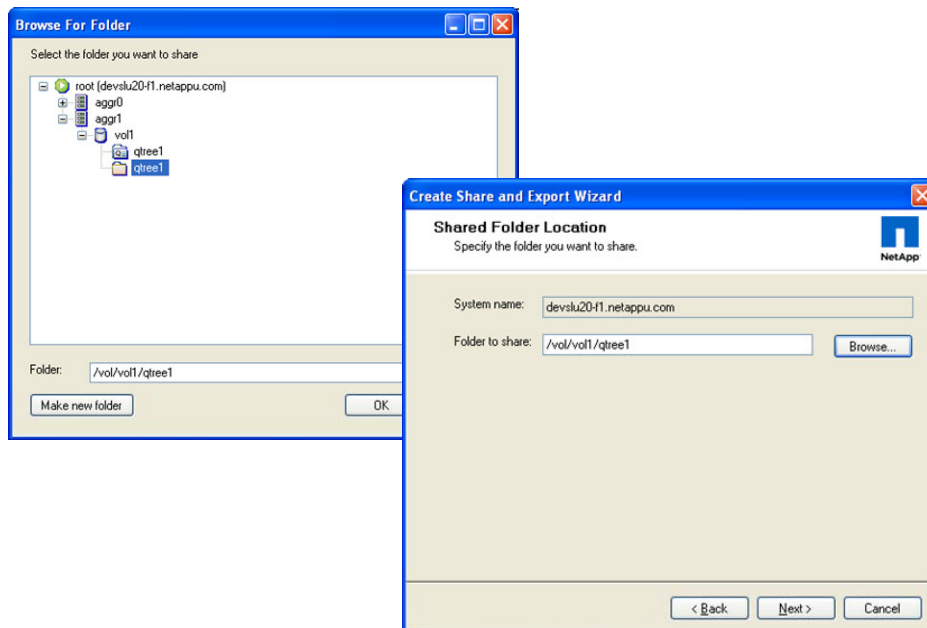


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING (CONT.)



NetApp System Manager: Exporting (Cont.)

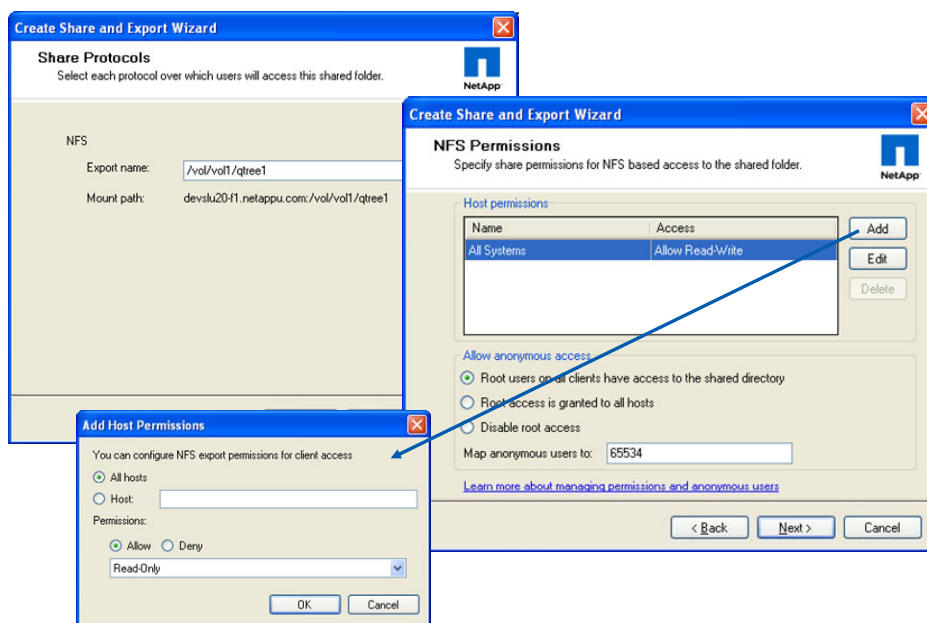


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING (CONT.)



NetApp System Manager: Exporting (Cont.)

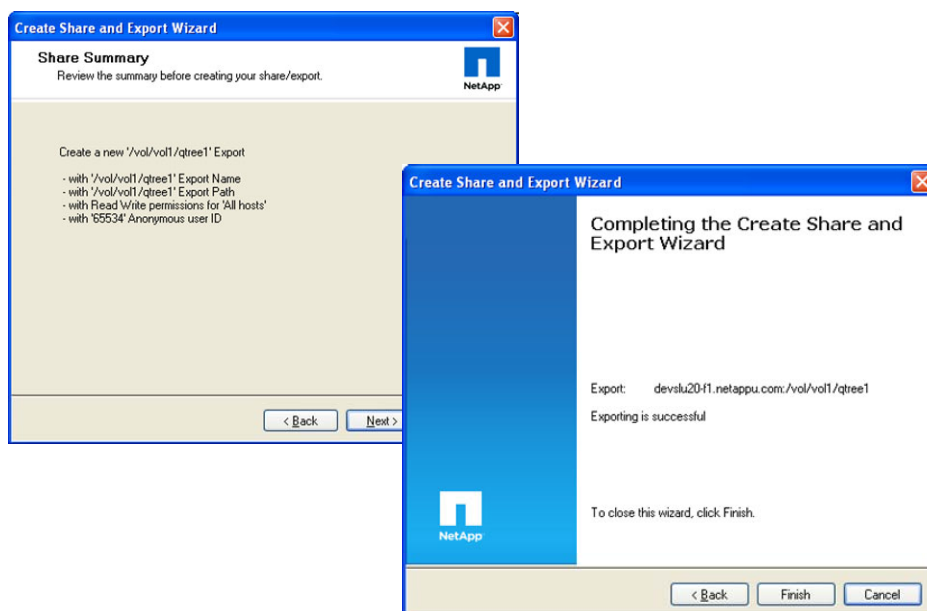


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING (CONT.)



NetApp System Manager: Exporting (Cont.)

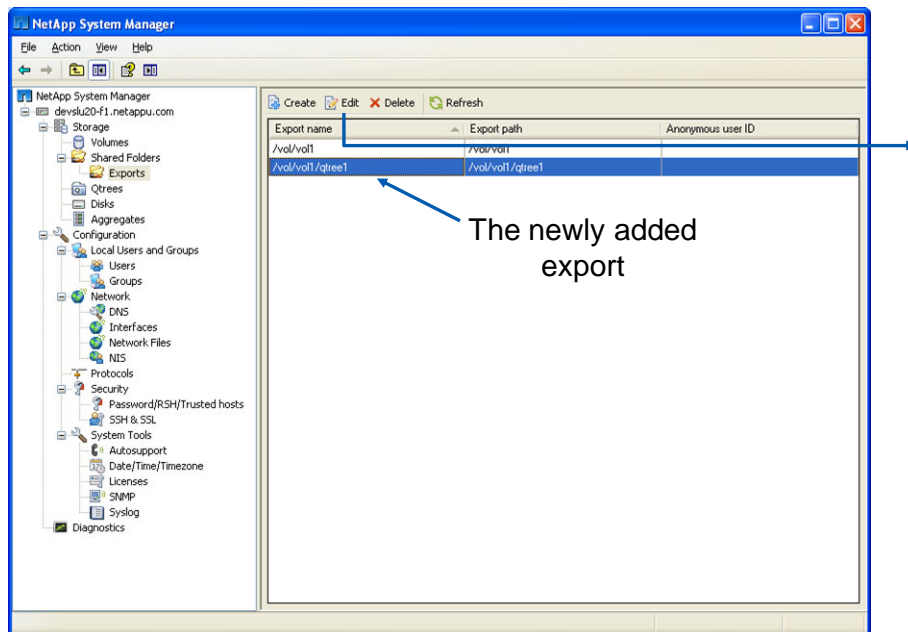


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING (CONT.)



NetApp System Manager: Exporting (Cont.)

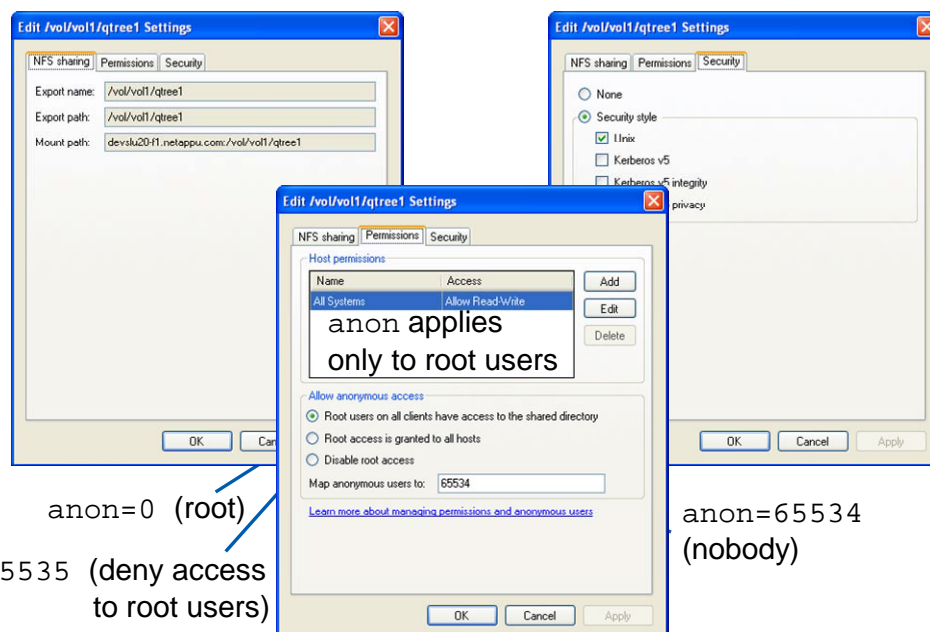


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING (CONT.)



NetApp System Manager: Exporting (Cont.)



© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: EXPORTING (CONT.)



Mounts

© 2010 NetApp, Inc. All rights reserved.

MOUNTS



Mounts

- Mounts are used to attach a storage system's exported hierarchy to the target's file system hierarchy
 - Requires a mountpoint—that is, a directory
 - Mounted by way of:
 - The `mount` command
 - Mount tables (for example, `/etc/vfstab` on Solaris)
 - Automounters

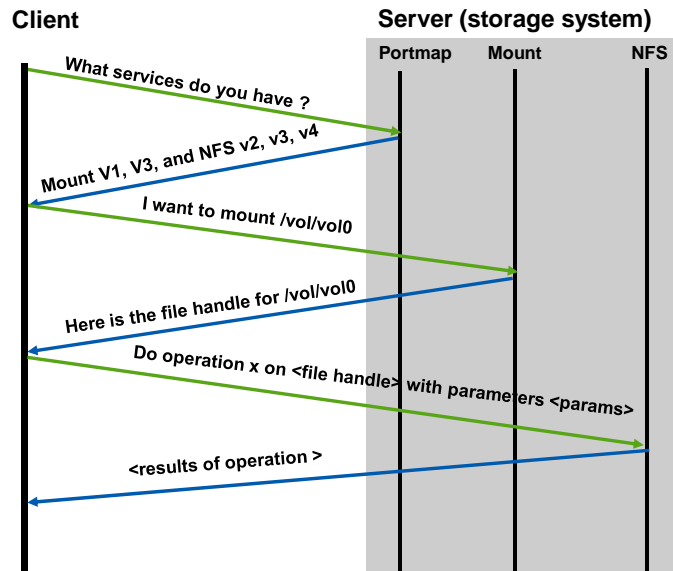
© 2010 NetApp, Inc. All rights reserved.

MOUNTS



The Mount Process

■ Abbreviated Mount Process



© 2010 NetApp, Inc. All rights reserved.

THE MOUNT PROCESS



Mount Configuration

- At the client
 - Create a directory (mountpoint)
`mkdir /nfsmount`
- To mount the storage system export, use the `mount` command from the command line as follows:
`mount <server>:/vol/vol0/home /nfsmount`
- Use `umount` command to unmount the export

© 2010 NetApp, Inc. All rights reserved.

MOUNT CONFIGURATION



Target Mount Tables

- To make the mounts persistently reload during boot up, edit the */etc/vfstab* (or equivalent) file to include the following entries:

```
Server:/vol/vol0/home - /nfsmount nfs rw 0 0;
```

- Storage system is mounted when client reboots
- Entry above automatically mounts the directory; user can only make changes to content in */vol/vol0/home* after mounting the directory

© 2010 NetApp, Inc. All rights reserved.

TARGET MOUNT TABLES

For persistent mounts across reboots, edit the *vfstab* file (for systems, such as Solaris™ host systems), or *fstab* file (for BSD systems, such as Linux host systems). To mount directories, use the mount command from the command line. Be aware that the mount instruction from the command line is not persistent across storage system reboots.



Automounter

- Automounter is an remote procedure call daemon on a target that automatically mounts NFS exports
 - Mount times out and disconnects when not in use
 - Useful with a large amount of resources to mount or unmount on startup or shutdown, for example, with home directories
- Uses an automounter map to store information
- Each automounter map contains:
 - Information about the server hosting the file system
 - The pathname of the file system in the server
 - The local pathname
 - Mount options that apply to the file system

© 2010 NetApp, Inc. All rights reserved.

AUTOMOUNTER

In the UNIX world, automounter is a program that mounts a resource when requested by a user, and when the user is done with it, the mount will time out and automatically disconnect itself. This can be useful when there is a large amount of resources to mount and unmount on startup and shutdown, like home directories.

There are two different types of automounter maps: direct and indirect. Direct maps are sets of unrelated mountpoints that can be spread out across the file system. An indirect map sets aside a directory and mounts everything in the map within that directory; users' home directories are a great example of this. Automounter maps may be distributed by using a directory service, such as NIS or NIS+, for centralized management.



Verifying Mounts

- To verify exports on a target:
 - `mount` (without option displays mounted files)
 - `showmount -a storagesystemx`
 - Displays list of clients mounting from the storage system
 - `showmount -e storagesystemx`
 - Prints list of shared file systems
- To verify mounts on a storage system:
 - `/etc/rmtab`
 - Refreshes when the storage system boots

© 2010 NetApp, Inc. All rights reserved.

VERIFYING MOUNTS

To verify the exported resources, use the `mount` command in UNIX systems. Use `showmount -e` or `nfsstat -m` or an equivalent command on the client to view output that verifies exported resources and also shows the mount options. With the `showmount` command, you can display what is being exported by the storage system and the clients mounting the storage system. `/etc/rmtab` stores the mount requests that the storage system has received since boot up. NOTE: `/etc/rmtab` may not be reliable to see which clients are actually mounted because a client may crash without unmounting the directory and then the mount entry will remain in the `/etc/rmtab` file.



Access Cache

© 2010 NetApp, Inc. All rights reserved.

ACCESS CACHE



Access Cache

- When an NFS target attempts to access an export, mountd must determine whether to grant or deny access
 - First checks the access cache
 - If no entry, determines a client's read, write and root access to an export path; possible values:
 - Positive = granted access
 - Negative = denied access
 - Delayed = Could not resolve
 - In Progress = Attempting to resolve
 - Uninitialized = Error condition
- Access cache is used after the mount when determining export permissions of a client

© 2010 NetApp, Inc. All rights reserved.

ACCESS CACHE

Whenever an NFS client attempts to access a file system path, Data ONTAP must determine whether to grant or deny access. The cache is a reservoir of access results. Each access result is cached for particular client address and export path for their read, write, and root access. Possible values for each of these three access types are: positive (that is, granted access), negative (that is, denied access), delayed (that is, could not be resolved--usually because netgroup or hostname could not be resolved), in progress, or uninitialized.

The access cache is persisted to disk every 15 minutes. The files are saved as: `/etc/exports_arc/cache` and `/etc/exports_arc/restore`. This helps improve the storage system's response time during the period after a restart. To turn off the access cache persistence, set `nfs.acache.persistence.enable` to off.



Cache Administration

- To check NFS target-to-access cache
 - `exportfs -c clientaddr path [options] [securitytype]`
- To remove entries from access cache
 - `exportfs -f [path]`
- To specify the cache expiration
 - `options nfs.export.neg.timeout`
 - Refresh time for denied access entries
 - `options nfs.export.pos.timeout`
 - Refresh time for granted access entries
 - `options nfs.export.harvest.timeout`
 - Sets idle expiration time for entries in cache

© 2010 NetApp, Inc. All rights reserved.

CACHE ADMINISTRATION

`Exportfs -c` parameters:

clientaddr

- The IP address of the NFS client

path

- The file system path

accesstype

- `ro` | `rw` | `root`

security type

- `sys` | `none` | `krb5` | `krb5i` | `krb5p`

Negative (that is, denied access) cache entries are purged from the access cache after the period of seconds set by the `nfs.export.neg.timeout` option. The negative timeout defaults to 1800 seconds or 30 minutes. The minimum value for this option is 60 seconds. The maximum value for this option is 7 days.

Positive (that is, granted access) cache entries are purged from the access cache after the period of seconds set by the `nfs.export.pos.timeout` option. The positive timeout defaults to 36000 seconds or 10 hours. The minimum value for this option is 60 seconds. The maximum value for this option is 7 days.

Delayed cache entries are purged from the access cache after 15 seconds. This timeframe cannot be changed.

A cache entry that has not been accessed will be purged from the access cache after the period of seconds set by the `nfs.export.harvest.timeout` option.



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Identify exportable resources
- Export and unexport resources to clients, subnets, and netgroups
- Administer storage system with the `exportfs` command
- Create mountpoints and mount exported resources on a client
- Monitor the usage of exported resources

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 4: Exports and Mounts
Estimated Time: 60 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- What is a target?
- What is a resource?
- Give examples of a target.
- Give examples of a resource.
- Can you specify `rw` and `ro` for the same host within a single export?
- What must be defined on the target to receive exported resources?
- The `/etc/exports` file typically resides on the target. True or false?
- What does `exportfs -c 10.0.0.1 /vol/vol2 rw` do?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

CIFS Overview

Module 5
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



CIFS OVERVIEW



Module Objectives

By the end of this module, you should be able to:

- Describe basic CIFS features
- Describe the following network environments:
 - Microsoft® Windows® workgroup
 - Non-Windows workgroup
 - Windows domains
- Describe how a storage system authenticates users in each server environment
- Explain the advantages and disadvantages of each server environment

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



CIFS Features

© 2010 NetApp, Inc. All rights reserved.

CIFS FEATURES



CIFS Definition

- Common Internet File System (CIFS)
 - A Microsoft network file-sharing protocol that evolved from the Server Message Block (SMB) protocol
 - Access and manipulate files and folders on remote servers as if they are on a local machine

© 2010 NetApp, Inc. All rights reserved.

CIFS DEFINITION

The Common Internet File System (CIFS) is a Microsoft network file-sharing protocol that evolved from the Server Message Block (SMB) protocol.

When using CIFS, any application that processes network I/O can access and manipulate files and folders (directories) on remote servers in a similar way that it accesses and manipulates files and folders on the local system.



CIFS: Basic Functions

- Network browsing to locate:
 - Machines within an environment (provided by a browse list)
 - Shared resources that are available on a given machine (provided by that machine)
- User authentication
- Authorization
 - Shared resource access
 - Folder and file access

© 2010 NetApp, Inc. All rights reserved.

CIFS: BASIC FUNCTIONS

The following are some CIFS features available in a Windows workgroup and domain:

- Network browsing to locate machines within a domain or workgroup (provided by a browse list) and shares that are available on each machine (provided by that machine).
- User authentication.
- Authorization at the share and folder or file level.



CIFS: Basic Functions (Cont.)

- Basic file attributes
 - Read-only
 - Archive
 - System
 - Hidden
- Extended NTFS file attributes of indexing, compression, and encryption
- Unicode support
- File locking (opportunistic locks)
- Dialect negotiation

© 2010 NetApp, Inc. All rights reserved.

CIFS: BASIC FUNCTIONS (CONT.)

EXTENDED ATTRIBUTES

Extended NTFS file attributes generally are not supported on a storage system. However, Encrypted File Systems (EFS) is supported with Open Systems SnapVault®.

UNICODE SUPPORT

The universal character encoding standard provides a unique number for every character, no matter what the platform, program, or language.

Characters are represented by more than eight bits.

OPPORTUNISTIC LOCKS

Guarantee to the client that file content is not allowed to be changed by the server or, if some change is imminent, the client is notified before the change proceeds.

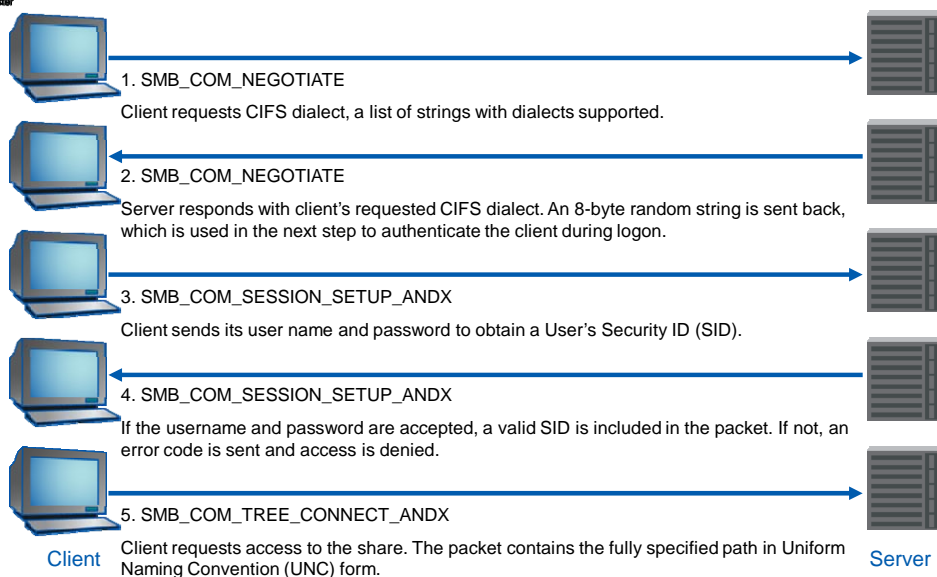
Oplocks are used to synchronize data and enhance performance.

DIALECT NEGOTIATION

Each protocol version is referred to as a “dialect” and assigned a unique string identifier.



Client-Server Communications



© 2010 NetApp, Inc. All rights reserved.

CLIENT-SERVER COMMUNICATIONS

This example demonstrates client-server communications for session, share access, and file authorization. The following are the basic steps:

The client contacts the server and requests CIFS dialect.

The server responds with the supported CIFS dialect and next logon step. Together, these two steps are called dialect negotiation.

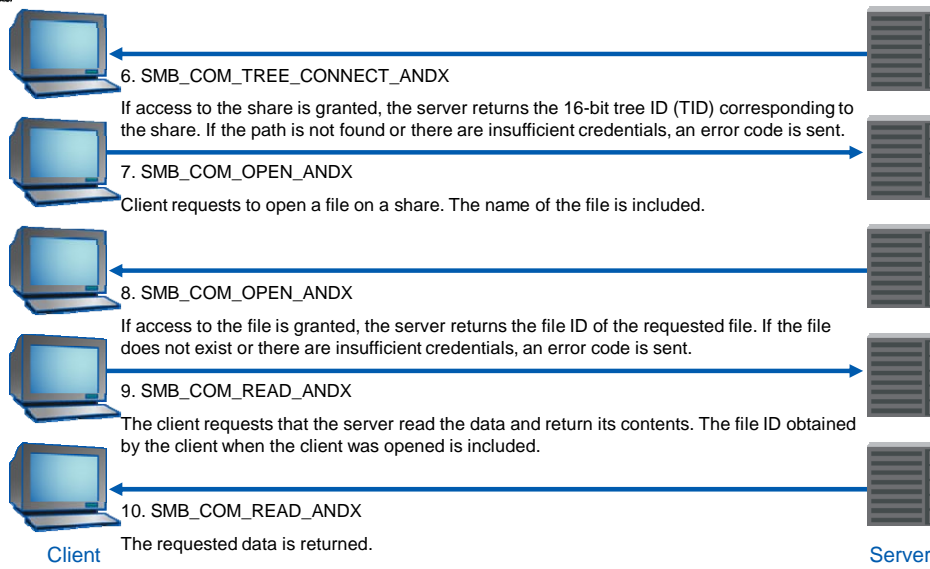
The client responds with the username and password.

The server sends a User ID (UID) if the username and password are accepted or an error message if not accepted. Together, these two steps are called user authentication.

The client requests access to a share. The storage system caches all security IDs (SIDs) and usernames received from the domain controller at boot time.



Client-Server Communications (Cont.)



© 2010 NetApp, Inc. All rights reserved.

CLIENT-SERVER COMMUNICATIONS (CONT.)

The server responds with a tree ID to the requested share (if access is allowed). Together, Steps 5 and 6 are called shared resource authorization.

The client requests to open a file on a share.

If access is allowed, the server responds with the ID of the requested file. Together, these two steps are called folder/file authorization.

The client requests that the server read the data and return its contents.

The server sends the requested data. During this process, the access control lists (ACLs) are checked for permissions. Together, these two steps are called folder/file I/O.



CIFS Environments

© 2010 NetApp, Inc. All rights reserved.

CIFS ENVIRONMENT



Network Environments

Storage systems can participate in:

- Workgroups
 - Windows workgroup
 - Non-Windows workgroup
- Domains
 - Windows NT® 4.0
 - Windows Active Directory

© 2010 NetApp, Inc. All rights reserved.

NETWORK ENVIRONMENTS



Client Requirements

Each client in a CIFS environment must:

- Locate other computers
- Request resources from a server
 - Requires user authentication
 - Requires resource authorization
 - Share permissions
 - File-level permissions

NOTE: Implementation differs depending on the CIFS environment

© 2010 NetApp, Inc. All rights reserved.

CLIENT REQUIREMENTS

In a network, a Windows client user requires the ability to:

- Find other machines (computers)
- Request resources from a server meaning any machine in the role of a server

Requesting resources requires user authentication (verification of a user's identity) to establish a session with a server and user authorization (permission) to access a share and resources (folders and files) in a share.



Windows Workgroups

© 2010 NetApp, Inc. All rights reserved.

WINDOWS WORKGROUPS



Windows Workgroup

- A Windows workgroup:
 - Logical grouping of networked machines
 - Shares resources, such as folders and files
- Each machine in the workgroup authenticates and authorizes users through a local security database

NOTE: Users must have an account on the machine they wish to access

© 2010 NetApp, Inc. All rights reserved.

WINDOWS WORKGROUP

A Windows workgroup is a simple, logical group of networked machines (computers) that share resources, such as folders and files.


- Each machine has its own Security Accounts Manager database (for Windows NT) or a local security database (for Windows 2000 Server or later) that is used to perform user authentication and user authorization.
- Each user that wants to access resources on a machine must have a user account on that machine.



Storage System Joins a Workgroup

For a storage system to “join” a Windows workgroup...

- It must broadcast its “name” to the network
- The master browser must update the master browse list
- It must broadcasts the browse list to all members of the domain
 - 15-minute delay possible
 - **NOTE:** Storage systems do not pull the master browse list.



© 2010 NetApp, Inc. All rights reserved.


STORAGE SYSTEM JOINS A WORKGROUP

When workgroup machines normally pull the updated master browse list to their local machine, the storage system does not. The browse list is a mechanism for members of the workgroup to find other members. The storage system always acts in a server role. Therefore, there is no need to discover other members in the workgroup.



Name Resolution in a Workgroup

- Machine name to IP resolution through NetBIOS resolution:
 - A user broadcasts a name query on the network
 - The requested machine responds to the name query by returning its IP address
- Machine name to IP resolution through DNS resolution is also available (discussed later)



© 2010 NetApp, Inc. All rights reserved.

NAME RESOLUTION IN A WORKGROUP

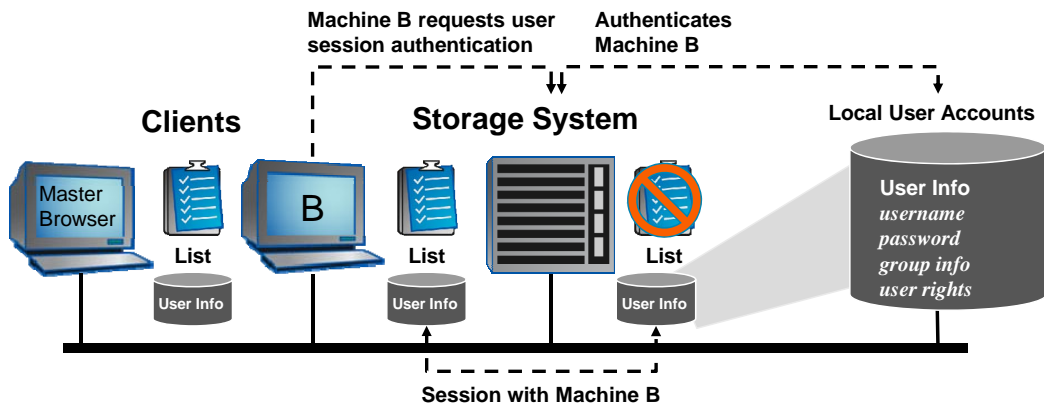
A machine broadcasts a name query to other machines in the network. For example, Machine A may broadcast a query for the IP address of the storage system. The storage system responds to the name query by broadcasting the IP address back to Machine A.



Storage System User Authentication

Storage system user authentication is performed locally:

- Users are added to a storage system
- Authentication is performed locally
- Authenticated users are provided with a session



© 2010 NetApp, Inc. All rights reserved.

STORAGE SYSTEM USER AUTHENTICATION

Users (local-user accounts) are added to a storage system and user authentication is performed locally on the storage system. User session authentication with a user name and password enables a user to establish a session with the storage system. Data access on a storage system requires a network logon to the storage system. A user can administer a storage system through the network (for example, a Telnet session) using a local account on the storage system; however, a user cannot log on locally to a storage system to access data. Machine-B's user requests user session authentication with the storage system. The storage system authenticates Machine-B's user by using the user name and password found in the storage system local-user account. After Machine-B's user successfully authenticates, a session is established with Machine-B's user and the storage system. Authenticated users can browse a storage system for available resources, but must be authorized to access a share and resources in a share.



Windows Workgroup Advantages

- Does not require running Windows Domain Controller
 - Advantageous for small organizations
- Simple to design and implement
- Convenient for a limited number of machines in close proximity
 - Limited to 96 local clients

© 2010 NetApp, Inc. All rights reserved.

WINDOWS WORKGROUP ADVANTAGES



Windows Workgroup Disadvantages

- Administrative overhead in maintaining a list of user accounts on multiple machines
 - Any changes to a user account (for example, passwords) must be made on each machine
- Joining or leaving a workgroup must be replicated by the master browse list
 - Delay up to 15 minutes
- Generally, a browse list cannot span subnets
 - Workgroup depends on subnet broadcasting

© 2010 NetApp, Inc. All rights reserved.

WINDOWS WORKGROUP DISADVANTAGES



Non-Windows Workgroups

© 2010 NetApp, Inc. All rights reserved.

NON-WINDOWS WORKGROUP



Non-Windows Workgroups

- A non-Windows workgroup:
 - Support for Windows client machines when there is no Windows workgroup or domain
 - Share resources with Windows client users
- This environment is also referred to as:
 - UNIX® password workgroup
 - /etc/passwd-style workgroup

© 2010 NetApp, Inc. All rights reserved.

NON-WINDOWS WORKGROUPS

A non-Windows workgroup is a logical group of networked machines that shares resources with Windows client users; the networked machines are members of neither a Windows workgroup nor a Windows domain.

This network environment also is called:

- UNIX password workgroup
- /etc/passwd-style workgroup



Non-Windows Workgroup Storage System

- Provides user authentication by one or more of the following:
 - Storage system local /etc/passwd file
 - Network Information Services (NIS) server
 - Lightweight Directory Access Protocol (LDAP) server
- Provides name-to-IP resolution by one or more of the following:
 - Storage system local /etc/hosts file
 - NIS server
 - Domain Name System (DNS) server

NOTE: /etc/nsswitch.conf sets the order of precedence for the mechanism used

© 2010 NetApp, Inc. All rights reserved.

NON-WINDOWS WORKGROUP STORAGE SYSTEM

When a storage system becomes a non-Windows workgroup server, it provides services to clients. An example is an all-UNIX work environment with many UNIX workstations and a few Windows clients with users that need CIFS resources. Note that any UNIX reference also includes Linux servers functioning in the role of a directory store for user information (user names, passwords, and group information):

- Storage system's local /etc/passwd file
- NIS server
- LDAP server

Servers that can provide machine (host) name resolution:

- Storage system's local /etc/hosts file
- NIS server
- Domain Name System (DNS) server



Non-Windows Workgroup Advantages

- In a mostly UNIX environment, CIFS shares are made available to the few Windows client users
- User authentication performed by existing:
 - NIS
 - LDAP server
 - /etc/passwd file
- Name-to-IP resolution performed by existing:
 - NIS
 - DNS server
 - /etc/hosts

© 2010 NetApp, Inc. All rights reserved.

NON-WINDOWS WORKGROUP ADVANTAGES



Non-Windows Workgroup Disadvantages

- Administrative overhead in maintaining a list of user accounts on multiple machines
 - Any changes to a user account (for example, passwords) must be made on each machine
 - Sends passwords in clear text
- Requires both NFS and CIFS licenses
- Generally, a browse list cannot span subnets
 - Workgroup depends on subnet broadcasting

© 2010 NetApp, Inc. All rights reserved.

NON-WINDOWS WORKGROUP DISADVANTAGES



Windows Domains

© 2010 NetApp, Inc. All rights reserved.

WINDOWS DOMAINS



Window Domains

- A Windows domain:
 - A logical grouping of networked machines
 - Share a central directory of resources
- A domain controller centralizes:
 - User/Group/Machine account management
 - User authentication
 - Group policy management across the domain

NOTE: In this course, we will consider NT-style security and Active Directory domains together

© 2010 NetApp, Inc. All rights reserved.

WINDOW DOMAINS



Typical Machines in a Domain

Type of machines in a domain:

- Clients
 - Clients requires resources from a server
- Member servers
 - Servers that provide resources to clients
- Domain controllers (DCs)
 - Servers that each maintain a copy of a centralized database
- Domain name resolution servers
 - Windows Internet Name Service (WINS) for Windows NT-style domains
 - Domain Name System (DNS) for Windows 2000 Server (or later) domains


© 2010 NetApp, Inc. All rights reserved.

TYPICAL MACHINES IN A DOMAIN



Storage System Joins a Domain

- When a storage system joins a domain:
 - Domain controller adds the storage system to a domain database
 - Becomes a member server



© 2010 NetApp, Inc. All rights reserved.

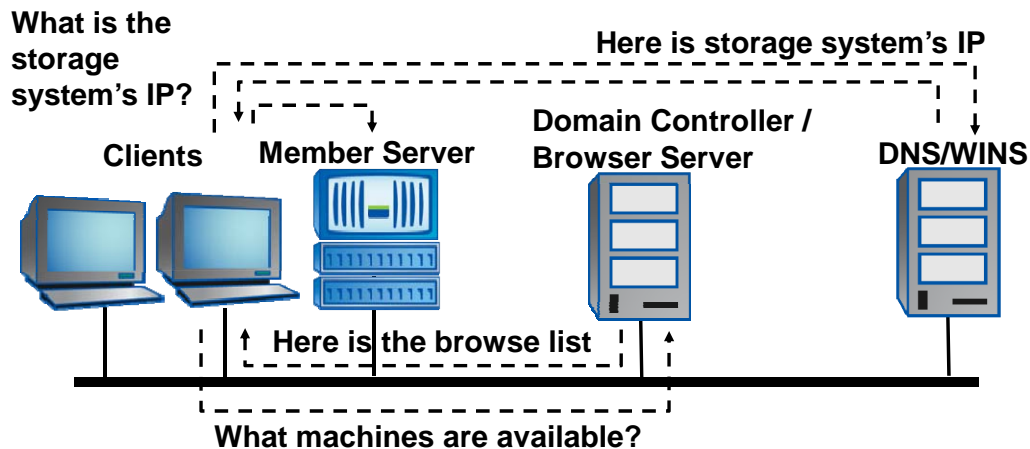
STORAGE SYSTEM JOINS A DOMAIN

When a storage system joins a domain, it becomes a member server that provides services to clients. The storage system (member server) goes to a domain controller and the domain controller adds the machine account to the directory database.



Domain Name to IP Resolution

- When a client accesses a storage system's resource:
 - Requests the browse list from the DC
 - Contacts DNS/WINS server for the IP address
 - Client communicates with storage system



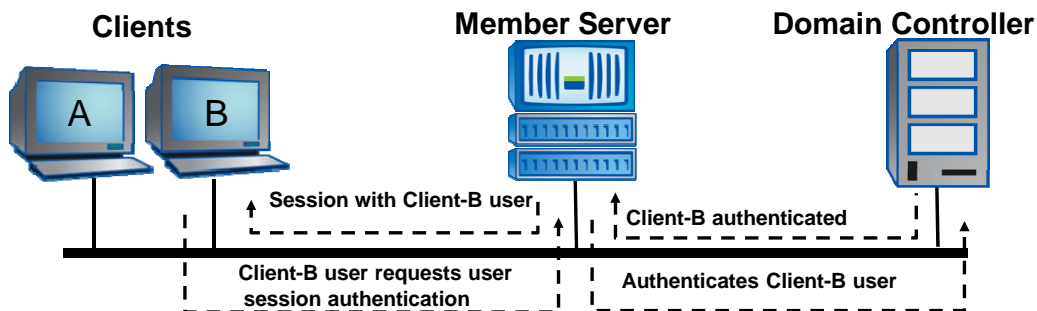
© 2010 NetApp, Inc. All rights reserved.

DOMAIN NAME TO IP RESOLUTION



User Authentication

- User authentication on a storage system in a domain:
 - Domain users created on domain controller (DC)
 - User session authentication occurs at the DC
 - Authenticated users must be authorized to access a share and resources



© 2010 NetApp, Inc. All rights reserved.

USER AUTHENTICATION

Domain users (already added to the domain controller) can browse the storage system for available shares and then request access to the storage system and its shares and resources in a share. User session authentication with a user name and password is performed centrally on the domain controller; this establishes a user session with the storage system. Users must be authorized to access a share and resources in a share. Data access on a storage system requires a network logon to the storage system. A user can administer a storage system through the network (for example, a Telnet session) using a local account on the storage system; however, a user cannot log on locally to a storage system to access data. Client-B's user requests user session authentication with the member server (storage system). The member server goes to the domain controller to authenticate Client-B's user. The domain controller authenticates Client-B's user and a session is established with Client-B's user and the member server (storage system).



Domain Advantages and Disadvantages

- Advantages:
 - Centralized administration of all user information
 - A centralized mechanism for authentication
 - Scalable
- Disadvantages:
 - Complexity of architecting Active Directory
 - Requires server license

© 2010 NetApp, Inc. All rights reserved.

DOMAIN ADVANTAGES AND DISADVANTAGES



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Describe basic CIFS features
- Describe the following network environments:
 - Microsoft Windows workgroup
 - Non-Windows workgroup
 - Windows domains
- Describe how a storage system authenticates users in each server environment
- Explain the advantages and disadvantages of each server environment

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 5: CIFS Overview
Estimated Time: 15 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- In a network, which two abilities does a Windows client user require?
- What is the difference between authentication and authorization?
- What are the three types of storage system CIFS service environments?
- What is the purpose of a name resolution server?
- What kind of information is kept in the directory that the domain controller stores and maintains?
- In a Windows domain, how does a storage system authenticate users?
- In a non-Windows workgroup, how does a storage system authenticate users?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

CIFS Workgroups

Module 6
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



CIFS WORKGROUP



Module Objectives

By the end of this module, you should be able to:

- License CIFS on a storage system
- Join a storage system to a Windows® workgroup environment using the `cifs setup` command
- Observe the results of `cifs setup`
- Manage newly created configuration files for the CIFS workgroup environment

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Setup Overview

© 2010 NetApp, Inc. All rights reserved.

SETUP OVERVIEW



CLI: CIFS Setup

- To prepare a storage system to support Windows client users, complete the following steps:
 1. License CIFS
 2. Perform the initial CIFS configuration by running the `cifs setup` program or using NetApp® System Manager
- If the setup is successful, the CIFS server starts automatically

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS SETUP



Configuring CIFS Using `cifs setup`

© 2010 NetApp, Inc. All rights reserved.

CONFIGURING CIFS USING CIFS SETUP



CLI cifs setup: WINS

■ During cifs setup

```
system> cifs setup
```

This process will enable CIFS access to the filer from a Windows system.

Note: Use "?" for help at any prompt and Ctrl C to exit without committing changes.

Your filer does not have WINS configured and is visible only to clients on the same subnet.

Do you want to make the system visible via WINS?

[n]:

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: WINS

Windows Internet Name Service (WINS) is Microsoft's implementation of NetBIOS Name Server on Windows. As of Windows 2000 Server, DNS is preferred over WINS, particularly for Active Directory. WINS servers are required for supporting pre-Windows 2000 Server and mixed Windows 2000 Server installs.



CLI cifs setup: Initial Questions

■ During cifs setup (Cont.)

A filer can be configured for multiprotocol access, or as an NTFS-only filer. Since NFS, DAFS, VLD, FCP, and iSCSI are not licensed on this filer, we recommend that you configure this filer as an NTFS-only filer

- (1) NTFS-only filer
- (2) Multiprotocol filer

Selection (1-2)? [1]:

This list varies depending on other licensed protocols

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: INITIAL QUESTIONS

If the storage system will be in a Windows-only environment, selecting the NTFS-only option configures the storage system to be most compliant with Microsoft® environments. **NOTE:** All existing volumes will be converted to NTFS, but qtrees are unaffected.

If the storage will participate in both Windows and non-Windows environments, the storage system should be configured as multiprotocol.

When storage system has both NFS and CIFS are licensed, option 1 will be Multiprotocol. If only CIFS is licensed, option 1 will be NTFS-only as shown.



Results of NTFS-Only

- NTFS-only security style changes as a result of `cifs setup`

Option	Defaults Before	Values After
<code>wafl.default_security_style</code>	unix	ntfs
<code>wafl.nt_admin_priv_map_to_root</code>	on	off

- Verify by `options wafl` command

NOTE: NTFS-only will change all existing volumes' security style

© 2010 NetApp, Inc. All rights reserved.

RESULTS OF NTFS-ONLY

After running the `cifs setup` command, the `options wafl` command is run. The option `wafl.default_security_style` is changed from UNIX® to NTFS. This causes all new volumes to default to NTFS security style.

Additionally, the `nt_admin_priv_map_to_root` option changes from **on** to **off**.



Switching Back to Multiprotocol

- To switch back to multiprotocol:
 - Use `cifs setup` and select multiprotocol
 - Switches `wapl.default_security_style` to `unix`
 - Switches `wapl.nt_admin_priv_map_to_root` to `on`
- Results of switching NTFS-only to multiprotocol:
 - Existing ACLs (if any) are unchanged
 - Security style of volumes and qtrees remains unchanged
 - New volumes have security style of UNIX

© 2010 NetApp, Inc. All rights reserved.

SWITCHING BACK TO MULTIPROTOCOL

Although you can change a storage system from NTFS-only to multiprotocol using the `cifs setup` command, you can achieve the same effect more easily by simply setting the `wapl.default_security_style` option to `unix`.

The results of changing an NTFS-only storage system to a multiprotocol storage system are as follows:

- Existing ACLs remain unchanged
- The security style of all volumes and qtrees remains unchanged
- When you create a volume, its default security style is UNIX
- The `wapl.default_security_style` option is set to UNIX



Switching Back to Multiprotocol (Cont.)

Root volume security style will remain `ntfs`

- UNIX root user might be denied access
- You can gain access:
 - Map of Windows user to UNIX root
(*Discussed in Module 10*)
 - `cifs.nfs_root_ignore_acl on`

© 2010 NetApp, Inc. All rights reserved.

SWITCHING BACK TO MULTIPROTOCOL (CONT.)

Because the security style of the root volume remains NTFS after you change the storage system from NTFS-only to multiprotocol, you might be denied access to the root volume when you connect from UNIX as root.

- You can gain access if the ACL for the root volume allows full control for the Windows user that maps to root
- You also can gain access by setting the `cifs.nfs_root_ignore_acl` option to on
- When this option is on, ACLs will not affect root access from the Network File System (NFS)



CLI `cifs` setup: Root User

■ During CIFS setup (Cont.):

CIFS requires local `/etc/passwd` and `/etc/group` files and default files will be created. The default `passwd` file contains entries for 'root', 'pcuser', and 'nobody'.

NOTE: These files are used during CIFS authentication processing when mapping Windows users to UNIX users even if it is NTFS-only security style.

Enter the password for the root user []: ← The password is entered, but it is not displayed
Retype the password: ←

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: ROOT USER

With respect to CIFS, the root user is used in a non-Windows (UNIX) workgroup only and when authentication is performed with the `/etc/passwd` file.

This is the root user that was created in the `/etc/passwd` file. With respect to CIFS, this root user is used in a non-Windows workgroup only and when authentication is performed with the `/etc/passwd` file.



CLI `cifs` setup: Server Name

■ During CIFS setup (Cont.):

The default name for this CIFS server is ' system '.
would you like to change this name? [n]:

System's name should be the
same as registered in the
infrastructures DNS server

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: SERVER NAME



CIFS Authentication Methods

■ During CIFS setup (Cont.):

Data ONTAP CIFS services support four styles of user authentication.

Choose the one from the list below that best suits your situation.

1. Active Directory domain authentication (Active Directory domains only)
2. Windows NT 4 domain authentication (Windows NT or Active Directory domains)
3. Windows Workgroup authentication using the filer's local user accounts
4. etc/passwd and/or NIS/LDAP authentication

Selection (1-4)? [1]:

© 2010 NetApp, Inc. All rights reserved.

CIFS AUTHENTICATION METHODS

If you plan to have the storage system join a Windows domain and make use of that domain's users and groups, you should choose Option 1 or 2. Options 3 and 4 are authentication methods that do not require the use of domain controllers, but may still require other systems for full functionality.

Option 1: Use this option if the storage system is joining an Active Directory Native Mode.

Option 2: Use this option if the storage system is joining a Windows NT® 4-based domain or an Active Directory Mixed Mode.

Option 3: Use this option if you want to join a Windows Workgroup and do not want to depend on external domain controllers. You will need to define a set of local users on the storage system.

Option 4: Use this option for a non-Windows Workgroup that uses UNIX style authentication. This style requires the use of cleartext passwords from Windows clients.



CLI: cifs setup Workgroup

■ Selecting Windows Workgroup:

1. Active Directory domain authentication
(Active Directory domains only)
2. Windows NT 4 domain authentication
(Windows NT or Active Directory domains)
3. Windows Workgroup authentication using
the filer's local user accounts
4. /etc/passwd and/or NIS/LDAP
authentication

Selection (1-4)? [1]: 3

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS SETUP WORKGROUP



CLI: cifs setup Workgroup (Cont.)

```
What is the name of the Workgroup? [WORKGROUP]: workgroup1  
Fri Jun 23 19:32:53 GMT [wafl.quota.sec.change:notice]:  
security style for /vol/vol0/ changed from unix to ntfs  
CIFS - Starting SMB protocol...
```

```
It is recommended that you create the local administrator  
account(system\administrator)for this filer.  
Do you want to create the system\administrator account? [y]: y  
Enter the new password for system\administrator:  
Retype the password:
```

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS SETUP WORKGROUP (CONT.)

The local administrator account can be locally authenticated by way of CIFS authentication and has privileges to administer CIFS on the storage system. The local users and passwords are stored in the storage system registry file.



CLI: cifs setup Workgroup (Cont.)

Workgroup completion continued

Welcome to the WORKGROUP1 Windows(R) workgroup

CIFS local server is running.

system> Fri Jun 23 19:33:18 GMT

[nbt.nbns.registrationComplete:info]: NBT: All CIFS name registrations have completed for the local server.

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS SETUP WORKGROUP (CONT.)



Configuring CIFS Using NetApp System Manager

© 2010 NetApp, Inc. All rights reserved.

CONFIGURING CIFS USING NETAPP SYSTEM MANAGER



System Manager: CIFS Setup

NetApp System Manager

devst020-f1.netapp.com

Storage

- Volumes
- Shared Folders
- Shares/Exports
- Sessions
- Oncores
- Disks
- Aggregates
- Configuration
- Local Users and Groups
 - Users
 - Groups
- Network
 - DNS
 - Interfaces
 - Network Files
 - NIS
- Protocols
 - CIFS
- Security
 - Password(RSH)/Trusted hosts
 - SSH & SSL
- System Tools
 - Autosupport
 - Date/Time/Timezone
 - Licenses
 - SNMP
 - Syslog
- Diagnostics
 - CIFS

Name	Type	Key	Expires on
Windows Shares(CIFS)	Evaluation	JMW/MR0A	7/17/2009 10:54:49 AM
Linux Exports(NFS)	Evaluation	TPQIZBN	7/17/2009 10:10:43 AM

The newly added CIFS license

Notice the new categories

License description

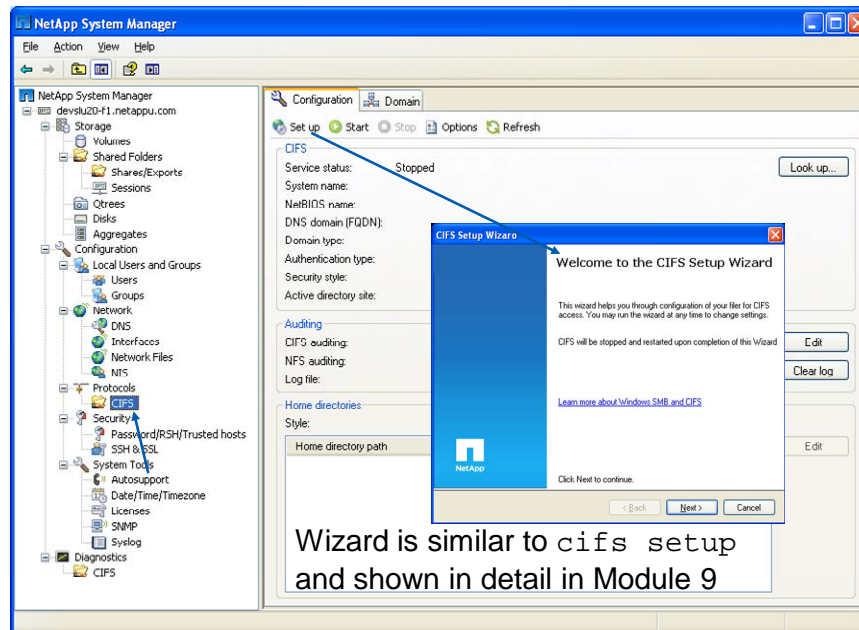
The storage system requires a software license to enable CIFS service. The license is installed on the storage system at the factory per your order, therefore, the initial setup of your storage system does not involve entering license codes.

© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP



System Manager: CIFS Setup (Cont.)



© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP (CONT.)



Results

© 2010 NetApp, Inc. All rights reserved.

RESULTS



CIFS Server Files Created During Setup

- `/etc/cifsconfig_setup.cfg`
 - Stores CIFS setup configuration
- `/etc/usermap.cfg`
 - Multiprotocol support for NFS and CIFS (*Discussed in Module 10*)
- `/etc/passwd`
 - Multiprotocol and UNIX workgroup
- `/etc/cifsconfig_share.cfg`
 - Default share definitions
- `/etc/lclgroups.cfg`
 - Local groups definitions

NetApp
recommends
not to edit
these files
directly

NOTE: Additional files are created depending on the environment

© 2010 NetApp, Inc. All rights reserved.

CIFS SERVER FILES CREATED DURING SETUP

During the CLI cifs setup script or NetApp System Manager's CIFS Setup Wizard, CIFS support and configuration files are created in the `/etc` directory. The number and content of the files are dependent on the environment.

The following are files that are common to all environments:

- `/etc/cifsconfig_setup.cfg` (stores the CIFS setup configuration)
- `/etc/usermap.cfg` (multiprotocol support for mapping users of NFS and CIFS)
- `/etc/passwd` (multiprotocol and UNIX workgroup)
- `/etc/cifsconfig_shares.cfg` (default shares definitions)

Additional files are created depending on the environment as in a workgroup (Windows or non-Windows) or a Windows domain.



/etc/cifsconfig_setup.cfg File

- /etc/cifsconfig_setup.cfg file
 - Contents are persistent across reboots
 - Runs each time the CIFS service is started

```
system> rdfile /etc/cifsconfig_setup.cfg
#Generated automatically by cifs commands
cifs setup -security unix -cp 0 -NTFSonly
```

NOTE: The content of the file varies depending on the environment that is selected

© 2010 NetApp, Inc. All rights reserved.

/ETC/CIFSCONFIG_SETUP.CFG FILE

The following shows the contents of an /etc/cifsconfig_setup.cfg file:

```
system> rdfile /etc/cifsconfig_setup.cfg
#Generated automatically by cifs commands
cifs setup -security unix -cp 0 -NTFSonly
```

The content of the file varies depending on the environment that is selected. This file is used each time CIFS service is started and persists across reboots.



/etc/passwd File

■ /etc/passwd file

```
system> rdfile /etc/passwd
```

An encrypted root user password is shown

```
root:_J9../ongnoStt3Ei79o:0:1::/  
pcuser::65534:65534::/  
nobody::65535:65535::/  
ftp::65533:65533:FTP Anonymous:/home/ftp:
```

- Is checked during CIFS authentication processing when mapping Windows users to UNIX UID and GID
- Can be used for authentication in a non-Windows (UNIX) workgroup environment

NOTE: Unless the Windows user is mapped to a specific UNIX user name, `pcuser` is the default (See *Module 10*)

© 2010 NetApp, Inc. All rights reserved.

/ETC/PASSWD FILE

NOTE: This root user was created during CIFS setup for the /etc/passwd file. This is not for the storage system user “root” that is used for system administration.



CIFS Default Shares

- Setup creates three default shares:
 - C\$ maps to /vol/<root volume>
 - ETC\$ maps to /vol /<root volume>/etc
 - HOME is /vol /<root volume>/home
 - Home directory is accessible to everyone
- \$ shares are hidden
- C\$ and ETC\$ are available only to administrators

© 2010 NetApp, Inc. All rights reserved.

CIFS DEFAULT SHARES

These are the three default share definitions:

- **C\$** is /vol/<root volume>. This is a hidden “admin share” to root of the root volume.
- **ETC\$** is /vol /<root volume>/etc. This is a hidden “admin share” to the /etc directory on the root volume. The /etc directory stores storage system configuration files, executables that are required to boot the system, and some log files.
- **HOME** is /vol /<root volume>/home. This share is the path to the /home directory on the root volume that is accessible to everyone.

A hidden share means that it is not visible when browsing. An “admin share” is only available to users who are members of an administrator group.

The storage system default root volume is /vol/vol0. You can change which volume on your storage system is used as the root volume or create a new one and in the process designate a different name for the root volume. The root volume contains special directories and configuration files for administering the storage system.



/etc/cifsconfig_share.cfg File

- /etc/cifsconfig_share.cfg file

```
system> rdfile /etc/cifsconfig_share.cfg
#Generated automatically by cifs commands
cifs shares -add "ETC$" "/" "Remote Administration"
cifs access "ETC$" S-1-5-32-544 Full Control
cifs shares -add "HOME" "/vol/vol0/home" "Default Share"
cifs access "HOME" S-NONE "nosd" ← nosd = No Security Descriptor
```

- The HOME share maps to the user who logs in
- The security descriptors on the user's home directory applies

```
cifs shares -add "C$" "/" "Remote Administration"
cifs access "C$" S-1-5-32-544 Full Control
```

- This file can be altered by way of command-line interface commands or GUIs

© 2010 NetApp, Inc. All rights reserved.

/ETC/CIFSCONFIG_SHARE.CFG FILE

The HOME share is unique because it maps to the user who is trying to log in. The security descriptors on the user's home directory apply.



/etc/lclgroups.cfg File

- The local administrator is added to lclgroups.cfg:

```
system> rdfile /etc/lclgroups.cfg
[ "Replicators" 552 ( "not supported" ) ]
[ "Backup Operators" 551 ( "Members can bypass file
security to backup files" ) ]
[ "Power Users" 547 ( "Members that can share
directories" ) ]
[ "Guests" 546 ( "Users granted Guest Access" ) ]
[ "Users" 545 ( "Ordinary Users" ) ]
[ "Administrators" 544 ( "Members can fully
administer the filer" ) ]
S-1-5-21-265246955-68147109-1151652928-500
```

Local Administrator

© 2010 NetApp, Inc. All rights reserved.

/ETC/LCLGROUPS.CFG FILE

The lclgroups.cfg maps groups to RIDs. A RID is a unique number that is generated when a group or user is created. For more information, check out this link in NOW:

http://now.netapp.com/NOW/knowledge/docs/ontap/rel732_vs/html/ontap/sysadmin/GUID-89A9ACCA-501C-42DB-949B-B57B9AFBB.B98.html.



SIDs

© 2010 NetApp, Inc. All rights reserved.

SIDS



CLI: cifs lookup

- Windows security identifiers (SIDs) can be converted to user and group IDs or the reverse
 - CLI: `cifs lookup` command
 - NetApp System Manager

```
system> cifs lookup administrator
SID = S-1-5-21-265246955-68147109-1151652928-500

system> cifs lookup S-1-5-32-544
name = BUILTIN\Administrators

system> cifs lookup S-1-5-21-265246955-68147109-1151652928-500
name = system\administrator
```

NOTE: SID might be listed in the `/etc/lclgroups.cfg` file

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS LOOKUP

Security IDs (SIDs) can be converted to user and group IDs using the command-line interface or NetApp System Manager.

The following examples demonstrate conversion using the command-line interface with the `cifs lookup` command.

```
system> cifs lookup S-1-5-32-544
name = BUILTIN\Administrators
```

The SID S-1-5-32-544 is the name “BUILTIN\Administrators.”

```
system> cifs lookup S-1-5-21-265246955-68147109-1151652928-500
name = system\administrator
```

This is the SID for the local administrator, system’s administrator, which also is listed in the `/etc/lclgroups.cfg` file.



SID Cache

To manage the SID Cache,

- **options cifs.sidcache.enable on**
 - Turns on SID Cache
- **options cifs.sidcache.lifetime *time***
 - Sets the normal life span of cached SIDs
- **cifs sidcache clear all**
 - Clears all CIFS SID-to-name map cache entries
- **cifs sidcache clear domain [*domain*]**
 - Clears CIFS SID-to-name map cache entries for a particular domain
- **cifs sidcache clear user [*user*]**
 - Clears CIFS SID-to-name map cache entries for a particular user
- **cifs sidcache clear sid [*sid*]**
 - Clears CIFS SID-to-name map cache entries for a particular SID

© 2010 NetApp, Inc. All rights reserved.

SID CACHE

CIFS frequently is required to map SIDs to user and group names and vice versa for user authentication, quota management, console command processing, and various remote procedure call responses. The SID-to-name map cache contains entries that map SIDs to pre-Windows 2000 Server user and group names.

The storage system obtains the SID-to-name mapping information by querying the domain controller. To minimize multiple lookups of the same names, SID-to-name information received from the domain controller is saved in the SID-to-name map cache on the storage system.

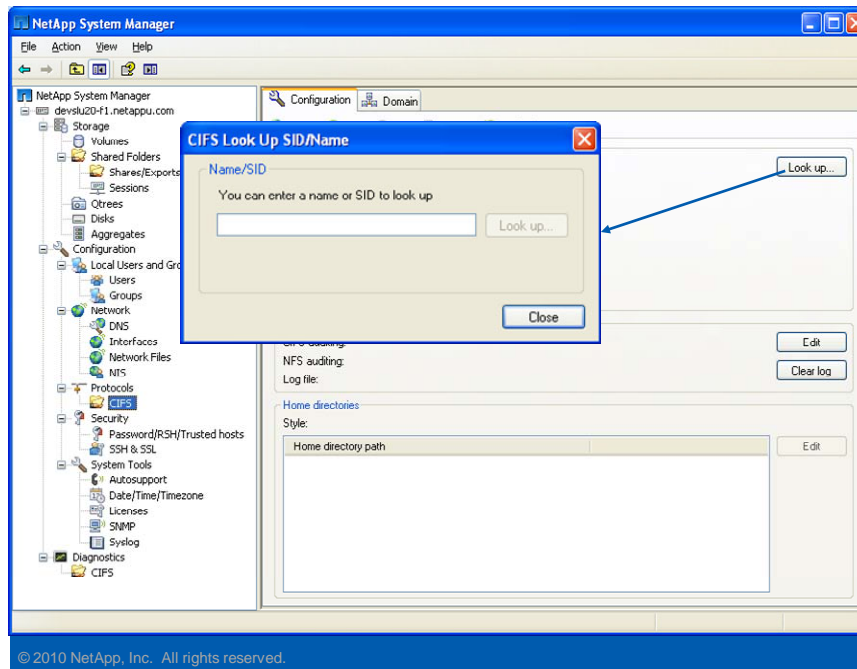
The SID-to-name map cache is enabled on the storage system by default. You can manually control the cache by changing the lifetime of the entries, clearing entries, or turning the SID-to-name map caching off or on. The cache persists if CIFS is terminated or restarted, but it does not persist across a reboot or a takeover and giveback.

When the storage system requires SID-to-name mapping information, it first looks for a matching entry in the SID-to-name map cache. If a matching entry is not found or if an expired matching entry is found, the storage system queries the appropriate domain controller for current mapping information. If the domain controller is not available, an expired mapping entry might be used by the storage system.

Defaults for the SID cache lifetime is 1440 minutes (24 hours) with a minimum value of 20 minutes and a maximum value of 100080 minutes (168 hours).



NetApp System Manager: SID Look Up



NETAPP SYSTEM MANAGER: SID LOOK UP



NetBIOS Aliases

© 2010 NetApp, Inc. All rights reserved.

NETBIOS ALIASES



NetBIOS

- NetBIOS over TCP/IP
 - Means “Network Basic Input/Output System”
 - Is an API that allows machines to be discovered by “name” over TCP/IP
 - A single machine can have multiple names
 - Is typically used by various applications such as Network Neighborhood and `net use`
- Windows clients use the LMHOSTS file to provide NetBIOS name resolution
- On the storage system, set NetBIOS name(s) using `nbalias` and the `cifs_nbalias.cfg` file

© 2010 NetApp, Inc. All rights reserved.

NETBIOS

Network Basic Input/Output System, or NetBIOS, is an application programming interface (API) that provides simple networking services, enabling users to share and use one another’s resources easily.

NetBIOS over TCP/IP (NBT or NetBT) is the standard protocol used for CIFS prior to Windows 2000 Server. The NetBIOS Name Server (NBNS) protocol is part of the NetBIOS over TCP/IP family of protocols. For more information about the NetBIOS over TCP/IP, see chapter 11 of *TCP/IP Fundamentals for Microsoft Windows* located at: www.microsoft.com/downloads/details.aspx?familyid=c76296fd-61c9-4079-a0bb-582bca4a846f&displaylang=en.



NetBIOS Aliases

```
system> rdfile /etc/cifs_nbalias.cfg
#
# This file contains NetBIOS aliases used by the filer.
# See the System Administrator's Guide for a full
# description of this file.
#
# There is a limit to the number of aliases that may be specified.
# Currently that limit is 200.
#
# Aliases must be entered one per line.
#
# After editing this file, use the console command "cifs nbalias load"
# to make the filer process the entries in this file.
#
# Note that the "#" character is valid in a CIFS NetBIOS alias.
# Therefore the "#" character is only treated as a comment in this
# file if it is in the first column.
#
LegacyEMC
Clariion1
Celerra2
OldHitachi1
Stumpy
```

© 2010 NetApp, Inc. All rights reserved.

NETBIOS ALIASES

The /etc/cifs_nbalias.cfg configuration file contains the NetBIOS aliases for the storage system. A NetBIOS alias allows the storage system to be accessed by a Windows client using an alternate name for the storage system. To list the current NetBIOS aliases, do the following:

```
system> cifs nbalias
No NetBIOS aliases
```

```
system> rdfile /etc/cifs_nbalias.cfg
# After editing this file, use the console command
# "cifs nbalias load"
# to make the filer process the entries in this file.
#
# Note that the "#" character is valid in a CIFS# NetBIOS alias.
# Therefore the "#" character is only treated as a# comment in this file if it
is in the first column.
grumpy
```

happy *[Edit and add the NetBIOS aliases.]*



NetBIOS Aliases (Cont.)

- List aliases

```
system> cifs nbalias
```

- Load file after making changes

```
system> cifs nbalias load
```

© 2010 NetApp, Inc. All rights reserved.

NETBIOS ALIASES (CONT.)

After the `/etc/cifs_nbalias.cfg` file has been edited with the proper NetBIOS aliases, use the `cifs nbalias load` command to register the update with the WINS server.



Terminating and Restarting CIFS

© 2010 NetApp, Inc. All rights reserved.

TERMINATING AND RESTARTING CIFS



Stopping and Restarting CIFS

- To terminate CIFS service (a complete shutdown) where all CIFS sessions are ended:
`system> cifs terminate [-t minutes]`
 - To stop a `cifs terminate` command if you have set a duration, click Control-C
- To restart CIFS service after terminating:
`system> cifs restart`
 - Or reconfigure CIFS services (will start automatically)`system> cifs setup`

© 2010 NetApp, Inc. All rights reserved.

STOPPING AND RESTARTING CIFS



CLI: Stopping and Restarting CIFS

- As an example, stop and restart CIFS services on the storage system called “system”

```
system> cifs terminate
```

```
CIFS local server is shutting down...
```

```
CIFS local server has shut down...
```

```
system> cifs restart
```

```
CIFS local server is running.
```

```
GMT[nbt.nbns.registrationComplete:info]: NBT:
```

```
All CIFS name registrations have completed for  
the local server.
```

© 2010 NetApp, Inc. All rights reserved.

CLI: STOPPING AND RESTARTING CIFS



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- License CIFS on a storage system
- Join a storage system to a Windows workgroup environment using the `cifs setup` command
- Observe the results of `cifs setup`
- Manage newly created configuration files for the CIFS workgroup environment

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 6: CIFS Workgroups
Estimated Time: 30 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- In `cifs setup`, what are the two security style choices for which a storage system can be configured?
- During the initial questions in `cifs setup`, for which root user can you enter a password?
- What are the three default share volumes created as a result of `cifs setup`?
- What is the name of the NetBIOS alias file?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

CIFS Shares and Sessions

Module 7
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



CIFS SHARES AND SESSIONS



Module Objectives

By the end of this module, you should be able to:

- Display all shares available on the storage system
- List the default shares
- Configure a client machine to access any share
- Identify the CIFS sessions established by accessing a share on the storage system
- Add, modify, and delete shares

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Share Administration

- Shares may be managed by way of:
 - The command-line interface
 - NetApp® System Manager
 - Microsoft® Management Console (MMC)
 - Computer Management
- Share administration includes:
 - Display shares
 - Add shares
 - Provide access to shares
 - Remove shares

© 2010 NetApp, Inc. All rights reserved.

SHARE ADMINISTRATION



Displaying Shares

© 2010 NetApp, Inc. All rights reserved.

DISPLAYING SHARES



CLI: Displaying CIFS Shares

- As a result of setting up the CIFS service, default shares are created
- To display all shares: `cifs shares`
- Example:

```
system> cifs shares
Name           Mount Point           Description
-----
ETC$           /etc                   Remote Administration
                BUILTIN\Administrators / Full Control

HOME           /vol/vol0/home         Default Share
                everyone / Full Control

C$             /                       Remote Administration
                BUILTIN\Administrators / Full Control
```

© 2010 NetApp, Inc. All rights reserved.

CLI: DISPLAYING CIFS SHARES



System Manager: Displaying CIFS Shares

The screenshot displays the NetApp System Manager interface. On the left, a navigation tree shows the 'Shares/Exports' section selected. The main pane is divided into two sections: 'Shares' and 'Exports'.

Shares Section:

Share name	Shared path	Description	Maximum users	Caching
C\$	/	Remote Administration		
ETC\$	/etc	Remote Administration		
HOME	/vol/vol0/home	Default Share		

Exports Section:

Export name	Export path	Anonymous user ID
/vol/vol0	/vol/vol0	
/vol/vol0/home	/vol/vol0/home	
/vol/vol1	/vol/vol1	

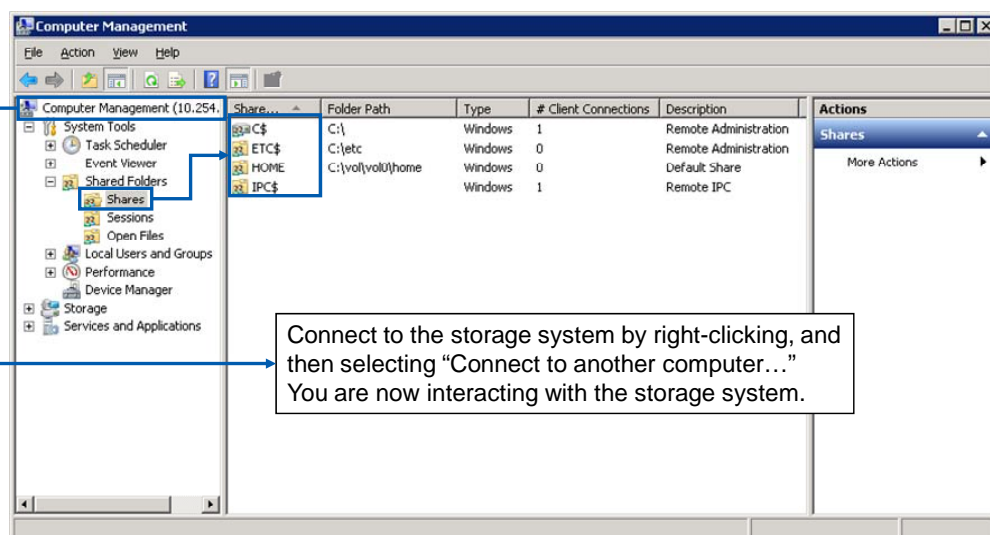
Annotations in the image include a blue box around 'Shares/Exports' in the navigation tree, a blue arrow pointing to the 'HOME' share in the 'Shares' table, and the text 'Default Shares' next to the arrow.

© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: DISPLAYING CIFS SHARES



MMC: Displaying Storage System Shares



NOTE: You must log in with a user account that is defined in the BUILTIN\Administrators group

© 2010 NetApp, Inc. All rights reserved.

MMC: DISPLAYING STORAGE SYSTEM SHARES

To display storage system shares, first connect to the storage system by right-clicking and then selecting “Connect... to another computer...” Click the **Shares** folder in the console tree. The three default shares—**C\$**, **ETC\$**, and **HOME**—display, as does the hidden **IPC\$** share.

The IPC\$ share is an inter-process communications mechanism for temporary connections between clients and servers. It is primarily used to administer network servers remotely. This share enables the communication between the Windows® Computer Management GUI and the storage system.



Accessing Shares

© 2010 NetApp, Inc. All rights reserved.

ACCESSING SHARES



Accessing a Share

- After the share has been created, it may be accessed from Windows by:
 - Microsoft's `net use` command

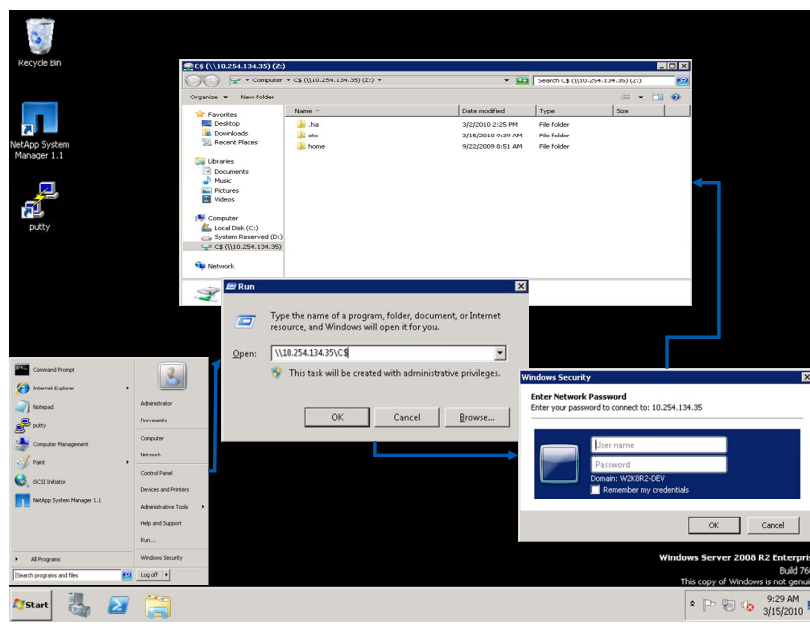
```
C:\> net use e: \\toaster\jdoe /user:marketing\jdoe
```
 - Using the run dialog
 - Mapping a drive from the GUI

© 2010 NetApp, Inc. All rights reserved.

ACCESSING A SHARE



Run Dialog



© 2010 NetApp, Inc. All rights reserved.

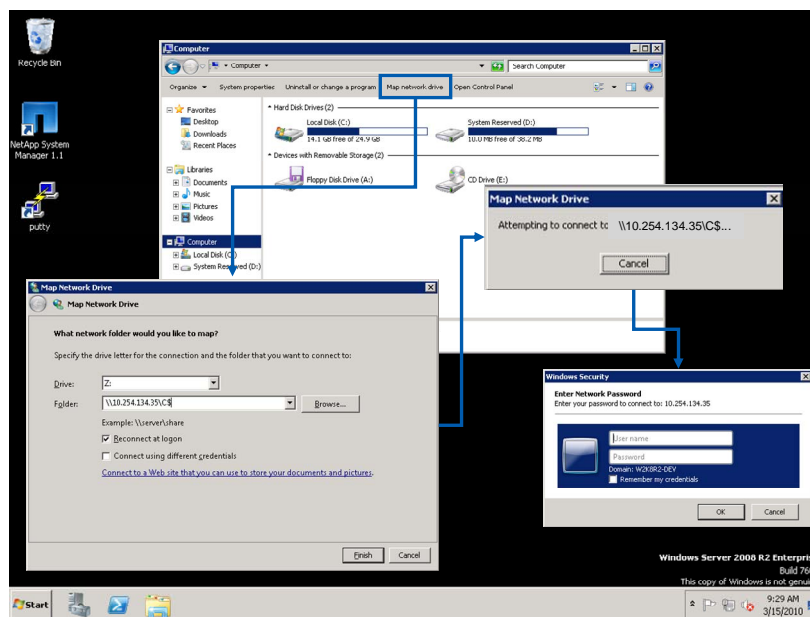
RUN DIALOG

On a Windows workstation using the Windows “run line,” access the C\$ share on the storage system “system” by performing the following steps:

- On the Windows desktop, click the **Start** menu and choose **Run**. The Run window appears.
- In the **Open** text box, type `\\storage_system_name\C$` (`\\system\C$`). **Note:** The storage system name can be the name or IP address. Click the **OK** button, and then the **Connect To** window appears.
- In the **Connect To** window, type the user name **administrator** and the password, and then click the **OK** button. The `\\system\C$` window appears with the share access to C\$ that displays the “etc” and “home” folders.



Mapping a Drive to a Share



© 2010 NetApp, Inc. All rights reserved.


MAPPING A DRIVE TO A SHARE

On a Windows workstation, map a network drive letter to a share by performing the following steps:

- Open **Windows Explorer** and go to **Tools > Map Network Drive**. The Map Network Drive window appears.
- In the **Drive** list box, select any unused letter. In the example, the letter K is selected.
- In the Folder list box, type **\\storage_system\C\$**. **Note:** The storage system name can be the name or IP address.
- Click the **Finish** button. The Map Network Drive attempts to connect to the storage system and share.
- When the Connect to window appears, in the **User name** text box, type administrator and in the **Password** text box, type the administrator's password.
- Click the **OK** button.



Mapping a Drive to a Share (Cont.)



© 2010 NetApp, Inc. All rights reserved.

MAPPING A DRIVE TO A SHARE (CONT.)

(The following continues the mapping of a network drive letter to a share.)

- The mapped network drive letter (K is shown in this example) displays the mapping to the C\$ share. Both the etc and home folders are in the C\$ share.



Encoding

- CIFS uses Unicode for its encoding
- If a volume is exclusively being accessed by CIFS, consider:

```
system> vol options vol create_unicode on
```

```
system> vol options vol convert_unicode on
```
- If the `unicode` options are not set, Data ONTAP® will transparently convert a non-Unicode directory when first accessed by CIFS
 - Time consuming
 - If read-only (that is, Snapshot™ copy), then access is refused

© 2010 NetApp, Inc. All rights reserved.

ENCODING

The CIFS protocol requires the Unicode encoding method. Unicode is an industry standard allowing computers to consistently represent text in most of the world's writing systems. Unicode provides a unique number for every character regardless of the language. See www.unicode.org for more information.

If a volume is exclusively being accessed by CIFS or NFS v4 or later, then consider setting the `create_unicode` and `convert_unicode` volume options. The `Create_unicode` option forces newly created directories to be Unicode directories for both NFS and CIFS. By default it is set to `off`, in which case all directories are created in a non-Unicode format and the first CIFS access will convert it to the Unicode format. The `convert_unicode on` option forces all directories to be converted to the Unicode format when accessed from both NFS and CIFS. By default, this option is set to `off`.

Unicode is not defaulted on a storage system because Unicode directories take up more space and are slower on some workloads.



Sessions

© 2010 NetApp, Inc. All rights reserved.

SESSIONS



CIFS Sessions

- A client establishes a session with a storage system upon the first share access
 - Access is based on user authentication and share access rules
- Display a CIFS session status by using these methods:
 - CLI: `cifs sessions` command
 - NetApp System Manager
 - Windows Computer Management: MMC > System Tools > Shared Folders > Sessions

© 2010 NetApp, Inc. All rights reserved.

CIFS SESSIONS

A client user establishes a session with a storage system upon the first share access. Access is based on user authentication and share access rules. The authentication method is defined by the environment into which the storage system is added.

You can display a CIFS session status by using these methods:

- CLI `cifs sessions` command
- NetApp System Manager
- Windows Computer Management GUI ->**SystemTools** -> **SharedFolders** -> **Sessions**



cifs sessions Command

With the `cifs sessions` command, you can display the following types of session information:

- A summary of session information, including the number of open shares and files opened by user
`system> cifs sessions`
- Share and file information about a specified connected user or all connected users, including shares and files opened
`system> cifs sessions [username|IPaddress|host]`
`system> cifs sessions * [all connected users]`
- Security information
`system> cifs sessions -s`

© 2010 NetApp, Inc. All rights reserved.

CIFS SESSIONS COMMAND

With the `cifs sessions` command, you can display the following types of session information:

- A summary of session information, including storage system information and the number of open shares and files opened by each connected user:
 - `cifs sessions`
- Share and file information about a specified connected user or all connected users, including:
 - The names of shares opened by a specified connected user or all connected users
 - The access levels of opened files
 - `cifs sessions user_name | IP_address |workstation_name`
 - `cifs sessions * [all connected users]`
- Security information about a specified connected user or all connected users, including the UNIX® user ID (UID) and a list of UNIX groups and Windows groups to which the user belongs:
 - `cifs sessions -s user_name | IP_address | workstation_name`
 - `cifs sessions -s [all connected users]`

NOTE: The number of open shares shown in the session information includes the hidden IPC\$ share.

The `cifs sessions` command can be used as a “status” command even when there is no session.

Example 1 is a storage system in a Windows workgroup. The storage system uses local user authentication.

```
system> cifs sessions
Server Registers as 'system' in workgroup 'WORKGROUP1'Root volume language is
not set. Use vol lang. Using Local Users authentication
Comment: This is a Windows workgroup server
=====
PC IP(PC Name) (user)                #shares  #files
```

Example 2 is a storage system in a Windows 2000 Server domain. The storage system uses the domain controller for authentication.

```
system> cifs sessions
Server Registers as 'system' in Windows 2000 domain 'DEVELOPMENT'
Root volume language is not set. Use vol lang.
Selected domain controller \\DEVDC01 for authentication
Comment: This is a Windows 2000 member server
=====
PC IP(PC Name) (user)          #shares  #files
```

Options:

- The -t option displays the total count of CIFS sessions, open shares, and open files.
- If you include the user argument, the command displays information about the specified user, along with the names and access level of files that user has opened. If you use * as the specified user, the command lists all users.
- Specifying the -c option with a user argument will display the names of open directories and the number of active change notify requests against the directory.
- The -s option displays security information for a specified connected user. If you do not specify a user or workstation name, the command displays security information for all users.

Here are examples using the *machine_name* and *machine_IP_address* arguments:

```
system> cifs sessions 192.168.228.4
users  shares/files opened
TORTOLA (nt-domain\danw - root)
HOME
```

```
system> cifs sessions tortola
users  shares/files opened
TORTOLA (nt-domain\danw - root)
HOME
```

Here is an example using the -t option:

```
system> cifs sessions -t
Using domain authentication. Domain type is Windows NT.
Root volume language is not set. Use vol lang.
Number of WINS servers: 2
CIFS sessions: 1
CIFS open shares: 1
CIFS open files: 3
CIFS sessions using security signatures: 0
```



cifs sessions Example

- The following example of the `cifs sessions` command shows a session with a storage system in a Windows domain

```
system> cifs sessions
Server Registers as 'system' in workgroup 'WORKGROUP'
Root volume language is not set. Use vol lang.
Using Local Users authentication
=====
PC IP(PC Name)  (user)                                #shares  #files
10.254.134.40() (system\administrator 1              0
               - root)
```

© 2010 NetApp, Inc. All rights reserved.

CIFS SESSIONS EXAMPLE

The following example of the `cifs sessions` command shows a session with a storage system in a Windows workgroup.

The **PC IP address 10.254.134.40** is the Windows workstation WIN.

The **system\administrator** user is the local administrator account on the storage system.

The user mapping for this account is root.

One share is currently being accessed.



CLI: cifs sessions Security Information

```
system> cifs sessions -s
users
Security Information
10.254.134.40() (system\administrator - root)
*****
UNIX uid = 0
user is a member of group daemon (1)
user is a member of group daemon (1)

NT membership
    system\administrator
    BUILTIN\Administrators
User is also a member of Everyone, Network Users,
Authenticated Users
*****
```

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS SESSIONS SECURITY INFORMATION

The following example of `cifs sessions -s` command shows security information for a user with a session with a storage system in a Windows workgroup.



NetApp System Manager: CIFS Sessions

NetApp System Manager

File Action View Help

NetApp System Manager

devslu20-f1.development.netapp

storage

Volumes

Shared Folders

Shares/Exports

Sessions

Quotas

Qtrees

Disks

Aggregates

Configuration

Diagnostics

Refresh

User	Computer	IP address	# Open shares	# Open directories	# Open files
(DEVSLU20-F1)\administrator - pcuser	10.254.147.54	10.254.147.54	1	1	0

Current Sessions

Local storage system's administrator account shown

Accessed volumes list:

vol0

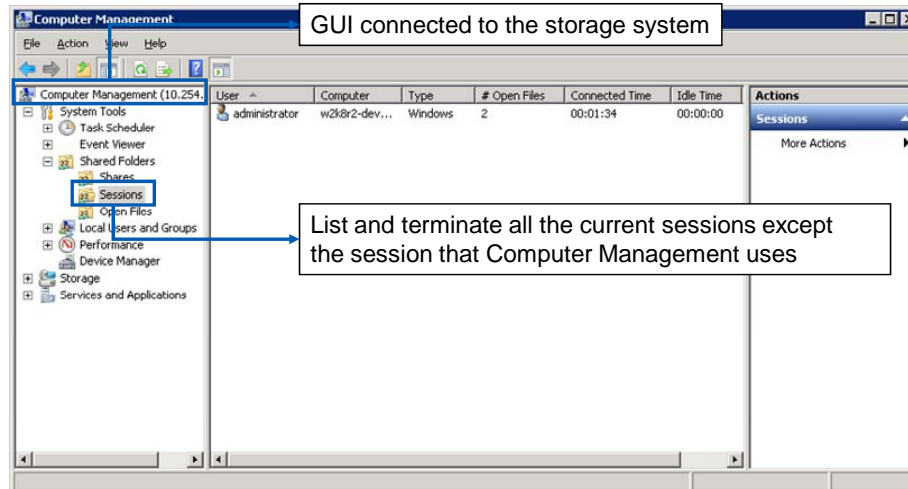
Highlighted session's access volume

© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: CIFS SESSIONS



MMC: CIFS Sessions



© 2010 NetApp, Inc. All rights reserved.

MMC: CIFS SESSIONS

With the Computer Management GUI, click the **System Tools->SharedFolders->Sessions** folders to display the CIFS sessions.

In this example, the **local administrator** has a session with the storage system **DEV270-1**, which is in a Windows workgroup.

- The name of the administrator's computer is **10.254.134.40**
- The number of **Open Files** is **3**
- This account is not a **Guest** account



Broadcasting a Message

- To display a message on Windows users' sessions:

```
system> cifs broadcast [workstation| -v  
volname] "message"
```

 - You can inform users about pending terminations or other important events
- The Messenger service on the Windows workstation must be enabled
 1. On your Windows workstation, go to: Start > Programs > Administrative Tools > Services > Messenger
 2. If the Messenger service is disabled, start the service

NOTE: The Messenger service is not available on Microsoft Windows Server 2008 R2 and therefore the `cifs broadcast` command is not available

© 2010 NetApp, Inc. All rights reserved.

BROADCASTING A MESSAGE

To display a message on Windows users' workstations, use the following command:

```
cifs broadcast {workstation | -v volname} "message"
```

You can inform users about pending terminations or other important events.

The Messenger service on the Windows workstation must be enabled. **NOTE:** It is disabled by default for security reasons.

To enable the Messenger service on your Windows workstation:


- Go to **Start->Programs->Administrative Tools->Services->Messenger**.
- If the Messenger service is disabled, start the service. (The default is disabled.)



Broadcasting a Message Example

- Example of broadcasting a message from a storage system:

```
system> cifs broadcast -v flexvol1 "The shutdown will start in 10 minutes."
```
- The following message displays on the Windows workstation:



© 2010 NetApp, Inc. All rights reserved.

BROADCASTING A MESSAGE EXAMPLE

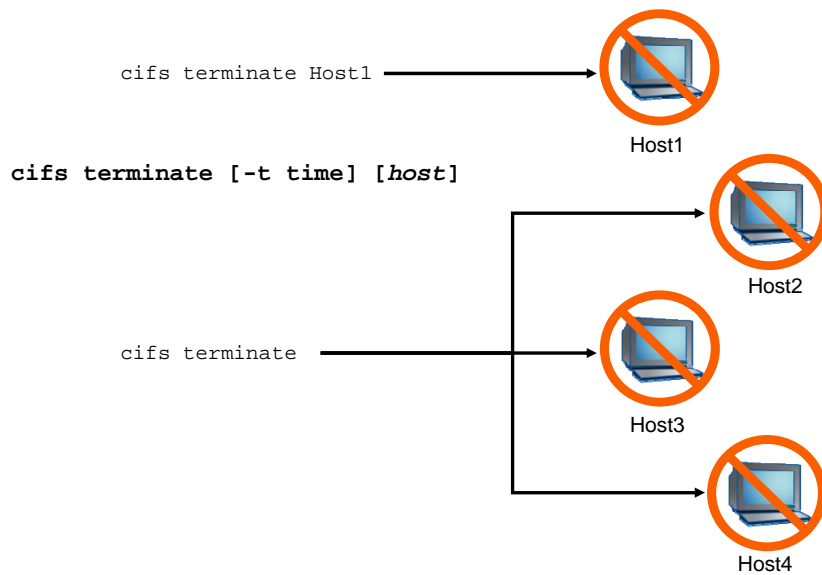
The following is an example of broadcasting a message using the volume option from a storage system:

```
system> cifs broadcast -v flexvol1 "The shutdown will start in 10 minutes."
```

The message "The shutdown will start in 10 minutes" will be broadcasted to all users that have sessions on the volume named flexvol1.



Terminating Sessions



© 2010 NetApp, Inc. All rights reserved.

TERMINATING SESSIONS

The `cifs terminate` command stops the CIFS service. If a single host is named, all CIFS sessions opened by that host are terminated. If a host is not specified, all CIFS sessions are terminated and the CIFS service is shut down.

If you run the `cifs terminate` command without specifying a time until shutdown and there are users with open files, you are prompted to enter the number of minutes to delay before terminating. If the CIFS service is terminated immediately on a host that has one or more files open, users will not be able to save changes. You can use the `-t` option to warn users of an impending service shutdown. If you execute `cifs terminate` from `rsh`, you must supply the `-t` option.

EXAMPLE	RESULT
<code>cifs terminate -t 10 gloriaswan</code>	Terminates a session in 10 minutes for the host gloriaswan. Alerts are sent periodically to the affected host(s).
<code>cifs terminate -t 0</code>	Terminates all CIFS sessions immediately for all clients.
<code>cifs restart</code>	Reconnects the storage appliance to the domain controller and restarts CIFS service.



Creating and Deleting Shares

© 2010 NetApp, Inc. All rights reserved.

CREATING AND DELETING SHARES



Default Shares

- As you recall, three default share definitions are created upon completion of `cifs` setup:
 - C\$
 - ETC\$
 - HOME
- But you can create new shares...

© 2010 NetApp, Inc. All rights reserved.

DEFAULT SHARES



Creating a Share

- When you create a share, you must provide:
 - Complete path name
 - Name of the share
 - Optionally, a description of the share
- Data ONTAP CLI also allows:
 - Group membership for files in the share
 - Support for wide symbolic links
 - Disabling or enabling of virus scanning when files in the share are first opened
- MMC also allows permissions for the share

© 2010 NetApp, Inc. All rights reserved.

CREATING A SHARE

When you create a share, you must provide these items:

- The complete path name of an existing volume or directory to be shared
- The name of the share entered by users when they connect to the share
- Optionally, a description of the share

When creating a share from the Data ONTAP command-line interface, you can specify a variety of share properties, including group membership for files in the share, support for wide symbolic links, and disabling of virus scanning when files in the share are first opened. Virus scanning occurs when files are opened, renamed, and closed after being modified.

Microsoft interfaces additionally allow the administrator to set permissions as the share is created.



Creating a Share (Cont.)

- Additional properties can be set or modified after creating a share:
 - Maximum number of users who can simultaneously access the share
 - If not specified, the limit is defined by the storage system's memory
 - Share-level access control list (ACL)

© 2010 NetApp, Inc. All rights reserved.

CREATING A SHARE (CONT.)

After you have created a share, you can specify these share properties:

- Maximum number of users who can simultaneously access the share
 - If you do not specify a number, the number of users is limited by storage system memory
- Share-level access control list (ACL)



CLI: Preparing to Create a Share

- You can create shares for folders, qtrees, or volumes
- For example:
 - To prepare for creating a share on a qtree, first create the following resources:
 - An aggregate (aggr1)
 - A flexible volume (flexvol1) on aggr1
 - A qtree (datatree1) on flexvol1

NOTE: This path example will be used throughout this module

© 2010 NetApp, Inc. All rights reserved.

CLI: PREPARING TO CREATE A SHARE

You can create shares for volumes or directories including qtrees.

For example, to prepare for creating a share on a qtree, first create the following resources:

- An aggregate (aggr1)
- A flexible volume (flexvol1) on aggr1
- A qtree (datatree1) on flexvol1



CLI: Adding a Share

- As an example, add a share called `datatree1` (for the `qtree datatree1`)

```
system> cifs shares -add datatree1
                        /vol/flexvol1/datatree1
                        -comment "Qtree for Windows Users"
```

The share name 'datatree1' will not be accessible by some MS-DOS workstations

Are you sure you want to use this share name? [n]: **y**

Name	Mount Point	Description
----	-----	-----
datatree1	/vol/flexvol1/datatree1	Qtree for Windows Users
	<div style="border: 1px solid blue; padding: 2px;">everyone / Full control</div>	

Default access control (*discussed later*)

© 2010 NetApp, Inc. All rights reserved.

CLI: ADDING A SHARE

For example, on a storage system, add a share called `datatree1` (for the `qtree datatree1`).

```
system> cifs shares -add datatree1 /vol/flexvol1/datatree1 -comment "Qtree for Windows Users"
```

The share name 'datatree1' will not be accessible by some MS-DOS workstations

Are you sure you want to use this share name? [n]:y

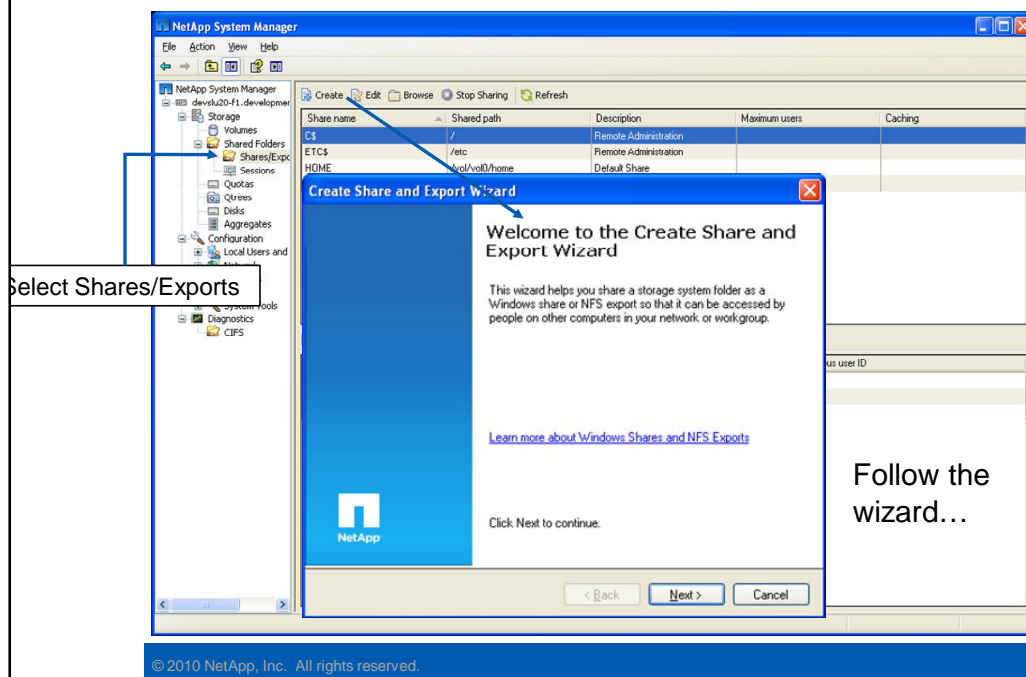
```
system> cifs shares datatree1
```

Name	Mount Point	Description
----	-----	-----
datatree1	/vol/flexvol1/datatree1	Qtree for Windows Users
	everyone / Full Control	

The default access control is full control for everyone.



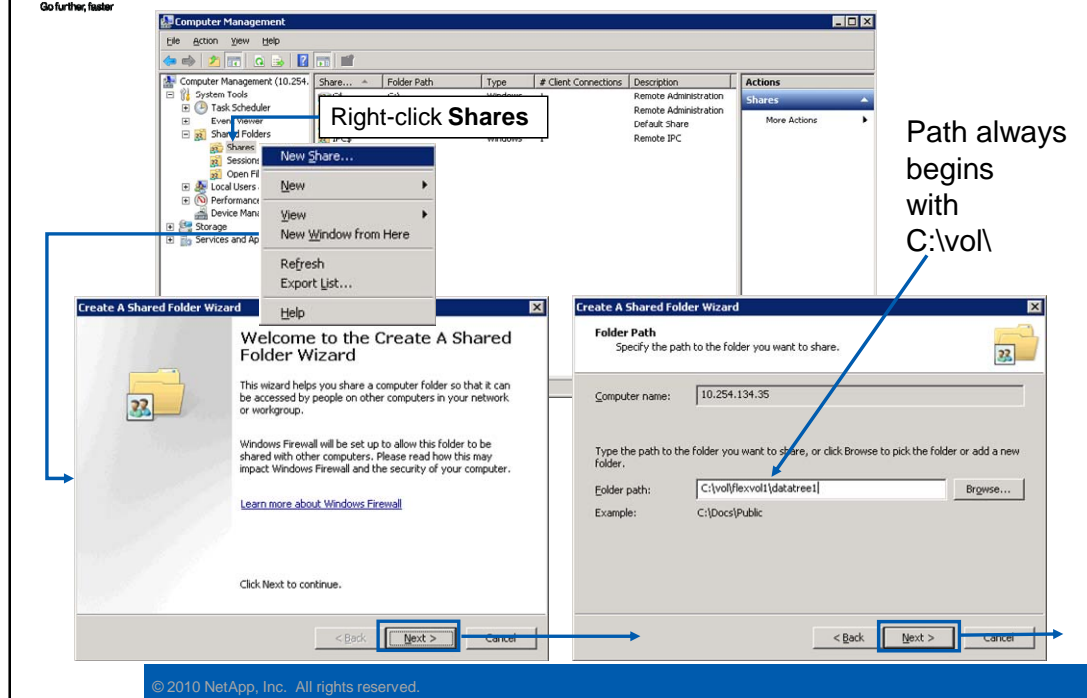
NetApp System Manager: Adding a Share



NETAPP SYSTEM MANAGER: ADDING A SHARE



MMC: Adding a Share

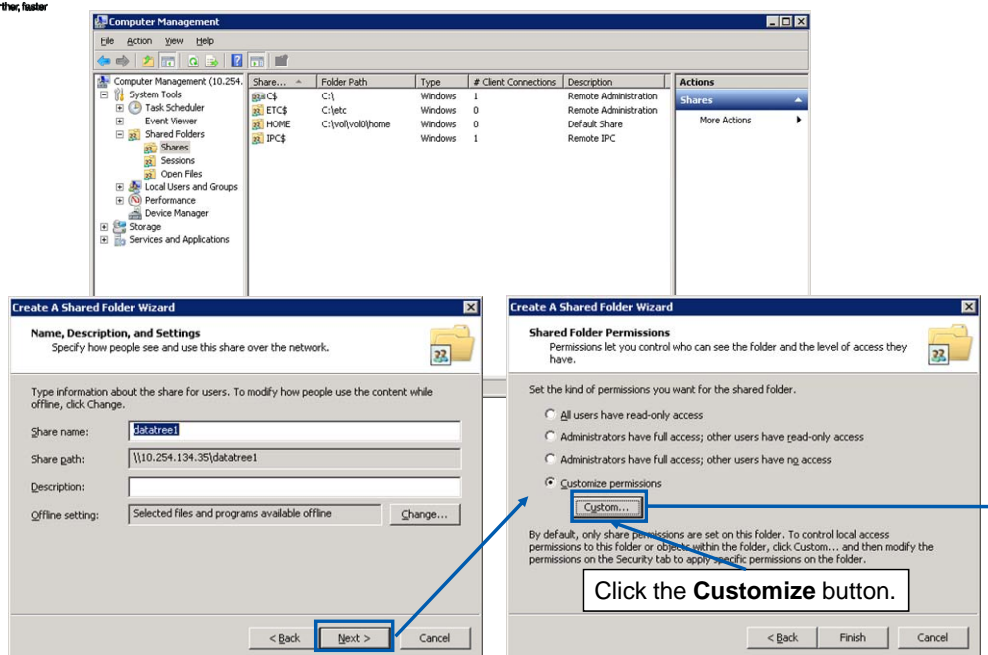


MMC: ADDING A SHARE

As an example with the Windows Computer Management GUI, add a new share called **datatree1** (for the qtree **datatree1**) on volume **flexvol1** by performing the following steps:

- In the console tree, right-click the **Shares** folder and choose **New Share....** The **Welcome to the Share a Folder Wizard** appears.
- Click the **Next** button to start the wizard, and the **Folder path** page displays with the **Computer name** text box showing your storage system name or IP address.
- In the **Folder path** text box, type the path **C:\vol\flexvol1\datatree1** for the **datatree1** share, and click the **Next** button.

MMC: Adding a Share (Cont.)



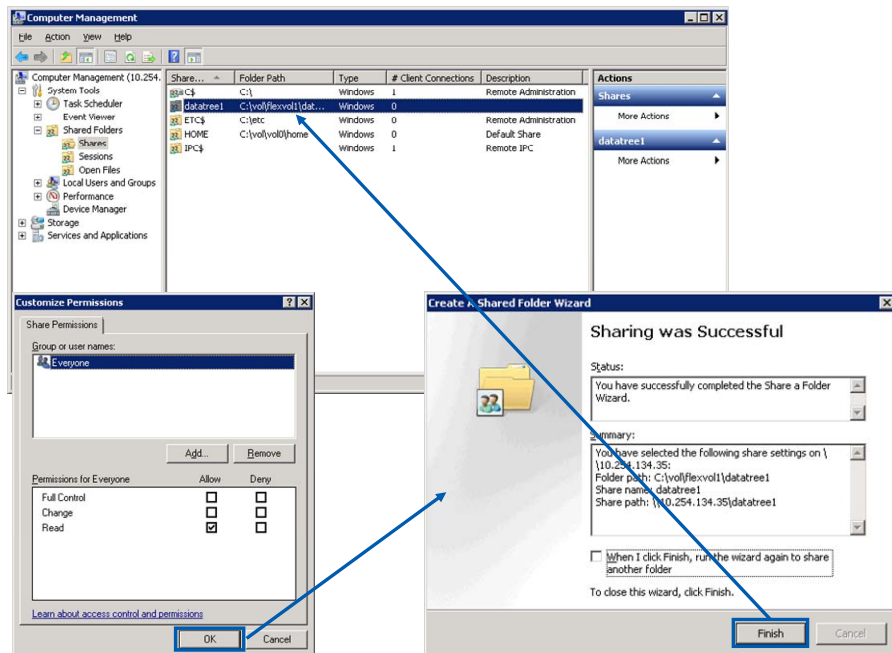
© 2010 NetApp, Inc. All rights reserved.

MMC: ADDING A SHARE (CONT.)

(The following continues the adding of a CIFS share.)

- On the “Name, Description, and Settings” page, in the **Share name** text box, enter **datatree1**.
- In the **Description** text box, type **Qtree for Windows Users**, and then click the **Next** button.
- In the “Permissions” page, mark the **Use custom share and folder permissions** radio button, and then click the **Customize** button.

MMC: Adding a Share (Cont.)



© 2010 NetApp, Inc. All rights reserved.

MMC: ADDING A SHARE (CONT.)

(The following continues the adding of a CIFS share.)

- In the Customize Permissions window, mark the check boxes for **Full Control**, **Change**, and **Read**, and click the **OK** button.
- In the Permissions page, click the **Finish** button.
- You receive a message stating that sharing was successful.
- Click the **Close** button to close the wizard.



CLI: Deleting a Share

- As an example, delete the share called `datatree1`

```
system> cifs shares -delete datatree1
```

```
system> cifs shares
```

Name	Mount Point	Description
----	-----	-----
ETC\$	/etc	Remote Administration BUILTIN\ Administrators / Full Control
HOME	/vol/vol0/home	Default Share everyone / Full Control
C\$	/	Remote Administration BUILTIN\ Administrators / Full Control

NOTE: The share `datatree1` is deleted not the underlying volume, `qtree`, or directory

© 2010 NetApp, Inc. All rights reserved.

CLI: DELETING A SHARE

For example, delete the share called `datatree1`:

```
system> cifs shares -delete datatree1
```

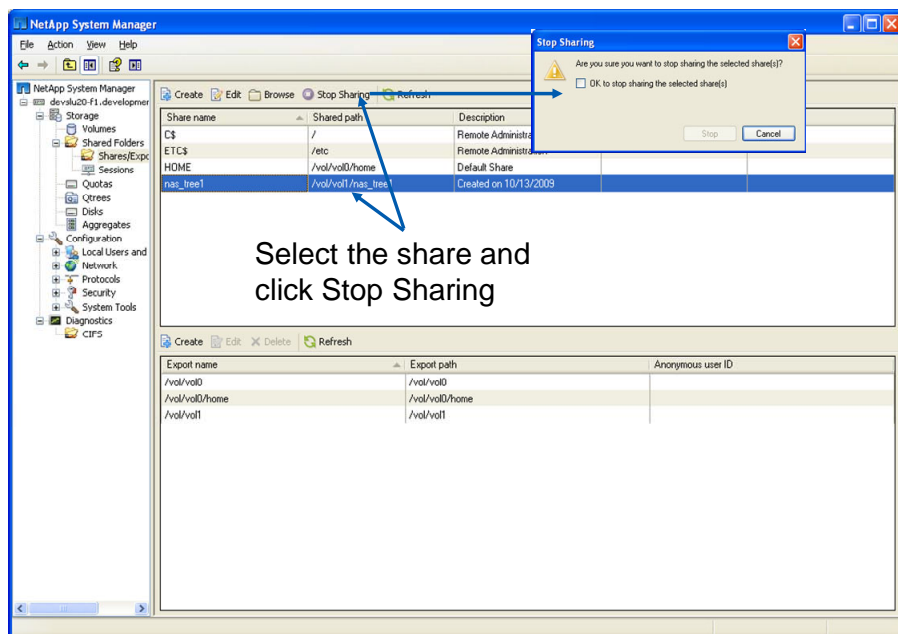
```
system> cifs shares
```

Name	Mount Point	Description
----	-----	-----
ETC\$	/etc	Remote Administration BUILTIN\Administrators / Full Control
HOME	/vol/vol0/home	Default Share everyone / Full Control
C\$	/	Remote Administration BUILTIN\Administrators / Full Control

NOTE: The share `datatree1` is deleted.



NetApp System Manager: Deleting a Share

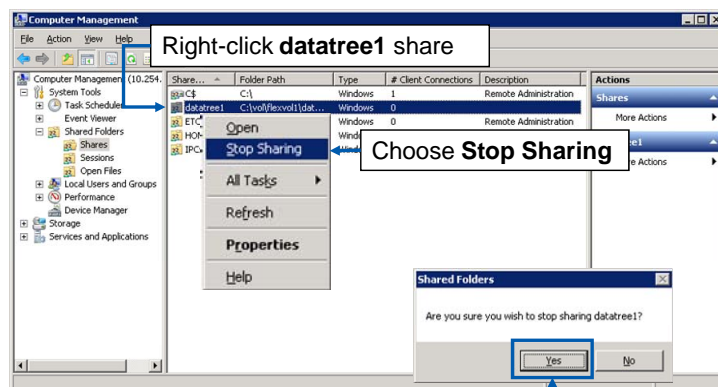


© 2010 NetApp, Inc. All rights reserved.

NETAPP SYSTEM MANAGER: DELETING A SHARE



MMC: Deleting a Share



Click the **Yes** button to confirm stop sharing datatree1

© 2010 NetApp, Inc. All rights reserved.

MMC: DELETING A SHARE

As an example with the Windows Computer Management GUI, delete the share called **datatree1** by performing the following steps:

- In the Computer Management window, right-click the **datatree1** share and choose **Stop Sharing**.
- In the Shared Folders window, when it asks if you are sure that you wish to stop sharing **datatree1**, click the **Yes** button.



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Display all shares available on the storage system
- List the default shares
- Configure a client machine to access any share
- Identify the CIFS sessions established by accessing a share on the storage system
- Add, modify, and delete shares

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 7: CIFS Shares and Sessions

Estimated Time: 15 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- For which storage objects can you create shares?
- What are three methods to manage CIFS shares?
- What command would you use to view the connected CIFS users?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

CIFS Access Control

Module 8
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



CIFS ACCESS CONTROL



Module Objectives

By the end of this module, you should be able to:

- Create and manage local users for a storage system
- Identify how to create a local group and make a local user a member of that group
- Use the command-line interface, NetApp® System Manager or Microsoft® tools to add, delete, and modify access permissions of shares
- Use Microsoft tools to add, delete, and modify access permissions of files and folders

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Local Users

© 2010 NetApp, Inc. All rights reserved.

LOCAL USERS



Local Users

Local users are:

- Accounts that are authenticated locally
- Associated with groups on the storage system
- Created and managed using the `useradmin` command or a text editor
- Saved in `/etc/registry` or `/etc/passwd`

© 2010 NetApp, Inc. All rights reserved.

LOCAL USERS

On the storage system, the domain administrators group and the local administrator account are part of the `BUILTIN\Administrators` group. They can do the following:

- Provide a text editor to edit configuration files. Data ONTAP® does not include an editor.
- Provide the ability to administer a storage system and hence have access to the root file system (`C$` and `ETC$`).
- Modify the share access for `C$` and `ETC$` to grant additional users access.
- The local administrator can set up local users on the storage system with the `useradmin user add` command.



Purpose of Local Users

Two main reasons for local user authentication:

1. Provides local administrators the ability to configure the storage system
 - Discussed in the *Data ONTAP 8.0 7-Mode Administration* course
2. Provides local client users access to the resources on the storage system for all environments
 - Windows® workgroup
 - Non-Windows workgroup
 - Windows domain

NOTE: You can create a maximum of 96 local user accounts

© 2010 NetApp, Inc. All rights reserved.

PURPOSE OF LOCAL USERS

Reasons for local user accounts include the following:

- Windows workgroup
 - You must create local user accounts so that the storage system can authenticate local users.
- Non-Windows workgroup (UNIX mode)
 - Do not create local user accounts because the storage system authenticates users with the UNIX password (/etc/passwd) database.
- Windows domain
 - The storage system can authenticate users (with the local user accounts) who try to connect to the storage system from an untrusted domain.
 - Local users can access the storage system when the domain controller is down or not available for domain authentication.

NOTE: You can create a maximum of 96 local user accounts.



Purpose of Local Users (Cont.)

When the CIFS server is configured for:

- Windows workgroup
 - You must create local user accounts so that the storage system can authenticate users
 - Use the `useradmin` command
 - User accounts are stored in `/etc/registry`
- Non-Windows workgroup (UNIX mode)
 - You must create local UNIX users
 - Use the `passwd` command
 - User accounts are stored in `/etc/passwd` and `/etc/shadow`

© 2010 NetApp, Inc. All rights reserved.

PURPOSE OF LOCAL USERS (CONT.)



Purpose of Local Users (Cont.)

When the CIFS server is configured for:

- Windows domain
 - Storage system can authenticate users (with the local user accounts) who try to connect to the storage system from an untrusted domain
 - Local users can access the storage system when the domain controller is down or not available for domain authentication
 - Use the `useradmin` command
 - User accounts are stored in `/etc/registry`

© 2010 NetApp, Inc. All rights reserved.

PURPOSE OF LOCAL USERS (CONT.)



Local Administrator

As you recall, during `cifs` setup, the local administrator account may be created

It is highly recommended that you create the local administrator account: `(system\administrator)` for this filer. This account allows access to CIFS from Windows when domain controllers are not accessible.

Do you want to create the `system\administrator` account? `[y]`:

Enter the new password for `system\administrator`:
Retype the password:

© 2010 NetApp, Inc. All rights reserved.

LOCAL ADMINISTRATOR



Local User Definitions

List the local users on the storage system

```
system> useradmin user list
```

```
Name: root
```

```
Info: Default system administrator.
```

```
Rid: 0
```

```
Groups:
```

} This is the storage
system root user
account

```
Name: administrator
```

```
Info: Built-in account for administering the filer
```

```
Rid: 500
```

```
Groups: Administrators
```

© 2010 NetApp, Inc. All rights reserved.

LOCAL USER DEFINITIONS

A local administrator is added to the user list if the response during `cifs setup` was to create a local administrator account for the storage system. Be sure to set an appropriate password for the administrator account.



Administrating Local Users

- Local users
 - Must provide a unique name
 - Associate user to a group
 - Created only by way of the command-line interface's `useradmin` command when the storage system is set to CIFS workgroup authentication

© 2010 NetApp, Inc. All rights reserved.

ADMINISTRATING LOCAL USERS



Local User Management

- Manage local users fully by using the command-line interface `useradmin` command

- To add a new local user:

```
system> useradmin user add user -g group
```

- To modify a local user :

```
system> useradmin user modify user -g group
```

- To list user information:

```
system> useradmin user list user
```

- To delete a local user:

```
system> useradmin user delete user
```

© 2010 NetApp, Inc. All rights reserved.

LOCAL USER MANAGEMENT



CLI: Adding a New Local User

- As an example, add a local user called Jane to the predefined Guests group

NOTE: User names are not case sensitive

```
system> useradmin user add jane -g Guests
New password:
Retype new password:
user <jane> added.
system> Mon Jul 31 01:13:18 GMT
[useradmin.added.deleted:info]:
The user 'jane' has been added.
```

← Password is typed but not displayed

© 2010 NetApp, Inc. All rights reserved.

CLI: ADDING A NEW LOCAL USER

As an example, add a local user called Jane to the predefined Guests group.

NOTE: User names are not case sensitive.

```
system> useradmin user add jane -g Guests
New password:Retype new password:User <jane> added.system> Mon Jul 31 01:13:18
GMT [useradmin.added.deleted:info]: The user 'jane' has been added.
```

NOTE: The password is typed but not displayed.



CLI: Adding a New Local User (Cont.)

- In the example, verify that the local user Jane has been added to the predefined Guests group

```
system> useradmin user list jane
Name: jane
Info:
Rid: 131075
Groups: Guests
Full Name:
Allowed Capabilities:
Password min/max age in days: 0/4294967295
Status: enabled
```

© 2010 NetApp, Inc. All rights reserved.

CLI: ADDING A NEW LOCAL USER (CONT.)

In the example, verify that the local user Jane has been added to the predefined Guests group.

```
system> useradmin user list jane
Name: jane
Info:
Rid: 131075
Groups: Guests
Allowed Capabilities:
Password min/max age in days: 0/4294967295
Status: enabled
```

NOTE: Jane has no allowed capabilities in the Guests group, but she can log in and be authenticated.



Local Groups

© 2010 NetApp, Inc. All rights reserved.

LOCAL GROUPS



Local Groups

- Local groups
 - Contain local and domain users
 - Created only by way of the command-line interface's `useradmin` command when the storage system is set to CIFS workgroup authentication

© 2010 NetApp, Inc. All rights reserved.

LOCAL GROUPS

MMC tools have some capabilities that are discussed in the next module because they only are available when the storage system is using CIFS domain authentication.



CLI: Group Management

- Manage local groups by using the command-line interface command `useradmin`
 - To add a new group:
`system> useradmin group add group -r role`
 - To modify an existing group:
`system> useradmin group modify group -g newName`
 - To list group information:
`system> useradmin group list group`
 - To delete a group:
`system> useradmin group delete group`

© 2010 NetApp, Inc. All rights reserved.

CLI: GROUP MANAGEMENT



CLI: Local Groups

- As an example, add a local group called Helpers with the predefined admin role

```
system> useradmin group add Helpers -r admin
Group <Helpers> added.
system > Mon Jul 31 02:02:43 GMT
[useradmin.added.deleted:info]: The group
'Helpers' has been added.
```

```
system > useradmin group list Helpers
Name: Helpers
Info:
Rid: 131076
Roles: admin
Allowed Capabilities: login-*, cli-*, api-*,
                      security-*
```

© 2010 NetApp, Inc. All rights reserved.

CLI: LOCAL GROUPS

As an example, add a local group called Helpers with the predefined admin role and verify the results.

```
system> useradmin group add Helpers -r admin
Group <Helpers> added.system> Mon Jul 31 02:02:43 GMT
[useradmin.added.deleted:info]: The group 'Helpers' has been added.

system> useradmin group list Helpers
Name: Helpers
Info:
Rid: 131076
Roles: admin
Allowed Capabilities: login-*,cli-*,api-*,security-*
```

NOTE: The admin role has full capabilities.

When groups are created, they are placed in the `lclgroups.cfg` file. Normally, this file is for administrative reference only; it is not used to reload groups into the system memory. However, sometimes you need Data ONTAP to reload this file—for example, when you migrate a storage system. Do not edit this file without direction from Technical Support.



Share Permissions

© 2010 NetApp, Inc. All rights reserved.

SHARE PERMISSIONS



Permissions

- Permissions can be set at:
 - Share level
 - Folder or file level
- Both permission levels must be satisfied to gain access to the resource

© 2010 NetApp, Inc. All rights reserved.

PERMISSIONS



Share Permissions

- Share permissions can be managed by:
 - Command-line interface: `cifs access` command
 - NetApp System Manager
 - MMC such as Computer Management
- Windows share permissions are the following:
 - Read-only
 - Full control
 - Change
- If all the permissions are denied, then there is no access

© 2010 NetApp, Inc. All rights reserved.

SHARE PERMISSIONS



cifs access Command

- The command-line interface `cifs access` command sets or modifies the share-level ACL to share definitions
 - To modify a share access:
`cifs access <share> [-g] [user_rights]`
 - To delete an ACL entry for a user on a share:
`cifs access -delete <share> [-g] [user]`
- The `-g` option specifies that the user is the name of a UNIX group; use this command when you have:
 - A UNIX group and a UNIX user or an NT user or group with the same name

© 2010 NetApp, Inc. All rights reserved.

CIFS ACCESS COMMAND

The command-line interface `cifs access` command sets or modifies the share-level access control list (ACL) to share definitions.

- To modify a share access:
`cifs access <sharename> [-g] [user_rights]`
- To delete an ACL entry for a user on a share:
`cifs access -delete <sharename> [-g] [user]`

The `-g` option specifies that the user is the name of a UNIX group. Use this command when you have a UNIX group and a UNIX user or an NT user or group with the same name.



CLI: Setting and Deleting Share Access

- As an example, on the datatree1 share, set the share access for the friends group to Full Control and delete the Everyone access

```
system> cifs access datatree1 friends Full Control
1 share(s) have been successfully modified
```

```
system> cifs access -delete datatree1 everyone
1 share(s) have been successfully modified
```

```
system> cifs shares datatree1
```

Name	Mount Point	Description
----	-----	-----
datatree1	/vol/flexvol1/datatree1	Windows Qtree

```
system\friends / Full Control
```

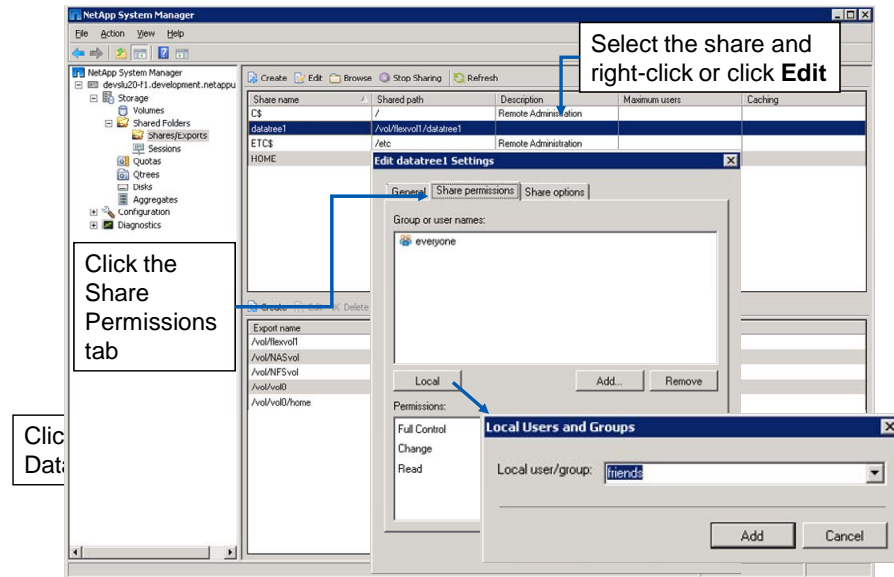
NOTE: This is the storage system local administrator

© 2010 NetApp, Inc. All rights reserved.

CLI: SETTING AND DELETING SHARE ACCESS



System Manager: Setting Share Access

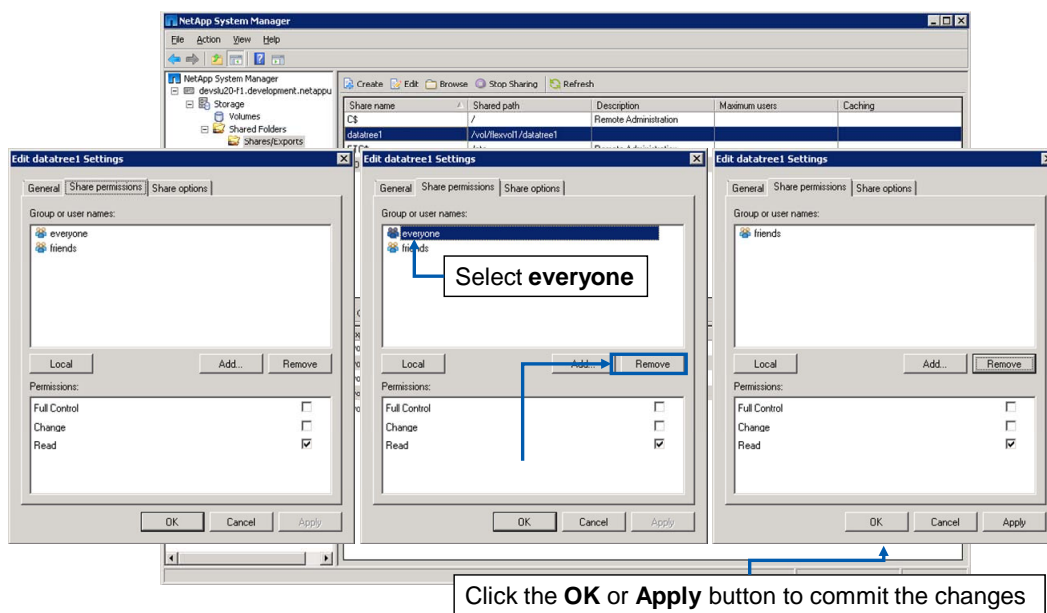


© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: SETTING SHARE ACCESS



System Manager: Deleting Share Access

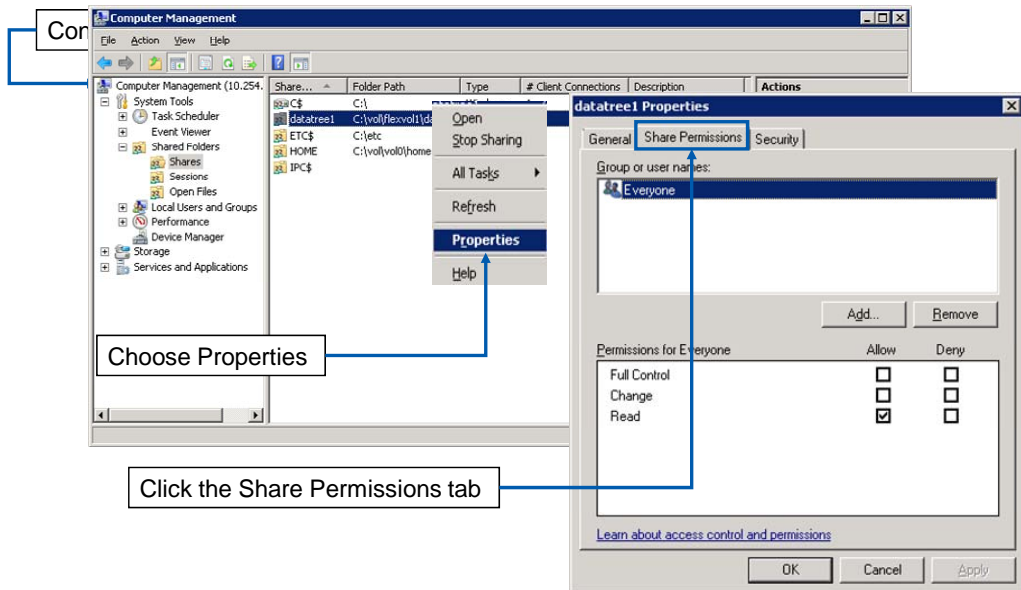


© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: DELETING SHARE ACCESS



MMC: Setting and Deleting Share Access



© 2010 NetApp, Inc. All rights reserved.

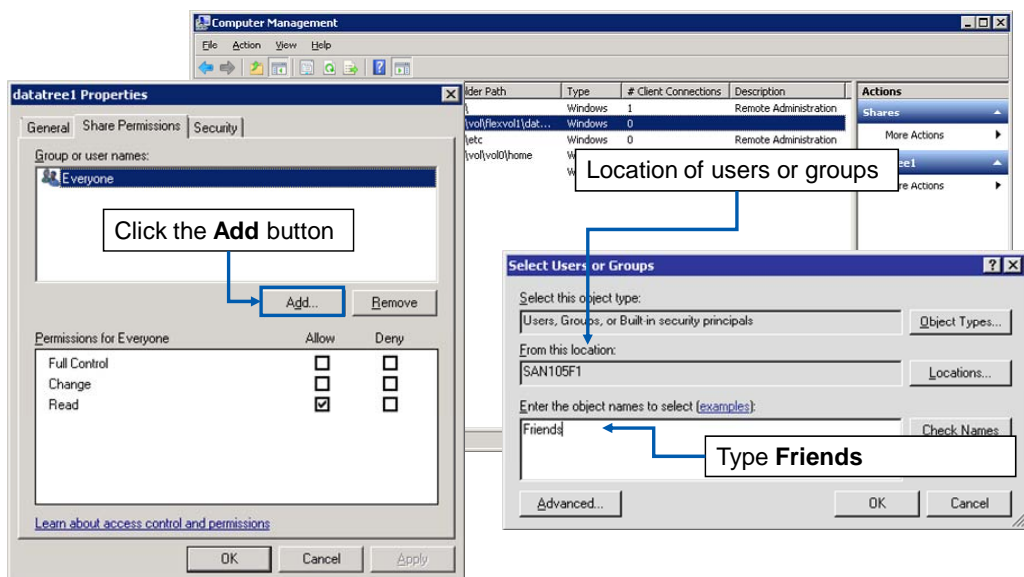
MMC: SETTING AND DELETING SHARE ACCESS

As an example with Windows Computer Management GUI, on the **datatree1** share, set the share access for the administrator to Full Control and delete the Everyone access by performing the following steps:

- Right-click the **datatree1** share and choose **Properties**.
- In the **datatree1 Properties** window, the **General** tab appears displaying the share name, folder path, and description for the **datatree1** share. Click the **Share Permissions** tab.



MMC: Managing Share Access (Cont.)



© 2010 NetApp, Inc. All rights reserved.

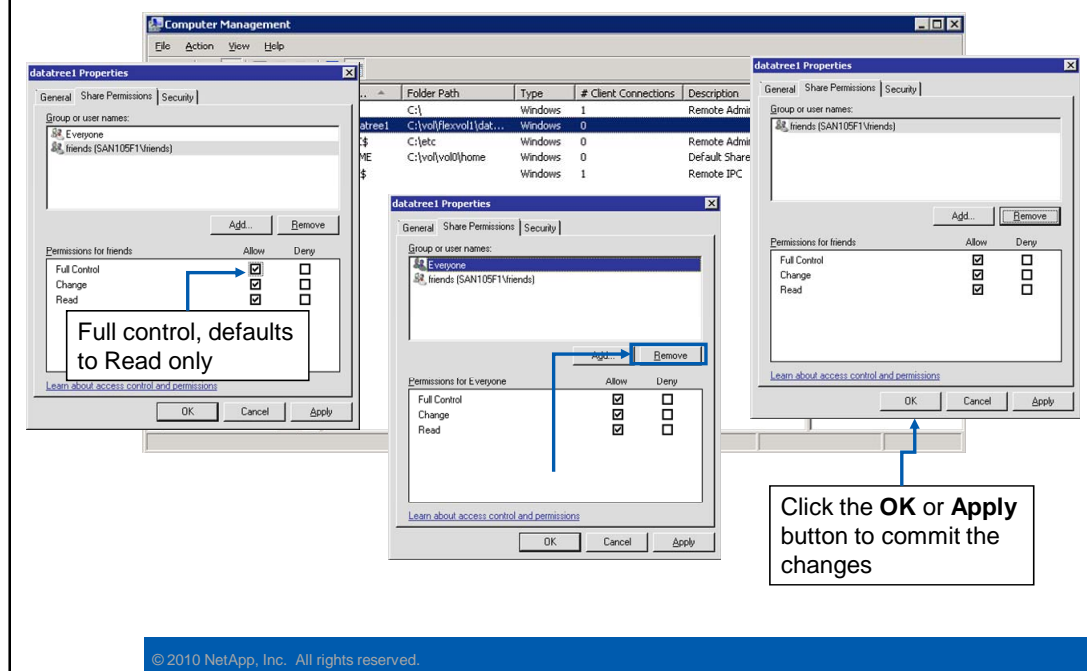
MMC: MANAGING SHARE ACCESS (CONT.)

(The following continues the setting and deleting of share access.)

- In the **Share Permissions** tab, click the **Add** button. The Select Users, Computers, or Groups window appears.
- In the **Enter the object names to select** text box, type **friends** group and click the **OK** button. The datatree1 Properties window appears, displaying the new share access for the friends group.



MMC: Managing Share Access (Cont.)



MMC: MANAGING SHARE ACCESS (CONT.)

(The following continues the setting and deleting of share access.)

- In the datatree1 Properties window, select **Everyone** and click the **Remove** button to delete share access for Everyone.
- The datatree1 Properties window displays that the Everyone share access is deleted.



File Permissions

© 2010 NetApp, Inc. All rights reserved.

FILE PERMISSIONS



Folder and File Permissions

- A storage system stores the NTFS file-level permissions for folders and files
 - Managed only from a Windows client or GPOs
- Standard Windows GUI tools display and set permissions
- Manage permissions as you would an NTFS file system on a Windows workstation or server

© 2010 NetApp, Inc. All rights reserved.

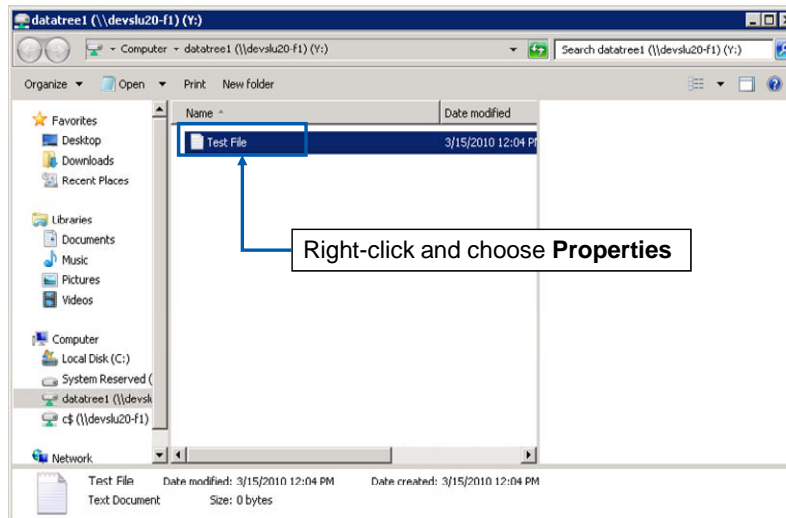
FOLDER AND FILE PERMISSIONS

A storage system stores the NTFS file-level permissions for folders and files. They can be managed from a Windows client only or Group Policy Objects (GPOs).

Standard Windows GUI tools display and set permissions. Manage permissions as you would an NTFS file system on a Windows workstation or server.



File Permissions of a Mapped Drive



© 2010 NetApp, Inc. All rights reserved.

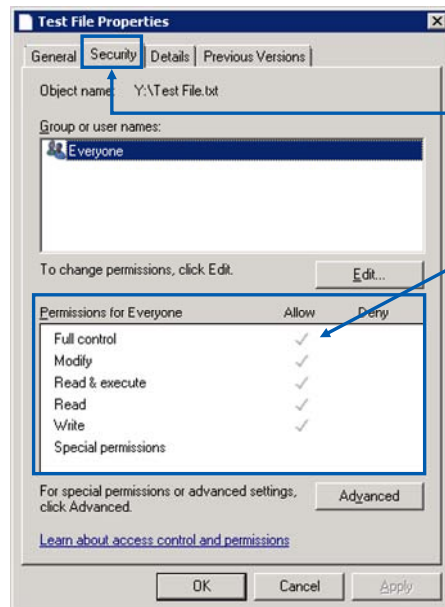
FILE PERMISSIONS OF A MAPPED DRIVE

To display the file permissions, perform the following steps:

- From a mapped network drive, right-click the file.
- Choose **Properties** from the shortcut menu.



Security Tab



Click the **Security** tab

NOTE: Grayed out permission is inherited from parent folders

The Everyone system group has full control for permissions, including Modify, Read & Execute, Read, Write, and Special Permissions

© 2010 NetApp, Inc. All rights reserved.

SECURITY TAB

- In the file Properties window, click the **Security** tab.
- **NOTE** the group and user names and the permissions for the group or user.
- Click the **OK** button.

In this example, the Everyone system group has full control for permissions including Modify, Read & Execute, Read, Write, and Special Permissions.



Access-Based Enumeration

© 2010 NetApp, Inc. All rights reserved.

ACCESS-BASED ENUMERATION



Access-Based Enumeration

- Share permissions conventionally allow users to view shared folders or files regardless of whether the users have access to them
 - Causes security risk
- Administrators can protect sensitive information using the Access-Based Enumeration (ABE) option

```
cifs shares -change share
[-accessbasedenum | -noaccessbasedenum]
```

 - May be set with `-add` switch when creating shares
 - No ABE is the default

© 2010 NetApp, Inc. All rights reserved.

ACCESS-BASED ENUMERATION

Conventional share properties allow you to specify which users (individually or in groups) have permission to view or modify shared resources. However, they do not allow you to control whether shared folders or files are visible to users who do not have permission to access them. This could pose problems, if the names of shared folders or files describe sensitive information, such as the names of customers or new products under development.

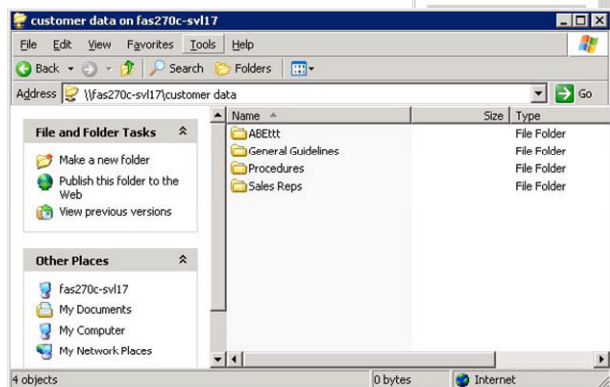
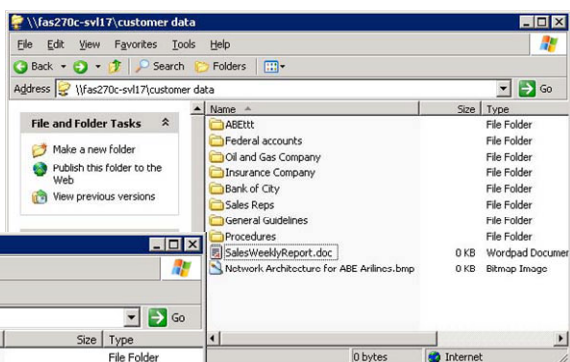
Access-Based Enumeration (ABE) extends share properties to include the enumeration of shared resources. When ABE is enabled on a CIFS share, users who do not have permission to access a shared folder or file underneath it (whether through individual or group permission restrictions) do not see that shared resource displayed in their environment. ABE therefore enables you to filter the display of shared resources based on user access rights.

ABE for a CIFS share on a NetApp storage system can be managed by the CIFS shares option `[-accessbasedenum | -noaccessbasedenum]`.



Access-Based Enumeration (Cont.)

Without ABE



With ABE

© 2010 NetApp, Inc. All rights reserved.

ACCESS-BASED ENUMERATION (CONT.)

The two figures illustrate how ABE affects the Data ONTAP directory listing. In the first figure, all the folders under the share “customer data” are visible to the user, even though the user does not have access to some of the folders containing sensitive information. In the bottom figure, after enabling ABE on this share, users can see only the folders to which they have access.



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Create and manage local users for a storage system
- Identify how to create a local group and make a local user a member of that group
- Use the command-line interface, NetApp System Manager or Microsoft tools to add, delete, and modify access permissions of shares
- Use Microsoft tools to add, delete, and modify access permissions of files and folders

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 8: CIFS Access Control
Estimated Time: 30 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- What is the purpose of a local administrator account on a storage system, and why does `cifs setup` recommend creating one?
- What does it mean when a storage system is configured for multiprotocol access?
- What command adds local users and groups to the storage system?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

CIFS Domains

Module 9
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



CIFS DOMAINS



Module Objectives

By the end of this module, you should be able to:

- Terminate the CIFS service to prepare for CIFS domain configuration
- Reconfigure the CIFS service for a Windows® domain
- Identify the resulting files
- Create domain users and add the domain users to a local storage system group
- Set up Preferred Domain Controllers (DCs)

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Reconfiguring CIFS Using `cifs setup`

© 2010 NetApp, Inc. All rights reserved.

RECONFIGURING CIFS USING CIFS SETUP



Reconfiguring CIFS

- To reconfigure CIFS on a storage system:
 1. Disconnect users and stop CIFS service:
 - `cifs terminate`
 2. Reconfigure CIFS service:
 - `cifs setup`
- CIFS server restarts with the new configuration
- Next we will investigate reconfiguring a storage system for an Active Directory domain

© 2010 NetApp, Inc. All rights reserved.

RECONFIGURING CIFS

To reconfigure CIFS on a storage system:

- Disconnect users and stop CIFS service:
 - `cifs terminate`
- Reconfigure CIFS service:
 - `cifs setup`

The storage system automatically attempts to restart the CIFS service with the new CIFS configuration.



CLI cifs setup: AD

```
(1) Active Directory domain authentication
(Active Directory domains only)
(2) Windows NT 4 domain authentication
(Windows NT or Active Directory domains)
(3) Windows Workgroup authentication using
the filer's local user accounts
(4) /etc/passwd and/or NIS/LDAP
authentication
```

```
Selection (1-4)? [1]:
```

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD

This is an example of the administrator configuring the storage system for an Active Directory (AD) domain.



CLI cifs setup: AD (Cont.)

What is the name of the Active Directory domain? []: **development.netappu.com**

In Active Directory-based domains, it is essential that the filer's time match the domain's internal time so that the Kerberos-based authentication system works correctly.

If the time difference between the filer and the domain controllers is more than 5 minutes, CIFS authentication will fail. Time services currently are not configured on this filer.

Would you like to configure time services? [y]:

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD (CONT.)

AD uses a time-based key mechanism. It is important for the domain controller and the storage system to be in sync by five (5) minutes or less.



CLI cifs setup: AD (Cont.)

CIFS Setup will configure basic time services. To continue, you must specify one or more time servers. Specify values as a comma or space separated list of server names or IPv4 addresses. In Active Directory-based domains, you can also specify the fully qualified domain name of the domain being joined (for example: ("DEVELOPMENT.NETAPPU.COM") and time services will use those domain controllers as time servers.

Enter the time server host(s) and/or address(es)
[DEVELOPMENT.NETAPPU.COM]:10.254.134.2

NOTE: The IP address is for the domain controller or a time server

Would you like to specify additional time servers? [n]:
Wed Jun 21 16:28:22 GMT [rc:ALERT]: timed: time daemon started

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD (CONT.)

The IP address is for the domain controller or a time server. It is best to enter the IP address of the main (root) domain controller for the domain.

The timed daemon allows the storage system to synchronize its time with external resources.

You need to configure the following:

- options timed.max_skew 30m
- options timed.protonntp
- options timed.sched hourly
- options timed.servers [server_ip_or_name,...]
- options timed.enable on
- options timed.log on



CLI cifs setup: AD (Cont.)

In order to create an Active Directory machine account for the filer, you must supply the name and password of a Windows account with sufficient privileges to add computers to the DEVELOPMENT.NETAPPU.COM domain.

Enter the name of the Windows user
[Administrator@DEVELOPMENT.NETAPPU.COM]:

[This Windows user is the domain administrator or any other account with privileges to add computer accounts to the domain.]

Password for Administrator@DEVELOPMENT.NETAPPU.COM:
CIFS -Logged in as Administrator@DEVELOPMENT.NETAPPU.COM.

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD (CONT.)

This Windows user is the domain account administrator that has privileges to join (add) the storage system to the domain controller.



CLI cifs setup: AD (Cont.)

The user that you specified has permission to create the filer's machine account in several (4) containers. Please choose where you would like this account to be created.

- (1) CN=computers
 - (2) OU=Domain Controllers
 - (3) OU=Additional_OU
 - (4) OU=sub_Additional_OU,OU=Additional_OU
 - (5) None of the above
- Selection (1-5)? [1]: 1

NOTE: CN means common name

The storage system is being registered in active computer as a computer under the default OU

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD (CONT.)

The container list displays Organizational Units (OUs) in which you have permission to create computer accounts. The list reflects your AD domain structure and may contain customized OUs.



CLI cifs setup: AD (Cont.)

```
Wed Jun 21 16:29:23 GMT [wafl.quota.sec.change:notice]:  
security style for /vol/vol0/ changed from unix to ntfs
```

```
CIFS - Starting SMB protocol...
```

```
It is highly recommended that you create the local  
administrator account (system\administrator) for this  
filer. This account allows access to CIFS from Windows  
when domain controllers are not accessible.
```

```
Do you want to create the system\administrator account?  
[y]:
```

```
Enter the new password for system\administrator:  
Retype the password:
```

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD (CONT.)

The local administrator account has privileges to administer CIFS on the storage system even if the domain controller is down. The local administration can set up local users on the storage system with the `useradmin user add` command.



CLI cifs setup: AD (Cont.)

Currently, the user "system\administrator" and members of the group "DEVELOPMENT\Domain Admins" have permission to administer CIFS on this filer. You may specify an additional user or group to be added to the filer's "BUILTIN\Administrators" group, thus giving them administrative privileges as well.

Would you like to specify a user or group that can administer CIFS? [n]:

Wed Jun 21 16:30:18 GMT

[nbt.nbns.registrationComplete:info]: NBT: All CIFS name registrations have completed for the local server.

Welcome to the DEVELOPMENT.NETAPPU.COM (DEVELOPMENT) Active Directory(R) domain.

CIFS local server is running.

© 2010 NetApp, Inc. All rights reserved.

CLI CIFS SETUP: AD (CONT.)



Reconfiguring CIFS Using NetApp System Manager

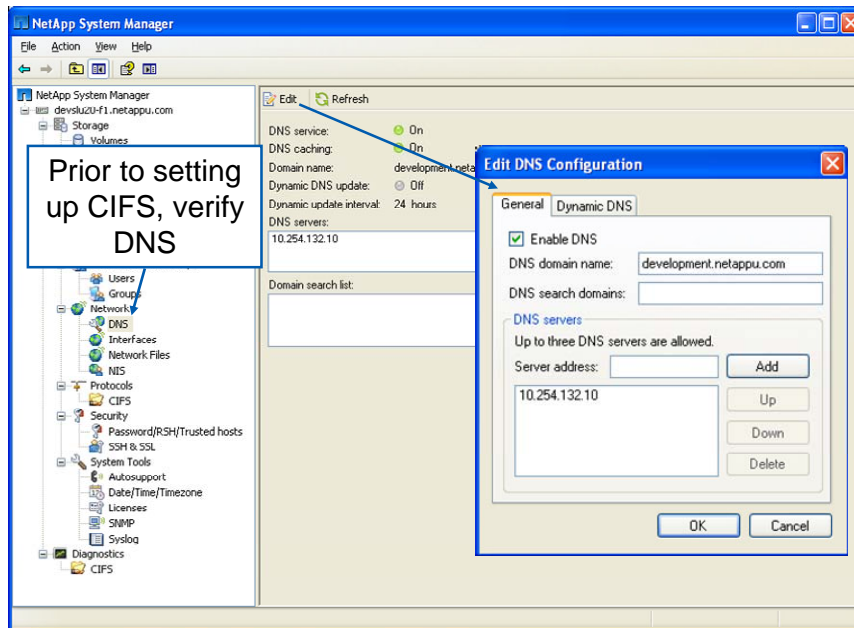


© 2010 NetApp, Inc. All rights reserved.

RECONFIGURING CIFS USING NETAPP SYSTEM MANAGER



System Manager: CIFS Setup

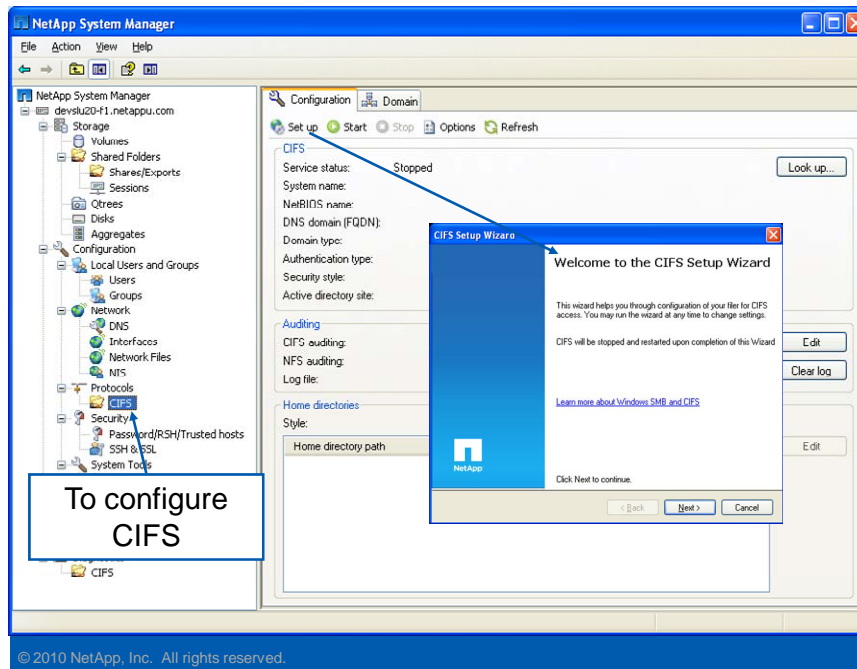


© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP



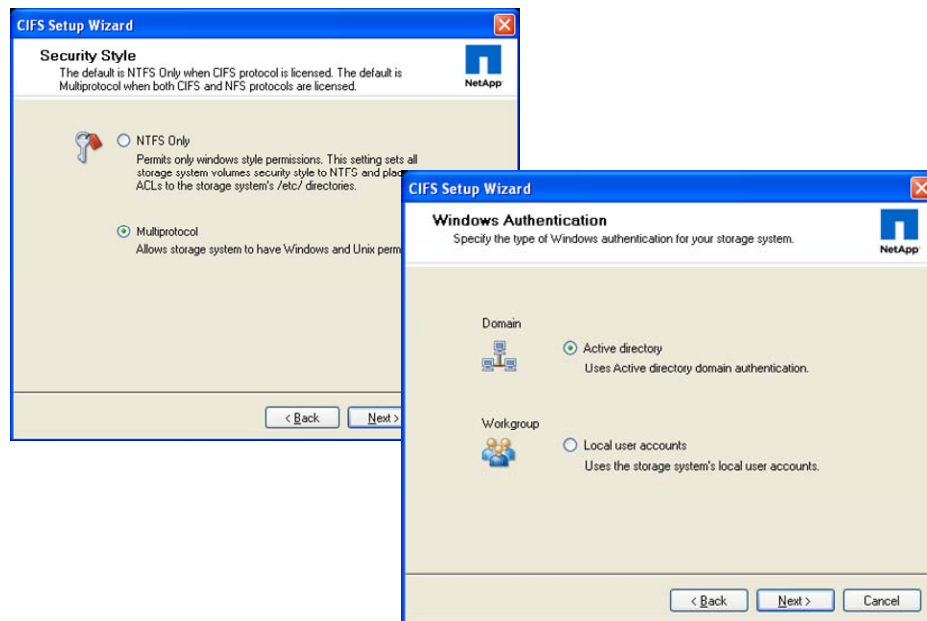
System Manager: CIFS Setup (Cont.)



SYSTEM MANAGER: CIFS SETUP (CONT.)



System Manager: CIFS Setup (Cont.)



© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP (CONT.)



System Manager: CIFS Setup (Cont.)

CIFS Setup Wizard

Active Directory Domain
Specify the Active Directory Domain name and Credentials.

Domain
Enter the active directory domain to join.
Domain name: development.netappu.com

Credentials
Enter the domain administrator credentials.
Administrator: administrator
Password: [masked]

< Back Next >

CIFS Setup Wizard

System Name, Description and WINS Servers
Specify a system name, description and WINS Servers for your storage system.

System name and description
Enter the storage system name and description as seen by Windows.
Name: DEVSLU20-F1
Description:

WINS servers
Enter up to four WINS server names.

WINS Server 1: 0 . 0 . 0 . 0
WINS Server 2: 0 . 0 . 0 . 0
WINS Server 3: 0 . 0 . 0 . 0
WINS Server 4: 0 . 0 . 0 . 0

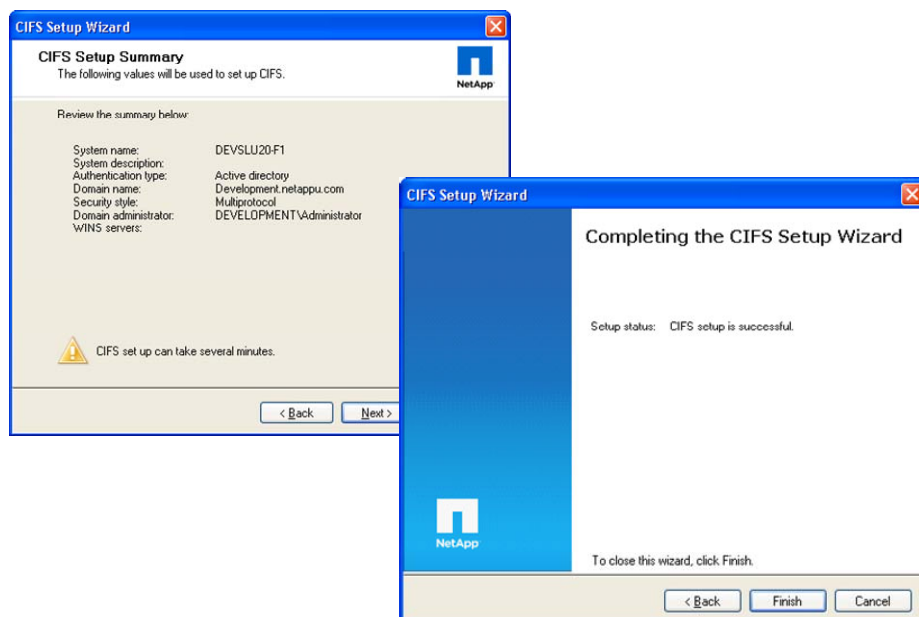
< Back Next > Cancel

© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP (CONT.)



System Manager: CIFS Setup (Cont.)

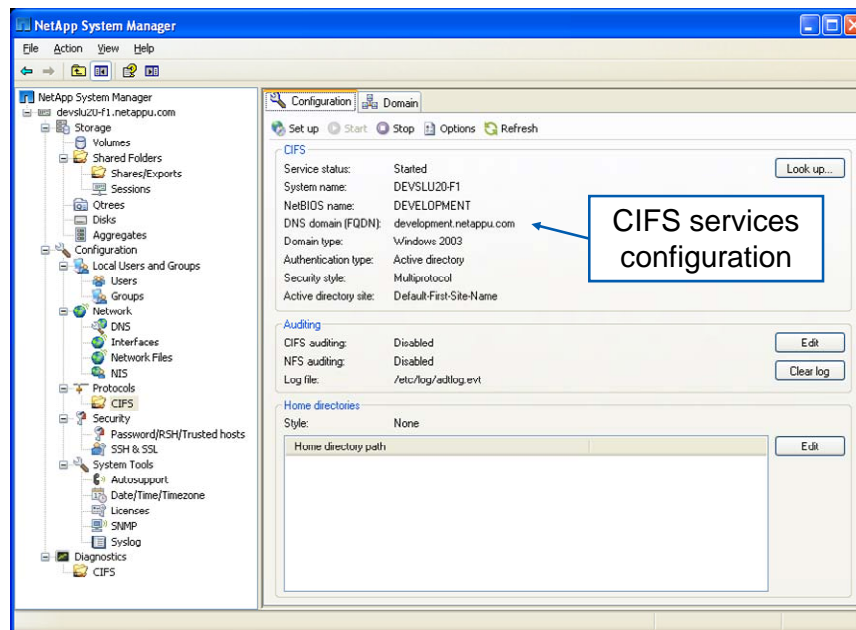


© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP (CONT.)



System Manager: CIFS Setup (Cont.)



© 2010 NetApp, Inc. All rights reserved.

SYSTEM MANAGER: CIFS SETUP (CONT.)



Results

© 2010 NetApp, Inc. All rights reserved.

RESULTS



Results

Additional files created in domain environment:

- `/etc/filersid.cfg`
 - Contains the storage system SID
- `/etc/cifssec.cfg`
 - Contains the Windows domain SID

NOTE: These files are not readable; do not edit the files

© 2010 NetApp, Inc. All rights reserved.

RESULTS

The `/etc/filersid.cfg` file is created in a domain environment and contains the storage system security ID (SID).

The `/etc/cifssec.cfg` file contains the Windows domain controller account information.

NOTE: These files are not readable; do not edit the files.



lclgroups.cfg Changes

- Domain administrators are added to lclgroups.cfg:

```
system> rdfile /etc/lclgroups.cfg
[ "Replicators" 552 ( "not supported" ) ]
[ "Backup Operators" 551 ( "Members can bypass
file security to backup files" ) ]
[ "Power Users" 547 ( "Members that can share
directories" ) ]
[ "Guests" 546 ( "Users granted Guest Access" ) ]
[ "Users" 545 ( "Ordinary Users" ) ]
[ "Administrators" 544 ( "Members can fully
administer the filer" ) ]
```

S-1-5-21-265246955-68147109-1151652928-500
S-1-5-21-3723512375-496415379-1150184651-512

Local Administrator
Domain Admins Group

- Remember to use `cifs lookup` to resolve SIDs

© 2010 NetApp, Inc. All rights reserved.

LCLGROUPS.CFG CHANGES



Domain-Specific Commands

After configuring the storage system for a domain environment, do the following:

- Display your domain information:
 - `cifs domaininfo`
- Test the storage system connection using NetBIOS over TCP/IP if used:
 - When CIFS has been successfully started and is operational:
 - `cifs testdc`
 - When the CIFS subsystem is not running:
 - `cifs testdc`
 `[WINSSvrIPAddress]domainname`
 `[storage_sys_name]`

© 2010 NetApp, Inc. All rights reserved.

DOMAIN-SPECIFIC COMMANDS



CLI: `cifs domaininfo` Command

- Example output from the `cifs domaininfo` command:

```
system> cifs domaininfo
NetBios Domain:          DEVELOPMENT
Windows 2000 Domain Name: Development.netappu.com
Type:                   Windows 2000
Filer AD Site:          none
```

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS DOMAININFO COMMAND



CLI: cifs domaininfo Command (Cont.)

- Example output from the `cifs domaininfo` command (cont.):

```
Current Connected DCs:      \\WIN2K3
Total DC addresses found: 2
Preferred Addresses:        None
Favored Addresses:          None
Other Addresses:            10.0.0.5 WIN2K3 PDC

Connected AD LDAP Server:  \\win2k3.netapp.com
Preferred Addresses:        None
Favored Addresses:          None
Other Addresses:            10.0.0.5 win2k3.netapp.com
                           10.0.0.6 win2k3-2.netapp.com
```

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS DOMAININFO COMMAND (CONT.)



CLI: `cifs testdc` Command

- The following example is output from the `cifs testdc` command on a storage system in a domain

```
system> cifs testdc
Using Established configuration
Current Mode of NBT is B Mode
Netbios scope ""
Registered names...
system      < 0> Broadcast
system      < 3> Broadcast
system      <20> Broadcast
GRUMPY      < 0> Broadcast
GRUMPY      < 3> Broadcast
GRUMPY      <20> Broadcast
HAPPY       < 0> Broadcast
HAPPY       < 3> Broadcast
HAPPY       <20> Broadcast
```

B Mode = Uses broadcast for name registration and resolution

These three names correspond to the Workstation, Server, and Messenger services, respectively

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS TESTDC COMMAND

For more information about NetBIOS over TCP/IP, see chapter 11 of TCP/IP Fundamentals for Microsoft® Windows: www.microsoft.com/downloads/details.aspx?familyid=c76296fd-61c9-4079-a0bb-582bca4a846f&displaylang=en.



CLI: `cifs testdc` Command (Cont.)

Output from the `cifs testdc` command (cont.):

```
SNEEZY          < 0> Broadcast
SNEEZY          < 3> Broadcast
SNEEZY          <20> Broadcast
DEVELOPMENT     < 0> Broadcast
```

```
Testing all Primary Domain Controllers
found 1 unique addresses
```

```
found PDC WIN2K3 at 10.0.0.5
```

```
Testing all Domain Controllers
found 1 unique addresses
```

```
found DC WIN2K3 at 10.0.0.5
```

© 2010 NetApp, Inc. All rights reserved.

CLI: CIFS TESTDC COMMAND (CONT.)



Preferred DCs

© 2010 NetApp, Inc. All rights reserved.

PREFERRED DCS



Preferred DCs

- Microsoft Active Directory members use a mechanism called “site awareness” to discover their closest domain controllers within AD
- A site is a physical, geographical, or subnet boundary of the network
- Storage system administrators accept the default and have `cifs.site_awareness.enable` turned on
- Storage system administrators can override this default mechanism by setting preferences for other domain controllers

```
system> options cifs.site_awareness.enable off
```

© 2010 NetApp, Inc. All rights reserved.

PREFERRED DCS

Site awareness, also called site discovery, is the process of automatically discovering the preferred domain controller. By default, a storage system is configured with `cifs.site_awareness.enable` turned on. A storage administrator can override this default mechanism by turning the `cifs.site_awareness.enable` option to off and setting preferred domain controllers using the `cifs prefcdc` command.



Configuring `prefdc` List

The `cifs prefdc` command configures and displays CIFS preferred domain controller information

- To display the preferred domain controller list:

```
system> cifs prefdc print [domain]
```

- To add a preferred domain controller list:

```
system> cifs prefdc add domain address [address]
```

- To delete a preferred domain controller list:

```
system> cifs prefdc delete domain
```

- Example:

```
system> cifs prefdc print
```

```
No preferred domain controllers configured.  
Domain controllers will be automatically  
discovered.
```

© 2010 NetApp, Inc. All rights reserved.

CONFIGURING PREFDC LIST

The `cifs prefdc` command configures and displays CIFS preferred domain controller information.

To display the preferred domain controller list:

- `cifs prefdc print [domain]`

To add a preferred domain controller list:

- `cifs prefdc add domain address [address...]`

To delete a preferred domain controller list:

- `cifs prefdc delete domain`


In the following example, there are no preferred domain controllers configured and domain controllers will be automatically discovered.

```
system> cifs prefdc print
```

```
No preferred Domain Controllers configured.  
DCs will be automatically discovered.
```



DC Ping Ordering



© 2010 NetApp, Inc. All rights reserved.

DC PING ORDERING

Most Windows server environments have multiple domain controllers. A NetApp® storage system contacts a domain controller in the following order:

- Preferred: Any domain controller(s) configured as preferred with the `cifs prefdc` command
- Favored: Any domain controller(s) that is determined by site awareness rules to be readily accessible
- Other: Any other domain controller(s) that is reachable

NOTE: DC ping occurs every time the CIFS service starts, every time `cifs prefdc` is executed, and every four hours.



Domain Users

© 2010 NetApp, Inc. All rights reserved.

DOMAINS USERS



Domain User

- Domain user is:
 - Created in a domain
 - Authenticated by the domain
 - Created with the Active Directory Users and Computers tool

© 2010 NetApp, Inc. All rights reserved.

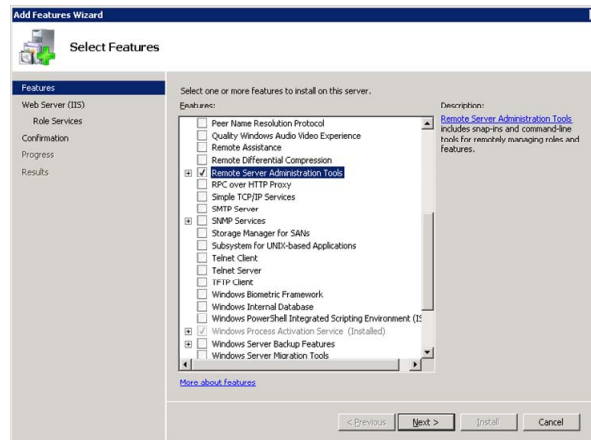
DOMAIN USER

A domain user is a non-local user who belongs to a Windows domain and is authenticated by the domain. This type of user also can be placed into storage system groups to grant them capabilities on the storage system. On the Windows workstation, you can create a domain user with the Active Directory Users and Computers tool. The Windows Active Directory Users and Computers tool allows management of users, groups, organizational units, and all other Active Directory objects. You can administer and publish information in the directory.



W2k8R2: Remote Administration Tools

- Within Windows Server 2008 R2, administrators must add the Remote Administration Tools to remotely administrate Active Directory
 - Same as the AdminPack for Windows Server 2003



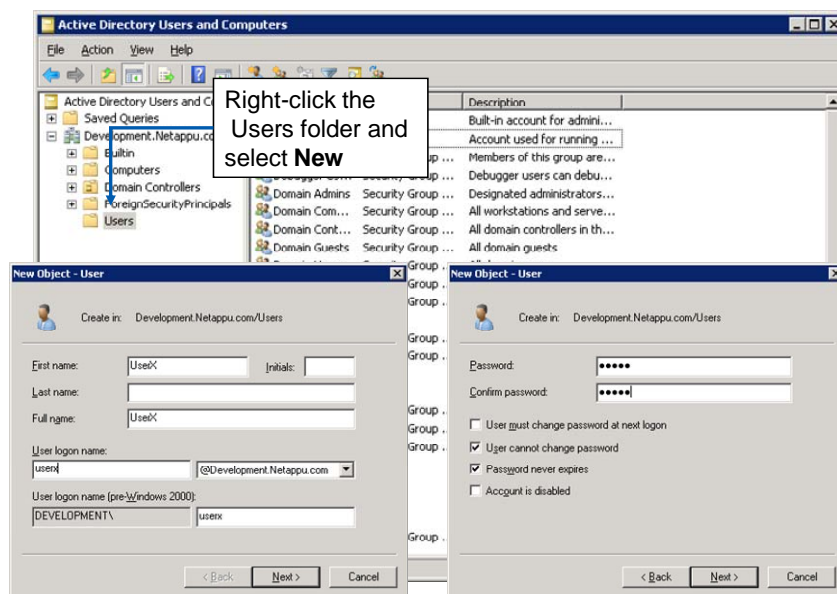
NOTE:
Reboot required

© 2010 NetApp, Inc. All rights reserved.

W2K8R2: REMOTE ADMINISTRATION TOOLS



Creating a Domain User



© 2010 NetApp, Inc. All rights reserved.

CREATING A DOMAIN USER

To create a domain user with the Active Directory Users and Computers Tool, perform the following steps:

1. To open the tool from your Windows workstation, go to Start > Control Panel> Administrative Tools > Active Directory Users and Computers.
2. To add a new domain user, right-click the **Users** folder and choose **New > User**
3. In the New Object – User window, type the name of the user in the **First name**, **Last name**, and **Full name** text boxes.

In this example, user_jdoe (for Jane Doe) is typed in the First name text box and repeated in the Full name text box.

4. In the **User logon name** text box, type the user logon of user_jdoe to add the domain user Jane Doe.
5. Click the **Next** button.
6. In the password window, type the password for Jane Doe and confirm the password.
7. Mark the **Password never expires** check box for this example.
8. Click the **Next** button.
9. Click the **Finish** button to complete adding user_jdoe to the domain.



Local User Authentication

When the storage system is using CIFS Domain authentication:

- Local user authentication is still possible
- Additional MMC functionality is available
 - Users:
 - Displays a current list of local users only
 - Cannot create, delete, or view properties of local users
 - Cannot administer passwords
 - Groups:
 - Can display, create, and delete a group, and add or delete users in the group
 - Cannot add or modify roles (and hence, capabilities) for the group

© 2010 NetApp, Inc. All rights reserved.

LOCAL USER AUTHENTICATION



Adding Domain Users to Groups

Assign a Windows domain user to a custom or predefined local group

- CLI: `useradmin domainuser`

- Syntax

```
system> useradmin domainuser add user  
          -g group | Administrators |  
          "Backup Operators" | Guests |  
          "Power Users" | Users
```

- To add an existing Windows domain user to a group:

```
system> useradmin domainuser add user -g group
```

- To list Windows domain users in a group:

```
system> useradmin domainuser list -g group
```

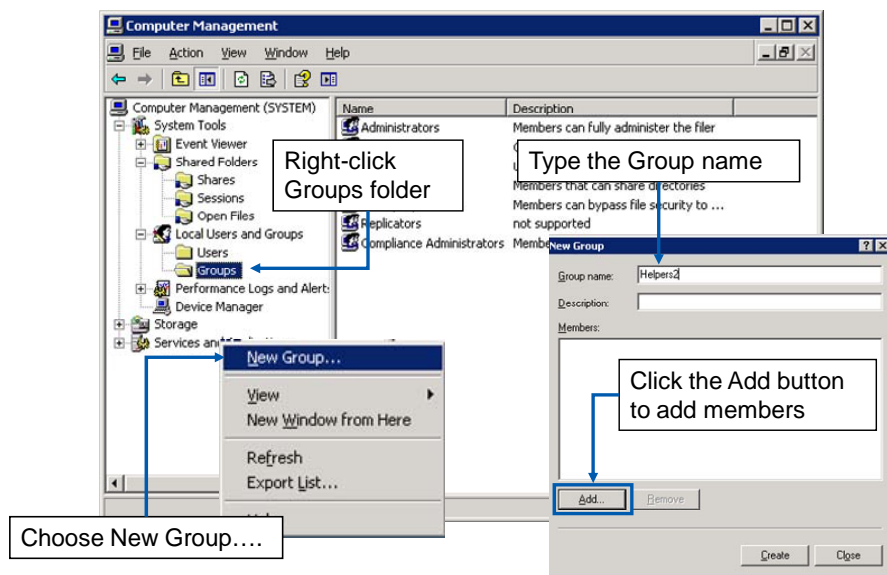
- Computer Management (MMC)

© 2010 NetApp, Inc. All rights reserved.

ADDING DOMAIN USERS TO GROUPS



MMC: Groups



© 2010 NetApp, Inc. All rights reserved.

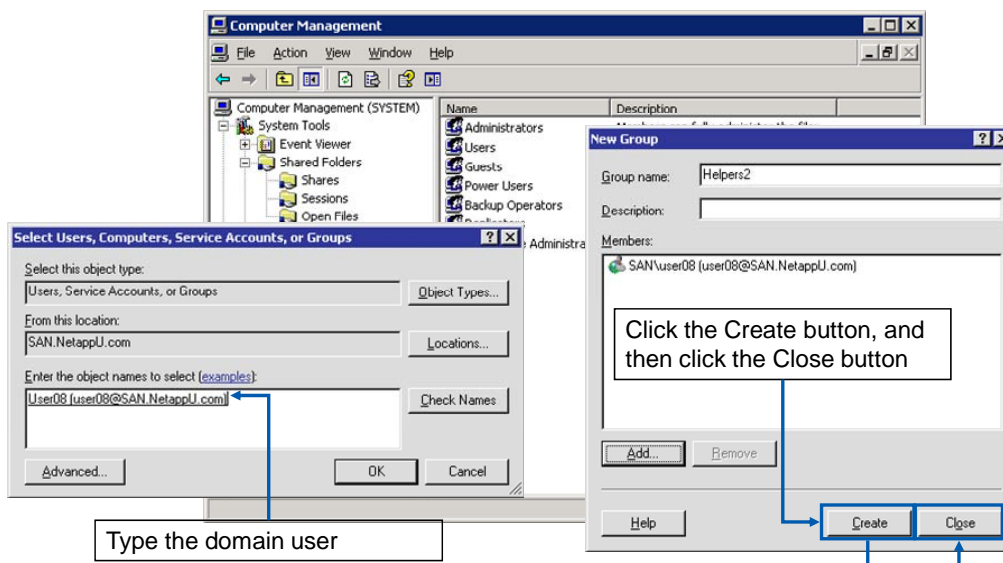
MMC: GROUPS

As an example with the Windows Computer Management GUI, in the Groups folder, add a new group Helpers2 and add local user Jane to the group by performing the following steps:

1. Go to System Tools > Local Users and Groups > Groups.
2. Right-click the Groups folder and choose New Group.
3. In the New Groups window, in the Group name text box, type the group name Helpers2.
4. Click the Add button to add members to the new group.



MMC: Groups (Cont.)



© 2010 NetApp, Inc. All rights reserved.

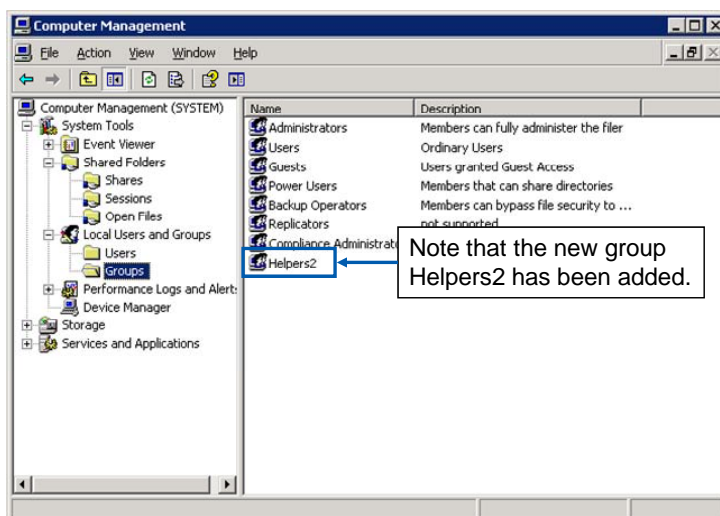
MMC: GROUPS (CONT.)

(The following continues the adding of a new local group.)

5. In the Select User, Computers, or Groups window, add the domain user.
6. Click the **OK** button. The New Group window displays the domain user as a member of a local storage system group.
7. In the New Group window, click the **Create** button, and then click the **Close** button.



MMC: Groups (Cont.)



© 2010 NetApp, Inc. All rights reserved.

MMC: GROUPS (CONT.)

(The following continues the adding of a new local group.)

8. Note that in the Computer Management GUI, the new group Helpers2 has been added.



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Terminate the CIFS service to prepare for CIFS domain configuration
- Reconfigure the CIFS service for a Windows domain
- Identify the resulting files
- Create domain users and add the domain users to a local storage system group
- Set up Preferred Domain Controllers (DCs)

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 9: CIFS Domains
Estimated Time: 60 minutes



EXERCISE

Please refer to your Exercise Guide for more instruction.



Check Your Understanding

- For which objects can you create shares?
- What are three methods used to manage CIFS shares?
- CIFS Kerberos-based authentication fails if the time difference between the storage system and the domain controller is more than how many minutes?
- Which command or commands allow you to configure the preferred domain controllers?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

NAS Multiprotocol

Module 10
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



NAS MULTIPROTOCOL



Module Objectives

By the end of this module, you should be able to:

- Determine and verify user mappings for CIFS users accessing UNIX® and MIXED volumes and qtrees
- Determine and verify user mappings for UNIX users accessing NTFS and MIXED volumes and qtrees

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



NAS Multiprotocol

© 2010 NetApp, Inc. All rights reserved.

NAS MULTIPROTOCOL



Multiprotocol

- Volumes and qtrees can have either:
 - NTFS-style ACL permissions
 - UNIX-style permissions
- Having UNIX-style permissions does not prevent Windows® (CIFS) users from accessing a volume or qtree if multiprotocol is correctly configured
- Having NTFS-style ACL permissions does not prevent UNIX (NFS) users from accessing a volume or qtree if multiprotocol is correctly configured

© 2010 NetApp, Inc. All rights reserved.

MULTIPROTOCOL

The following describes the three qtree security styles:


- NTFS
 - For CIFS clients, security is handled using Windows NTFS ACLs.
 - For NFS clients, the NFS username is mapped to a Windows username which is then associated with a Windows security identifier (SID) and its groups. These mapped credentials are used to determine file access based on the NTFS ACL.
- UNIX
 - Just like UNIX, files and directories have UNIX permissions.
 - For CIFS client, the Windows username is mapped to a UNIX username. This mapped account is then used to determine file access based on the UNIX security.
- Mixed
 - Both NTFS and UNIX security is allowed. A file or directory can have either NTFS ACLs or UNIX permissions.
 - For NTFS ACLs and NFS clients, the NFS username is mapped to a Windows username and its associated groups. These mapped credentials are used to determine file access based on the NTFS ACL.
 - For UNIX permissions and CIFS clients, the Windows username is mapped to a NFS user. These mapped credentials are used to determine file access based on the UNIX security.
 - The default file security style is the style most recently used to set permissions on that file.



Security Style Interaction

For a Windows user to access:

- An NTFS-style volume or qtree
 - Windows user is tested against NTFS-style ACLs
- A UNIX-style volume or qtree
 - Windows user must be mapped to a UNIX user (and associated UNIX group)



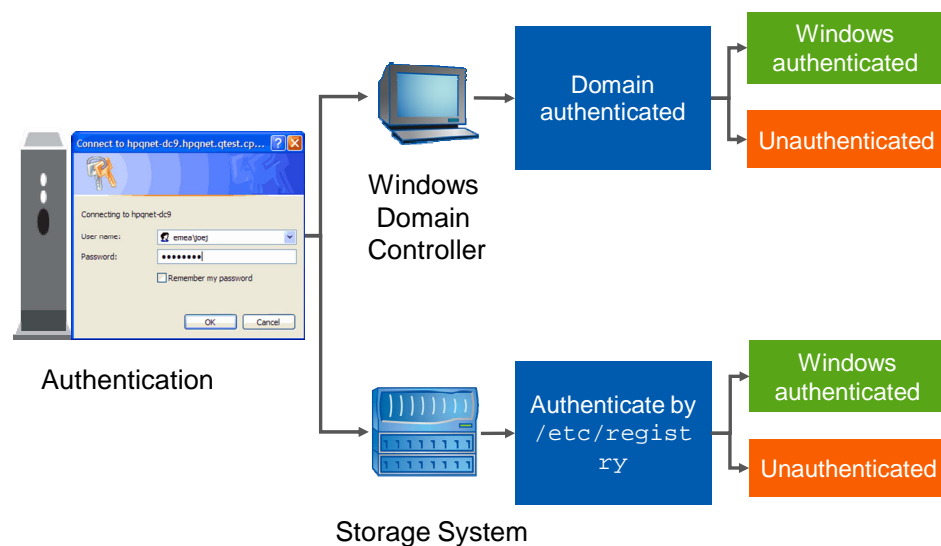
© 2010 NetApp, Inc. All rights reserved.

SECURITY STYLE INTERACTION

NOTE: There is always a user mapping (UNIX user -> NTFS user) whether the chosen security style is NTFS or multiprotocol. Even when a Windows client user is accessing data through an NTFS qtree on a storage system with NTFS security style, a user mapping occurs for the Windows client user. Both NTFS and UNIX users are always mapped.



Windows-to-UNIX User Resolution



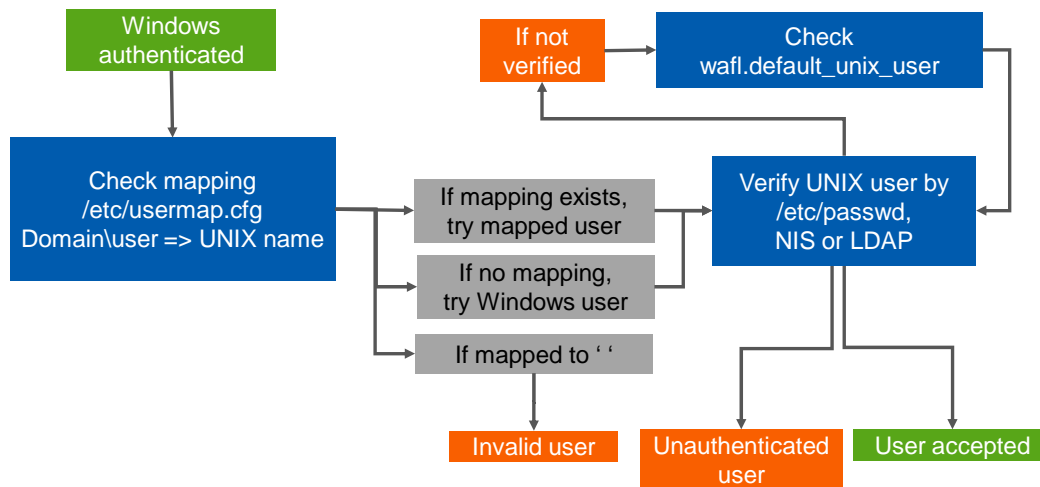
© 2010 NetApp, Inc. All rights reserved.

WINDOWS-TO-UNIX USER RESOLUTION

When a CIFS user attempts to access a volume or qtree that has UNIX permissions, the user is authenticated with the method by which the CIFS server has previously been configured. If the storage system has been configured for domain authentication, the storage system passes the credentials to the domain controller for proper authentication. The credentials are either authenticated or not. If the storage system has been configured for workgroup authentication, then the storage system will authenticate the user by way of the /etc/registry.



Windows-to-UNIX User Resolution (Cont.)



© 2010 NetApp, Inc. All rights reserved.

WINDOWS-TO-UNIX USER RESOLUTION (CONT.)

A Windows authenticated user then is looked up in the `/etc/usermap.cfg` file. Three possibilities are available. The user may be mapped to a UNIX user, not mapped at all, or mapped to an empty string. If the user is mapped, then the mapped UNIX user is passed to verification. If the user is not mapped, then the authenticated CIFS user's name is tried for UNIX verification with all letters lowercased. If the user is mapped to an empty string "", then the user is invalid.

VERIFICATION

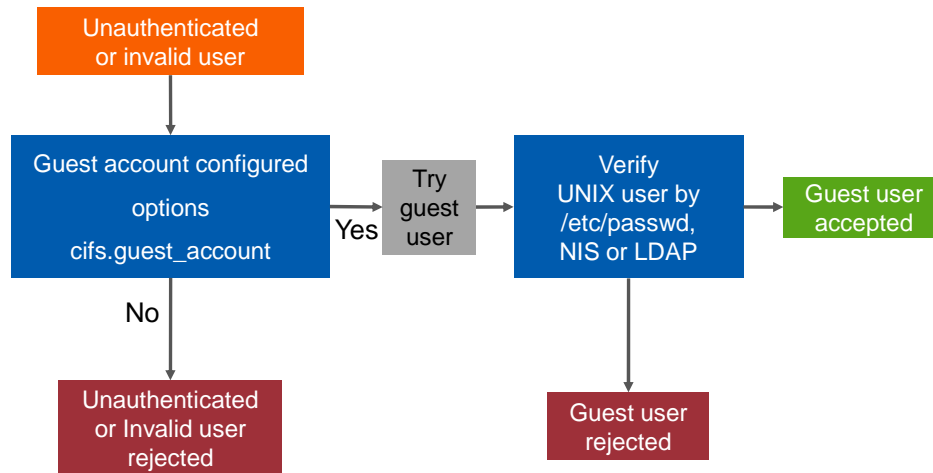
The storage system will attempt to verify a UNIX user by employing the mechanism as stated in the `/etc/nsswitch.conf` file. These mechanisms are using `/etc/passwd`, NIS, and/or LDAP. If verification is unsuccessful, then the option `wapl.default_unix_user` is tried as a generic user account. A typical default UNIX user is "pcuser" with UID=65534 and GID=65534, which is stored in `/etc/passwd` file by default. If verification is successful, the CIFS user is properly associated with a UNIX account. If verification is unsuccessful, the CIFS user is invalid.

WINDOWS ADMINISTRATOR

The Windows Administrator user is a special case. The administrator is mapped to the UNIX user name "root" with UID=0 and GID= if the `wapl.nt_admin_priv_map_to_root` option is set "on."



Windows-to-UNIX User Resolution (Cont.)



© 2010 NetApp, Inc. All rights reserved.

WINDOWS-TO-UNIX USER RESOLUTION (CONT.)


Unauthenticated or invalid users still may be allowed access to the resource if `options cifs.guest_account` is configured. **NOTE:** Windows guest account is not a default, unlike in the Windows operating system. It must be specifically set.

The guest account then is passed to the storage system for UNIX verification that is specified by the `/etc/nsswitch.conf` file.



UNIX User Access to Files

- For a UNIX user to access:
 - A UNIX-security style volume or qtree
 - The UNIX user is tested against the UNIX files permissions
 - An NTFS-security style volume or qtree:
 - The UNIX user and group must be mapped to a Windows user (and associated Windows groups)



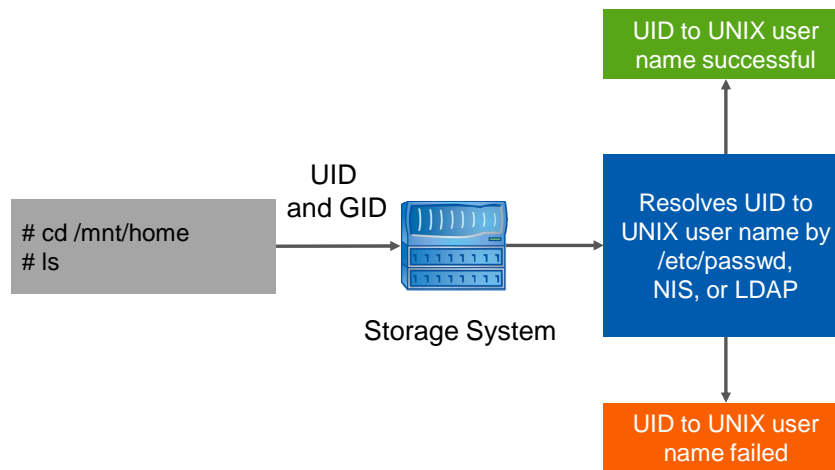
© 2010 NetApp, Inc. All rights reserved.

UNIX USER ACCESS TO FILES

This section explains the default mechanism (`/etc/usermap.cfg`) for mapping UNIX user names to Windows accounts. This mapping can also be accomplished by using LDAP, Active Directory, or NIS servers as described in www.netapp.com/library/tr/3458.pdf.



UNIX-to-Windows User Resolution



NOTE: UNIX UID (and GID) were assigned at user login when user name and password were authenticated

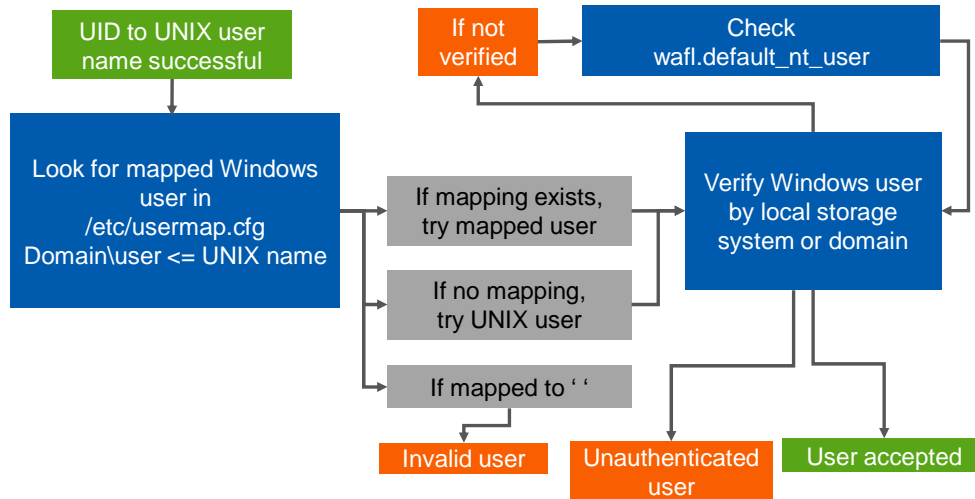
© 2010 NetApp, Inc. All rights reserved.

UNIX-TO-WINDOWS USER RESOLUTION

For the sake of this example, we are assuming that NFS v2 or v3 is being used. When an NFS user attempts to access a volume or qtree that has NTFS ACLs, the user's UID is passed from the client to the storage system, where the storage system attempts to resolve the user's name by the normal UNIX methods as defined in `/etc/nsswitch.conf`.



UNIX-to-Windows User Resolution (Cont.)



© 2010 NetApp, Inc. All rights reserved.

UNIX-TO-WINDOWS USER RESOLUTION (CONT.)


A valid user name is then looked up in the `/etc/usermap.cfg` file. Three possibilities are available. The user may be mapped to a Windows user, not mapped at all, or mapped to an empty string. If the user is mapped, then the mapped Windows user is passed to verification. If the user is not mapped, then the UNIX user's name is tried for CIFS verification. If the user is mapped to an empty string "", then the user is automatically invalid.

VERIFICATION

The storage system will attempt to verify a Windows user by using the mechanism as configured by the CIFS server. These mechanisms are either using the local accounts defined in the `/etc/registry` or passing verification to a domain controller. If verification is unsuccessful, then the option `wapl.default_nt_user` is tried as a generic user account. There is no default setting for this value, so it must be configured. If verification is successful, the NFS user is properly associated with a Windows account. If verification is unsuccessful, the NFS user is invalid.



UNIX-to-Windows User Resolution (Cont.)



© 2010 NetApp, Inc. All rights reserved.

UNIX-TO-WINDOWS USER RESOLUTION (CONT.)

Unlike in the Windows-to-UNIX resolution, there is no guest user account for NFS users. If the user is invalid, the user is rejected.



Security Styles

Security Styles			
Security Style	Hosts that can change Security/Permissions	CIFS Client Access Determined by	NFS Client Access Determined by
unix	NFS clients	UNIX permissions (Windows user names mapped to UNIX account)	UNIX permissions
mixed	NFS and CIFS clients	Depends on the last client to set security settings (permissions)	
ntfs	CIFS clients	Windows NTFS ACLs	Windows NTFS ACLs (UNIX user names mapped to Windows account)

© 2010 NetApp, Inc. All rights reserved.

SECURITY STYLES

NOTE: A CIFS user can access the file without disrupting UNIX permissions by using one of the following techniques:

- Prior to Data ONTAP® 7.2, the CIFS user must have a Windows add-on from the NOW™ site called the SecureShare®.
- With Data ONTAP 7.2 and later, the CIFS user can manage security directly with `cifs.preserve_unix_security`.

For more information, please see the *CIFS Administration on Data ONTAP* courses.



Setting Security Styles

- To set a security style for a volume:

```
system> qtree security /vol/vol0 ntfs
```

- To set a security style for a qtree:

```
system> qtree security /vol/vol0/q1 ntfs
```

- Changing a security resets all security permissions within a volume or qtree to default

- NTFS: Everyone has read-write access
- UNIX: Has user/group/world having rwx

```
drwxrwxrwx 2 root root 4096 cifs_tree1
```

© 2010 NetApp, Inc. All rights reserved.

SETTING SECURITY STYLES



Verify Mappings

- A Windows-to-UNIX user mapping is kept as part of the CIFS session credential
 - A fresh Windows-to-UNIX user mapping is required only when a new CIFS session is established for a user
 - Use `cifs session -s` command to verify mapping

© 2010 NetApp, Inc. All rights reserved.

VERIFY MAPPINGS



Multiprotocol Options

- A CIFS user can access the file without disrupting UNIX permissions
- A CIFS user might then attempt to set security restrictions on a file or folder
 - Prior to Data ONTAP 7.2, the CIFS user must have an add-on from the NOW site called the SecureShare file locking system
 - Data ONTAP 7.2 and later, the CIFS user can manage security directly with `cifs.preserve_unix_security`

© 2010 NetApp, Inc. All rights reserved.

MULTIPROTOCOL OPTIONS



Preserving UNIX Permissions

- The `cifs.preserve_unix_security` option preserves UNIX permissions as files are edited and saved by Windows applications that perform the following steps:
 1. Read the security properties of the file.
 2. Create a new temporary file.
 3. Apply those properties to the temporary file.
 4. Rename temporary file with original file name.
- Windows clients that perform a security query receive a constructed ACL that exactly represents the UNIX permissions

© 2010 NetApp, Inc. All rights reserved.

PRESERVING UNIX PERMISSIONS

Enabling the `cifs.preserve_unix_security` option preserves UNIX permissions as files are edited and saved by Windows applications that perform the following steps:

- Read the security properties of the file.
- Create a new temporary file.
- Apply those properties to the temporary file.
- Give the temporary file the original file name.

Windows clients that perform a security query receive a constructed ACL that exactly represents the UNIX permissions.



Preserving UNIX Permissions (Cont.)

- The `cifs.preserve_unix_security` option allows manipulation of UNIX permissions by using the Security tab on a Windows client
 - When enabled, UNIX qtrees appear as NTFS volumes
 - The default for this option is “off”

NOTE: You cannot change the owner and group from the Windows Security tab

© 2010 NetApp, Inc. All rights reserved.

PRESERVING UNIX PERMISSIONS (CONT.)

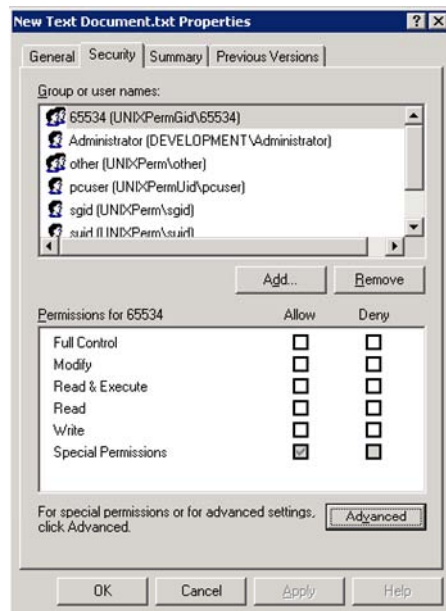
Enabling the `cifs.preserve_unix_security` option also allows you to manipulate the UNIX permissions by using the **Security** tab on a Windows client, or using any application that can query and set Windows ACLs.

When enabled, UNIX qtrees appear as NTFS volumes. The default for this option is **off**.

NOTE: You cannot change the owner and group from the Windows Security tab.



File Permissions with Mapped UNIX User



UNIX credentials are used when evaluating access requests by comparing Windows credentials against the file or folder's permissions

© 2010 NetApp, Inc. All rights reserved.

FILE PERMISSIONS WITH MAPPED UNIX USER

In this example, a Windows user is accessing a UNIX file. The **Security** tab in the file Properties window displays the user's mapped UNIX credentials.

The UNIX credentials are used when evaluating the user's access requests by comparing the user's credentials against the file or folder's UNIX access permissions.




WAFL Credential Cache

© 2010 NetApp, Inc. All rights reserved.



WAFL Credential Cache

- The WAFL® Credential Cache (WCC)
 - Caches user mappings for the UNIX UIDs and GIDs to Windows SIDs for users and groups
 - Use the `wcc` command to manage the cache
- The `wcc -u unixname` command
 - Displays the Windows user mappings for the UNIX account
- The `wcc -s ntname` command
 - Displays the UNIX user mappings for the Windows account
- Timeout value for WCC
 - `options wafl.wcc_minutes_valid 20` 
Default value

© 2010 NetApp, Inc. All rights reserved.

WAFL CREDENTIAL CACHE

The WAFL Credential Cache (WCC) contains the cached user mappings for the UNIX user identities (UIDs and GIDs) to Windows identities (SIDs for users and groups). After a UNIX-to-Windows user mapping is performed (including group membership), the results are stored in the WCC.

A Windows-to-UNIX user mapping is not stored in the WCC, but instead is kept as part of the CIFS session credential. A fresh Windows-to-UNIX user mapping is required only when a new CIFS session is established for a user.

The `wcc` command does not look in the WCC, but performs a current user mapping operation and displays the results. This command is useful for troubleshooting user mapping issues.

The `wcc -s ntname` command, where *ntname* can be a Windows user name or SID, displays the current user mappings for the Windows account.



wcc Command (root)

```
system> wcc -u root
(NT - UNIX) account name(s):(system\administrator -
root)
*****
UNIX uid = 0
user is a member of group daemon (1)
user is a member of group daemon (1)

NT membership
    system\administrator
    BUILTIN\Administrators
User is also a member of Everyone,
    Network Users, Authenticated Users
*****
```

© 2010 NetApp, Inc. All rights reserved.

WCC COMMAND (ROOT)



wcc Command (Local Administrator)

```
system> wcc -s administrator
(NT - UNIX) account name(s):(DEVSLU10-F1\administrator -
pcuser)
*****
UNIX uid = 65534

NT membership
DEVSLU10-F1\administrator
BUILTIN\Administrators
User is also a member of Everyone,
Network Users,
Authenticated Users
*****
```

© 2010 NetApp, Inc. All rights reserved.

WCC COMMAND (LOCAL ADMINISTRATOR)

The following example displays the user mapping for a local administrator.

```
system> wcc -s administrator
(NT - UNIX) account name(s):
(system\administrator - pcuser)
*****
UNIX uid = 65534
NT membership                system\administrator
                              BUILTIN\Administrators
User is also a member of Everyone, Network Users,      Authenticated Users
*****
```



wcc Command (Domain Administrator)

```
system> wcc -s development\administrator
(NT - UNIX) account name(s): (DEVELOPMENT\Administrator
- pcuser)
*****
UNIX uid = 65534
NT membership
    DEVELOPMENT\Administrator
    DEVELOPMENT\Group Policy Creator Owners
    DEVELOPMENT\Domain Users
    DEVELOPMENT\Domain Admins
    DEVELOPMENT\Enterprise Admins
    DEVELOPMENT\Schema Admins
    DEVELOPMENT\Debugger Users
    BUILTIN\Users
    BUILTIN\Administrators
User is also a member of Everyone, Network
Users,
Authenticated Users
*****
```

© 2010 NetApp, Inc. All rights reserved.

WCC COMMAND (DOMAIN ADMINISTRATOR)

The following example displays the user mapping for a domain administrator.

```
system> wcc -s
Development\administrator(NT - UNIX)
account name(s): (DEVELOPMENT\Administrator - pcuser)
*****
UNIX uid = 65534
NT membership
Owners
    DEVELOPMENT\Administrator
    DEVELOPMENT\Group Policy Creator
    DEVELOPMENT\Domain Users
    DEVELOPMENT\Domain Admins
    DEVELOPMENT\Enterprise Admins
    DEVELOPMENT\Schema Admins
    DEVELOPMENT\Debugger Users
    BUILTIN\Users
    BUILTIN\Administrators
User is also a member of Everyone, Network Users, Authenticated Users
*****
```



Module Summary

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Module Summary

In this module, you should have learned to:

- Determine and verify user mappings for CIFS users accessing UNIX and MIXED volumes and qtrees
- Determine and verify user mappings for UNIX users accessing NTFS and MIXED volumes and qtrees

© 2010 NetApp, Inc. All rights reserved.

MODULE SUMMARY



Go further, faster®

Exercise

Module 10: NAS Multiprotocol
Estimated Time: 30 minutes



PLEASE REFER TO YOUR EXERCISE GUIDE FOR MORE INSTRUCTION.



Check Your Understanding

- A UNIX user cannot ever access a file with a Windows ACL but a Windows user can access a file with UNIX permissions. True or false?
- Windows users are only associated and resolved to a UNIX user if the Windows user is attempting to access a file with UNIX permissions. True or false?
- Which file is used to associate Windows users and UNIX users?
- Which command allows administrators to display a cached UNIX user's mapping to a given Windows user?

© 2010 NetApp, Inc. All rights reserved.

CHECK YOUR UNDERSTANDING



Go further, faster®

NAS Troubleshooting

Module 11
Accelerated NCDA Boot Camp
Data ONTAP 8.0 7-Mode



NAS TROUBLESHOOTING



Module Objectives

By the end of this module, you should be able to:

- Locate options and configuration files that can be misconfigured on the storage system
- Test for Domain Name System (DNS) resolution on both the storage system and the client
- Use client-side tools to test the client configuration
- Use storage system and client tools to isolate network system blockages
- Recognize typical error messages and list commands to identify the source of the error messages

© 2010 NetApp, Inc. All rights reserved.

MODULE OBJECTIVES



Troubleshooting Overview

■ Initial Configuration



■ Problems arise on:

- Storage systems
- Clients
- The network
- Infrastructure support
 - DNS servers
 - NIS servers
 - LDAP servers

© 2010 NetApp, Inc. All rights reserved.

TROUBLESHOOTING OVERVIEW



NFS Troubleshooting: Storage System

© 2010 NetApp, Inc. All rights reserved.

NFS TROUBLESHOOTING: STORAGE SYSTEM



NFS Storage System Configuration

- Verify NFS service
 - NFS licensed
 - NFS properly configured
 - Interfaces properly configured
- Verify exports
 - `exportfs -v`
 - */etc/exports*

© 2010 NetApp, Inc. All rights reserved.

NFS STORAGE SYSTEM CONFIGURATION



NFS Troubleshooting: Client

© 2010 NetApp, Inc. All rights reserved.

NFS TROUBLESHOOTING: CLIENT



Client Troubleshooting Tools

- `ping, dig, host, getent`
- `ypwhich, ypcat, domainname`
- `showmount -e|-a`
- `/etc/init.d/autofs start|stop`
- `nfsstat -m`
- Check:
 - `/etc/nsswitch.conf`
 - `/etc/vfstab` or `/etc/fstab`
 - `/etc/resolv.conf`
 - `/etc/mtab`

© 2010 NetApp, Inc. All rights reserved.

CLIENT TROUBLESHOOTING TOOLS

`dig` (domain information groper) is used to gather information from the Domain Name System (DNS) servers.

`host` is a simple utility for performing DNS lookups. It is normally used to convert names to IP addresses and vice versa.

`getent` gets a list of entries from the administrative databases, for example:

- `# getentpasswd` or
- `# getent hosts v210-inst`

`ypwhich` returns the name of the NIS server that supplies the NIS name services to the NIS client.

`ypcatmapname` prints the values of all keys from the NIS database specified by map name.

`domainname` shows or sets the system NIS domain name.

`showmount` queries the mount daemon on a remote host for information about the state of the NFS server on that machine.

`Autofs` controls the operation of the automount daemons.

`Nfsstat` displays statistical information about the NFS and remote procedure call interfaces to the kernel. Verify:

- `nsswitch.conf` is using the name services that you think it is
- `vfstab` or `fstab`, as appropriate, is controlling the mounts as intended
- `resolv.conf` points to valid name servers

`mtab` shows the mount parameters you expect.



UNIX Networking Configuration Files

Main network settings and their file locations

Settings	HP-UX	Solaris	Suse Linux	Red Hat Linux
IP Mapping	/etc/hosts	/etc/hosts or /etc/inet/ipnodes	/etc/sysconfig/ network/config	/etc/hosts
DNS Domain and Name Servers	/etc/resolv.conf	/etc/resolv.conf	/etc/resolv.conf	/etc/resolv.conf
Interface's Network Address	/etc/rc.config.d/ netconf	/etc/hostname	/etc/hosts	network-scripts/ifcfg- interfacename
Hostname	/etc/rc.config.d/ netconf	/etc/nodename	/etc/HOSTNAME	/etc/sysconfig/ network
Default Route	/etc/rc.config.d/ netconf	/etc/defaultrouter	/etc/sysconfig/ network/routes	network-scripts/ifcfg- interfacename

© 2010 NetApp, Inc. All rights reserved.

UNIX NETWORKING CONFIGURATION FILES

To automatically configure networking at boot time in UNIX®, you need to:

- Set up the networking hardware
- Configure name resolution
- Activate the machine's NIC
- Specify any routing settings
- IP address and netmask

At boot time, the system executes a number of startup scripts files. You usually specify network settings in these files to ensure a standard automatic configuration of the system.

Different versions of UNIX use different files to initialize the settings for networking properties such as interfaces, domain names, and IP address mappings.

For a cross-referenced list of tasks mapped to the commands for various Linux or UNIX systems see:
<http://bhami.com/rosetta.html>.



Mount Process

- Client resolves name to IP
 - Remote procedure calls
 - Ports must be open on clients and storage systems
 - portmap TCP 111
 - nfsd TCP 2049
 - mountd TCP 4046
- } Verify with `rpcinfo`
- Storage system looks at exports in memory
 - Who can access this path?
 - Performs reverse name lookup
 - Grants or denies access
 - From the client, `showmount -e`

© 2010 NetApp, Inc. All rights reserved.

MOUNT PROCESS

The mount command verifies that the mountpoint is a full pathname and then passes arguments and options to `/usr/lib/fs/nfs/mount`, which takes control of the process as follows:

- `mount` opens `/etc/mnttab` and verifies that the file system was not already mounted.
- `mount` parses the argument storage system:/vol/volname/path into host storage system and remote directory /vol/volname/path.
- `mount` calls the storage system `rpcbind` to get the port number of the storage system's mountd.
- `mount` calls the storage system mountd daemon and passes it to /vol/volname/path, requesting it to send a file handle for that path.
- The storage system's mountd daemon handles the client's mount requests. If the directory /vol/volname/path is available to the client, the mountd daemon does a `NFS_GETFH` system call on /vol/volname/path to get the file handle, and then it sends it to the client's mount process. /usr/lib/fs/nfs/mount does a regular mount system call with the file handle and the mountpoint directory.
- The client kernel looks up the given mountpoint directory. If OK, it binds the file handle to the hierarchy in a mount record.
- The client kernel looks up the directory /vol/volname/path on the storage system.
- The client kernel does a `statvfs` call to the storage system NFS server `nfsd`.
- The mount system call.
- Mount opens `/etc/mnttab` and adds an appropriate entry to the end with the mounted file system and mountpoint directory information.



Mounting Options

- UNIX clients can mount either:
 - Soft
 - Clients try to mount export a few times and then return an error
 - Hard
 - Clients will continue to send out mount requests indefinitely until the server responds
- An example of a hard mount with reasonable defaults:

```
mount -o  
rw,bg,vers=3,tcp,timeo=600,  
rsize=32768,wsiz=32768,hard,intr...
```

© 2010 NetApp, Inc. All rights reserved.

MOUNTING OPTIONS

The following options are used for NFS mounts:

- Hard or soft: specifies whether the program using a file by way of an NFS connection should stop and wait for the server to come back online if the host serving the exported file system is unavailable (hard), or if it should report an error (soft).

If hard is specified, the user cannot terminate the process waiting for the NFS communication to resume unless the `intr` option is also specified. If they have mounted file systems with the hard option, they continue to send out mount requests indefinitely until the server responds. If soft is specified, the user can set an additional `timeo=<value>` option, where `<value>` specifies the number of seconds to pass before the error is reported.

Mount Option Examples

On older Linux® systems, if you do not specify any mount options, the Linux mount command (or the Automounter) automatically chooses these defaults:

```
mount -o rw,fg,vers=2,udp,rsize=4096,wsiz=4096,hard,intr, timeo=7,retrans=5
```

These default settings are designed to make NFS work right out of the box in most environments. Almost every NFS server supports NFS v2 over UDP. Rsize and wsize are relatively small because some network environments fragment large UDP packets, which can hurt performance if there is a chance that fragments can be lost. The remote procedure call retransmit timeout is set to 0.7 seconds by default to accommodate slow servers and networks. The example on the slide is reasonable mount options. Bg option causes the mount attempts to be run in the background. In fact, on many newer Linux distributions, these are the default mount options.



Problem: Stale NFS File Handle

Error code 70: Stale file handle

- What would you do?
- Resolution tips:
 - Check connectivity to the storage system (server)
 - Check mountpoint
 - Check client `vfstab` or `fstab` as relevant
 - Check `showmount -e filerx` from client
 - Check `exportfs` from command line of the storage system
 - Check storage system `/etc/exports` file

© 2010 NetApp, Inc. All rights reserved.

PROBLEM: STALE NFS FILE HANDLE

Sample Error Messages - NFS Error 70

The probable cause of this problem is that a file or directory that was opened by an NFS client was either removed or replaced on the NFS file server. To determine possible cause(s) of stale file handles after an NFS server reboot, check for the following:

- Remove “qtree security” lines in the `/etc/rc` file, if they exist. Qtree security entries are not required, and if you manually change a qtree security to NFS or MIXED, and a reboot causes the security to become NT file system (NTFS) again, stale file handles may occur.
- Check for IP address changes. Did existing mounts work through all interfaces? Did machines with the failed existing mounts have working mounts to other mountpoints?
- Check if the `exportfs` list changed. It is possible to make command-line additions through `exportfs` for a qtree to have its own mountpoint, but not add the entry to the `/etc/exports` file making the change persistent.

Possible Solution

For this scenario, there is currently a workaround:

If you experience a stale file handle while editing a file, write the file to a local file system instead. Try remounting the file system. If problems persist, consult your NFS client support to determine if you should shut down the NFS client processes that access stale file handles or, as a last resort, reboot the NFS client.

If the stale file handle number = 20, an opened file or directory has been destroyed or re-created. You can resolve the problem by unmounting and remounting the file system.



Problem: No Space Left on Disk

No space left on disk error

- What would you do?
- Resolution tips:
 - Check `df` for available disk space
 - Check for Snapshot™ copy overruns
 - Check quota report for exceeded quotas

© 2010 NetApp, Inc. All rights reserved.

PROBLEM: NO SPACE LEFT ON DISK



CIFS Troubleshooting: Storage System

© 2010 NetApp, Inc. All rights reserved.

CIFS TROUBLESHOOTING: STORAGE SYSTEM



CIFS Troubleshooting Checklist

- Verify the following:
 - CIFS service is licensed on the storage system
 - What CIFS configuration are you working with?
 - Windows® workgroup
 - Non-Windows workgroup
 - Windows domain

© 2010 NetApp, Inc. All rights reserved.

CIFS TROUBLESHOOTING CHECKLIST



Problem: DC Connectivity

- Problem: Communication from storage system to domain controller fails or trust across multiple domains fails
 - Perform the following steps:
 - a) `system> cifs domaininfo`
 - This provides information about domain and known domain controllers
 - If you receive an error and want more verbose output, then go to Step b)

© 2010 NetApp, Inc. All rights reserved.

PROBLEM: DC CONNECTIVITY

- **Potential Issue:** “Communication from storage system to domain controller fails or trust across multiple domains fails.”
- Perform the following steps:
- a) `system> cifs domaininfo`

This provides information about domain and known domain controllers.

If you receive an error and want a more verbose output, then go to Step b).



Problem: DC Connectivity (Cont.)

- b) Set the following option **on**:
`system> options cifs.trace_dc_connection on`
 - When this option is on, the storage system logs all DC address discovery and connection activities
- c) `system> cifs resetdc`
 - This command tells the storage system to disconnect from the domain controller and then establish a new CIFS connection with the DC (The steps are being logged with the `cifs_trace_dc_connection` option)
- d) Check the trace output on the console or logged output in `/etc/messages` file to find the problem

© 2010 NetApp, Inc. All rights reserved.

PROBLEM: DC CONNECTIVITY (CONT.)

- b) Set the following option **on**:
`system> options cifs.trace_dc_connection on`

When this option is on, the storage system logs all DC address discovery and connection activities.

- c) `system> cifs resetdc`

This command tells the storage system to disconnect from the domain controller and then establish a new CIFS connection with the DC. (The steps are being logged with the `cifs_trace_dc_connection` option.)

- d) Check the trace output on the console or logged output in `/etc/messages` file to find the problem.

The following is sample output from running the **cifs resetdc** command with the `cifs.trace_dc_connection` option set on.

```
system> options cifs.trace_dc_connection on
```

```
system> cifs resetdc
```



CIFS Troubleshooting: Client

© 2010 NetApp, Inc. All rights reserved.

CIFS TROUBLESHOOTING: CLIENT



Client Troubleshooting

- Most Windows client troubleshooting involves either:
 - Network communication issues (discussed later)
 - Infrastructure issues (discussed later)
- Other issues might also arise within Windows regarding which tracing and debuggers can be used:
 - Windows user and kernel debuggers (windbg)
 - Time Traveling Tracing
 - ETW Tracing
 - Windows Resource Kit and Sysinternals

© 2010 NetApp, Inc. All rights reserved.

CLIENT TROUBLESHOOTING

Windows user/kernel debuggers (windbg) is the most common debugger in use for customer issues.

Time traveling tracing can identify hard to find issues. Time traveling tracing traces a program's flow and then is analyzed internal at Microsoft®. This tool is currently available only through Microsoft's support.

Event Tracing for Windows (ETW) provides a mechanism to monitor, log, and troubleshoot SMB.

Sysinternals and Windows Resource Kit are available at <http://technet.microsoft.com>.

There is an excellent presentation by Hongwei Sun, a Microsoft Escalation Engineer, which was given at the 2009 File Sharing Windows Protocols Plug-fest. The presentation can be found at:
<http://channel9.msdn.com/posts/Darryl/Troubleshooting-Windows-SMBSMB2-Issues/>.



Problem: Client Communication

- Potential Issue: “Network failed or is slow.”
 - Check the following:

```
system> ifstat
system> netdiag
system> ping
C:\> tracert
```
- Potential Issue: Firewall prevents communications between storage system and DC
 - If using SMB over TCP/IP
 - Windows 2000 Server and later
 - Requires TCP port 445

© 2010 NetApp, Inc. All rights reserved.

PROBLEM: CLIENT COMMUNICATION

In a domain environment, a Windows client user requests user session authentication with the storage system.

Potential Issue: “Network failed or is slow.”

Check the following:

- `system > ifstat`
The `ifstat` command displays statistics about packets received and sent on all or a specified network interface.
- `system > netdiag`
The `netdiag` command analyzes the statistics continuously gathered by the network protocol code, performs various tests (if required), displays the results of analysis, and suggests remedial actions if problems are encountered.
- `system > ping`
The `ping` command sends ICMP ECHO_REQUEST packets to a network hosts to elicit an ICMP ECHO_RESPONSE from the specified host or gateway.
- `C:\> tracert`
The Windows `tracert` command visually displays a network packet being sent and received and the number of hops required for the packet to reach its destination.



Problem: Network Connectivity

- Problem: Windows client cannot 'find' the storage system
 - If using DNS, try pinging the storage system by name
- ```
C:\> ping system_name
```
- Have routes been configured correctly?

© 2010 NetApp, Inc. All rights reserved.

## PROBLEM: NETWORK CONNECTIVITY



## NAS Troubleshooting: Network

© 2010 NetApp, Inc. All rights reserved.

### NAS TROUBLESHOOTING: NETWORK



## Network

- The “cleaner” the better
  - Matched parameters all the way through
  - Not saturated (no Quality of Service in Ethernet)
  - Auto versus half or full duplex
- Use TCP instead of UDP
  - TCP is acknowledged
  - UDP is not acknowledged
- Are there firewalls (network or client) in the way?
  - Remember remote procedure call ports should not be blocked

© 2010 NetApp, Inc. All rights reserved.

## NETWORK

Data ONTAP® 7.2.1 and later introduced a new multi-threaded mount process. Clients that are still mounting the file systems from the storage system using UDP cannot benefit from the new multi-threaded mount processing. UDP requests still use single-threaded operations. Clients mounting with TCP benefit greatly from this enhancement.



# Packet Traces

## ■ Overview

- Data ONTAP utility for packet capture
- Captures data for further analysis by support personnel

## ■ Syntax

- `pktt start <if>|all [-d dir] [-m pklen] [-b bsize] [-i ipaddr -i ...]`
  - Starts packet tracing
- `pktt dump [<if>|all [-d dir]] | [<if> [-f file]]`
  - Writes data from memory to file (disk)
- `pktt stop <if>|all`
  - Stops packet tracing

## ■ Optional commands

- `pktt pause <if>|all`
- `pktt status [<if>|all] [-v]`
- `pktt delete [filename.trc]+`
- `pktt list`

© 2010 NetApp, Inc. All rights reserved.

## PACKET TRACES

### PKTT DUMP

The dump subcommand causes the contents of the packet trace buffer to be written to a file. If the “-d [dir]” option is used, the file will be written to that directory, otherwise it will be written to the root directory of the root volume. The name of the file is always .trc and the contents are in “tcpdump” format. If a file by that name already exists it will be overwritten.

### PKTT STOP

This causes all tracing to stop on the named interface or all interfaces. If any unwritten data is in the trace buffer it will be flushed to disk. If you have not dumped the trace data, and you were not tracing to a disk file, the trace data will be lost. This action is not confirmed, so be careful when using this command.

### PKTT STATUS

This can be used to display the buffer and file status of an existing trace. Using `pktt status -v` will give you full tracing status for all interfaces. This can be used to display the buffer and file status of an existing trace. Using `pktt status -v` will give you full tracing status for all interfaces.





## Reading Packet Traces

- Pktt trace saved in *tcpdump* format
  - Reference [www.tcpdump.org](http://www.tcpdump.org)
- Use a *tcpdump*-compliant program to review the packet trace, such as Wireshark ([www.wireshark.org/](http://www.wireshark.org/))
- Alternatively, convert pktt trace to *Netmon*-compliant format using
  - “*capconv*” utility—reference <http://now.netapp.com/NOW/download/tools/capconv/>
  - Netmon-compliant packet analyzers such as Windows Netmon

© 2010 NetApp, Inc. All rights reserved.

## READING PACKET TRACES



## **NAS Troubleshooting: Infrastructure Support**

© 2010 NetApp, Inc. All rights reserved.

### **NAS TROUBLESHOOTING: INFRASTRUCTURE SUPPORT**



## Problem: Hostname-to-IP Resolution

- Cannot resolve hostnames to IP addresses
  - Look at:
    - *nsswitch.conf*
    - *hosts file*
    - *resolv.conf*
    - On the storage system, DNS or NIS options
    - Changing the order of DNS or NIS servers
    - Consider circumventing DNS or NIS by temporarily entering hosts into the hosts file
  - Remember:
    - Data ONTAP caches NIS maps in slave mode
    - Data ONTAP caches DNS

© 2010 NetApp, Inc. All rights reserved.

## PROBLEM: HOSTNAME-TO-IP RESOLUTION

Name resolution is critical to a working NFS system. Make sure both the storage system and host can resolve names – and that they get the same results.

The *nsswitch.conf* file is the place to start when troubleshooting name-resolution issues. Make sure that you are using the name services you intend to be using. If that file is correct, move to the services listed; files = */etc/hosts*, DNS = */etc/resolv.conf*, NIS = *domainname*, and *ypwhich* for starters.

Remember, there are several options in Data ONTAP used to configure and manage DNS:

Option:

`dns.cache.enable` is used to enable/disable DNS name resolution caching.

`dns.domainname` is the storage system DNS domain name.

`dns.enable` is used to - enable or disable DNS name resolution.

`dns.update.enable` is used to dynamically update the storage system 'A' record (CIFS).

`dns.update.ttl` is the time-to-live for a dynamically inserted 'A' record.

One troubleshooting method when managing name-resolution problems is to enter hostnames or addresses in the */etc/hosts* file of the storage system or host, thereby eliminating external name resolution services. This is not a fix, but a workaround to assist in fault isolation.

Remember that NIS maps in slave mode are cache as well as DNS. You can flush the DNS cache at any time by entering the `dns flush` command.



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Locate options and configuration files that can be misconfigured on the storage system
- Test for Domain Name System (DNS) resolution on both the storage system and the client
- Use client-side tools to test the client configuration
- Use storage system and client tools to isolate network system blockages
- Recognize typical error messages and list commands to identify the source of the error messages

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 11: NAS Troubleshooting  
Estimated Time: 0 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- When troubleshooting a NAS protocol on the storage system, what is one of the first steps you should do to verify the appropriate service is available?
- How do you configure the order of the hostname-to-IP resolution mechanism on the storage system?
- Which command is used to capture network packet traces on the storage system?
- What third-party application can be used to read the native packet traces created on the storage system?

© 2010 NetApp, Inc. All rights reserved.

### CHECK YOUR UNDERSTANDING



Go further, faster®

# SAN Overview

Module 12  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## SAN OVERVIEW





## Module Objectives

By the end of this module, you should be able to:

- Describe the differences between network-attached storage (NAS) and storage area network (SAN)
- List the methods to implement a SAN environment
- Define logical unit number, initiator, and target
- Describe ports, worldwide node names, and worldwide port names
- List the basic steps to implement a SAN

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES




## SAN Introduction

© 2010 NetApp, Inc. All rights reserved.

### SAN INTRODUCTION



## SAN Versus NAS



© 2010 NetApp, Inc. All rights reserved.

### SAN VERSUS NAS

Operating systems and applications request data either at the block level or the file level. Network-attached storage (NAS) provides file-level access to data on a storage system. Access is by way of a network, using Data ONTAP® services such as CIFS and NFS. Storage area networks (SANs) provide block-level access to data on a storage system. SAN solutions can be any mixture of iSCSI or Fibre Channel (FC) protocols. When both SAN and NAS storage are present on the same storage system, it is referred to as “unified storage.”



## SCSI

- SAN uses Small Computer System Interface (SCSI) protocol over a distance
- SCSI features:
  - Block-level access
  - Efficiency
  - Lower overhead
  - Resiliency


© 2010 NetApp, Inc. All rights reserved.

## SCSI

Small Computer System Interface (SCSI) is a set of standards that define commands, protocols, and interfaces used to transmit data. SCSI allows low-level “block” access to data in units of 512-byte blocks. This is highly efficient and has low overhead compared to NAS or “file” level access. SCSI has a high level of resiliency that makes it perfect for an enterprise-level protocol.



## SCSI on Host and Controller



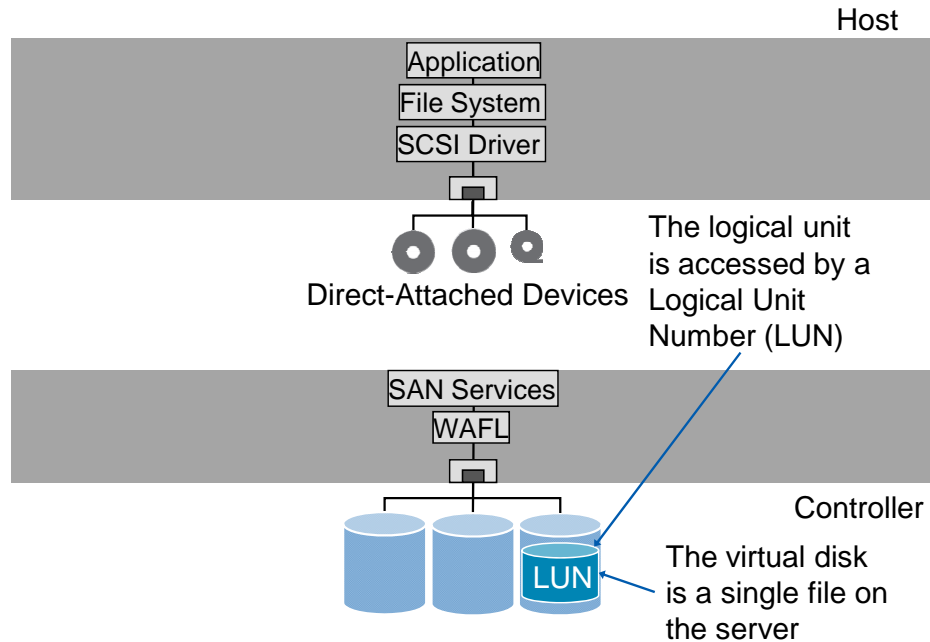
© 2010 NetApp, Inc. All rights reserved.

### SCSI ON HOST AND CONTROLLER

Traditionally, storage is attached to a local machine. SCSI is used for transmitting data between a host and peripheral devices either through SCSI adapters or other adapters that communicate using SCSI commands.



# Logical Unit



© 2010 NetApp, Inc. All rights reserved.

## LOGICAL UNIT



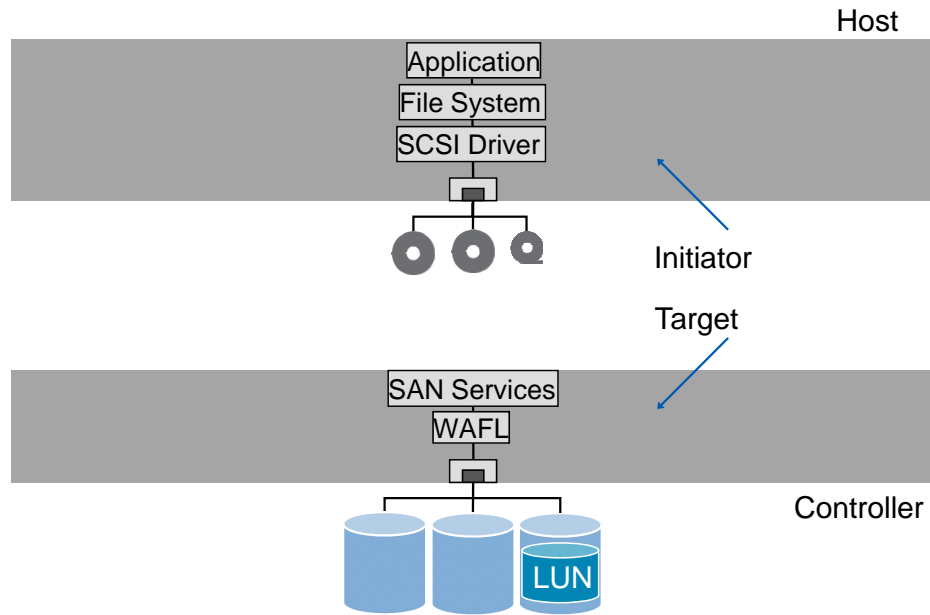
## Terms

© 2010 NetApp, Inc. All rights reserved.

## TERMS



## Initiator and Target



© 2010 NetApp, Inc. All rights reserved.




### INITIATOR AND TARGET

Initiators, including Windows® and UNIX®-type hosts, are consumers or clients within a SCSI relationship. Targets, including NetApp controllers and storage arrays, present data as logical units and are the servers within a SCSI relationship.





## SAN Types

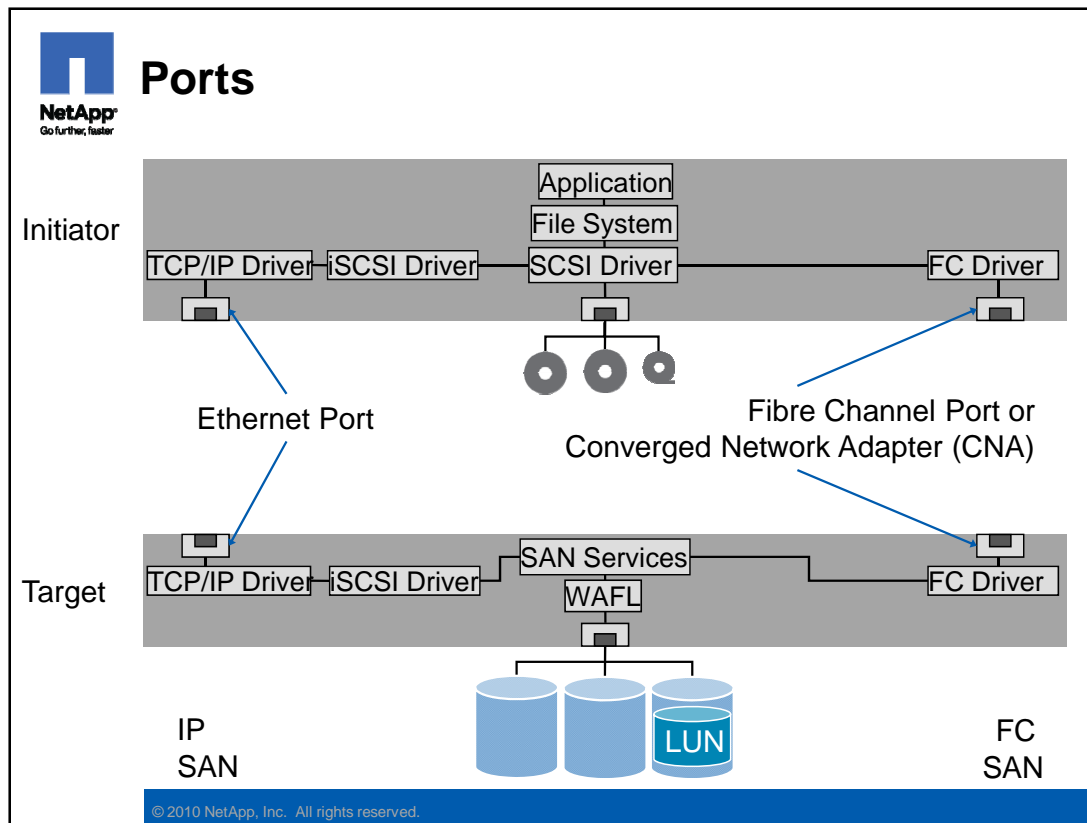
- A SAN may be implemented using either:
  - Fibre Channel (FC)
    - Referred to as FC SAN
    - Uses Fibre Channel Protocol to communicate
    - 
    - Uses Fibre Channel over Ethernet (FCoE) to communicate
    - 
  - Internet Protocol (IP)
    - Referred to as IP SAN
    - Uses Internet SCSI (iSCSI) to communicate
    - 

© 2010 NetApp, Inc. All rights reserved.

## SAN TYPES

LUNs on a NetApp storage system can be accessed through either a Fibre Channel (FC SAN) fabric using Fibre Channel Protocol or an Ethernet network using the Fibre Channel over Ethernet (FCoE) or Internet SCSI (iSCSI) protocols. In all cases, the transport portals (FC, FCoE or iSCSI) carry encapsulated SCSI commands as the data transport mechanism.

iSCSI is an IETF standard found here: [www.ietf.org/rfc/rfc3720.txt?number=3720](http://www.ietf.org/rfc/rfc3720.txt?number=3720).

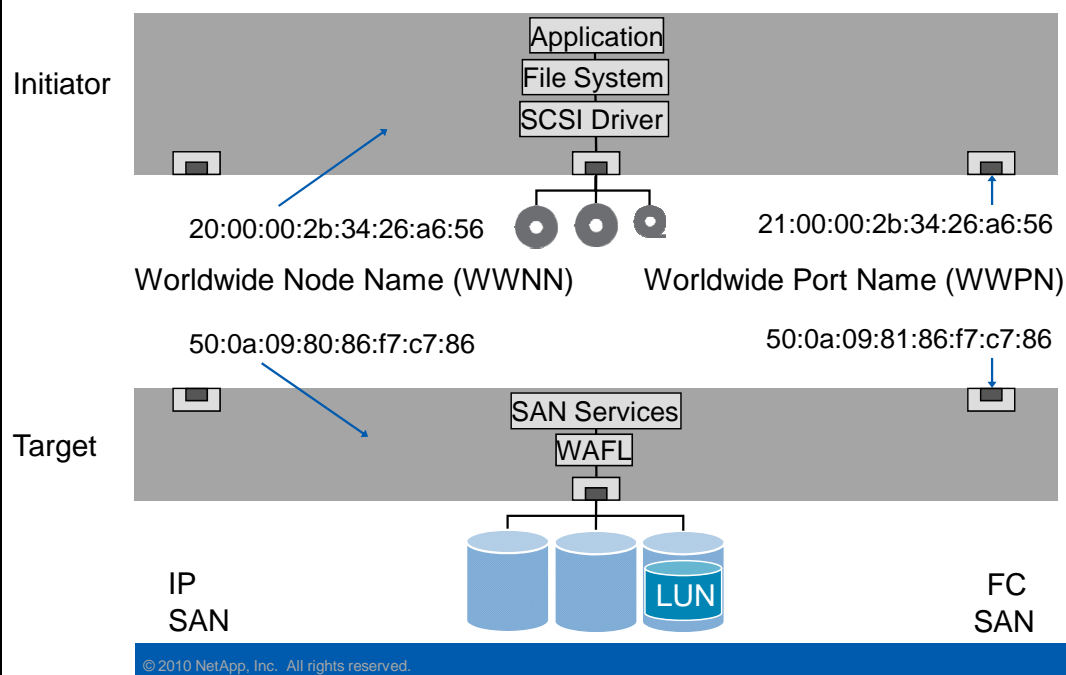


## PORTS

Data is communicated over ports. In an IP SAN, the data is communicated by way of Ethernet ports. In an FC SAN, the data is communicated over Fibre Channel ports.



## Node and Port Names in Fibre Channel

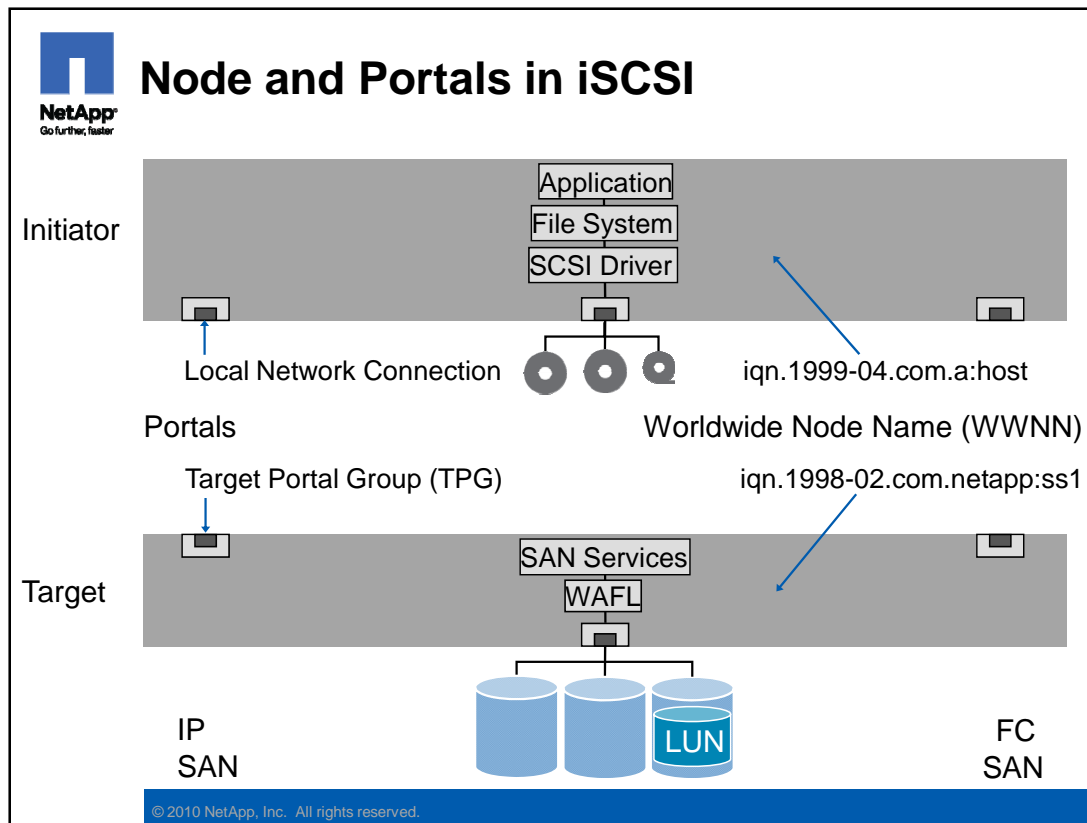


### NODE AND PRT NAMES IN FIBRE CHANNEL

In FC SAN, a worldwide node name (WWNN) describes a machine while a worldwide port name (WWPN) describes a physical port attached to that machine.

The FC specification for the naming of nodes and ports on those nodes can be fairly complicated. Each device is given a globally unique WWNN and an associated WWPN for each port on the node. WWNNs and WWPNS are 64-bit names made up of 16 hexadecimal digits grouped together in twos with a colon separating each pair (for example, 21:00:00:2b:34:26:a6:54).

The first number in the address defines what the other numbers in the address represent, according to the FC specification. The first number is generally a 1, 2, or 5. In the example of QLogic® initiator host bus adapters (HBAs), the first number is generally a 2. For Emulex® initiator HBAs, the first number is generally a 1. Finally, a NetApp storage system is assigned with a 5.



## NODE AND PORTALS IN ISCSI

In IP SAN, the worldwide node name (WWNN) describes a machine while the portal describes a physical port. Each iSCSI node must have a node name. There are two possible node name formats.

### IQN-TYPE DESIGNATOR

iSCSI Qualified Name or IQN node name is conventionally “iqn.yyyy-mm.backward\_naming\_authority:unique\_device\_name.” This is the most popular node name format and is the default used by a NetApp storage system. The components of the logical name are the following:

- Type designator, IQN, followed by a period (.)
- The date when the naming authority acquired the domain name, followed by a period
- The name of the naming authority, optionally followed by a colon (:)
- A unique device name

### EUI-TYPE DESIGNATOR

The Extended Unique Identifier or EUI node name is “eui.nnnnnnnnnnnnnnnnn.” The components of the logical name are the following:

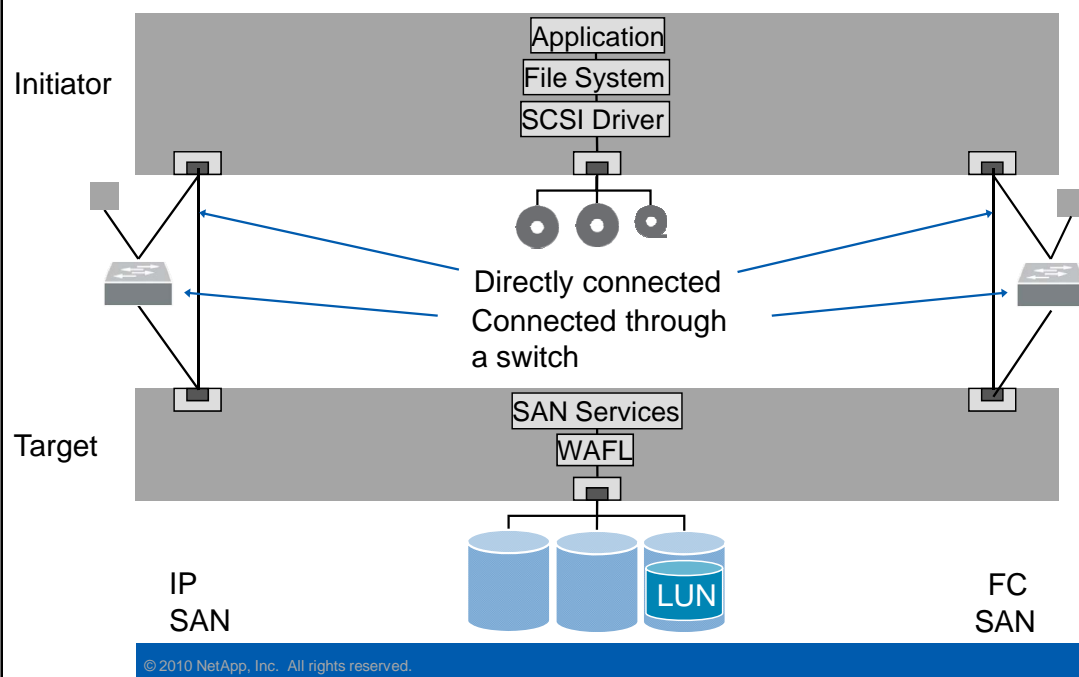
The type designator itself, “eui,” followed by a period (.)

Sixteen hexadecimal digits

Example: “eui.123456789ABCDEF0”



## Connectivity Between Initiator and Target



### CONNECTIVITY BETWEEN INITIATOR AND TARGET



## Implementing SAN

© 2010 NetApp, Inc. All rights reserved.

### IMPLEMENTING SAN



## Steps to Implement a SAN

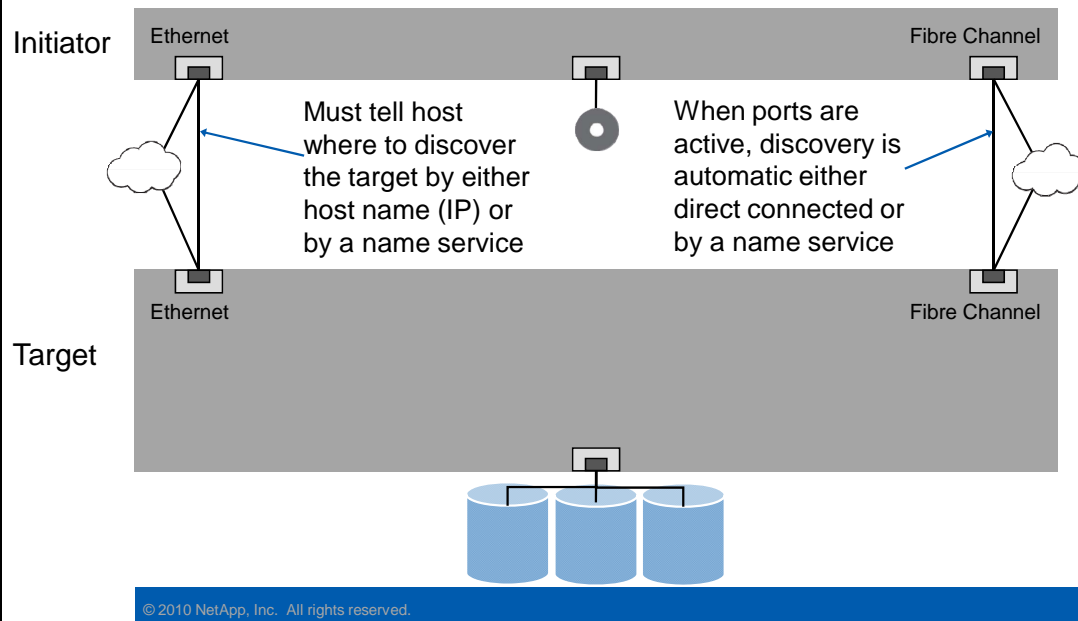
1. Have an initiator discover a target portal
2. Create a session between the initiator with the target and make the bindings persistent between reboots if required
3. Create an igroup on the storage system if necessary
4. Create a LUN on the storage system
5. Map the LUN to an igroup on the storage system
6. Find the LUN on the host
7. Prepare the disk for the host OS if necessary

© 2010 NetApp, Inc. All rights reserved.

## STEPS TO IMPLEMENT A SAN



# 1. Initiator Discovery of a Target




## 1. INITIATOR DISCOVERY OF A TARGET

The first step in a SAN implementation is discovery of a target. In IP SAN, an administrator must tell the client where to discover a target by either directly connecting by host name (IP) or discovery through a name service (iSNS). In FC SAN, if the ports are active, discovery is automatic--either directly connected or by a name service. For more information about IP SAN discovery, see module 14. For more information concerning FC SAN discovery, see module 13.





## 1. Initiator Discovery of a Target (Cont.)




© 2010 NetApp, Inc. All rights reserved.

## 1. INITIATOR DISCOVERY OF A TARGET (CONT.)



## 1. Initiator Discovery of a Target (Cont.)

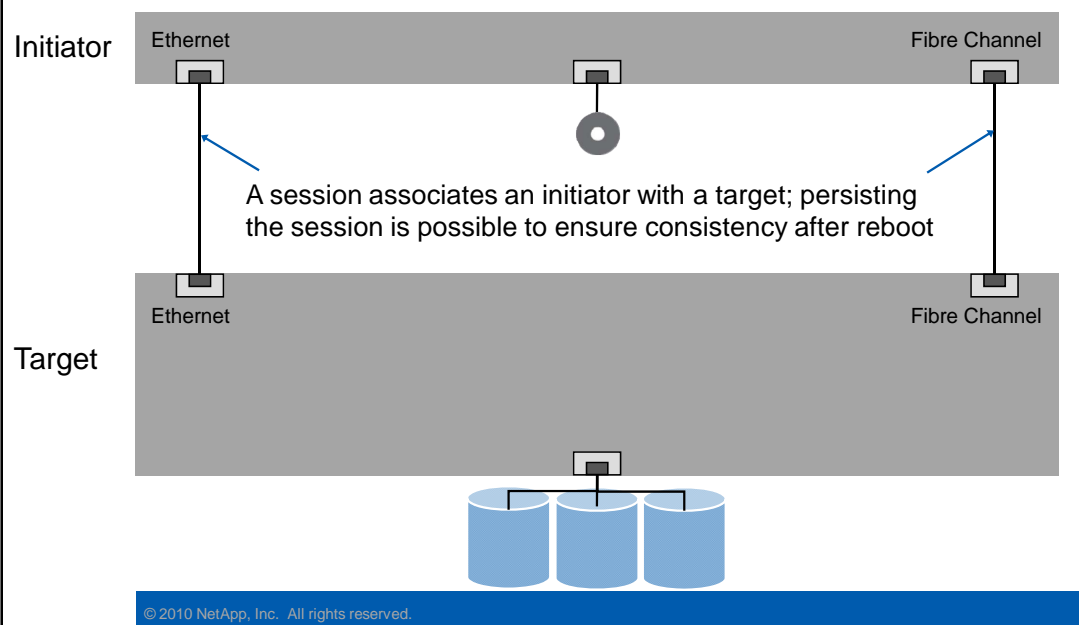


© 2010 NetApp, Inc. All rights reserved.

## 1. INITIATOR DISCOVERY OF A TARGET (CONT.)



## 2. Create a Session

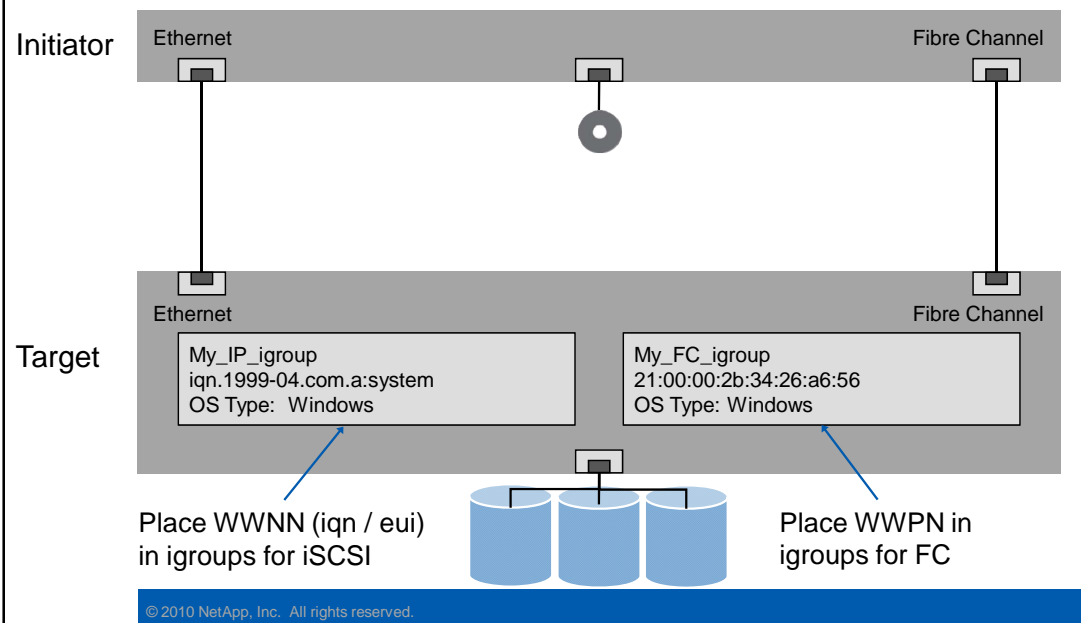


## 2. CREATE A SESSION

Sessions associate the initiators with targets. A session may be persisted to ensure availability after a host reboots.



### 3. Create an igroup

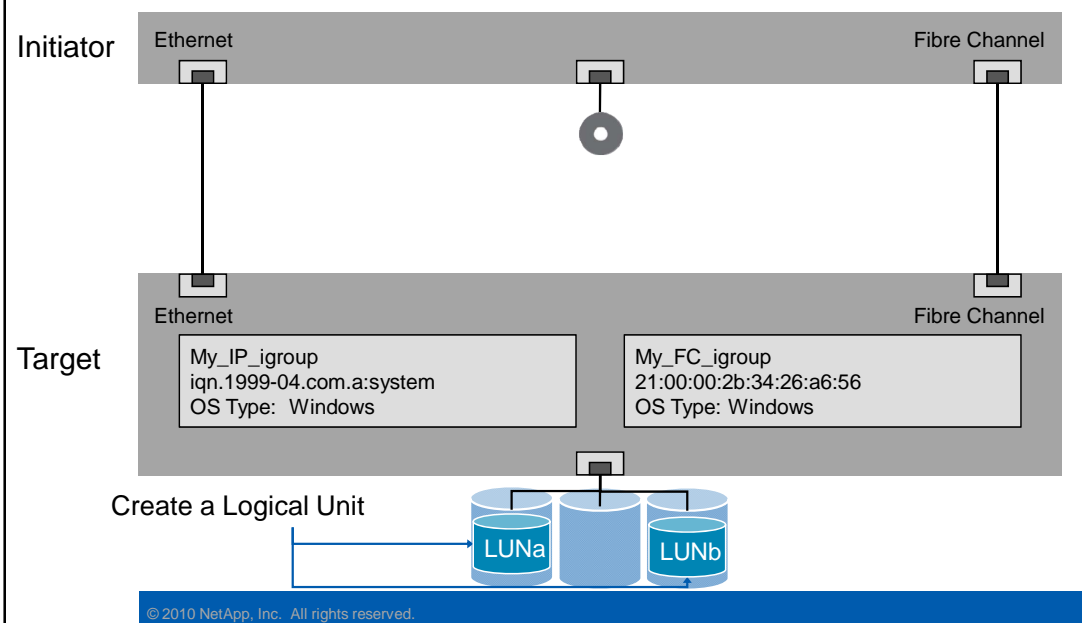


### 3. CREATE AN IGROUP

An igroup is a group of one or more initiators that have access to a target. In IP SAN, an administrator references an initiator by WWNN. In FC SAN, an administrator references an initiator by WWPN.



## 4. Create a Logical Unit



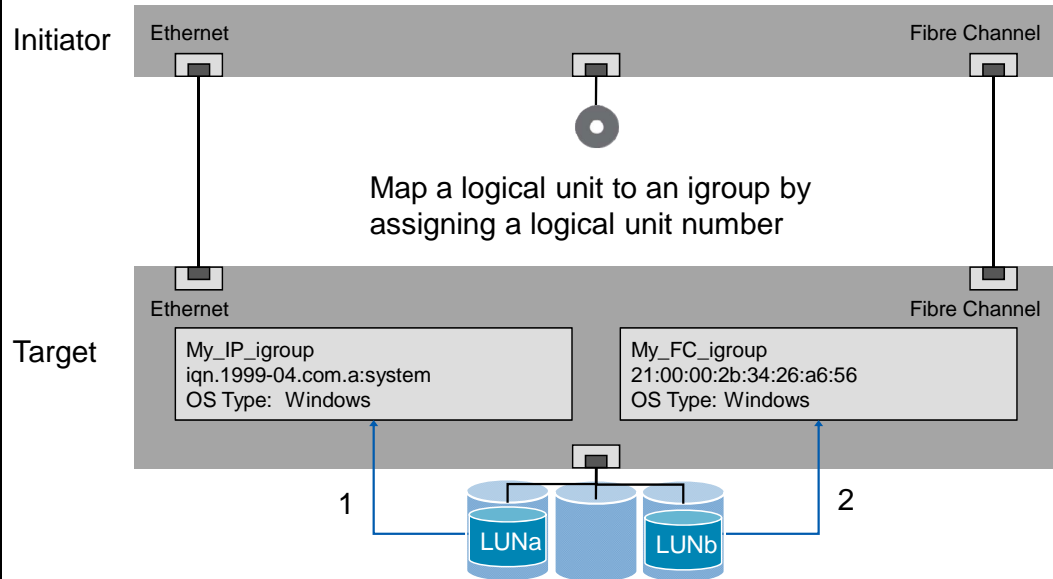
### 4. CREATE A LOGICAL UNIT

The next step in implementing a SAN is to create a logical unit. A logical unit is a logical representation of a physical unit of storage. It is a collection of, or a part of, physical or virtual disks configured as a single disk. When you create a logical unit, it is automatically striped across many physical disks. Data ONTAP manages logical units at the block level, so it cannot interpret the file system or data in a logical unit.



## 5. Map a Logical Unit to an igroup

**NOTE:** This step is also called LUN masking



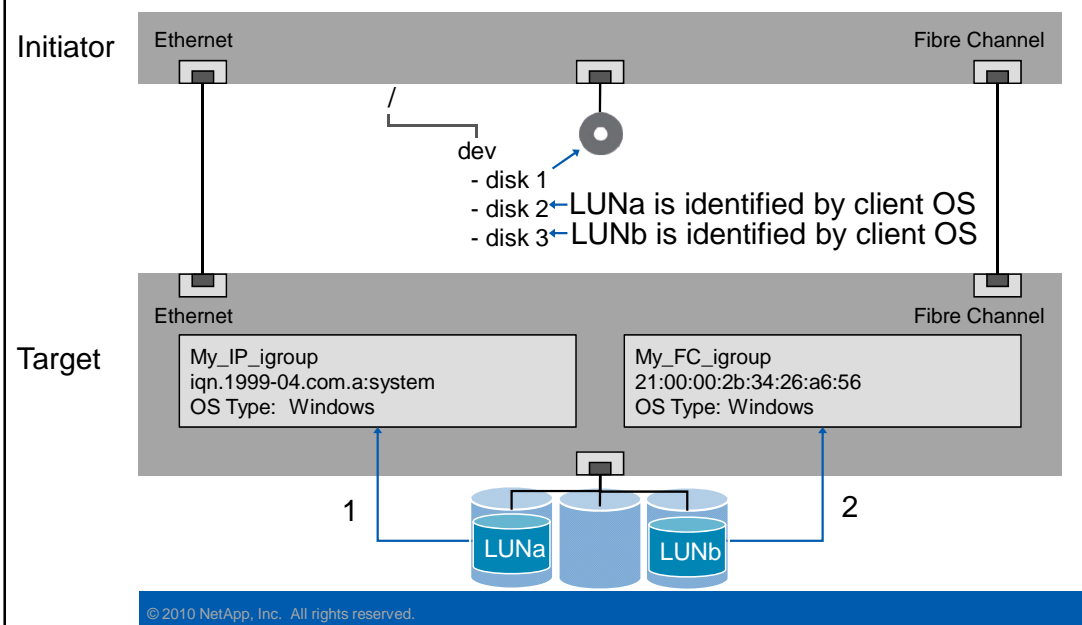
© 2010 NetApp, Inc. All rights reserved.

### 5. MAP A LOGICAL UNIT TO AN IGROUP

The logical unit is mapped to an igroup and referenced by an ID. The logical unit is then referred to as a logical unit number, or LUN.



## 6. Find the Disk

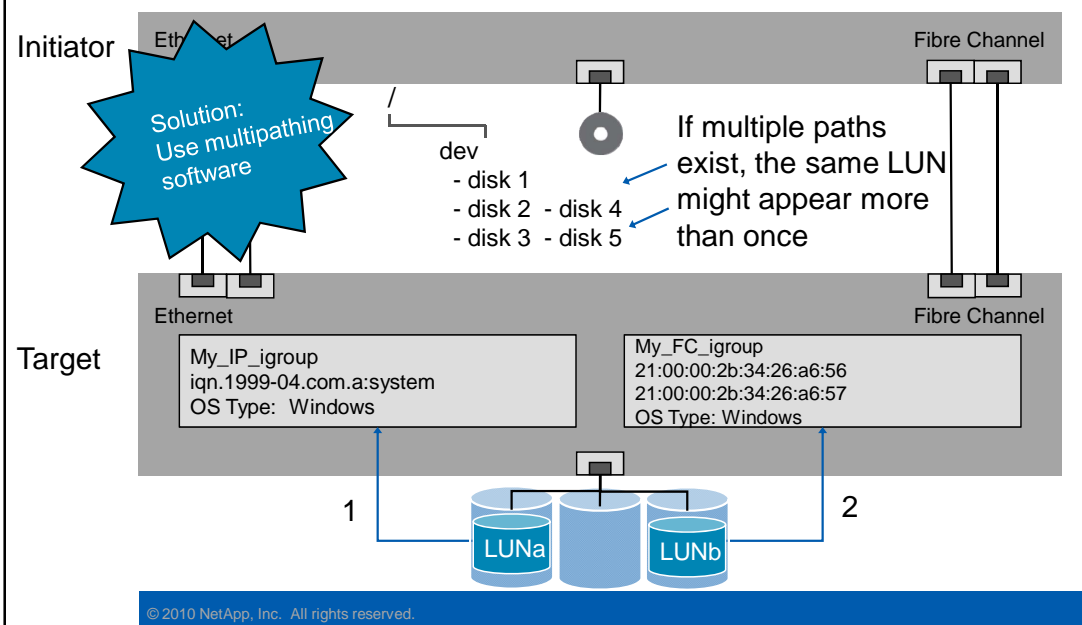


## 6. FIND THE DISK

The LUN is then identified by the client operating system. From the host, LUNs appear as local disks, allowing you to format and store data on them.



## 6. Find the Disk (Cont.)



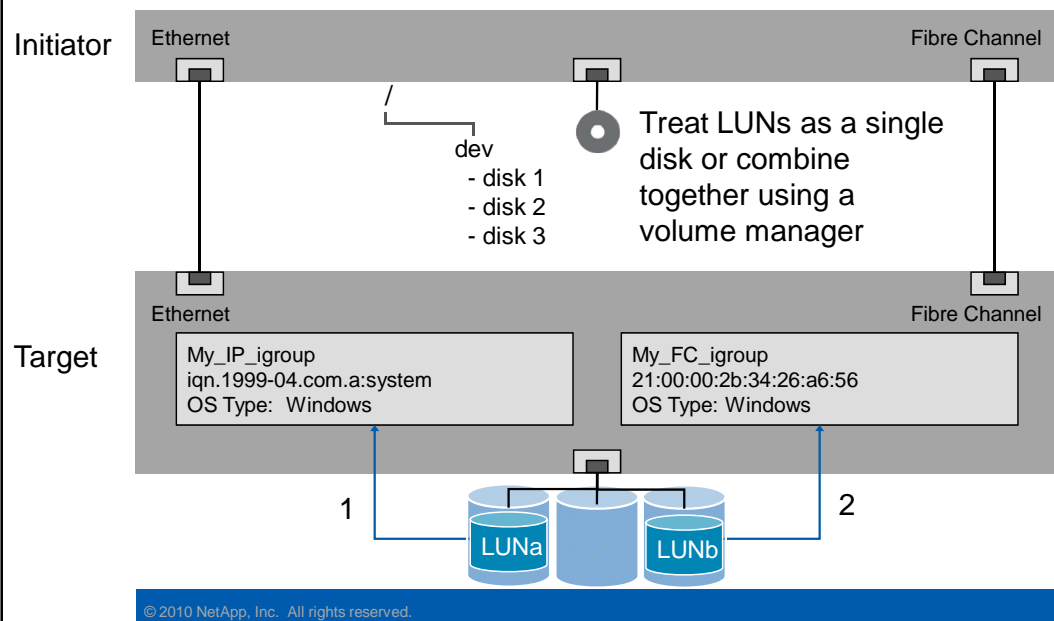
## 6. FIND THE DISK (CONT.)

If multiple paths exist, the LUN will appear more than once unless multipathing software is used.





## 7. Prepare the Disk for the Host OS

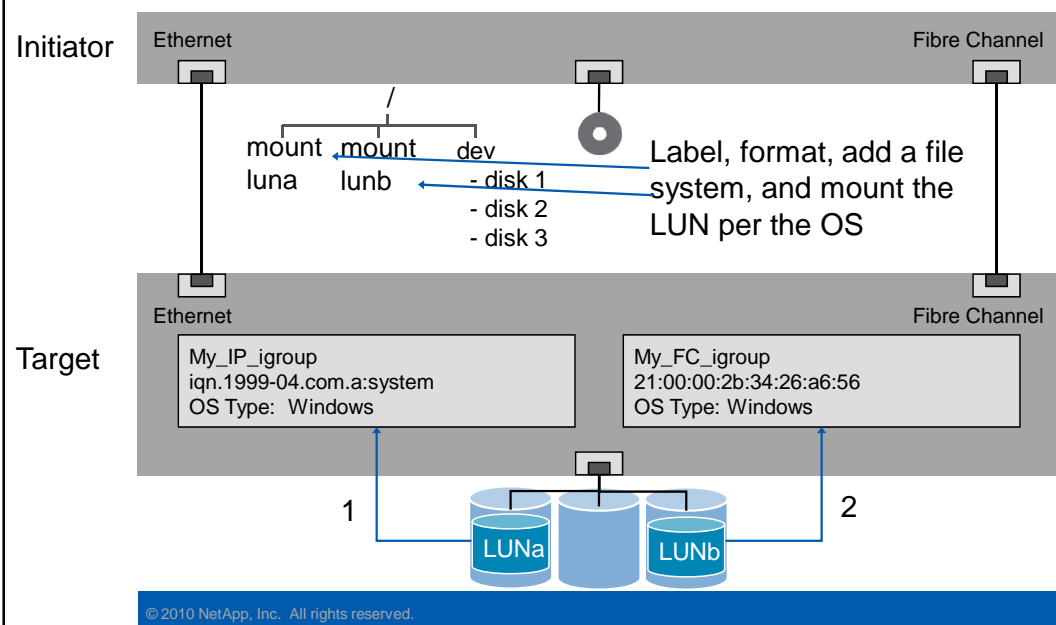


## 7. PREPARE THE DISK FOR THE HOST OS

LUNs may be used as single disk or combined together using a host-based volume manager.



## 7. Prepare the Disk for the Host OS (Cont.)

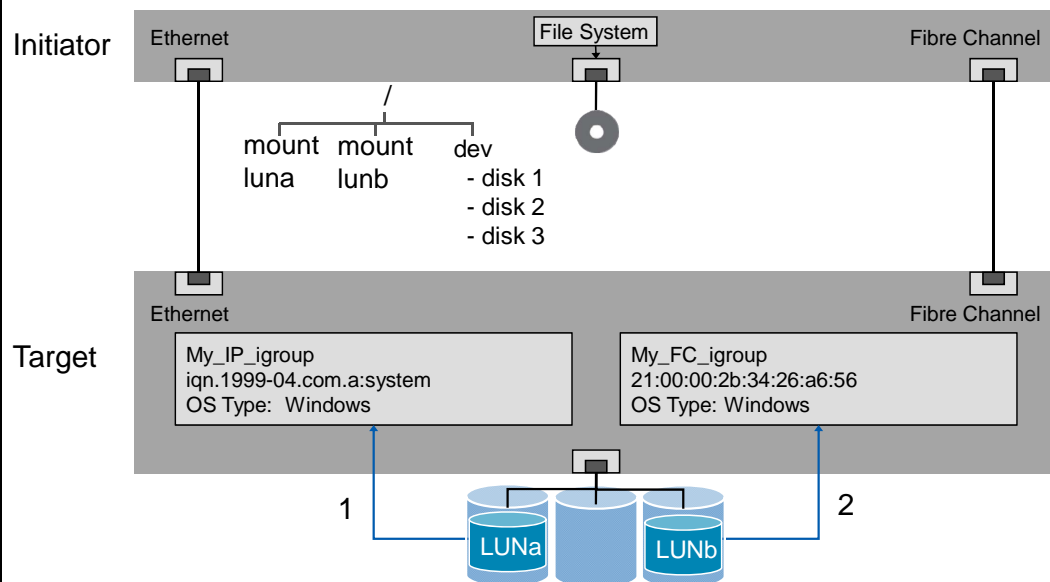


## 7. PREPARE THE DISK FOR THE HOST OS (CONT.)

Finally, the logical unit must be labeled, formatted, a file system added, and finally mounted by the client OS.



# LUN Setup Complete



© 2010 NetApp, Inc. All rights reserved.

## LUN SETUP COMPLETE



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe the differences between network-attached storage (NAS) and storage area network (SAN)
- List the methods to implement a SAN environment
- Define logical unit number, initiator, and target
- Describe ports, worldwide node names, and worldwide port names
- List the basic steps to implement a SAN

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 12: SAN Overview  
Estimated Time: 15 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- If NAS provides file-level access, then SAN provides what?
- What is it called when NAS and SAN are both implemented on a NetApp storage system?
- The initiator is on the host and the target is on the storage system. True or false?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# FC Connectivity

Module 13  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## FC CONNECTIVITY





## Module Objectives

By the end of this module, you should be able to:

- Describe multiple path implementation with Fibre Channel (FC) connectivity
- Describe how to configure FC ports on Windows®, for Red Hat®, and NetApp® systems
- Describe commands and utilities to identify the worldwide node name (WWNN) and worldwide port name (WWPN) on Windows, Red Hat, and NetApp systems

© 2010 NetApp, Inc. All rights reserved.

### MODULE OBJECTIVE



## FC Connectivity Configuration

The following are the steps to configure FC SAN:

1. Determine the FC topology
2. Verify host HBA configuration, drivers, firmware, cables, and multipathing software
3. Configure the switch (if in the topology)
4. Configure the target(s)
5. Configure the initiator(s)
6. Cable the devices together
7. Implement FC zoning (if required)

© 2010 NetApp, Inc. All rights reserved.

### FC CONNECTIVITY CONFIGURATION




## Data ONTAP

© 2010 NetApp, Inc. All rights reserved.

## DATA ONTAP



## High-Availability



- The exercise environment's storage systems have been configured as a high-availability pair
- To configure controller failover:  
system and system2> **license add xxxxxxx**  
system and system2> **reboot**  
system> **cf enable**

© 2010 NetApp, Inc. All rights reserved.

## HIGH-AVAILABILITY



## Data ONTAP as an FC Target

- Data ONTAP® 6.3 and later has support for FC
- Data ONTAP features:
  - Built-in FC functionality
  - Simple LUN creation and management
- To properly configure Data ONTAP for FC connectivity:
  - Confirm FC HBA(s) and port(s)
  - Configure and verify Fibre Channel protocol
  - Configure the FC Target HBA(s)
  - Identify the worldwide node name (WWNN)
  - Identify the worldwide port name (WWPN)

© 2010 NetApp, Inc. All rights reserved.

### DATA ONTAP AS AN FC TARGET



## HBA Confirmation

- Confirm the current HBAs:

```
system> sysconfig -a
```

- To identify the type of the on-board FC ports:

```
system> fcadmin config
```

|         |           | Local      |        |
|---------|-----------|------------|--------|
| Adapter | Type      | State      | Status |
| -----   |           |            |        |
| 0a      | initiator | CONFIGURED | online |
| 0b      | initiator | CONFIGURED | online |
| 0c      | initiator | CONFIGURED | online |
| 0d      | initiator | CONFIGURED | online |

**NOTE:** Add-on cards are configured to be either an initiator or target and cannot be changed

© 2010 NetApp, Inc. All rights reserved.

## HBA CONFIRMATION

The fcadmin utility manages the Fibre Channel adapters used by the storage subsystem. Use these commands to show link-level and loop-level protocol statistics, list the storage devices connected to the storage system, and configure the personality of embedded adapters.



## HBA Confirmation (Cont.)

- To change an onboard interface from an initiator to a target:

```
system> fcadmin config -d 0b ← Down the interface
system> fcadmin config -t target 0b ← Reconfigure the interface as a target
system> reboot ← A reboot is required
```

- To change an onboard interface from a target to an initiator:

```
system> fcadmin config -d 0b
system> fcadmin config -t initiator 0b
system> reboot
```

© 2010 NetApp, Inc. All rights reserved.

## HBA CONFIRMATION (CONT.)



## Configuring FC Services in Data ONTAP

1. Verify Fibre Channel protocol service is running:
  - `fcv status`
2. Verify FC is licensed (license the FC services if needed):
  - `license`
  - `license add XXXXXX`
3. Start the FC service:
  - `fcv start`

© 2010 NetApp, Inc. All rights reserved.

### CONFIGURING FC SERVICES IN DATA ONTAP





## Configuring FC HBA in Data ONTAP

1. List the installed target HBAs:
  - `fcv show adapters`
2. Take an HBA offline:
  - `fcv config adapter down`
3. Set target HBA speed to the FC switch port's speed to improve takeover and giveback performance:
  - `fcv config adapter speed [auto|1|2|4]`
4. Bring an HBA online:
  - `fcv config adapter up`

© 2010 NetApp, Inc. All rights reserved.

### CONFIGURING FC HBA IN DATA ONTAP



## Identify WWNN in Data ONTAP

- WWNN uniquely identifies the storage system
  - The default WWNN is generated by a serial number in its NVRAM and stored on disk
    - Normally doesn't need to be changed
- To identify the WWNN:

```
system> fcp nodename
```

```
Fibre Channel nodename: 50:0a:09:80:86:f7:c7:86
(500a098086f7c786)
```

- To change the WWN:

```
system*> fcp nodename new_nodename
```

← Data ONTAP 7.3.1.1  
and later

© 2010 NetApp, Inc. All rights reserved.

## IDENTIFY WWNN IN DATA ONTAP



## Identify WWPN in Data ONTAP

- WWPN uniquely identifies an FC HBA port
- WWPNs are determined by:
  - WWNN
  - Controller failover mode (cfmode)
  - Internal port index

- To verify the default WWPN:

```
system> fcp portname show
```

| Portname                | Adapter |                  |
|-------------------------|---------|------------------|
| -----                   | -----   |                  |
| 50:0a:09:81:86:f7:c7:86 | 0c      | Use fcp portname |
| 50:0a:09:82:86:f7:c7:86 | 0d      | show -v to list  |
|                         |         | available WWPNs  |

- To change a WWPN:

```
system> fcp portname set adapter
```

© 2010 NetApp, Inc. All rights reserved.

## IDENTIFY WWPN IN DATA ONTAP

Within Fibre Channel (FC) SAN, worldwide port names (or WWPNs) uniquely identify each Fibre Channel port. Each 64-bit name is determined by three factors. The first factor is the worldwide node name (WWNN), which is the unique identifier for the NetApp storage system running as an FC target device server. The second factor is the controller failover mode (cfmode) currently set on the NetApp storage system. The third and final factor is that each FC target port has an internal port index range that assists in assigning the WWPNs.



## FC cfmodes Defined

- Controller failover mode (cfmode) determines how HBAs of storage systems in an high-availability configuration:
  - Log in to the fabric
  - Provide access to local and partner LUNs
- Both storage systems in the active-active configuration must have the same settings
  - `fcv show cfmodes` to verify current setting
  - `fcv set cfmodes` to set the cfmodes
    - Requires advanced mode to set cfmodes  
`priv set advanced`

© 2010 NetApp, Inc. All rights reserved.

## FC CFMODES DEFINED

Cfmodes only applies to Fibre Channel environments in an high-availability NetApp storage controller configuration. The cfmodes determine how target ports do the following:

- Log in to the fabric
- Handle local and partner traffic for a cluster
- Provide access to local and partner LUNs in a cluster

In the original release of Data ONTAP 6.3, which included SAN support for Fibre Channel, **cfmodes standby** was the implied default. There was not a setting for cfmodes in that release, and it was not called `cfmodes standby`. However, when Data ONTAP 6.5 was released, four cfmodes were introduced. One of these modes was standby. The others were partner, mixed, and dual fabric. In Data ONTAP 7.1, a new cfmodes called single system image (SSI) became available. SSI is the default cfmodes for new installations with Data ONTAP 7.2. The availability of standby, partner, mixed, and dual fabric modes is dependent on the storage controller model, Data ONTAP version, and/or the use of 2-Gb or 4-Gb FC ports. With Data ONTAP 7.3 and later, the only configurable cfmodes is single system image.



## FC cfmode Types

- Prior to Data ONTAP 7.3, there were five types of cfmodes:

| cfmode       | Supported Systems                                                                                                                     | Benefits and Limitations                                                                                                                                                                               |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| standby      | All systems except the FAS270c, FAS20x0, FAS31x0, FAS6040, FAS6080, and FAS6030 and FAS6070 with a 4-Gb or 8-Gb FC adapter            | <ul style="list-style-type: none"><li>■ Requires more switch ports</li><li>■ Supports only Windows and Solaris™ hosts</li></ul>                                                                        |
| partner      | All systems except the FAS270c, FAS20x0, FAS3040, FAS3070, FAS31x0, FAS60x0, and any FAS system with a 4-Gb or 8-Gb target FC adapter | <ul style="list-style-type: none"><li>■ Supports all host OS types</li><li>■ Supports all switches</li></ul>                                                                                           |
| dual_fabric  | FAS270c only                                                                                                                          | <ul style="list-style-type: none"><li>■ Supports all host OS types</li><li>■ Requires fewer switch ports</li><li>■ Does not support all switches; requires switches that support public loop</li></ul> |
| mixed        | All systems except the FAS270c, FAS20x0, FAS30x0, FAS31x0, FAS6040, FAS6080, and FAS6030 and FAS6070 with a 4-Gb or 8-Gb FC adapter   | <ul style="list-style-type: none"><li>■ Supports all operating systems</li><li>■ Does not support all switches; requires switches that support public loop</li></ul>                                   |
| single_image | All systems                                                                                                                           | <ul style="list-style-type: none"><li>■ Supports all host OS types</li><li>■ Supports all switches</li><li>■ Makes all LUNs available on all target ports</li></ul>                                    |

© 2010 NetApp, Inc. All rights reserved.

## FC CFMODES TYPES

There are five possible cfmodes on the storage controller. Only one cfmode can be set per each storage controller, and in a cluster situation the cfmode must be the same for both systems.

### STANDBY

The standby mode is supported on all systems except the FAS270c. It supports only Windows and Solaris operating systems. In addition, this mode requires additional switch ports.

### PARTNER

The partner mode is supported on all systems except the FAS270c and the FAS6000 series. All switches and host operating systems are supported.

### DUAL-FABRIC

The dual-fabric mode is only supported on a FAS270c. All host operating systems are supported by this mode. This mode requires a switch that supports a public loop.

### MIXED

The mixed mode is supported on all systems except the FAS270c and the FAS6000 series. Mixed mode supports all host operating systems, but requires a switch that supports a public loop.

### SINGLE IMAGE

The single image mode is supported on all systems, switches, and host operating systems. This mode makes all LUNs available on all target ports.



## FC cfmodes Types (Cont.)

- Data ONTAP 7.3 and later supports only `single_image` cfmode
  - Partner and standby are supported only on upgrades of existing systems that currently support and use these cfmodes modes
  - After an upgrade to Data ONTAP 7.3 or later, a storage system pair may only be (re)configured to `single_image`
- `single_image` cfmode
  - Each high-availability pair has a single WWNN, allowing both storage systems in the active-active configuration to function as a single Fibre Channel storage system
  - All LUNs in a high-availability configuration are available on all ports in the high-availability pair
  - Multipathing software is required

© 2010 NetApp, Inc. All rights reserved.

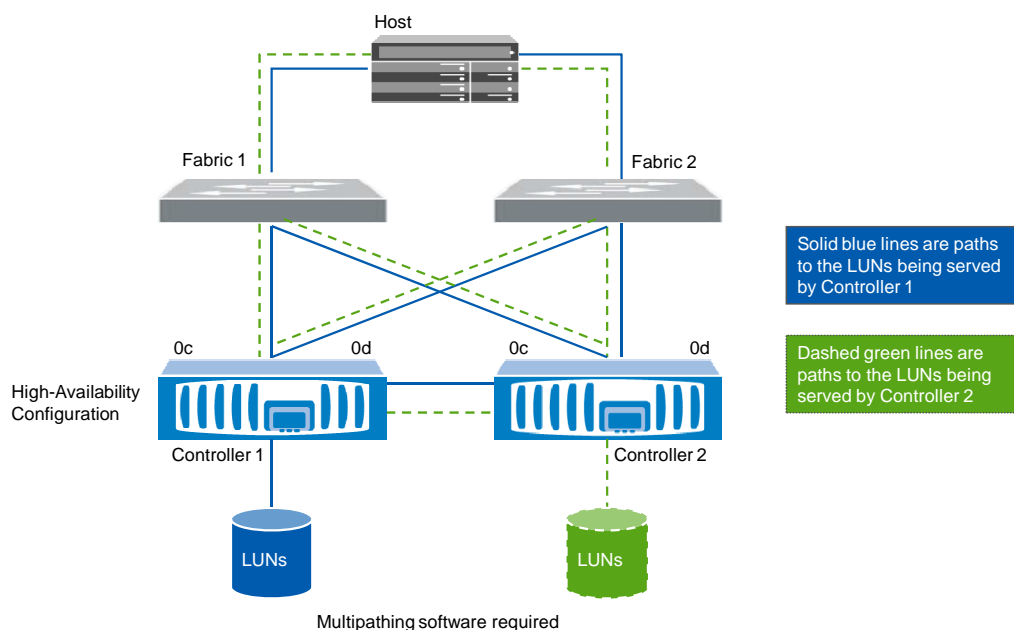
## FC CFMODES TYPES (CONT.)

The `single_image` cfmode setting is available in Data ONTAP 7.1. This cfmode setting is the default for new installs Data ONTAP 7.2 and later. Upgrades to Data ONTAP 7.2 will retain the cfmode from the previous version. In SSI cfmode, an high-availability storage controller configuration has a single WWNN, and both systems in the configuration function as a single Fibre Channel node. Each node in the cluster shares the partner node's LUN map information.

All LUNs in the cluster are available on all ports in the cluster by default. As a result, more paths to each LUN are stored on the cluster. Any port can provide access to both local and partner LUNs. You can specify the LUNs available on a subset of ports by defining port sets and binding them to an igroup. Any host in the igroup can then access the LUNs only by connecting to the target ports in the port set.



## Single Image Example



© 2010 NetApp, Inc. All rights reserved.

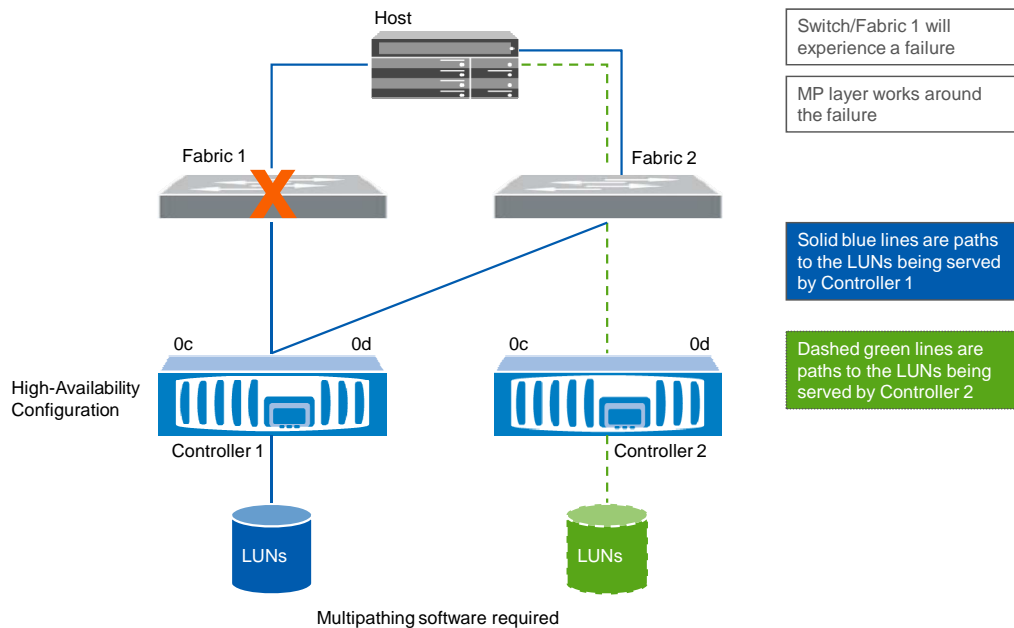
### SINGLE IMAGE EXAMPLE

LUNs from both controllers are visible through a single physical (and logical) port.

A single FC port is a primary path for a LUN served by that controller and a secondary path for a LUN on the partner controller.



## Single Image Example (Cont.)



© 2010 NetApp, Inc. All rights reserved.

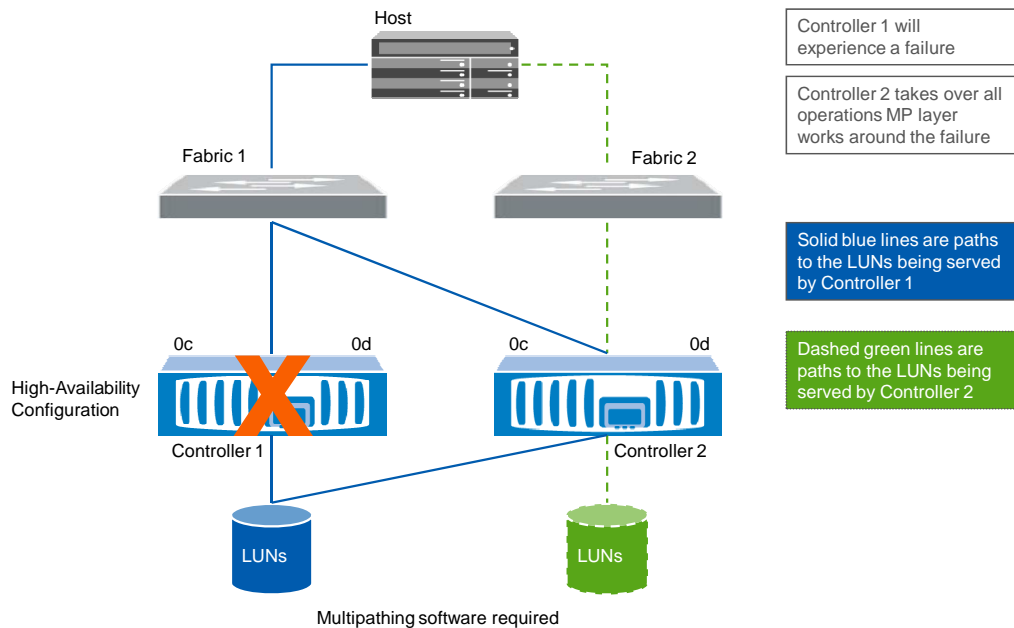
### SINGLE IMAGE EXAMPLE (CONT.)

Switch 1 experiences a failure. The host multipathing software layer works around that failure to reroute the I/O through Fabric 2 to the LUN.





## Single Image Example (Cont.)



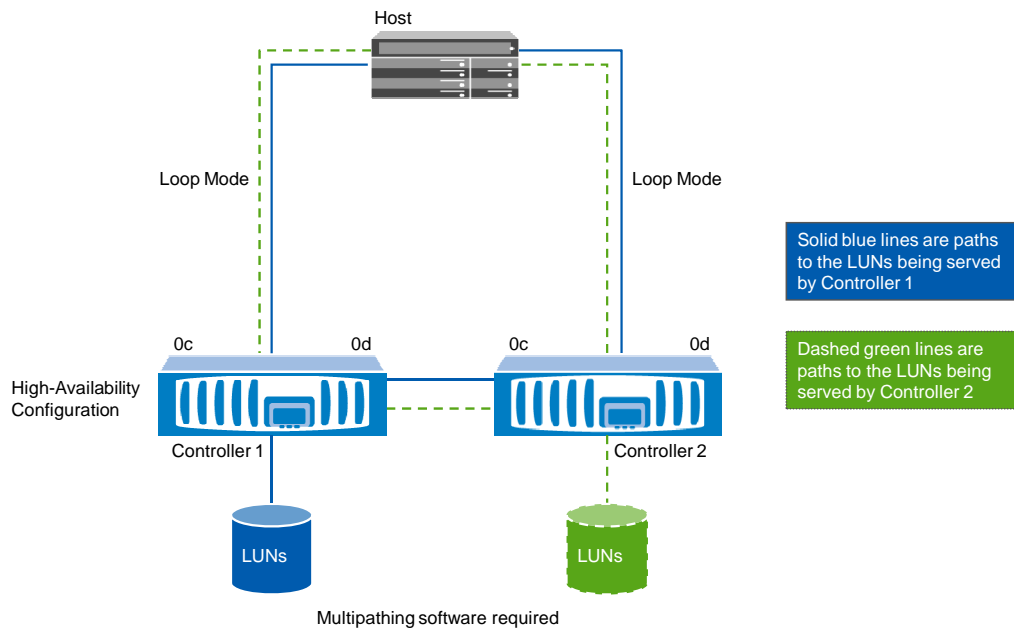
© 2010 NetApp, Inc. All rights reserved.

### SINGLE IMAGE EXAMPLE (CONT.)

When Controller 1 experiences a failure, Controller 2 takes over all operations. The host multipathing software layer works around the failure by rerouting I/O through the Controller 2 path to the LUN.



## Single Image Example (Cont.)



### SINGLE IMAGE EXAMPLE (CONT.)

Single system image does support this configuration. SSI allows ports to alternate between fabric point-to-point login and individual loop. All LUNs are available on a single port in the event that one of the links fails.



## fcg config (Single System Image Mode)

```
system or system2> fcp nodename
Fibre Channel nodename: 50:0a:09:80:86:f7:c7:86 (500a098086f7c786)
system or system2> fcp show cfmode
fcp show cfmode: single_image
system> fcp config
0c: ONLINE <ADAPTER UP> PTP Fabric
 host address 011000
 portname 50:0a:09:81:96:f7:c7:86 nodename 50:0a:09:80:86:f7:c7:86
 mediatype auto speed auto
0d: ONLINE <ADAPTER UP> PTP Fabric
 host address 011100
 portname 50:0a:09:82:96:f7:c7:86 nodename 50:0a:09:80:86:f7:c7:86
 mediatype auto speed auto
system2> fcp config
0c: ONLINE <ADAPTER UP> PTP Fabric
 host address 011200
 portname 50:0a:09:81:86:f7:c7:86 nodename 50:0a:09:80:86:f7:c7:86
 mediatype auto speed auto
0d: ONLINE <ADAPTER UP> PTP Fabric
 host address 011300
 portname 50:0a:09:82:86:f7:c7:86 nodename 50:0a:09:80:86:f7:c7:86
 mediatype auto speed auto
```

© 2010 NetApp, Inc. All rights reserved.

### FCG CONFIG (SINGLE SYSTEM IMAGE MODE)

This is an example of the adapter settings for `single_image` cfmode. Notice that all node names are identical and that the media type is set to auto. This means that the ports log in to the fabric using point-to-point mode. If point-to-point mode fails, then the ports will try loop mode.



## Windows


© 2010 NetApp, Inc. All rights reserved.

## WINDOWS



## Windows as an FC Initiator

- NetApp has supported Windows as an FC initiator OS since Windows 2000 Server
- Windows Server 2008 has many advantages over previous versions
  - New tools
    - Storage Explorer
    - Storage Manager for SANs
  - Multipath I/O (MPIO)
  - Built-in FC drivers
- Windows must be properly configured for FC connectivity



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS AS AN FC INITIATOR

Windows Server 2008 provides many new features that make configuring an FC SAN easier. Storage Explorer provides a one-stop interface for investigating local HBAs as well as the FC switches, if present. Storage Manager for SANs is an additional tool that is available in Windows Server 2003 R2 and later. This tool allows the configuration of a SAN environment. Storage Manager for SANs requires the Virtual Disk Service add-in provided by NetApp at the NOW™ (NetApp on the Web) site.



## Windows FC Design and Installation

1. Verify host operating system releases, required patches, and NetApp Windows Host Utility Kit with Interoperability Matrix:
  - Use System Properties dialog to verify release
2. Install compatible host bus adapters (HBAs).
3. Install and configure required HBA drivers and utilities.
4. Verify an HBA:
  - Emulex®: Use HBAnyware®
  - QLogic®: Use SANsurfer
  - All HBA Types: Device Manager Dialog
5. Install compatible NetApp Windows Host Utility Kit:
  - Provides Perl scripts to diagnose and troubleshoot Windows
  - Example: `windows_info` provides diagnostic information

© 2010 NetApp, Inc. All rights reserved.

### WINDOWS FC DESIGN AND INSTALLATION

Host utilities contain software tools and documentation that allow you to configure a host in a NetApp SAN environment.

**NOTE:** Host utilities were formerly called Host Attach and Support Kits. Kits that were released before this naming convention changed are still called Host Attach and Support Kits. The term host utilities will be used in this course, but be aware that NetApp is in the process of transitioning to this name.

Host utilities are available from the Download Software page on the NOW site at:  
[now.netapp.com/NOW/cgi-bin/software](http://now.netapp.com/NOW/cgi-bin/software).



## Windows Implementation

After installation, to configure a Windows Emulex or QLogic implementation:

- Verify the HBA is enabled
- Identify the WWNN on the host HBA(s)
- Identify the WWPN on the host HBA(s)
- Verify connectivity between the initiator(s) and target

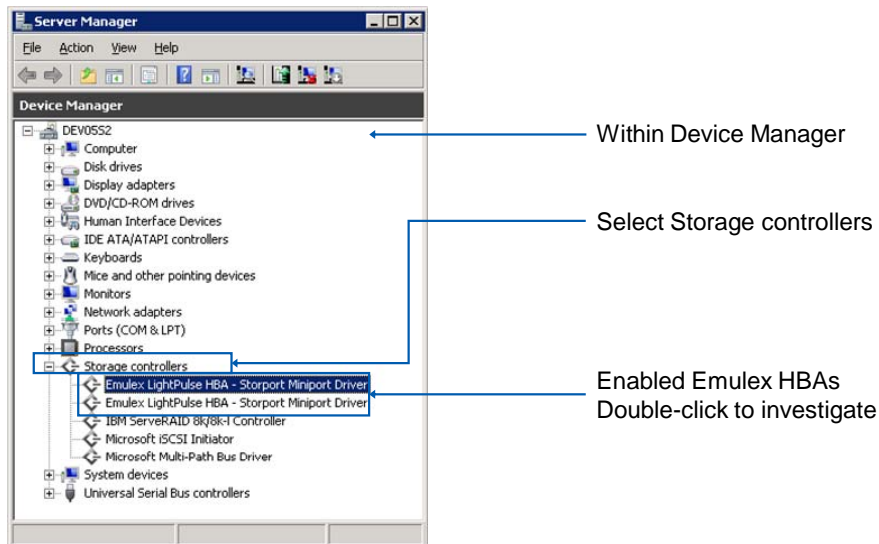
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS IMPLEMENTATION



# Windows/Emulex Implementation

- Verify that Windows Server 2008 has identified the HBA(s)



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/EMULEX IMPLMENTATION






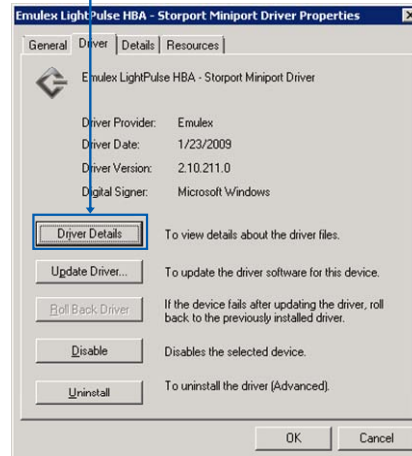
## Windows/Emulex Implementation (Cont.)

- Identify the driver associated with the HBA(s)

Other HBAs will be at a different location



For more information



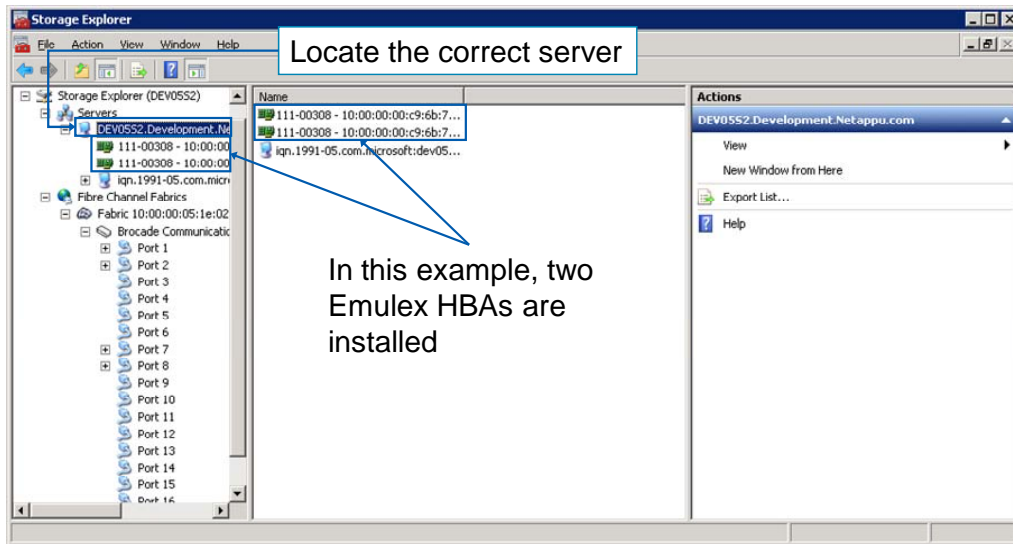
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/EMULEX IMPLEMENTATION (CONT.)



## Windows/Emulex Implementation (Cont.)

- Verify the HBA(s) is connected on Windows Server 2008



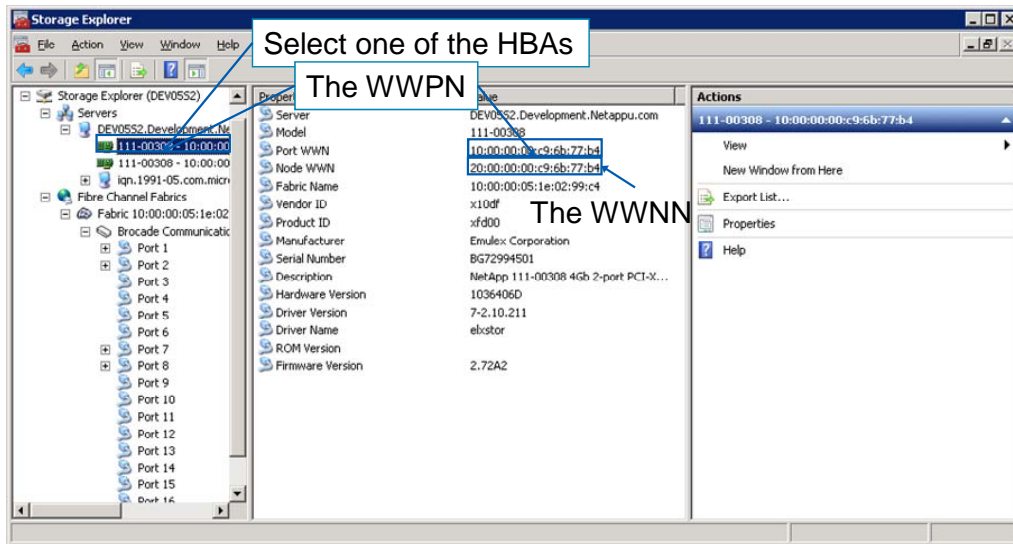
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/EMULEX IMPLMENTATION (CONT.)



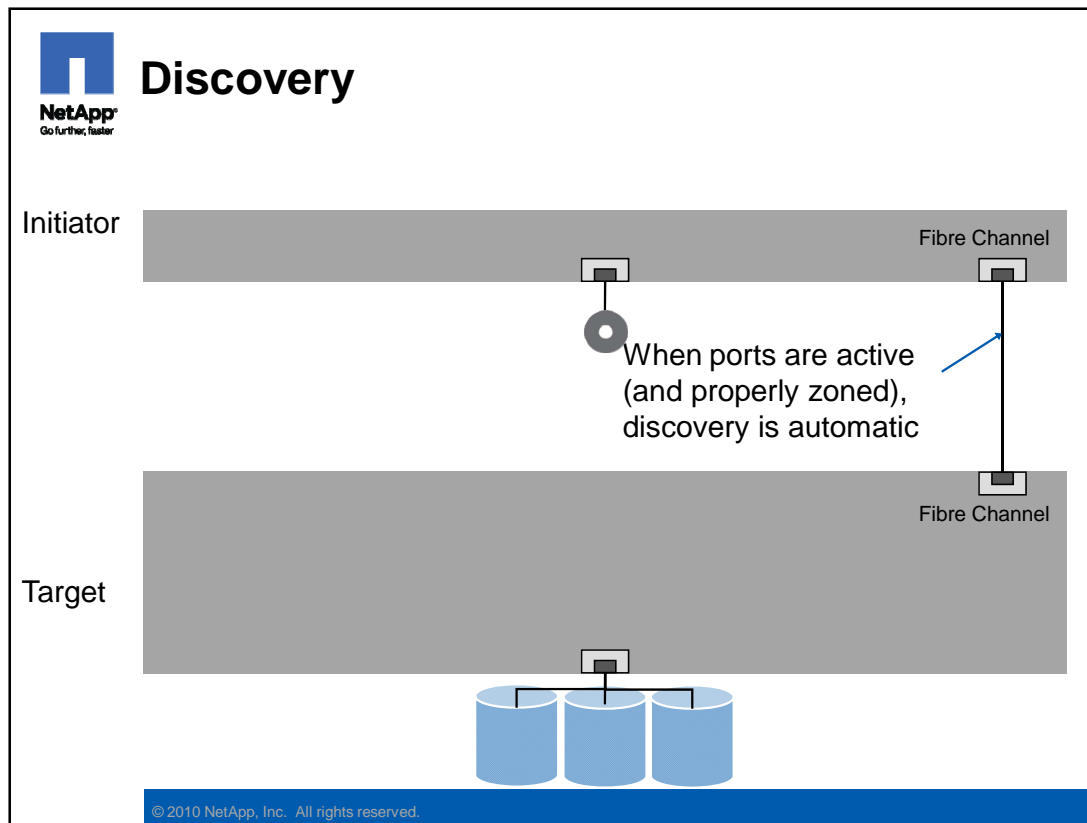
## Windows/Emulex Implementation (Cont.)

- Identify the local WWPN(s) on Windows Server 2008



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/EMULEX IMPLMENTATION (CONT.)



## DISCOVERY

Within FC SAN, discovery is automatic unless switch zoning prevents it. See Appendix 1 for a discussion about switch zoning.



## Data ONTAP Discovery of Initiators

- Verify connectivity from the storage system:

```
system> fcp show initiators
```

```
Initiators connected on adapter 0c:
```

```
Portname
```

```
Group
```

```

```

```

```

```
10:00:00:00:c9:6b:77:b4
```

← Windows WWPN

**NOTE:** For convenience, you may assign an alias to the Windows WWPN

© 2010 NetApp, Inc. All rights reserved.

### DATA ONTAP DISCOVERY OF INITIATORS



## WWPN Aliases

- In large FC installations, it can be difficult to identify WWPNs because of their cryptic 64-bit name
  - Example: 10:00:00:00:c9:6b:77:b4
- For convenience, WWPN may be assigned a name or “alias” within Data ONTAP
- Both target and initiator ports may be aliased

© 2010 NetApp, Inc. All rights reserved.

### WWPN ALIASES

One common problem administrators face in large Fibre installations is determining how to distinguish between WWPNs due to their cryptic 64-bit naming conventions. Now with Data ONTAP 7.3 and later, administrators can rename or “alias” a WWPN with a more convenient name to assist in easy identification. Aliases may be used for both target and initiator ports.



## fcpx wwpn-alias Command

### ■ To alias a WWPN:

- Use `fcpx wwpn-alias set [-f] alias wwpn`

- Example:

```
system> fcpx wwpn-alias set WIN1-FC
10:00:00:00:c9:6b:77:b4
```

### ■ To remove an alias:

- Use `fcpx wwpn-alias remove {-a alias...|-w wwpn}`

- Example:

```
system> fcpx wwpn-alias remove -a WIN1-FC WIN2-FC
```

*(removes the aliases)*

```
system> fcpx wwpn-alias remove -w
10:00:00:00:c9:6b:77:b4
```

*(removes all aliases for a particular WWPN)*

### ■ To show aliases:

- Use `fcpx wwpn-alias show [-a alias | -w wwpn]`

© 2010 NetApp, Inc. All rights reserved.

## FCP WWPN-ALIAS COMMAND

To alias a WWPN, an administrator must use the new `fcpx wwpn-alias set` command. The `-f` switch can be used to force or reassign an existing alias to a new WWPN. The new WWPN may have multiple aliases, but only one alias can be assigned to a single WWPN.

To remove an alias, an administrator may use either the `-a` switch to remove one or more particular aliases or the `-w` switch to remove all aliases for a given WWPN with the new `fcpx wwpn-alias remove` command.

To verify all aliases, an administrator may use the `fcpx wwpn-alias show` command. The administrator can then limit the alias return by requesting to see only a particular alias with the `-a` switch or a particular WWPN with the `-w` switch.



## Alias Rules

- A storage system can have up to 1,024 aliases
- An alias can have the following characters: A-Z, a-z, 0-9, '\_', '-', '.', '{', '}' and no spaces
- Many aliases may be associated with a single WWPN, but each alias will be assigned to only one WWPN
- Use `fcpx wwpn-alias help subcommand` for more information on the subcommand

© 2010 NetApp, Inc. All rights reserved.

## ALIAS RULES

The following rules apply to WWPN aliases:

- A storage system can have up to 1,024 aliases.
- An alias can have the following characters: A-Z, a-z, 0-9, '\_', '-', '.', '{', and '}', but no spaces.
- Many aliases may be associated with a single WWPN, but each alias will be assigned to only one WWPN.

Use the `fcpx wwpn-alias help <subcommand>` function for more information on a particular subcommand.





## WWPN Aliases Example

```
system> fcp wwpn-alias set WIN1-FC
10:00:00:00:c9:6b:77:b4

system> fcp wwpn-alias show

WWPN Alias
---- -
10:00:00:00:c9:6b:77:b4 WIN1-FC

system> fcp show initiators
Initiators connected on adapter 0c:
Portname Group
----- -
10:00:00:00:c9:6b:77:b4
WWPN Alias(es): WIN1-FC
```

© 2010 NetApp, Inc. All rights reserved.

### WWPN ALIASES EXAMPLE



# Windows Discovery of Targets

- Verify connectivity on Windows Server 2008

The screenshot shows the Storage Explorer application window. On the left, the tree view shows 'Storage Explorer (DEV05S2)' expanded, with 'Fibre Channel Fabrics' selected, and 'Fabric 10:00:00:05:1e:02' selected. A blue arrow points from the text 'Select the Brocade® fabric' to this fabric. The main pane displays a table of storage systems and their WWPNs. A blue box highlights the 'Attached Port WWN' column, with the text 'Storage systems' WWPNs show up' next to it. Another blue box highlights the 'Attached Node' column, with the text 'Windows WWPNs show up' next to it. The table has columns: Switch Port, Type, State, Port WWN, Attached Port WWN, and Attached Node. The 'Attached Port WWN' column shows values like '50:0a:09:81:96:88:37:5d' and '10:00:00:00:c9:6b:77:b3'. The 'Attached Node' column shows values like 'Other' and '111-00308-10:...'. The 'Actions' pane on the right shows 'Brocade Communicati...' expanded, with options like 'Manage telnet://10...', 'Manage http://10.25...', 'Manage telnet://10...', and 'Manage http://10.25...'.

| Switch Port | Type    | State   | Port WWN                | Attached Port WWN       | Attached Node    |
|-------------|---------|---------|-------------------------|-------------------------|------------------|
| Port 1      | F_Port  | Online  | 20:00:00:05:1e:02:99:c4 | 50:0a:09:81:96:88:37:5d | Other            |
| Port 2      | F_Port  | Online  | 20:01:00:05:1e:02:99:c4 | 50:0a:09:82:96:88:37:5d | Other            |
| Port 3      | Unkn... | Offline | 20:02:00:05:1e:02:99:c4 |                         |                  |
| Port 4      | Unkn... | Offline | 20:03:00:05:1e:02:99:c4 |                         |                  |
| Port 5      | Unkn... | Offline | 20:04:00:05:1e:02:99:c4 |                         |                  |
| Port 6      | Unkn... | Offline | 20:05:00:05:1e:02:99:c4 |                         |                  |
| Port 7      | F_Port  | Online  | 20:06:00:05:1e:02:99:c4 | 10:00:00:00:c9:6b:77:b3 | 111-00308-10:... |
| Port 8      | F_Port  | Online  | 20:07:00:05:1e:02:99:c4 | 10:00:00:00:c9:6b:77:b4 | 111-00308-10:... |
| Port 9      | Unkn... | Offline | 20:08:00:05:1e:02:99:c4 |                         |                  |
| Port 10     | Unkn... | Offline | 20:09:00:05:1e:02:99:c4 |                         |                  |
| Port 11     | Unkn... | Offline | 20:0a:00:05:1e:02:99:c4 |                         |                  |
| Port 12     | Unkn... | Offline | 20:0b:00:05:1e:02:99:c4 |                         |                  |
| Port 13     | Unkn... | Offline | 20:0c:00:05:1e:02:99:c4 |                         |                  |
| Port 14     | Unkn... | Offline | 20:0d:00:05:1e:02:99:c4 |                         |                  |
| Port 15     | Unkn... | Offline | 20:0e:00:05:1e:02:99:c4 |                         |                  |
| Port 16     | Unkn... | Offline | 20:0f:00:05:1e:02:99:c4 |                         |                  |

© 2010 NetApp, Inc. All rights reserved.

## WINDOWS DISCOVERY OF TARGETS



## Binding

- Binding or mapping in FC SAN is the process of associating an OS device name to a target's worldwide port name
- Persistent binding in FC SAN is the process of ensuring that the same binding occurs even after a host OS reboot

**NOTE:** Binding is done automatically in most modern OS; therefore, it does not need to be manually configured

© 2010 NetApp, Inc. All rights reserved.

## BINDING



## Window/Emulex Binding

The screenshot shows the HBAnyware interface with several annotations:

- Select**: Points to the 'General' tab in the 'Target Mapping' section.
- From Device Properties ...**: Points to the 'NETAPP LUN SCSI Disk Device Properties' window.
- Current Bindings**: Points to the 'Current Bindings' section in the 'General' tab.
- Current SCSI IDs**: Points to the 'SCSI ID' column in the 'Current Bindings' table.
- Do not use persistent binding**: Points to the 'Persistent Binding Configuration' section.
- SCSI ID = Bus Number, Target ID**: Points to the 'Location' field in the 'NETAPP LUN SCSI Disk Device Properties' window.

The 'Current Bindings' table shows the following data:

| W/WPN                   | W/WNN                   | D_ID  | SCSI ID | Type |
|-------------------------|-------------------------|-------|---------|------|
| 50:0A:09:81:86:F7:C7:86 | 50:0A:09:80:86:F7:C7:86 | 11000 | (0, 0)  | Auto |
| 50:0A:09:81:86:F7:C7:86 | 50:0A:09:80:86:F7:C7:86 | 11200 | (0, 1)  | Auto |

The 'Persistent Binding Configuration' section shows the following data:

| Target W/WPN            | SCSI ID |
|-------------------------|---------|
| 50:0A:09:81:86:F7:C7:86 |         |
| 50:0A:09:81:96:F7:C7:86 |         |

© 2010 NetApp, Inc. All rights reserved.

## WINDOW/EMULEX BINDING

There are no native Windows Server 2003 or 2008 tools for verifying binding of an initiator and one or more targets. Therefore, an administrator must use third-party tools such as HBAnyware from Emulex.

**NOTE:** FC-persistent binding is not supported by NetApp.



## Command-Line Interface for Server Core

- Command-line interface commands for FC management is available for the HBA vendors
  - Emulex: `hbacmd`
  - QLogic: `scli`
- Commands available to:
  - Verify connectivity
  - Interrogate the fabric
  - Manage bindings
  - Verify configuration
  - Administrate VPORTs

© 2010 NetApp, Inc. All rights reserved.

### COMMAND-LINE INTERFACE FOR SERVER CORE




## Multipath I/O

© 2010 NetApp, Inc. All rights reserved.

### MULTIPATHING I/O



## Multiple Paths to a LUN



© 2010 NetApp, Inc. All rights reserved.

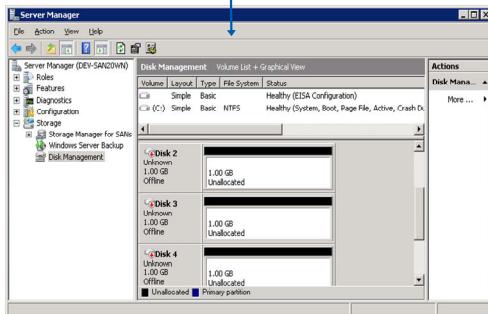
## MULTIPLE PATHS TO A LUN



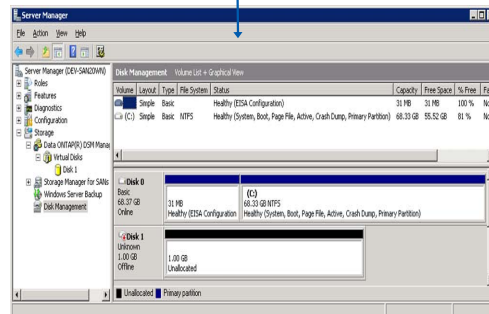
## Multiple Paths

- When multiple paths are present to a LUN, the same LUN would appear multiple times
  - The same LUN would appear a single instance for each path available

Without multipathing software



With multipathing software



© 2010 NetApp, Inc. All rights reserved.

## MULTIPLE PATHS





## Optimized or Non-Optimized Paths

- Not all paths are necessarily equal
  - Optimized or primary or favored = active path between initiator and target
    - Same latency level
  - Non-optimized or secondary or unfavored = inactive path between initiator and target
    - Latency differs between path

© 2010 NetApp, Inc. All rights reserved.

### OPTIMIZED OR NON-OPTIMIZED PATHS




## Multipath Access

### ■ Symmetric

- All paths are favored or optimized

### ■ Asymmetric

- Only certain paths are favored or optimized



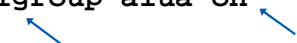
© 2010 NetApp, Inc. All rights reserved.

## MULTIPATH ACCESS



## Windows Multipath FC Implementation

- Storage management software available for Windows:
  - Native Windows Disk Management
  - Veritas™ Storage Foundation
- When possible, use the NetApp DSM and do not enable ALUA on the igroup  
`system> igroup set my_igroup alua off`
- Native Disk Management in W2K8 will use ALUA by default with Microsoft® native Device Specific Module (DSM); enable ALUA on the igroup  
`system> igroup set my_igroup alua on`  



Pre-defined igroup    Default is off
- This course focuses on native Disk Management, NetApp DSM, and Emulex HBAs

© 2010 NetApp, Inc. All rights reserved.

## WINDOWS MULTIPATH FC IMPLEMENTATION



## ALUA Output

### ■ Example:

```
system> igroup show -v
iWIN1_fcp (FCP) (ostype: windows):
 10:00:00:00:c9:6b:77:b3 (logged in on: 0b, 0d)
 WWPN Alias(es): dev05s1-fc1
 ↖ No ALUA
iWIN2_fcp (FCP) (ostype: windows):
 10:00:00:00:c9:6b:77:b4 (logged in on: 0d, 0b)
 WWPN Alias(es): dev05s2-fc1
ALUA: YES ↖
 ALUA enabled
```

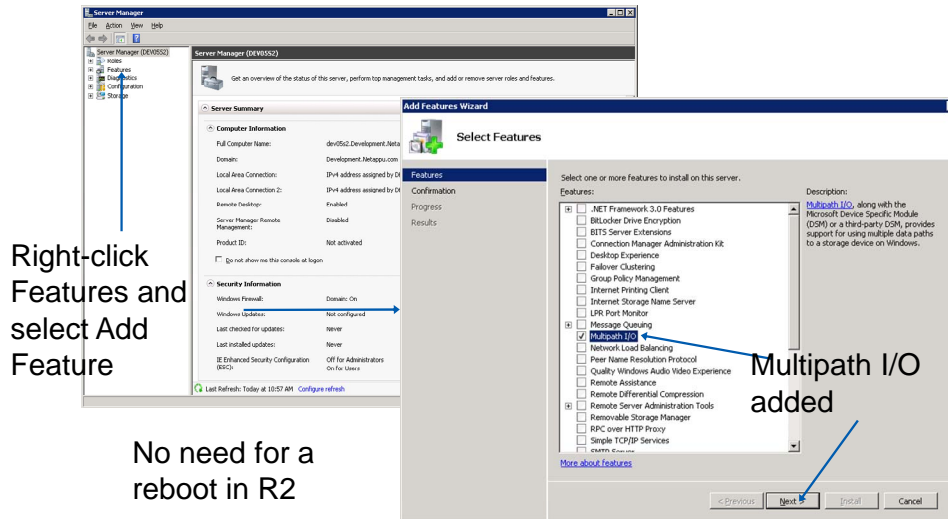
© 2010 NetApp, Inc. All rights reserved.

## ALUA OUTPUT



# Windows Native Multipath I/O

- Windows Server 2008 can be configured to support Multipath I/O



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS NATIVE MULTIPATH I/O



## Windows Native Multipath I/O (Cont.)

- Out of the box, Windows multipath I/O comes with a generic DSM
- NetApp recommends installing the NetApp Data ONTAP DSM
  - Verify with Interoperability Matrix for the recommended version

© 2010 NetApp, Inc. All rights reserved.

## WINDOWS NATIVE MULTIPATH I/O (CONT.)



## NetApp Data ONTAP DSM 3.3.1

### ■ Features:

- Supports Windows Server 2008 R2
- Supports multiple load-balancing policies
- Support for claiming iSCSI LUNs
- Coexists with other DSMs
- Multiprotocol LUN support (simultaneous iSCSI and FC paths to the same LUN)

### ■ Requirements:

- Windows Server 2003 or 2008 or 2008 R2
- Data ONTAP 7.2.2+

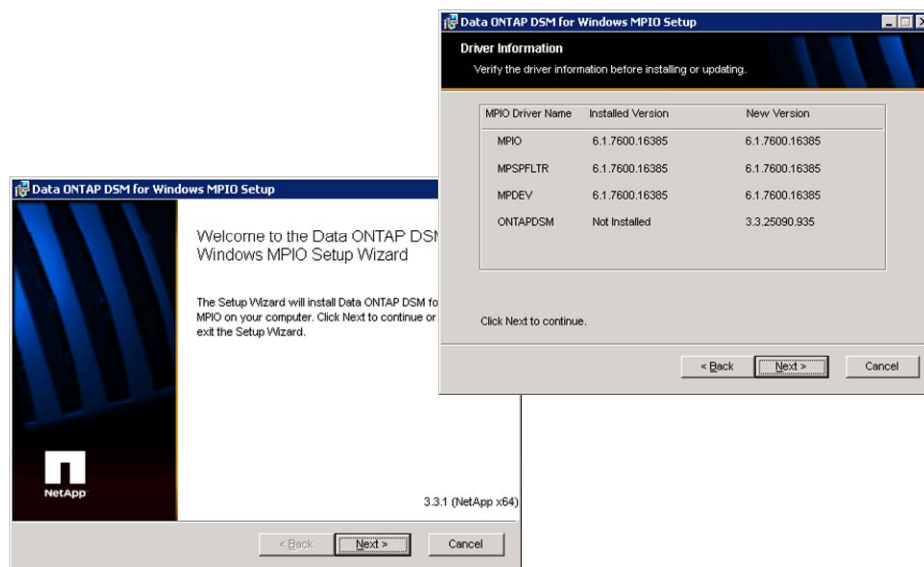
© 2010 NetApp, Inc. All rights reserved.

## NETAPP DATA ONTAP DSM 3.3.1

The Data ONTAP DSM for Windows MPIO is a Device Specific Module (DSM) that works with the Microsoft Windows MPIO drivers (mpdev.sys, mpio.sys, and mpspltr.sys) to manage multiple paths between NetApp & IBM® N series storage systems and Windows host computers. The DSM includes the storage-system-specific intelligence needed to correctly identify paths and to manage path failure and recovery.



# Install NetApp DSM



© 2010 NetApp, Inc. All rights reserved.

## INSTALL NETAPP DSM





## Red Hat

© 2010 NetApp, Inc. All rights reserved.

**RED HAT**



## Red Hat as an FC Initiator

- NetApp has supported Red Hat as an FC initiator OS since Red Hat 3
- Red Hat 5.3 has many advantages over previous versions:
  - Updated FC and iSCSI drivers
  - Better scalability
- Red Hat must be properly configured for FC connectivity

© 2010 NetApp, Inc. All rights reserved.

### RED HAT AS AN FC INITIATOR



## Red Hat 5.3 Design and Installation

1. Verify host operating system releases, required patches, and NetApp Linux® Host Utility Kit with Interoperability Matrix:
  - Use `/etc/redhat-release` and `uname -a` to verify Red Hat version
  - Interoperability Matrix can be found on the NOW site
2. Install compatible host bus adapters (HBAs).
3. Install and configure required HBA drivers and utilities if needed.
4. Verify an HBA:
  - All HBA Types: `lspci`
  - Emulex: Use `/usr/sbin/lpfc/lputil` or HBAnyware
  - QLogic: Use `/usr/bin/scli` or SANsurfer

© 2010 NetApp, Inc. All rights reserved.

### RED HAT 5.3 DESIGN AND INSTALLATION



## Red Hat/Emulex Implementation

After installation, to configure a Red Hat/Emulex implementation:

1. Identify the correct HBA driver
2. Verify that Red Hat has correctly identified the HBA and loaded the driver
3. Identify the WWNN and WWPN on the host
4. Verify connectivity between the initiator and target

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/EMULEX IMPLEMENTATION



## Red Hat/Emulex Implementation (Cont.)

- Verify that Red Hat has identified the HBA(s)

```
cd /sys/class/scsi_host
```

```
ls
```

```
host0 host1 host2
```

Emulex HBA(s) or in this case  
one HBA with two ports

- Identify the driver associated with the HBA(s)

```
cd host1
```

```
cat fwrev
```

```
ls
```

```
2.72A2 (B3F2.72A2), sli-3
```

```
... fwrev
```

```
... modeldesc
```

```
... modelname
```

```
... npiv_info
```

```
... portnum
```

```
... serialnum
```

```
... lpfc_drvr_version
```

NetApp part number  
of NetApp sold HBAs

```
cat lpfc_drvr_version
```

```
Emulex LightPulse
```

```
Fibre Channel SCSI
```

```
driver 8.2.0.33.3p
```

Current installed driver

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/EMULEX IMPLEMENTATION (CONT.)



## Red Hat/Emulex Implementation (Cont.)

- *Interoperability Matrix* requires an update to driver based upon existing firmware:

```
tar xzf lpfc_2.6_driver_kit-8.2.0.39-1.tar.gz
cd lpfc_2.6_driver_kit-8.2.0.39-1
./lpfc-install ...
reboot
```

The kernel build  
number of Linux

- The driver can be found here:

```
ls /lib/modules/2.6.18-128.el5/kernel/drivers/scsi/lpfc
lpfc.ko
```

- To verify the driver when loaded:

```
modprobe -c | grep lpfc
```

- To load the driver (if needed):

```
modprobe -v lpfc
```

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/EMULEX IMPLEMENTATION (CONT.)



## Red Hat/Emulex Implementation (Cont.)

- Install compatible Linux Host Utility Kit (HUK)
  - Provides Perl scripts to configure and tune Red Hat
    - Example: `sanlun` application to manage LUNs from Red Hat
  - HUK requires packages to be installed:
    - `libnl.so = libnl-1.0-0.10.pre5.5.x86_64.rpm` ← Use 64-bit version
    - `libnl.so = libnl-1.0-0.10.pre5.5.i386.rpm` ← Use 32-bit version
    - `HBAware = elxlinuxapps-4.0a31-8.2.0.39-1-1.tar`

**NOTE:** You must install these libraries in this order

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/EMULEX IMPLEMENTATION (CONT.)



## Red Hat/Emulex Implementation (Cont.)

### ■ Install HBAnyware:

```
tar xvf elxlinuxapps-4.0a31-8.2.0.39-1-1.tar
cd elxlinuxapps-4.0a31-8.2.0.39-1-1
./install
```

Select desired mode of operation for HBAnyware

- 1 Local Mode : HBA's on this Platform can be managed by HBAnyware clients on this Platform Only.
- 2 Managed Mode: HBA's on this Platform can be managed by local or remote HBAnyware clients.
- 3 Remote Mode : Same as '2' plus HBAnyware clients on this Platform can manage local and remote HBA's.

...

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/EMULEX IMPLEMENTATION (CONT.)





## Red Hat/Emulex Implementation (Cont.)

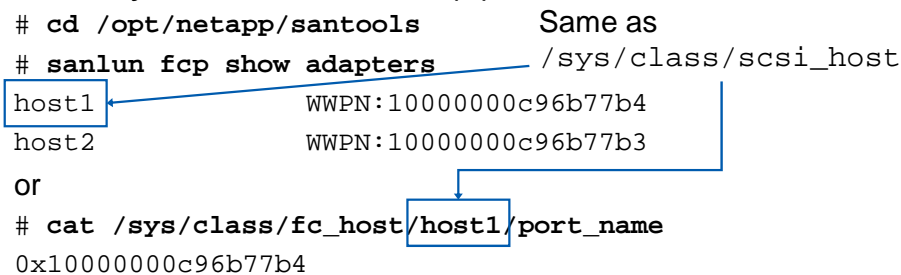
### ■ Install HUK:

```
tar xzf netapp_linux_host_utilities_5_0.tar.gz
cd netapp_linux_host_utilities_5_0
./install
cd /opt/netapp/santools
./san_version
NetApp Linux Host Utilities version 5.0
```

### ■ Identify WWPN of HBA(s):

```
cd /opt/netapp/santools Same as
sanlun fcp show adapters /sys/class/scsi_host
host1 WWPN:10000000c96b77b4
host2 WWPN:10000000c96b77b3

or
cat /sys/class/fc_host/host1/port_name
0x10000000c96b77b4
```

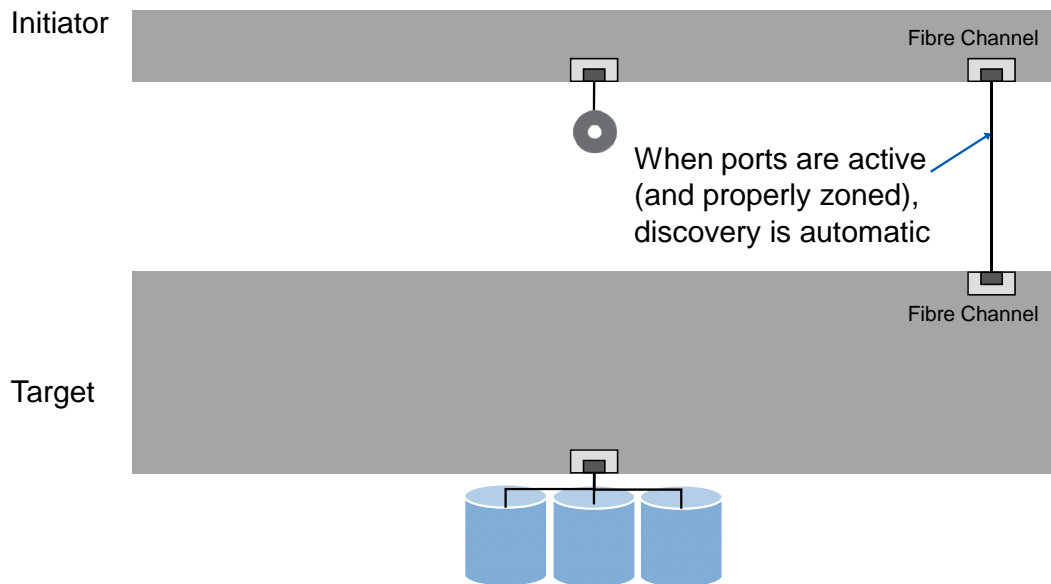


© 2010 NetApp, Inc. All rights reserved.

## RED HAT/EMULEX IMPLEMENTATION (CONT.)



## Discovery and Session Creation



© 2010 NetApp, Inc. All rights reserved.

## DISCOVERY AND SESSION CREATION



## Data ONTAP Discovery of Initiators

- Verify connectivity from the storage system:

```
system> fcp show initiators
```

```
Initiators connected on adapter 0c:
```

| Portname                | Group           |
|-------------------------|-----------------|
| -----                   | -----           |
| 10:00:00:00:c9:6b:77:b3 |                 |
| 10:00:00:00:c9:6b:77:b4 | ← Red Hat WWPNS |

**NOTE:** For convenience, alias the Red Hat WWPNS

© 2010 NetApp, Inc. All rights reserved.

### DATA ONTAP DISCOVERY OF INITIATORS



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe multiple path implementation with Fibre Channel (FC) connectivity
- Describe how to configure FC ports on Windows, Red Hat, and NetApp systems
- Describe commands and utilities to identify the worldwide node name (WWNN) and worldwide port name (WWPN) on Windows, Red Hat, and NetApp systems

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 13: FC Connectivity  
Estimated Time: 0 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- Using Data ONTAP 8.0 7-Mode, which cfmode(s) can you configure a new storage system?
- What two FC multipathing software stacks are available for Windows?
- Which three software libraries are required for the `sanlun` command function in the Linux Host Utilities Kit?
- Administrators should use what tool, provided by NetApp and available on the NOW site, to verify their Red Hat configuration?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# iSCSI Connectivity

Module 14  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## ISCSI CONNECTIVITY





## Module Objectives

By the end of this module, you should be able to:

- Describe multiple path implementation with iSCSI connectivity
- Configure network ports on Windows®, Red Hat®, and NetApp® systems
- Identify the worldwide node name (WWNN) on Windows, Red Hat, and NetApp systems
- Set up and verify multiple path iSCSI connectivity between Windows, Red Hat, and NetApp systems

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Data ONTAP

© 2010 NetApp, Inc. All rights reserved.

## DATA ONTAP



## Data ONTAP as an iSCSI Target

- Data ONTAP® 6.4 and later has support for iSCSI
- Data ONTAP features:
  - Built-in iSCSI service
  - Simple LUN creation and management
- Data ONTAP must be properly configured for iSCSI connectivity
  1. Configure IP interfaces
  2. Configure iSCSI services
  3. Configure the iSCSI interfaces
  4. Identify the worldwide node name (WWNN)

© 2010 NetApp, Inc. All rights reserved.

### DATA ONTAP AS AN ISCSI TARGET



## Configuring Interfaces

1. List the available interfaces:
  - `ifconfig -a`
2. Take an interface offline:
  - `ifconfig interface_name down`
3. Configure the interface:
  - `ifconfig interface_name ipaddress`
4. Bring an interface online:
  - `ifconfig interface_name up`

**NOTE:** Virtual interfaces (interface groups) may also be configured to be used with the iSCSI service

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURING INTERFACES



## Configuring iSCSI Services in Data ONTAP

1. Verify the iSCSI service is running:
  - `iscsi status`
2. Verify iSCSI is licensed (license it if needed):
  - `license`
  - `license add XXXXXX`
3. Start the iSCSI service:
  - `iscsi start`

© 2010 NetApp, Inc. All rights reserved.

### CONFIGURING ISCSI SERVICES IN DATA ONTAP



## Verify Interfaces

- Verify the interface is enabled for iSCSI:
  - `iscsi interface show`
  - By default, all interfaces are enabled
- To enable the interface for iSCSI traffic:
  - `iscsi interface enable interface_name`
- To disable iSCSI traffic for a particular interface:
  - `iscsi interface disable interface_name`

© 2010 NetApp, Inc. All rights reserved.

## VERIFY INTERFACES



## Interface Access List

- Administrators may force initiators to access a storage system through certain interfaces:
  - `iscsi interface accesslist add initiator_name {-a | interface_name}`
  - By default, all initiators may use any interface that is enabled for iSCSI traffic
  - To display the current access list, use:  
`iscsi interface accesslist show`
  - To remove an entry from the access list, use:  
`iscsi interface accesslist remove initiator_name {-a | interface_name}`

© 2010 NetApp, Inc. All rights reserved.

## INTERFACE ACCESS LIST

The `accesslist` feature controls initiator access to network interfaces. By default, an initiator does not have an access list and can access the storage system through any network interface. The administrator can create an access list using the `add` subcommand:

```
iscsi interface accesslist add initiator {-a | interface...}
```

This creates an access list for the initiator with a list of network interface names.

The specified initiator can only log in through the network interfaces in the access list.

If the specified initiator sends a `SendTargets` request, it will receive a list of target addresses. These target addresses are associated with only those network interfaces that are included in the access list.

An existing access list can be edited by means of the `add` and `remove` subcommands:

```
iscsi interface accesslist remove <initiator> [-a | <interface>...]
```

The `-a` parameter adds or removes all network interfaces. When the last network interface is removed from an access list, the access list itself is removed.



## Identifying WWNN in Data ONTAP

- To identify the WWNN:

```
system> iscsi nodename
```

```
iSCSI target nodename: iqn.1992-08.com.netapp:system
```

- Remember WWPNNs are not used within iSCSI

© 2010 NetApp, Inc. All rights reserved.

## IDENTIFYING WWNN IN DATA ONTAP





## Interfaces on the Storage System

- Verify interface status:

```
system> ifconfig -a

...
e0b: flags=108042<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
ether 00:a0:98:03:28:8f (auto-unknown-down) flowcontrol full

e0c: flags=108042<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
ether 00:a0:98:03:28:8f (auto-unknown-down) flowcontrol full

...
```

- Enable the interface:

```
system> ifconfig e0b 10.254.134.75 up
...[system: netif.linkUp:info]: Ethernet e0b: Link up.
system> ifconfig e0c 10.254.134.81 up
...[system: netif.linkUp:info]: Ethernet e0c: Link up.
```

- Ensure that the iSCSI service may use the interface:

```
system> iscsi interface enable e0b e0c
```

**NOTE:** The default is all interfaces are enabled for iSCSI traffic

© 2010 NetApp, Inc. All rights reserved.

## INTERFACES ON THE STORAGE SYSTEM



## Windows

© 2010 NetApp, Inc. All rights reserved.

## WINDOWS



## Windows as an iSCSI Initiator

- NetApp has supported Windows as an iSCSI initiator OS since Windows 2000 Server
  - Always check the Interoperability Matrix Tool for current supported operating systems
- Windows Server 2008 has many advantages over previous versions:
  - New tools such as Storage Explorer and Storage Manager for SANs
- Windows Server 2008 R2 has many advantages over Windows Server 2008:
  - User interface enhancement and redesign
  - iSCSI boot support for up to 32 paths
- Windows must be properly configured for iSCSI connectivity over a standard network interface

© 2010 NetApp, Inc. All rights reserved.

## WINDOWS AS AN ISCSI INITIATOR



## Windows iSCSI Design and Installation

1. Verify host operating system releases, required patches, and NetApp iSCSI Host Utility Kit with the Interoperability Matrix Tool and FC and iSCSI configuration guides:
  - Use System Properties Dialog
2. Identify and verify a network interface is properly configured or install a supported iSCSI HBA or TOE.
3. In Windows Server 2008, the software initiator is preinstalled.
4. Install compatible NetApp Windows Host Utility Kit:
  - Provides Perl scripts to monitor and diagnose iSCSI on Windows

© 2010 NetApp, Inc. All rights reserved.

### WINDOWS ISCSI DESIGN AND INSTALLATION



## Windows/NIC Implementation

After installation, to configure a Windows standard NIC software initiator implementation:

1. Identify the local network interface(s) to use
2. Verify iSCSI Initiator driver is enabled and the service is started
3. Identify the WWNN for the local Windows host
4. Identify which method of discovery to use and enter the storage system's portal IP address or iSNS address
5. Configure authentication security if necessary
6. Verify discovery and log on to the storage system

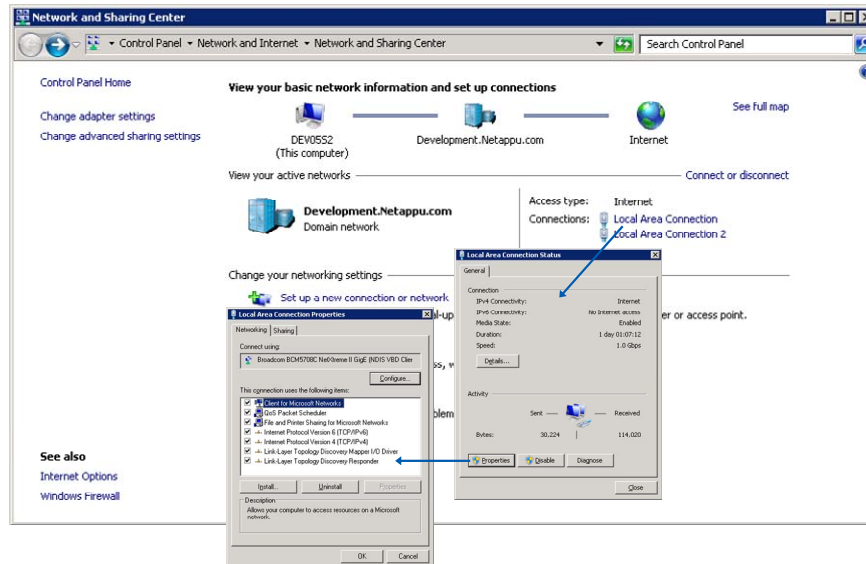
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION



# Windows/NIC Implementation (Cont.)

## 1. Identify and configure the local interfaces



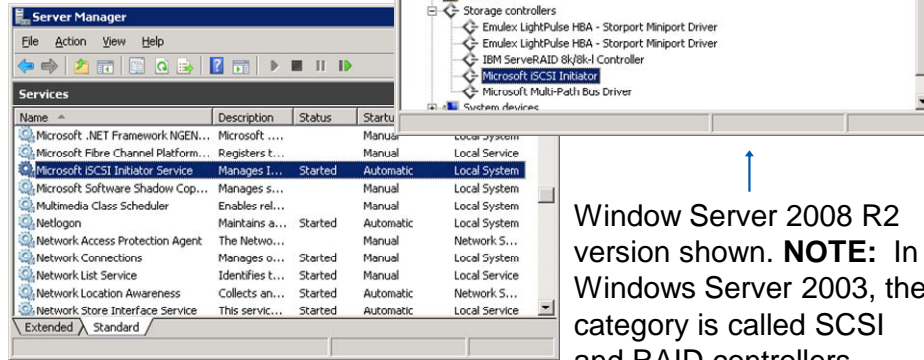
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)



## Windows/NIC Implementation (Cont.)

2. Verify iSCSI Initiator driver is enabled and verify iSCSI Initiator service is started



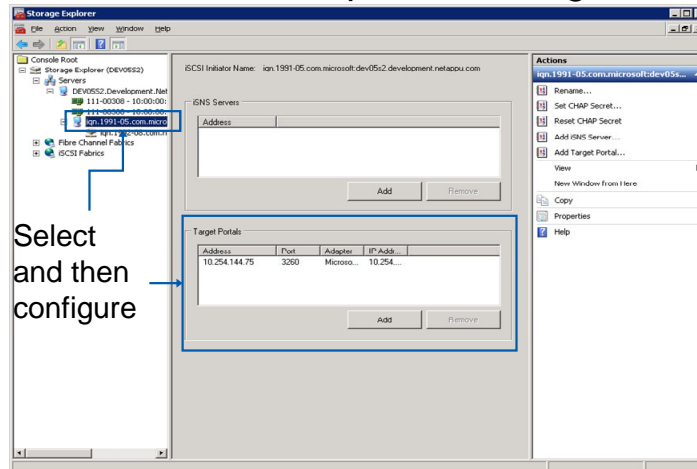
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)



## Windows/NIC Implementation (Cont.)

- iSCSI initiator may be configured through:
  - Storage Explorer
  - iSCSI Initiator Properties Dialog



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)

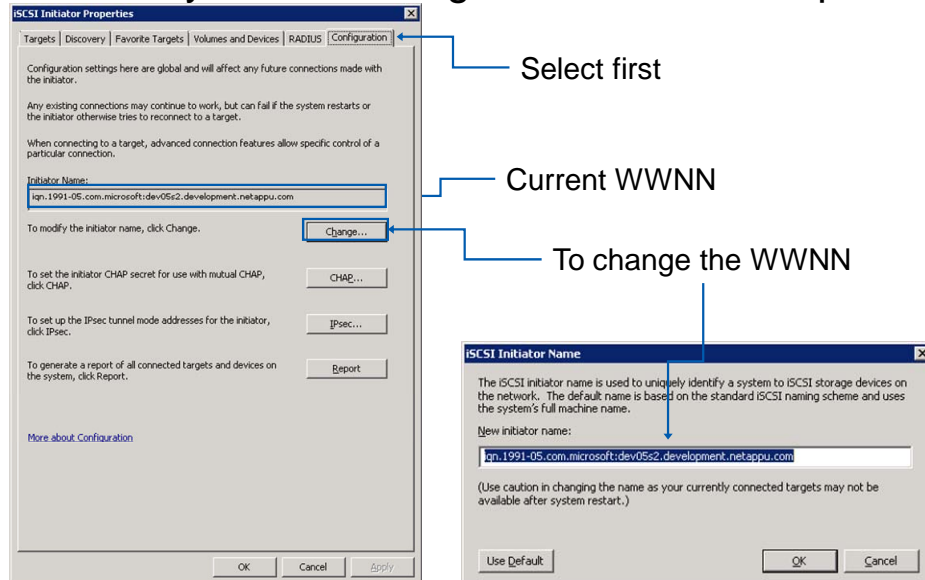
The Microsoft® iSCSI software initiator may be configured through either Storage Explorer in Windows Server 2008 or the iSCSI Initiator Properties dialog in Windows Server 2003 or 2008.





## Windows/NIC Implementation (Cont.)

### 3. Identify WWNN using iSCSI Initiator Properties



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)

The local host WWNN appears on the Configuration tab of the iSCSI Initiator Properties dialog. It generally does not need to be changed.



## WWNN Spoofing


- iSCSI node names are:
  - Spoof-able
  - Sniff-able
  - Can be attacked
- NetApp recommends using authentication methods and other security techniques (discussed later)

© 2010 NetApp, Inc. All rights reserved.

## WWNN SPOOFING



## Discovery



© 2010 NetApp, Inc. All rights reserved.

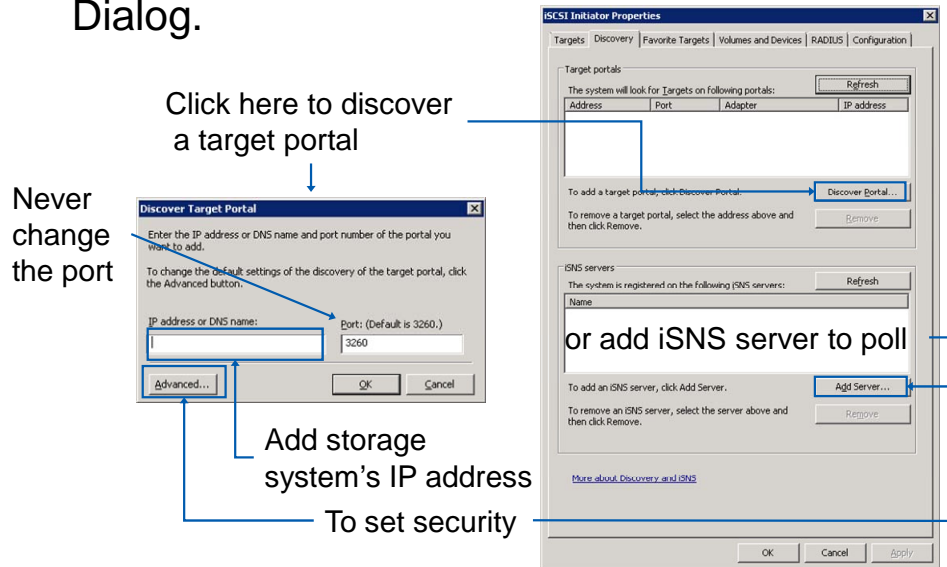
## DISCOVERY

Unlike FC discovery, iSCSI discovery is not automatic.



## Windows/NIC Implementation (Cont.)

### 4. Discovery with iSCSI Initiator Properties Dialog.



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)

From the Discovery tab on the iSCSI Initiator Properties dialog, an administrator may set the method of discovery of targets. Discovery may be accomplished either through add-a-target's portal address or by polling an Internet Storage Name Service (iSNS) server.

When adding a target portal address, if authentication is required, an administrator will set this by clicking the Advanced button. Next, iSCSI authentication is discussed in detail.



## iSCSI Authentication in Windows

5. To increase security, iSCSI may be configured to require authentication.

– Authentication methods:

■ CHAP

- Unidirectional: targets will authenticate initiators
- Bidirectional: initiators and targets will authenticate each other

■ RADIUS

- IPsec could also be used to increase security
- This course will discuss using CHAP authentication, but will not use it in the exercise

© 2010 NetApp, Inc. All rights reserved.

## ISCSI AUTHENTICATION IN WINDOWS

There are two methods of authenticating systems in iSCSI:

### CHALLENGE HANDSHAKE AUTHENTICATION PROTOCOL (CHAP)

CHAP requires a known secret that is shared by both target and initiator. There are four basic steps of unidirectional authentication:

After the completion of the link establishment phases, the target sends a “challenge” message to the initiator.

The initiator responds with a one-way hash function of the shared secret.

The target checks the response against its own calculation of the expected hash value.

At random intervals, the target will send a new challenge to the initiator and repeat Steps 1, 2, and 3.

**NOTE:** In bidirectional authentication the process is also implemented in the reverse.

### REMOTE AUTHENTICATION DIAL-IN USER SERVICE (RADIUS)

RADIUS is a networking protocol that uses access servers to provide centralized management of access to large networks.

Proper authentication resists man-in-the-middle attacks as well as other attacks.

IPsec can also be used to increase security.

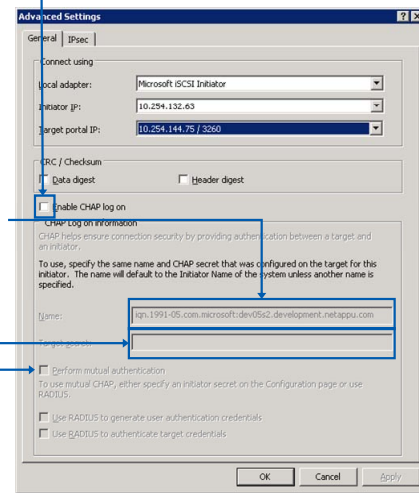


# iSCSI Unidirectional CHAP Authentication

- Configure the discovery to use the CHAP authentication

```
system> iscsi security add
-i iqn.1991-05.com.microsoft:win
-s CHAP
-n iqn.1991-05.com.microsoft:win
-p thisismysecret
```

Check to configure



To configure bidirectional, check here and then...

© 2010 NetApp, Inc. All rights reserved.

## ISCSI UNIDIRECTIONAL CHAP AUTHENTICATION

To configure unidirectional (inbound) CHAP authentication for the Microsoft Software Initiator, enter the `iscsi security add` command on the storage system and enter the user name and shared secret as appropriate. By convention, the user name is generally the same as the WWNN.


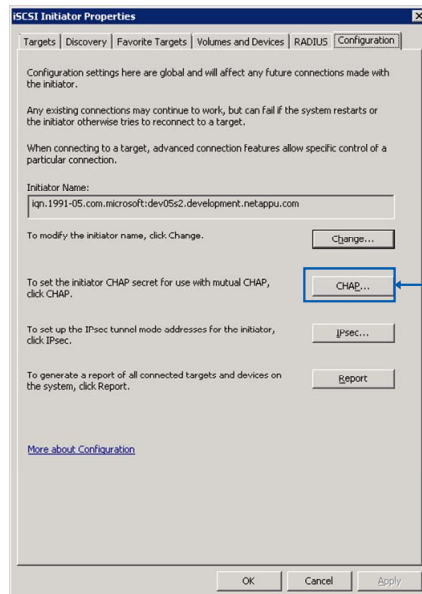
The Microsoft iSCSI Software Initiator requires both the initiator and target CHAP passwords to be at least 12 bytes if IPsec encryption is not being used. The maximum password length is 16 bytes regardless of whether IPsec is used.

**NOTE:** Data ONTAP provides an `iscsi security generate` command that creates a random 128-bit key that may, in some cases, be used as the password.



# iSCSI Bidirectional CHAP Authentication

## ■ Set Windows CHAP secret



Set CHAP secret  
from switch -o

## ■ On the storage system:

```
system> iscsi security add
-i iqn.1991-05.com.microsoft:win
-s CHAP
-n iqn.1991-05.com.microsoft:win
-p thisismysecret
-o thisismysecret2
-m iqn.1991-05.com.microsoft:win
```

© 2010 NetApp, Inc. All rights reserved.

## ISCSI BIDIRECTIONAL CHAP AUTHENTICATION

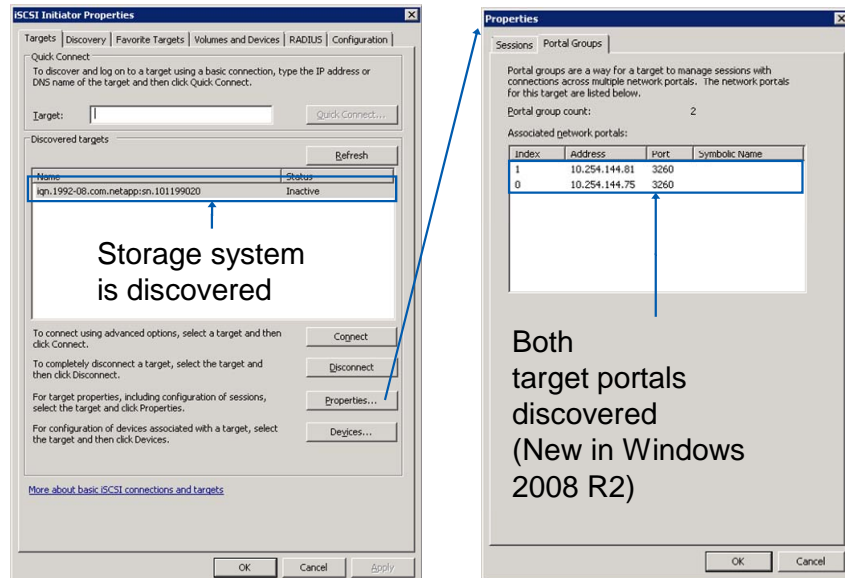
To add bidirectional (inbound and outbound) authentication to the Microsoft software initiator, check the mutual authentication on the Advanced Setting (see previous page) and then add the CHAP secret from switch -o of the `iscsi security add` command.

**NOTE:** The user name is the same as the WWNN but the shared secret is different. This is because the user name and password combination cannot be the same for inbound and outbound settings on a storage system.



## Windows/NIC Implementation (Cont.)

### 6. Discovered in iSCSI Initiator Properties Dialog



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)





## Binding

- iSCSI binding or logging on is the process of creating a session between an initiator and a target
- Persistent binding ensures that an initiator binds to a target after a reboot of the initiator OS

© 2010 NetApp, Inc. All rights reserved.

## BINDING



## Windows/NIC Implementation (Cont.)

### ■ Targets in the iSCSI Initiator Properties Dialog

The screenshot shows the 'iSCSI Initiator Properties' dialog box with the 'Targets' tab selected. A table lists discovered targets, including 'iqn.1992-08.com.netapp:sn.101199020'. A blue arrow points to the 'Quick Connect' button, with the annotation 'Storage system is discovered'. Another blue arrow points to the 'Connect' button, with the annotation 'Click here to connect'. A third blue arrow points to the 'Advanced...' button in the 'Connect To Target' sub-dialog, with the annotation 'To change the interface that is used with which to connect'. The 'Connect To Target' sub-dialog is also shown, with a blue arrow pointing to the 'Add this connection to the list of Favorite Targets' checkbox, with the annotation 'Best practice: check both'. The 'Advanced...' button in the sub-dialog is also highlighted with a blue arrow.

Storage system is discovered

Click here to connect

Best practice: check both

To change the interface that is used with which to connect

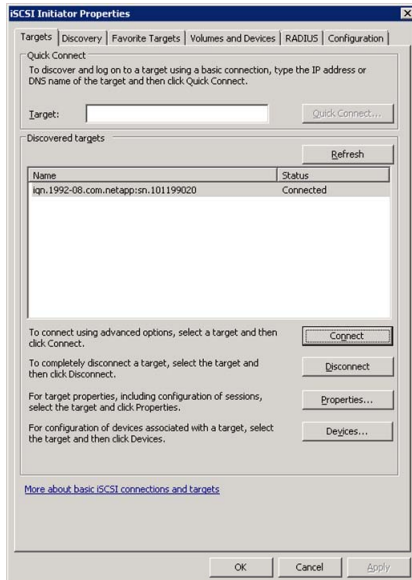
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)



## Windows/NIC Implementation (Cont.)

### ■ Connection in iSCSI Initiator Properties Dialog



NOTE: Console message appears

```
[system: iscsi.notice:notice]:
ISCSI: New session from initiator
iqn.1991-05.com.microsoft:dev05s2.
development.netappu.com at IP addr
10.254.144.75
```

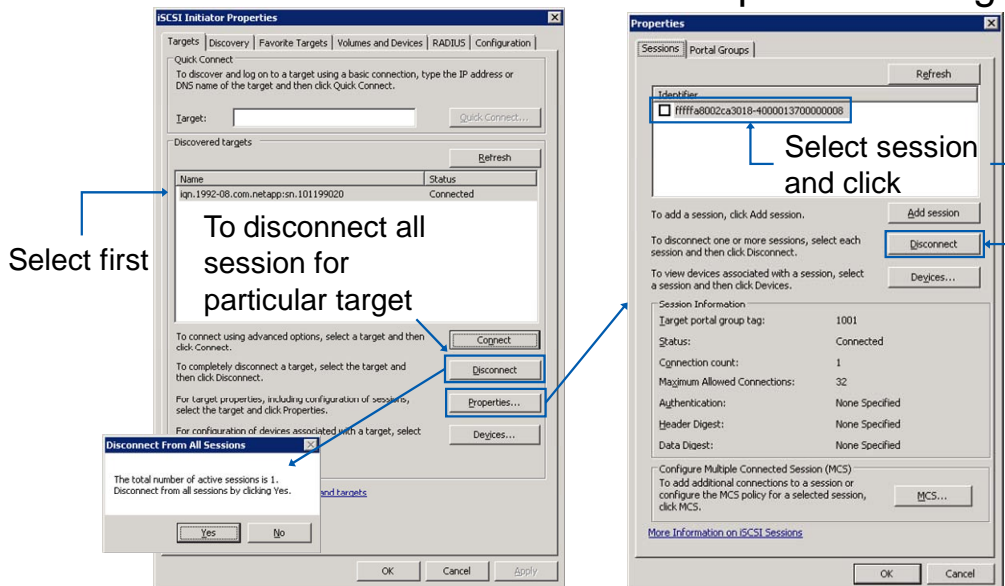
© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)



## Windows/NIC Implementation (Cont.)

### ■ Disconnect in iSCSI Initiator Properties Dialog



© 2010 NetApp, Inc. All rights reserved.

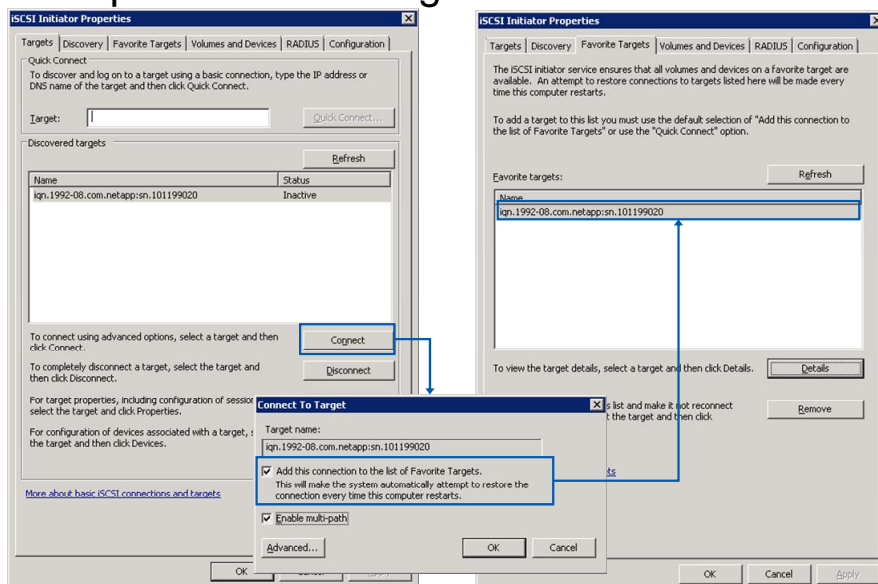
## WINDOWS/NIC IMPLEMENTATION (CONT.)

To disconnect, select the session from the target Properties dialog and click Disconnect.



## Windows/NIC Implementation (Cont.)

### ■ iSCSI persistent binding



© 2010 NetApp, Inc. All rights reserved.

## WINDOWS/NIC IMPLEMENTATION (CONT.)

While logging into a target, an administrator may optionally set the target as a favorite target that automatically logs in when the local host is booted.



## Red Hat

© 2010 NetApp, Inc. All rights reserved.

**RED HAT**



## Red Hat as an iSCSI Initiator

- NetApp has supported Red Hat as an iSCSI Initiator OS since Red Hat 4
- Red Hat 5.3 has many advantages over previous versions with:
  - Packages providing iSCSI device drivers and utilities
  - iSCSI software initiator for standard network interfaces
- Red Hat must be configured properly for iSCSI connectivity over a standard network interface

© 2010 NetApp, Inc. All rights reserved.

### RED HAT AS AN ISCSI INITIATOR



## Red Hat 5.3 Design and Installation

1. Verify host operating system releases, required patches, and NetApp Linux® Host Utility Kit with Interoperability Matrix:
  - Use `/etc/redhat-release` and `uname -a` to verify Red Hat version
  - Interoperability Matrix can be found on the NOW™ site
2. Install iSCSI software packages and patches:  
# `rpm -ivh iscsi-initiator-utils-6.2.0.868-0.18.el5.x86_64.rpm`
3. Identify and verify a network interface is properly configured.
4. Install compatible Linux Host Utility Kit:
  - Provides Perl scripts to configure and tune Red Hat for iSCSI
  - Example: `sanlun` application to manage LUNs from Red Hat

© 2010 NetApp, Inc. All rights reserved.

### RED HAT 5.3 DESIGN AND INSTALLATION





## Red Hat/NIC Implementation

After installation, to configure a Red Hat standard NIC software initiator implementation:

1. Identify the local network interface(s) to use
2. Verify the iSCSI service and WWNN for host
3. Configure authentication security if necessary
4. Identify which method of discovery to use and enter the storage system's portal IP address
5. Verify discovery and sessions with the target

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLMENTATION



## Red Hat/NIC Implementation (Cont.)

1. Using the software initiator, Red Hat supports iSCSI over standard network interface.

– Investigate and configure interfaces:

```
ifconfig -a
```

```
eth0
```

```
Link encap:Ethernet HWaddr 00:21:5E:6F:18:C4
```

```
inet addr:10.254.132.63 Bcast:10.254.135.25
```

```
Mask:255.255.252.0
```

```
inet6 addr: fe80::221:5eff:fe6f:18c4/64
```

```
Scope:Link UP BROADCAST RUNNING MULTICAST
```

```
MTU:1500 Metric:1
```

Unconfigured adapter; configure  
adapter with ifconfig and  
add to rc daemon

```
...
```

```
eth1
```

```
Link encap:Ethernet HWaddr 00:21:5E:6F:18:C6
```

```
BROADCAST MULTICAST MTU:1500 Metric:1 ...
```

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLMENTATION (CONT.)



## Red Hat/NIC Implementation (Cont.)

### 2. View the iSCSI Service and initiator node name.

- Start iSCSI service:

```
service iscsi start
```

- Verify the status of iSCSI service:

```
service iscsi status
```

```
iscsid (pid 5236 5235) is running...
```

- Identify the initiator's node name:

```
cat /etc/iscsi/initiatorname.iscsi
```

```
InitiatorName=ign.1994-05.com.redhat:rhel
```



Red Hat WWNN

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLMENTATION (CONT.)



# Discovery

Initiator


Ethernet



Discovery is not automatic

Target

Ethernet



© 2010 NetApp, Inc. All rights reserved.

## DISCOVERY



## iSCSI Authentication in Red Hat 5

3. To increase security, iSCSI may be configured to require authentication.
  - Authentication methods:
    - CHAP session authentication
      - Unidirectional: targets will authenticate initiators
      - Bidirectional: initiators and targets will authenticate each other
    - CHAP discovery (send targets) authentication
      - Unidirectional: targets will authenticate initiators
      - Bidirectional: initiators and targets will authenticate each other
  - This course will discuss using CHAP session authentication, but will not use it in the exercise

© 2010 NetApp, Inc. All rights reserved.

### ISCSI AUTHENTICATION IN RED HAT 5



## iSCSI Unidirectional CHAP Authentication

- Configure initiator's user name and password:

```
vi /etc/iscsi/iscsid.conf
```

```
...
```

```
#node.session.auth.username = username
#node.session.auth.password = thisismysecret
```

- To configure CHAP, enable CHAP authentication:

```
vi /etc/iscsi/iscsid.conf
```

```
...
```

```
#node.session.auth.authmethod = CHAP
```

Uncomment

```
service iscsi restart
```

- On the storage system, register the CHAP secret:

```
system> iscsi security add
-i iqn.1994-05.com.redhat:rhel
-s CHAP
```

```
-p thisismysecret
-n username
```

Same user name and  
password

© 2010 NetApp, Inc. All rights reserved.

## ISCSI UNIDIRECTIONAL CHAP AUTHENTICATION



## iSCSI Bidirectional CHAP Authentication

- Configure unidirectional CHAP and then configure the reverse direction:

```
system> iscsi security add
 -i iqn.1994-05.com.redhat:rhel
 -s CHAP
 -p thisismysecret
 -n username
 -o thisismysecret2
 -m username2
```

User name and password of inbound and outbound cannot be the same

- On Red Hat, register the storage system's CHAP secret:

```
vi /etc/iscsi/iscsid.conf
...
#node.session.auth.username = username2
#node.session.auth.password = thisismysecret2
```

Don't forgot to uncomment

- On Red Hat, restart the iSCSI service:

```
service iscsi restart
```

© 2010 NetApp, Inc. All rights reserved.

## ISCSI BIDIRECTIONAL CHAP AUTHENTICATION



## Red Hat/NIC Implementation

4. Discovery is possible through either:
- Static discovery
    - iSCSI targets added manually
  - Send-targets discovery
    - IP address of the target is added
    - Initiator communicates to target over port 3260
  - Internet Storage Name Service (iSNS)
    - Centralized management of discovery and configuration of iSCSI networks

This course will focus on the send-target discovery method

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLEMENTATION





## Red Hat/NIC Implementation (Cont.)

### ■ Set up iSCSI interface (for multiple paths):

```
iscsiadm -m iface -I iface0 --op=new
iscsiadm -m iface -I iface0 --op=update
 -n iface.net_ifacename -v eth0
```

Red Hat's eth0 and eth1 interface name

```
iscsiadm -m iface -I iface1 --op=new
iscsiadm -m iface -I iface1 --op=update
 -n iface.net_ifacename -v eth1
```

### ■ Verify iSCSI interfaces:

```
iscsiadm -m iface
iface0 tcp,default,eth0
iface1 tcp,default,eth1
ls /var/lib/iscsi/ifaces
iface0 iface1
```

Use vi or iscsiadm to configure

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLEMENTATION (CONT.)



## Red Hat/NIC Implementation (Cont.)

- Set up send targets discovery with interfaces:

```
iscsiadm -m discovery -t sendtargets -p
10.254.133.239 -I iface0 -I iface1
10.254.133.239:3260,1001 iqn.1992-08.com.netapp:system
10.254.133.239:3260,1001 iqn.1992-08.com.netapp:system
10.254.133.240:3260,1002 iqn.1992-08.com.netapp:system
10.254.133.240:3260,1002 iqn.1992-08.com.netapp:system
```

IP address of iSCSI-enabled  
interface on the storage system

Target portals discovered  
by each Red Hat interface

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLEMENTATION (CONT.)



## Red Hat/NIC Implementation (Cont.)

### 5. Explore the iSCSI targets discovered:

```
iscsiadm -m node --op=show
#BEGIN RECORD 2.0-868
node.name = iqn.1992-08.com.netapp:system
node.tpgt = 1002
node.startup = automatic
iface.hwaddress = 00:21:5E:6F:18:C6
iface.iscsi_ifacename = iface1
iface.net_ifacename = default
iface.transport_name = tcp
node.discovery_address = 10.254.133.239
...
#BEGIN RECORD 2.0-868
node.name = iqn.1992-08.com.netapp:system
node.tpgt = 1002
node.startup = automatic
iface.hwaddress = 00:21:5E:6F:18:C4
iface.iscsi_ifacename = iface0
...
```

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLEMENTATION (CONT.)



## Red Hat/NIC Implementation (Cont.)

### ■ Observe discovered targets:

```
iscsiadm -m node
10.254.133.239:3260,1001 iqn.1992-08.com.netapp:system
10.254.133.239:3260,1001 iqn.1992-08.com.netapp:system
10.254.133.240:3260,1002 iqn.1992-08.com.netapp:system
10.254.133.240:3260,1002 iqn.1992-08.com.netapp:system
```

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLMENTATION (CONT.)



## Red Hat/NIC Implementation (Cont.)

### ■ Create a session with discovered targets:

```
iscsiadm -m node -l
Logging in to [iface: iface1, target: iqn.1992-
08.com.netapp:system, portal: 10.254.133.240,3260]
Logging in to [iface: iface0, target: iqn.1992-
08.com.netapp:system, portal: 10.254.133.240,3260]
Logging in to [iface: iface1, target: iqn.1992-
08.com.netapp:system, portal: 10.254.133.239,3260]
Logging in to [iface: iface0, target: iqn.1992-
08.com.netapp:system, portal: 10.254.133.239,3260]
...
```

© 2010 NetApp, Inc. All rights reserved.

## RED HAT/NIC IMPLEMENTATION (CONT.)



## Red Hat iSCSI Sessions

### ■ View the current sessions:

```
iscsiadm -m session
tcp: [1] 10.254.133.240:3260,1002 iqn.1992-08.com.netapp:system
tcp: [2] 10.254.133.240:3260,1002 iqn.1992-08.com.netapp:system
tcp: [3] 10.254.133.239:3260,1001 iqn.1992-08.com.netapp:system
tcp: [4] 10.254.133.239:3260,1001 iqn.1992-08.com.netapp:system
```

### ■ View the sessions on the storage system:

```
system> iscsi session show
Session 30
Initiator Information
Initiator Name: iqn.1994-05.com.redhat:rhel
ISID: 00:02:3d:01:00:00
Initiator Alias: dev05s2.development.netappu.com

Session 31
Initiator Information
...
```

© 2010 NetApp, Inc. All rights reserved.

## RED HAT ISCSI SESSIONS



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe multiple path implementation with iSCSI connectivity
- Configure network ports on Windows, Red Hat, and NetApp systems
- Identify the worldwide node name (WWNN) on Windows, Red Hat, and NetApp systems
- Set up and verify multiple path iSCSI connectivity between Windows, Red Hat, and NetApp systems

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY





Go further, faster®

## Exercise

Module 14: iSCSI Connectivity  
Estimated Time: 90 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What is the format for the IQN model of WWNNs naming?
- What two iSCSI multipathing techniques are available for Windows Server 2003 and 2008?
- What command is used to configure the iSCSI software initiator on Red Hat Enterprise Linux?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# LUN Access

Module 15  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## LUN ACCESS



## Module Objectives

By the end of this module, you should be able to:

- Describe the steps to allow a Windows® Server 2008 R2 initiator to access a LUN on a storage system
- Describe the steps to allow a Red Hat® initiator to access a LUN on a storage system

© 2010 NetApp, Inc. All rights reserved.

### MODULE OBJECTIVE



## LUN Access Overview

© 2010 NetApp, Inc. All rights reserved.

### LUN ACCESS OVERVIEW



## LUN Access

To connect an initiator to a target's LUN:

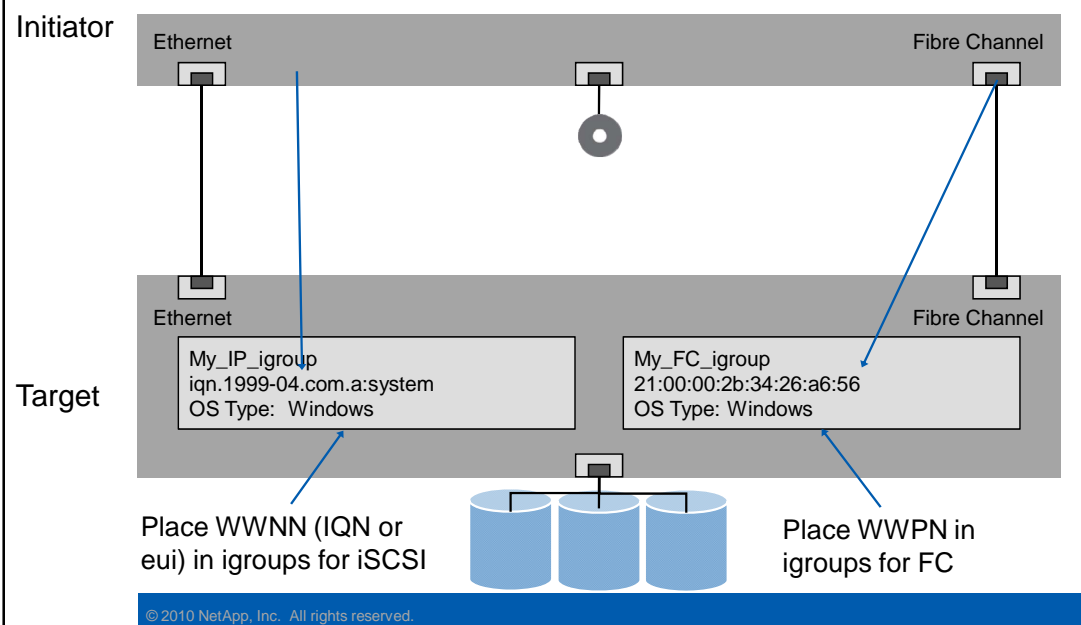
1. Create an igroup if necessary
2. Create the LUN
3. Map the LUN to the igroup
4. Find the LUN on the initiator
5. Prepare the LUN as a new disk on the initiator

© 2010 NetApp, Inc. All rights reserved.

## LUN ACCESS



# 1. Create an Igroup



## 1. CREATE AN IGROUP

If necessary, create an igroup to provide access to a LUN.

Initiator groups (igroups) are tables of host identifiers (FC WWPNs or iSCSI WWNs) that are used to control access to LUNs. Typically, you want all of the host's host bus adapters (HBAs) or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each HBA or software initiator of each clustered host needs redundant paths to the same LUN.

You can create igroups that specify which initiators have access to the LUNs either before or after you create LUNs, but you must create igroups before you can map a LUN to an igroup.

Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. However, you cannot map a LUN to multiple igroups that have the same initiator.

**NOTE:** An initiator cannot be a member of igroups of differing operating systems types (ostypes). Also, a given igroup can be used for FC or iSCSI, but not both.



## Steps to Create an Igroup

1. Optionally, verify initiators connectivity:
  - `fcplib show initiators`
  - `iscsi initiators show`
2. Create the igroup and place the initiators into the igroup:
  - `igroup create {-i|-f} -t ostype  
                                            igroup_name[node, node...]`
    - *i* = iSCSI igroup
    - *f* = FC igroup
    - *ostype*= solaris, windows, hpux, aix, linux, netware, vmware, hyper\_v, xen
    - *node*
      - iSCSI type has worldwide node (WWNN - IQN or eui)
      - FC type has worldwide port name (WWPN - may be aliased)
3. Verify the igroup:
  - `igroup show`

© 2010 NetApp, Inc. All rights reserved.

### STEPS TO CREATE AN IGROUP

Use the `igroup create` command to configure an igroup on a storage system. Note that you add nodes to the igroup and, therefore, the optional step of listing the currently connected initiators is provided in the first step.





## Data ONTAP Configuration

### ■ Add WWPNs to igroup:

```
system> igroup create -f -t windows iWIN_fcp
```

```
system> igroup add iWIN_fcp WIN1-FC WIN2-FC
```

### ■ Verify igroup:

```
system> igroup show -v
```

```
iWIN_fcp (FCP) (ostype: windows):
```

```
10:00:00:00:c9:6b:77:b3 (logged in on: 0d, 0c)
```

```
WWPN Alias(es): WIN1-FC
```

```
10:00:00:00:c9:6b:77:b4 (logged in on: 0d, 0c)
```

```
WWPN Alias(es): WIN2-FC
```

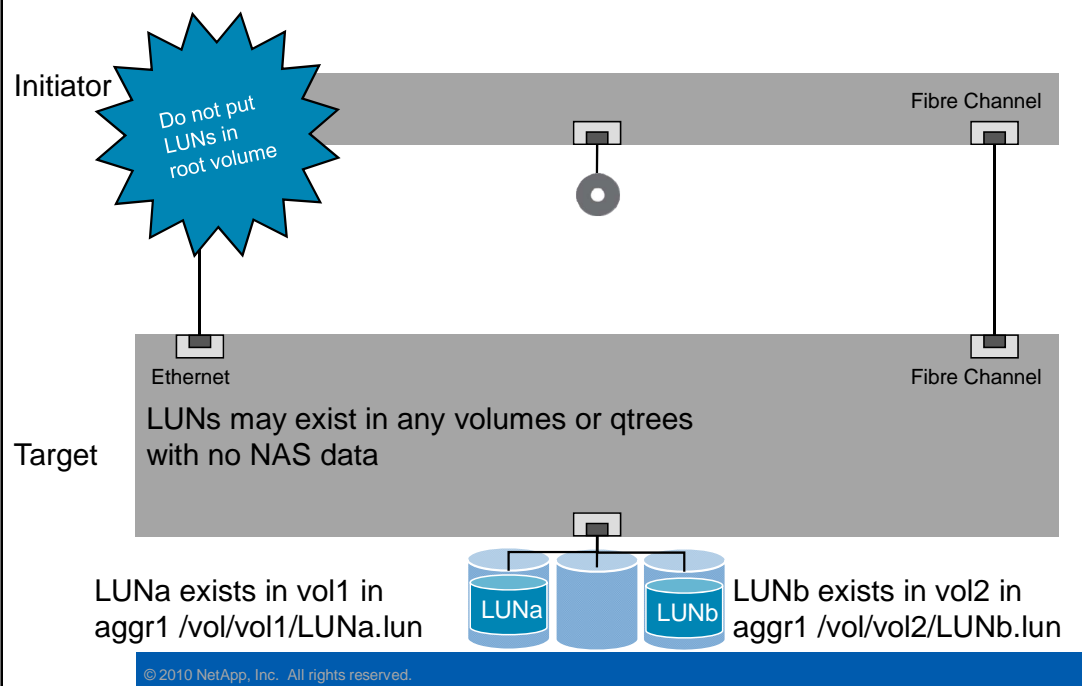
**NOTE:** Connected  
using paths displayed

© 2010 NetApp, Inc. All rights reserved.

## DATA ONTAP CONFIGURATION



## 2. Create a Logical Unit



## 2. CREATE A LOGICAL UNIT

LUNs may exist in any volume or qtree. When creating traditional or flexible volumes that contain LUNs, follow these guidelines:

- Do not create any LUNs in the system's root volume. Data ONTAP® uses this volume to administer the storage system. The default root volume is /vol/vol0.
- Ensure that no other files or directories exist in a volume that contains a LUN. If this is not possible and you are storing LUNs and files in the same volume, use a separate qtree to contain the LUNs.
- If multiple hosts share the same volume, create a qtree on the volume to store all LUNs for the same host. This is a recommended best practice that simplifies LUN administration and tracking.
- To simplify management, use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.



## Steps to Create a LUN

- Create the aggregate for the LUN:

```
system> aggr create aggr_SAN 7
```

- Create the volume for the LUN:

```
system> vol create vol_SAN0 aggr_SAN 10g
```

- Set Snapshot™ policy for the volume:  
(more on this in Module 13)

```
system> snap reserve vol_SAN0 0
```

```
system> vol options vol_SAN0 nosnap on
```

- Optional: Create a qtree for the LUN:

```
system> qtree create /vol/vol_SAN0/qtSAN0
```

© 2010 NetApp, Inc. All rights reserved.

## STEPS TO CREATE A LUN



## Steps to Create a LUN (Cont.)

### 1. Create a LUN:

```
–lun create -s size -t ostype lun_path
```

- *size* = in bytes by default

- Use m for megabytes

- Use g for gigabytes

NOTE: LUN sizing is discussed in detail in Module 13

- *ostype* = solaris, vld, windows, hpux, aix, linux, netware, vmware, windows\_gpt, windows\_2008, openvms, xen, hyper\_v, and solaris\_efi

- *lun\_path*

- LUN path begins with /vol/{VolumeName}/[qtreeName]

- Last portion of path is the LUN Name

- Example: /vol/vol\_SAN0/qtSAN0/lun0

### 2. Verify the LUN:

```
–lun show
```

© 2010 NetApp, Inc. All rights reserved.

## STEPS TO CREATE A LUN (CONT.)

Use the `lun create` command to create a LUN.

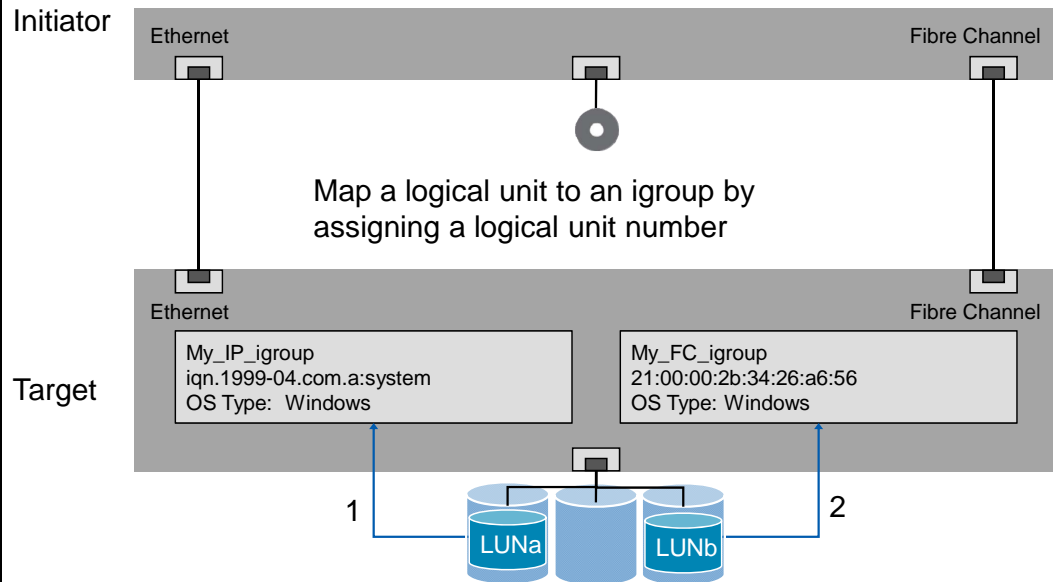
The host `ostype` indicates the type of operating system running on the host that accesses the LUN, which also determines the following:

- Geometry used to access data on the LUN
- Minimum LUN sizes
- Layout of data for multiprotocol access



### 3. Map a Logical Unit to an Igroup

**NOTE:** This step is also called LUN masking



© 2010 NetApp, Inc. All rights reserved.

### 3. MAP A LOGICAL UNIT TO AN IGROUP

When you map the LUN to the igroup, you grant the initiators in the igroup access to the LUN. If you do not map a LUN, the LUN is not accessible to any hosts. Data ONTAP maintains a separate LUN map for each igroup to support a large number of hosts and to enforce access control.



## Steps to Mask the LUN

### 1. Map a LUN to an igroup:

- `lun map lun_path igroup_name [lun_id]`
  - `lun_path` = path name of a LUN
  - `igroup_name` = name of an initiator group
  - `lun_id` = unique identification number that the initiator uses when the LUN is mapped to it
    - If not entered, automatically assigned
  - Example:  
`lun map /vol/voll/qtree1/luna My_IP_igroup 1`

### 2. Verify the LUN mapping:

- `lun show -m`

© 2010 NetApp, Inc. All rights reserved.

## STEPS TO MASK THE LUN

Use the `lun map` command to map an igroup to a LUN.

You map a LUN to an igroup by specifying the following attributes:

### LUN NAME

Specify the path name of the LUN to be mapped.

### INITIATOR GROUP

Specify the name of the igroup that contains the hosts that will access the LUN.

### LUN ID

Assign a number for the LUN ID, or accept the default LUN ID. Typically, the default LUN ID begins with 0 and increases in increments of one as each additional LUN is created. The host associates the LUN ID with the location and path name of the LUN. The range of valid LUN ID numbers depends on the host. For detailed information, see the documentation provided with your host utilities kit.



## Windows Setup

© 2010 NetApp, Inc. All rights reserved.

### WINDOWS SETUP



## Windows Steps

To connect an initiator to a target's LUN:

1. Create an igroup
2. Create the LUN
3. Map the LUN to the igroup
4. Find the LUN on the initiator
5. Prepare the LUN as a new disk on the initiator

© 2010 NetApp, Inc. All rights reserved.

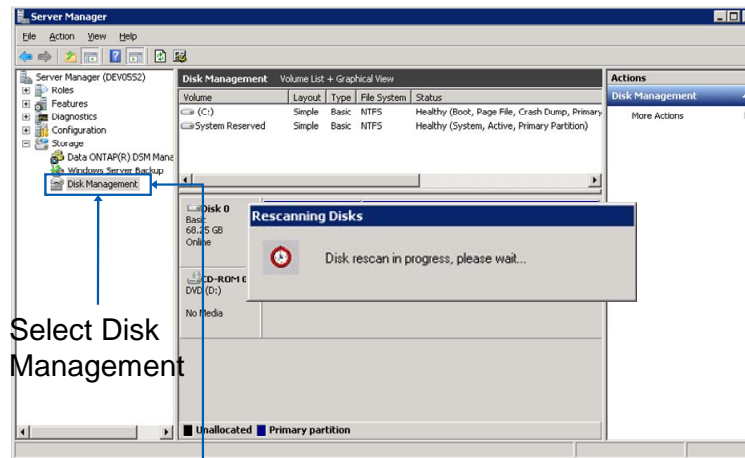
## WINDOWS STEPS





## 4. Find the LUN on Windows

- Using the Disk Management Tool, “Rescan Disks”



Select Disk  
Management

Right-click and choose  
“Rescan Disks”

© 2010 NetApp, Inc. All rights reserved.

## 4. FIND THE LUN ON WINDOWS

On a Windows host, you must partition and format any new LUN. To perform these tasks on Windows, use the Disk Management tool. First, access Computer Management by right-clicking My Computer and selecting Manage. Select Disk Management.

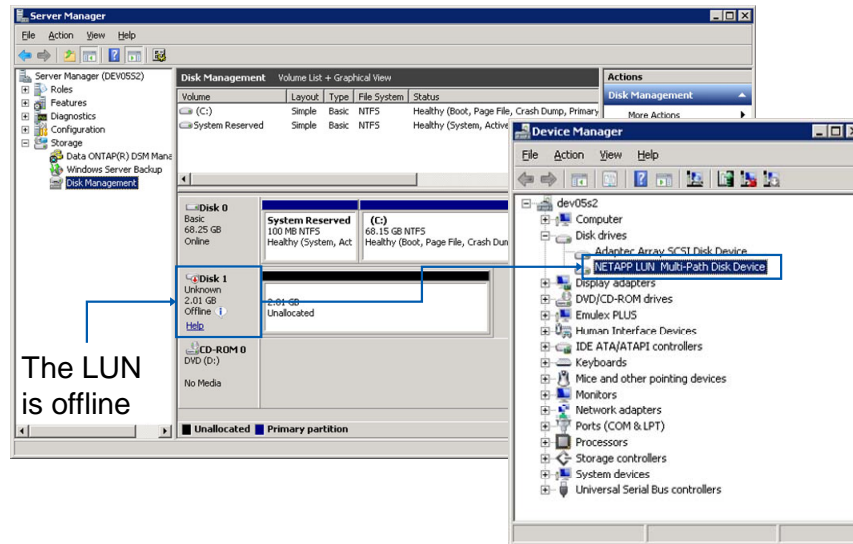
**NOTE:** The Disk Management tool will see and treat the LUN as though it is a local disk.

Within Computer Management, expand Storage and double-click Disk Management. In order for Disk Management to discover the new LUNs (virtual disks), select Action > Rescan Disks. From the Action menu, you can see tasks that may be performed on the new disk(s).



## 4. Find the LUN on Windows (Cont.)

- The LUN appears



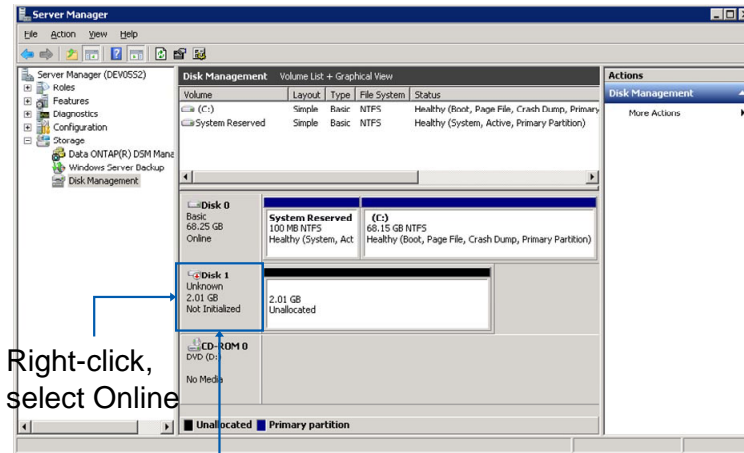
© 2010 NetApp, Inc. All rights reserved.

## 4. FIND THE LUN ON WINDOWS (CONT.)



## 5. Preparing the LUN for Windows

- Make the LUN online



**NOTE:** The LUN is not initialized

© 2010 NetApp, Inc. All rights reserved.

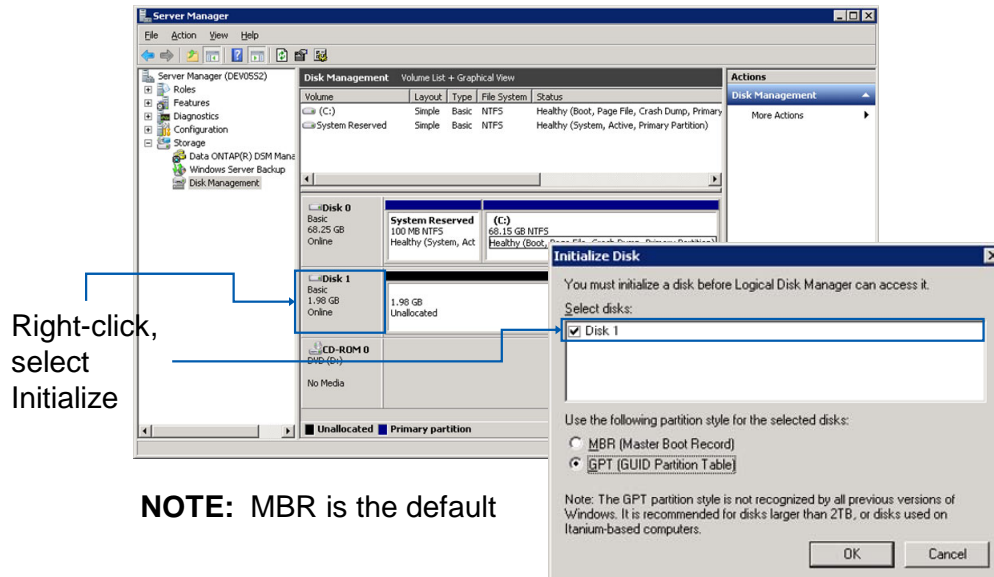
## 5. PREPARING THE LUN FOR WINDOWS

The disk must be brought online before we can use it. To bring the disk online, right-click “Disk #” and select Online.



## 5. Preparing the LUN for Windows (Cont.)

- Initialize the LUN in Windows Server 2008



© 2010 NetApp, Inc. All rights reserved.

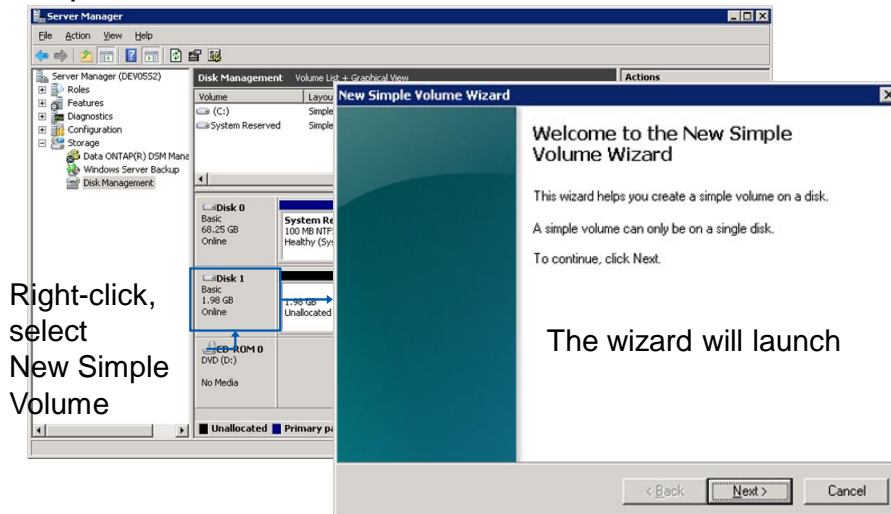
## 5. PREPARING THE LUN FOR WINDOWS (CONT.)

The disk must be initialized before it is formatted and partitioned. To initialize a disk, right-click “Disk #” and select Initialize. The administrator may select either a Master Boot Record (MBR) or a GUID Partition Table (GPT).



## 5. Preparing the LUN for Windows (Cont.)

- Windows Server 2008 has Disk Management's New Simple Volume Wizard



© 2010 NetApp, Inc. All rights reserved.

## 5. PREPARING THE LUN FOR WINDOWS (CONT.)

In Windows Server 2003, to partition the disk, select Action > All Tasks > Create Partition or right-click the unallocated box and select Create Partition.

In Windows Server 2008, to partition the disk, select Action > All Tasks > New Simple Volume or right-click the unallocated box and select New Simple Volume.

After the disk is initialized, partitioned, and formatted, you should test access to the disk by navigating to it and creating a file.

**NOTE:** The presentation displays the Windows Server 2008 version of the wizard. There are only cosmetic changes between the Windows Server 2003 and Windows Server 2008 version.



## 5. Preparing the LUN for Windows (Cont.)

### ■ New Simple Volume Wizard (Cont.)

Specify the volume size

**New Simple Volume Wizard**

**Specify Volume Size**  
Choose a volume size that is between the maximum and minimum disk space.

Maximum disk space in MB: 1025  
Minimum disk space in MB: 8  
Simple volume size in MB: 1025

**Assign Drive Letter or Path**  
For easier access, you can assign a drive letter or drive path to your partition.

☒ Assign the following drive letter: E  
☐ Mount in the following empty NTFS folder: Browse...  
☐ Do not assign a drive letter or drive path

< Back Next > Cancel

Specify the method to mount

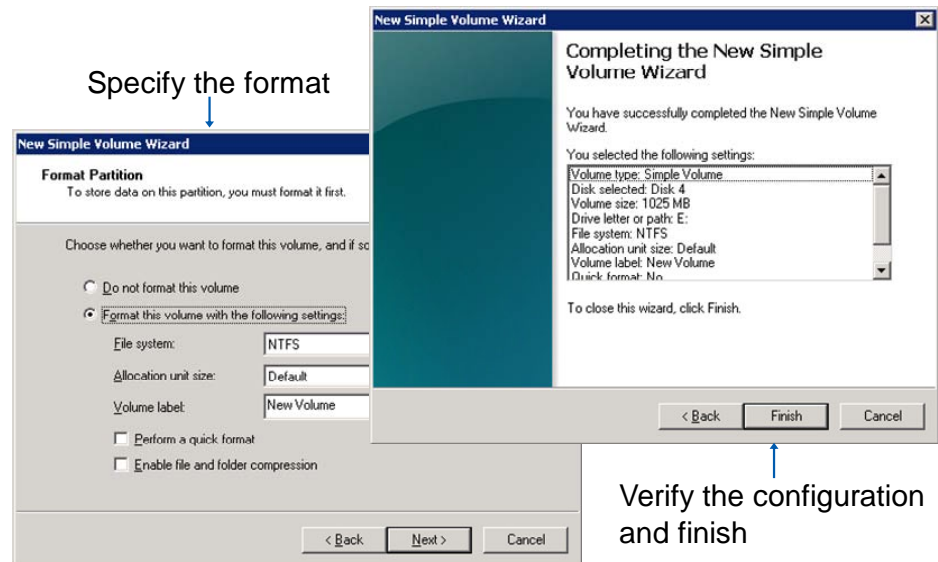
© 2010 NetApp, Inc. All rights reserved.

## 5. PREPARING THE LUN FOR WINDOWS (CONT.)



## 5. Preparing the LUN for Windows (Cont.)

### ■ New Simple Volume Wizard (Cont.)



© 2010 NetApp, Inc. All rights reserved.

## 5. PREPARING THE LUN FOR WINDOWS (CONT.)



## Red Hat Setup

© 2010 NetApp, Inc. All rights reserved.

## RED HAT SETUP





## Red Hat Steps

To connect an initiator to a target's LUN:

1. Create an igroup
2. Create the LUN
3. Map the LUN to the igroup
4. Find the LUN on the initiator
5. Prepare the LUN as a new disk on the initiator

© 2010 NetApp, Inc. All rights reserved.

## RED HAT STEPS



## 4. Find the LUN and 5. Prepare the LUN

### ■ Rescan HBA

```
cd /usr/sbin/lpfc
```

```
./lun_scan
```

### ■ Use the `fdisk -l` command

```
fdisk -l
```

```
...
```

```
Disk /dev/sdb: 2147 MB, 2147483648 bytes
```

```
67 heads, 62 sectors/track, 1009 cylinders
```

```
Units = cylinders of 4154 * 512 = 2126848 bytes
```

```
Disk /dev/sdb doesn't contain a valid partition table
```

```
Disk /dev/sdc: 2147 MB, 2147483648 bytes
```

```
67 heads, 62 sectors/track, 1009 cylinders
```

```
Units = cylinders of 4154 * 512 = 2126848 bytes
```

```
Disk /dev/sdc doesn't contain a valid partition table
```

```
...
```

The LUN shows up eight times... sdb - sdi; therefore use multipath

© 2010 NetApp, Inc. All rights reserved.

## 4. FIND THE LUN AND 5. PREPARE THE LUN



## Single Disk Configuration

### ■ Format a device:

```
fdisk /dev/sda
```

```
Command (m for help): n
```

```
Command action
```

```
 e extended
```

```
 p primary partition (1-4)
```

```
p
```

```
Partition number (1-4): 1
```

```
First cylinder (1-261, default 1): 1
```

```
Last cylinder or +size or +sizeM or +sizeK
(1-261, default 261): 261
```

```
Command (m for help): t
```

```
Hex code (type L to list codes): 83
```

```
Command (m for help): w
```

© 2010 NetApp, Inc. All rights reserved.

## SINGLE DISK CONFIGURATION



## More Information

- Additional resources are available in the *SAN Implementation Workshop* instructor-led course available from NetApp
  - Using dynamic disks within Microsoft Windows 2008 R2
  - Configure FC and iSCSI accessed LUN in VMware® vSphere™
  - DM-Multipath configuration in Red Hat Enterprise Linux®
  - LUN Provisioning Guidelines, including thin provisioning best practices

© 2010 NetApp, Inc. All rights reserved.

## MORE INFORMATION



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe the steps to allow a WindowsServer 2008 R2 initiator to access a LUN on a storage system
- Describe the steps to allow a Red Hat initiator to access a LUN on a storage system

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 15: LUN Access  
Estimated Time: 45 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What is an igroup?
- How does an FC igroup differ from an iSCSI igroup?
- Can you add a LUN to a different igroup using the LUN ID?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING





Go further, faster®

# Availability Overview

Module 16  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## AVAILABILITY OVERVIEW



## Module Objectives

By the end of this module, you should be able to:

- List the methods to back up and recover data for data availability
- Describe the methods to ensure system availability

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Availability

- Data storage providers need to be able to guarantee availability of user data
- To accomplish this, storage administrators need to develop a disaster recovery plan including:
  - A backup strategy for data availability
  - A system availability plan

© 2010 NetApp, Inc. All rights reserved.

## AVAILABILITY



## Disaster Recovery Plan

- Before you deploy a backup strategy, you need to set up a disaster recovery plan to:
  - Ensure continuity of operations
  - Guarantee availability of critical resources
  - Quickly recover vital business data
- The disaster recovery solution must be well-planned, rehearsed, tested, and documented before a disaster occurs

© 2010 NetApp, Inc. All rights reserved.

### DISASTER RECOVERY PLAN

Disaster recovery by definition includes the process, policies, and procedures related to the recovery of data, the continuation of data, as well as preparing for a data loss disaster. There are several approaches to protect data and maintain data availability in the face of hardware, software, or even site failures. However, before you deploy a backup strategy, purchase redundant hardware components, or replicate data between sites, you need to set up a disaster recovery plan.

A disaster recovery plan ensures the continuity of operations, the availability of critical resources and lists how to recover vital business operations successfully and quickly in the event of a disaster. The disaster recovery solution must be well-planned, rehearsed, tested, and documented before a disaster occurs.



## Standard Disaster Recovery Metrics

- Two basic recovery measures
  - **Recovery Point Objective (RPO)**— a measure backward in time from when the data was last protected
  - **Recovery Time Objective (RTO)**— a measure forward in time from the data loss event plus the accumulated time required to bring an application back online and re-create any lost transactions
- The amount of data replicated over a bandwidth size is used to predict the RPO
- The choice of a solution architecture helps determine the RTO

© 2010 NetApp, Inc. All rights reserved.

### STANDARD DISASTER RECOVERY METRICS

There are two basic recovery measures: recovery point and recovery time. Recovery point is a measure backwards in time from the data loss event that shows when data was last backed up or protected off-site. Recovery time is a measure forward in time from the data loss event; it is the accumulated time required to bring an application back online and re-create any lost transactions.

Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs) are the key elements in any recovery design. The traditionally accepted RPO and RTO are 24 hours, but this is changing with new technology and business requirements, including laws for some industries that require a near zero RPO.

When architecting a solution, administrators will need to identify the RPO and RTO targets. To do this, you need to assess the economic impact of incidents and disruptions that result from denial of access by way of systems, services, and facilities. This process might already be part of an organization's ongoing business impact analysis or risk analysis process.

Keep in mind that not all data has the same value or needs the same level of protection.



## Considerations for Disaster Recovery

- When deploying a disaster recovery solution, consider:
  - The distance between the data center and the disaster recovery site
  - Which data needs protection
  - The RPO
    - The maximum age of the data to be recovered
  - The RTO
    - The length of time required to recover

© 2010 NetApp, Inc. All rights reserved.

## CONSIDERATIONS FOR DISASTER RECOVERY

When deciding which type of disaster recovery solution to implement, you have to consider the following:

How far away from the primary storage site will the disaster recovery site be located? The distance between the main production center and the disaster recovery site depends on the likely geographic impact of the disaster.

Which applications need protection?

What are the RPO and the RTO in case of a disaster?

**RPO** is the maximum age of the data that the company must be able to recover in the event of a disaster. This parameter defines the amount of time for which work may be lost in the event of an unplanned outage at the primary site. It is used to define how frequently data is replicated to the disaster recovery site. A zero RPO implies no data loss, which can only be addressed by a synchronous replication solution. For higher RPO, an asynchronous replication solution with RPO in minutes or hours is flexible and cost-effective.

**RTO** is the length of time required to recover from a disaster. This parameter defines how quickly you need to fail over to a disaster recovery site—in seconds, minutes, hours, or days. Customers have to typically make tradeoffs between RPO, RTO, and cost. The target storage, the networking infrastructure, the software license and the deployment costs all contribute to the total cost of the overall disaster recovery solution.



## Disaster Recovery Solutions

|                                |                                        |                           |                        |                  |
|--------------------------------|----------------------------------------|---------------------------|------------------------|------------------|
| Data Recovery Techniques (RPO) | Tape / VTape                           | Data Vault w/ Tier Copy   | Data Vault             | Data Replication |
| Sample Data Loss               | 8-24 Hours                             | Near 0                    | Near 0                 | Near 0           |
| Equipment Provisioning (RTO)   | Tape or Disk-to-Disk                   | Hot Site                  | Re-Leveraged Resources | Standby          |
| Sample Recovery Times          | 30-76 Hours                            | 12-48 Hours               | 4-12 Hours             | 0-4 Hours        |
| Hosting Options                | Mobile                                 | Hot Site                  | Co-Location            | Internal         |
| NetApp Solution                | NDMP with 3 <sup>rd</sup> party or VTL | SnapVault with SnapMirror | SnapVault              | SnapMirror       |

© 2010 NetApp, Inc. All rights reserved.

## DISASTER RECOVERY SOLUTIONS



# Availability

- We will discuss the following:
  - Data availability overview
    - through backup and recovery techniques
  - System availability overview
    - through business continuous techniques

© 2010 NetApp, Inc. All rights reserved.

## AVAILABILITY





## Data Availability Overview

© 2010 NetApp, Inc. All rights reserved.

### DATA AVAILABILITY OVERVIEW



## Backup and Recovery Spectrum

Back up to Local Storage  
- Recover files

Back up to Local Storage  
- Recover single file, volume, or aggregate

Back up to Local/Remote Storage  
- Recover files

Back up to Remote Storage  
- Recover qtrees, directories, or files

Back up to Local or Remote Tape  
- Recover directories or files

Back up with Third-Party Tools  
- Recover qtrees, directories, or files

© 2010 NetApp, Inc. All rights reserved.

### BACKUP AND RECOVERY SPECTRUM




# Snapshot Copies

Back up to Local Storage  
- Recover files



- Administrators may back up and recover files quickly--almost instantaneously--with Snapshot™ copies
- NOTE: Snapshot copies do not replace standard backups to another media location



© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT COPIES



## SnapRestore

Back up to Local Storage  
- Recover single file, volume, or aggregate

Covered  
in  
Module 18

- SnapRestore® reverts a file system back to any specified Snapshot copy or restores single file from a Snapshot copy
  - Fast online restores of files, volumes, and aggregate
  - Multiple recovery points
  - Easy recovery process based on a single command input
  - Requires the `snaprestore` license code
- Use SnapRestore to recover from data corruption or to revert a file system

© 2010 NetApp, Inc. All rights reserved.

## SNAPRESTORE

SnapRestore enables you to quickly revert a local volume or a file on a storage system to the state it was in when a particular Snapshot copy was taken. In most cases, reverting a file or volume is much faster than restoring files from tape or copying files from a Snapshot copy to the active file system.

You use SnapRestore to recover from data corruption. If a primary storage system application corrupts data files in a volume, you can revert the volume or specified files in the volume to a Snapshot copy taken before the data corruption. You can also use SnapRestore if you are testing a volume or file and want to restore that volume or file to pretest conditions.

You must purchase and install the `snaprestore` license code to enable and use the SnapRestore service.




## ndmcopy Command

Back up to Local/Remote Storage  
- Recover files

Covered  
in  
DOTA

- The `ndmcopy` command:
  - Used to transfer data between storage systems that support NDMP v3 or v4  
`system> ndmpd on`
  - Can carry out full and incremental transfers
  - Limits incremental transfers to a maximum of two levels (one full and up to two incremental)
  - Applies NetApp® to NetApp only
  - Syntax:  
`system> ndmcopy [options]  
source_host:source_path  
destination_host:destination_path`



© 2010 NetApp, Inc. All rights reserved.

## NDMPCOPY COMMAND

The `ndmcopy` command enables you to transfer file system data between storage systems that support NDMP v3 or v4, and the UNIX® file system (UFS) dump format.

Using the `ndmcopy` command, you can carry out both full and incremental data transfers. However, incremental transfers are limited to a maximum of two levels (one full and no more than two incremental). You can transfer full or partial volumes, qtrees, or directories, but not individual files.

To copy data within a storage system or between storage systems using `ndmcopy`, use the following command from the source or the destination system, or from a storage system that is not the source or the destination:

- `system> ndmcopy [options] source_hostname:source_path  
destination_hostname:destination_path`
  - Where *source\_hostname* and *destination\_hostname* can be host names or IP addresses. If *destination\_path* does not specify a volume (or specifies a nonexistent volume), the root volume is used.




# SnapVault

Back up to Remote Storage  
- Recover qtrees, directories, or files

Covered  
in  
Module 19

- SnapVault® is the NetApp embedded disk-to-disk backup and archival software
- Administrators may back up and recover:
  - Qtrees
  - Directories on non NetApp storage



© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT

SnapVault is a disk-based storage backup feature of Data ONTAP®. SnapVault enables data stored on multiple storage systems to be backed up to a central, secondary storage system quickly and efficiently as read-only Snapshot copies.

In event of data loss or corruption on a storage system, backed-up data can be restored from the SnapVault secondary with less downtime and uncertainty than is associated with conventional tape backup and restore operations.

Additionally, users who wish to perform a restore of their own data may do so without the intervention of a system administrator. The SnapVault secondary may be configured with NFS exports and CIFS shares to let users copy the file from the Snapshot copy to the correct location.




# Tape Dumps and Restores

Back up to Local or Remote Storage  
- Recover directories or files

Covered  
in  
DOTA

- NetApp provides support for:
  - Classical tape backup and recovery
  - Virtual Tape Library (VTL)



© 2010 NetApp, Inc. All rights reserved.

## TAPE DUMPS AND RESTORES



## NDMP

Back up with Third-Party Tools  
- Recover qtrees, directories, or files

Covered  
in  
DOTA

- NDMP is an open standard that allows backup applications to control native backup and recovery functions on NetApp storage systems and other NDMP servers
- NDMP-compliant backup applications interact with the `ndmpd` process on the storage system
- NDMP requests from backup applications prompt the storage system to invoke native `dump` and `restore` commands to initiate backups and restores

© 2010 NetApp, Inc. All rights reserved.

## NDMP

The NDMP is an open standard for centralized control of data management across the enterprise. NDMP enables backup software vendors to provide support for NetApp storage systems without having to port client code.

An NDMP-compliant solution separates the flow of backup and restore control information from the flow of data to and from the backup media. These solutions invoke the Data ONTAP® operating system's native `dump`, and `restore` to back up data from, and restore data to a NetApp storage system.

NDMP also provides low-level control of tape devices and media changers.

Using data protection services through backup applications that support NDMP offers a number of advantages:

- Provides sophisticated scheduling of data protection operations across multiple storage systems
- Provides media management and tape inventory management services to eliminate or minimize manual tape handling during data protection operations
- Supports data catalogue services that simplify the process of locating specific recovery data; Direct Access Recovery optimizes the access of specific data from large backup tape sets
- Supports multiple topology configurations, allowing efficient sharing of secondary storage resources (tape library) through the use of three-way network data connections





## System Availability Overview

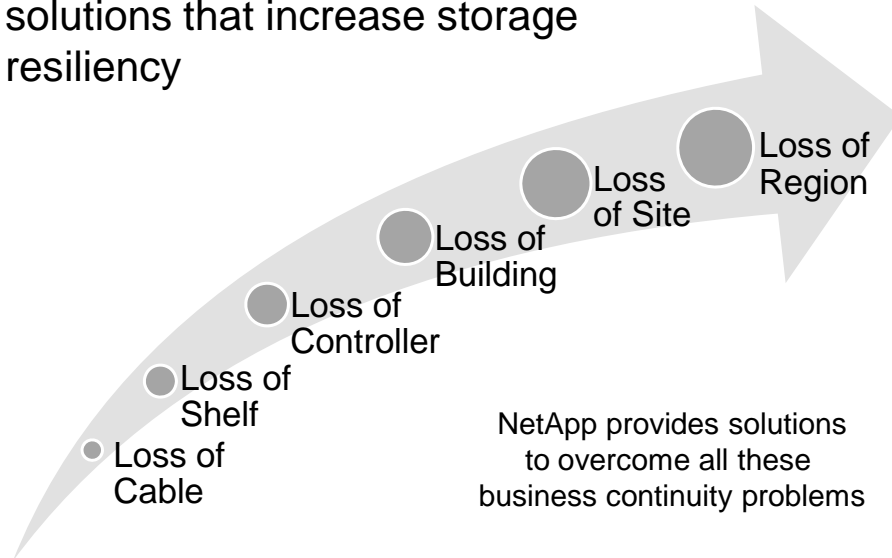
© 2010 NetApp, Inc. All rights reserved.

### SYSTEM AVAILABILITY OVERVIEW



## System Availability Spectrum

- High availability is the process of providing solutions that increase storage resiliency



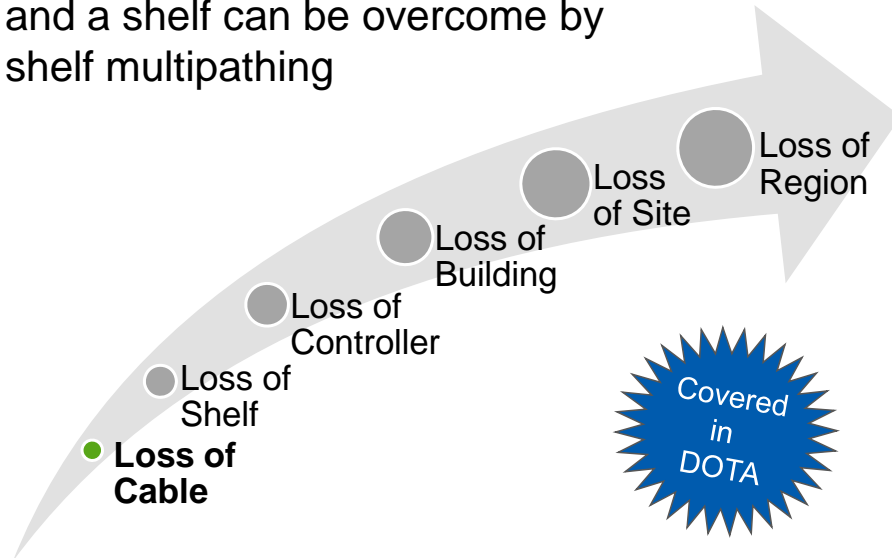
© 2010 NetApp, Inc. All rights reserved.

## SYSTEM AVAILABILITY SPECTRUM



## Loss of Cable

- Loss of a cable between the storage system and a shelf can be overcome by shelf multipathing



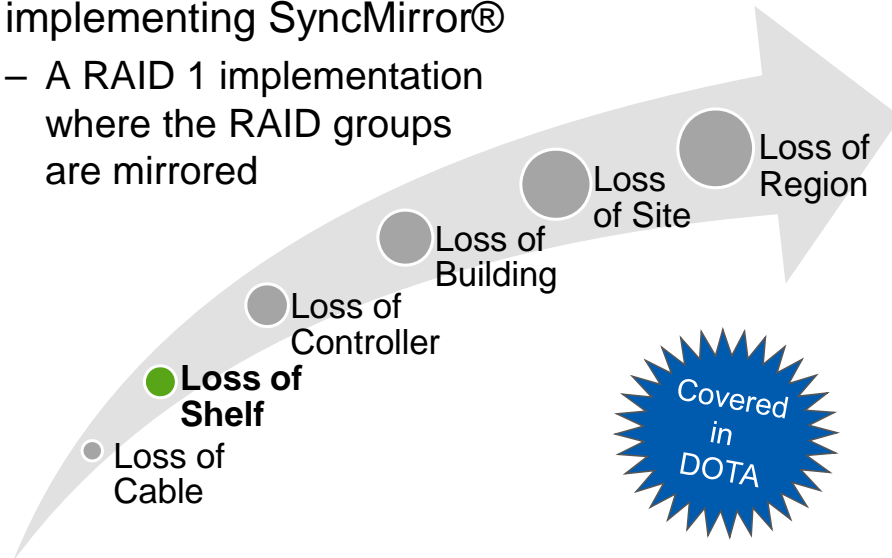
© 2010 NetApp, Inc. All rights reserved.

### LOSS OF CABLE



## Loss of Shelf

- Loss of a shelf can be overcome by implementing SyncMirror®
  - A RAID 1 implementation where the RAID groups are mirrored



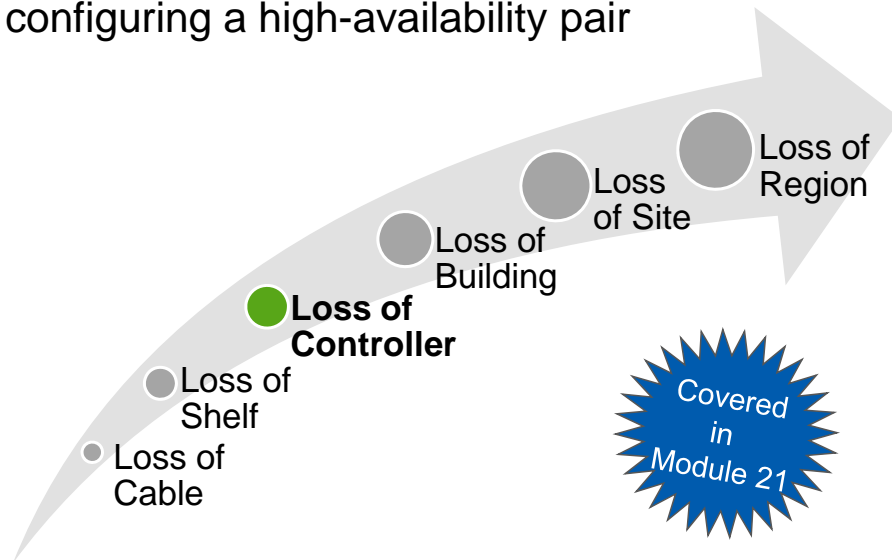
© 2010 NetApp, Inc. All rights reserved.

## LOSS OF SHELF



## Loss of Controller

- Loss of a controller may be overcome by configuring a high-availability pair



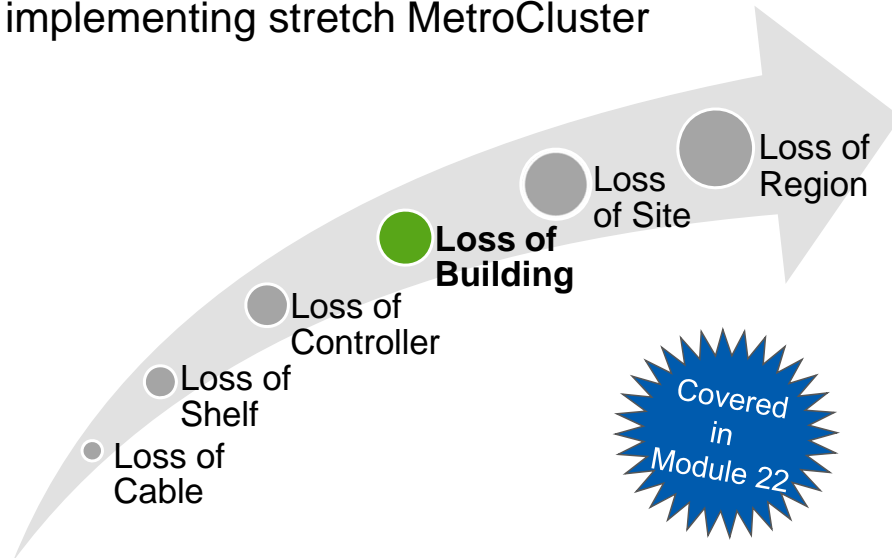
© 2010 NetApp, Inc. All rights reserved.

## LOSS OF CONTROLLER



## Loss of Building

- Loss of a entire building can be overcome by implementing stretch MetroCluster



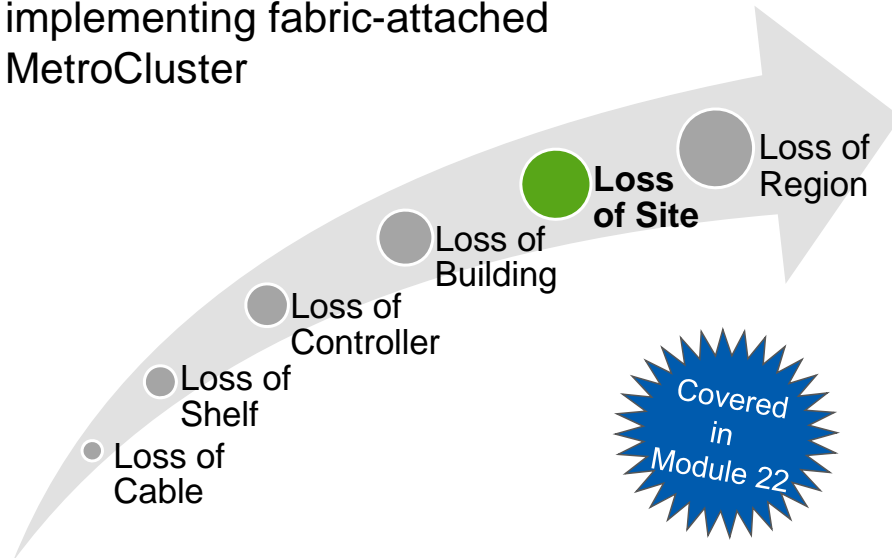
© 2010 NetApp, Inc. All rights reserved.

## LOSS OF BUILDING



## Loss of Site

- Loss of a site can be overcome by implementing fabric-attached MetroCluster



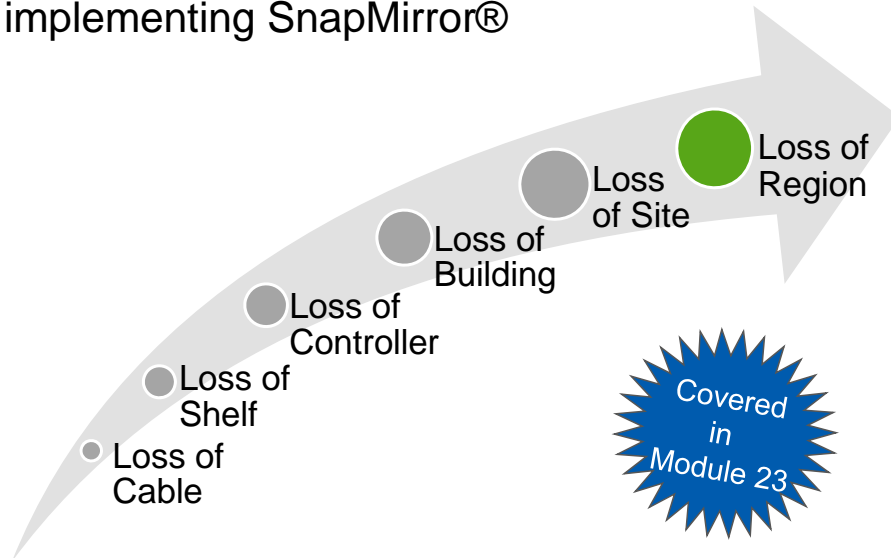
© 2010 NetApp, Inc. All rights reserved.

## LOSS OF SITE



## Loss of Region

- Loss of a region can be overcome by implementing SnapMirror®



© 2010 NetApp, Inc. All rights reserved.

## LOSS OF REGION





## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- List the methods to back up and recover data for data availability
- Describe the methods to ensure system availability

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 16: Availability Overview  
Estimated Time: 15 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- Which Data ONTAP solutions would you use for backup and rapid recovery?
- Which Data ONTAP solutions would you use for system availability?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# Snapshot Copies

Module 17  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## SNAPSHOT COPIES



## Module Objectives

By the end of this module, you should be able to:

- Describe the function of Snapshot™ copies
- Explain the benefits of Snapshot copies
- Identify and execute Snapshot commands
- Create and delete Snapshot copies
- Configure and modify Snapshot options
- Explain the importance of the `.snapshot` directory
- Describe how disk space is allocated by a Snapshot copy for volumes and aggregates
- Schedule Snapshot copies
- Configure and manage the Snapshot reserve

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Overview

© 2010 NetApp, Inc. All rights reserved.

## OVERVIEW



# Snapshot Technology

Back up to Local Storage  
- Recover files

- A Snapshot copy is a read-only image of the active file system at a point in time
- The benefits of Snapshot technology are:
  - Nearly instantaneous application data backups
  - Fast recovery of data lost due to:
    - Accidental data deletion
    - Accidental data corruption
- Snapshot technology is the foundation for:
  - SnapRestore®
  - SnapManager®
  - SnapDrive®
  - SnapMirror®
  - FlexClone®
  - SnapVault®

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT TECHNOLOGY

Snapshot technology is a key element in the implementation of the WAFL® (Write Anywhere File Layout) file system:

- A Snapshot copy is a read-only, space-efficient, point-in-time image of data in a volume or aggregate.
- A Snapshot copy is only a “picture” of the file system and does not contain any data file content.
- Snapshot copies are used for backup and error recovery.

Data ONTAP® automatically creates and deletes Snapshot copies of data in volumes to support commands related to Snapshot technology.





# How Snapshot Technology Works



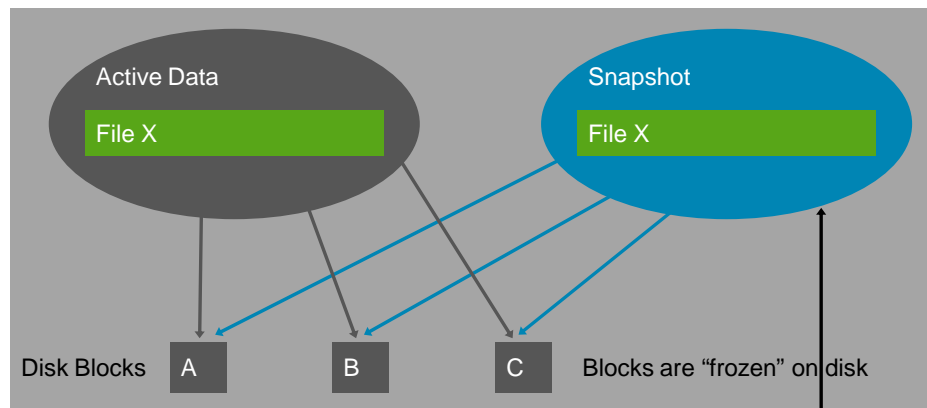
© 2010 NetApp, Inc. All rights reserved.

## HOW SNAPSHOT TECHNOLOGY WORKS

Before a Snapshot copy is created, there is a file system tree pointing to data blocks that contain content. When the Snapshot copy is created, a copy of the file structure metadata is created. The Snapshot copy points to the same data blocks.



## How Snapshot Technology Works (Cont.)



- Consistent (flushes NVRAM), point-in-time copy
- Ready to use (read-only)
- Consumes no space\*

\* With the exception of a 4-KB replicated root inode block that defines the Snapshot copy

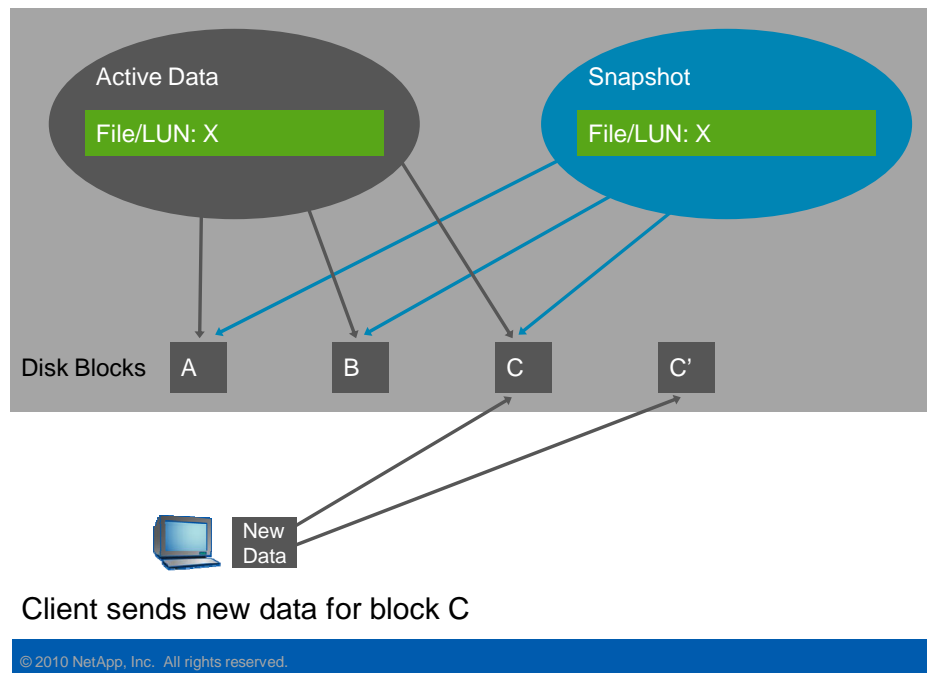
© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPSHOT TECHNOLOGY WORKS (CONT.)

There is no significant impact on disk space when a Snapshot copy is created. Because the file structure takes up little space, and no data blocks must be copied to disk, a new Snapshot copy consumes *almost* no additional disk space. In this case, the phrase “consumes no space” really means no *appreciable* space. The so-called “top-level root inode,” which is necessary to define the Snapshot copy, is 4 KB.



## How Snapshot Technology Works (Cont.)



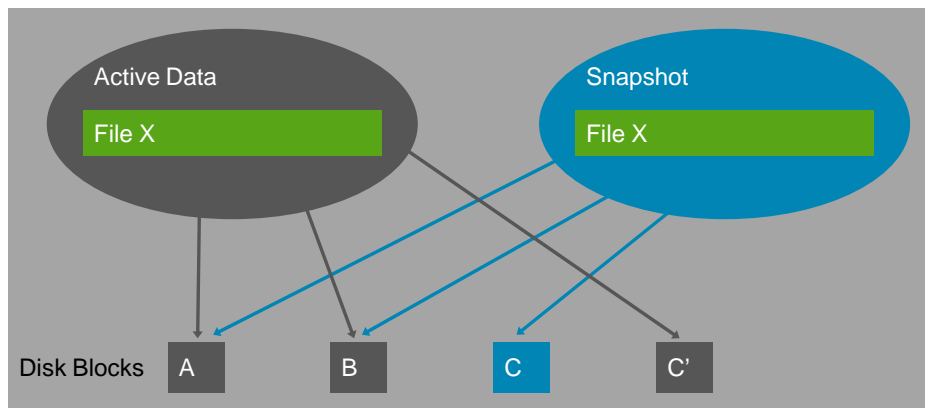
### HOW SNAPSHOT TECHNOLOGY WORKS (CONT.)

Snapshot copies begin to use space when data is deleted or modified. WAFL writes the new data to a new block (C') on the disk and changes the root structure for the active file system to point to the new block.

Meanwhile, the Snapshot copy still references the original block C. As long as there is a Snapshot copy referencing a data block, the block remains unavailable for other uses. This means that Snapshot copies start to consume disk space only as the file system changes after a Snapshot copy is created.



## How Snapshot Technology Works (Cont.)



- Active version of X is now comprised of blocks A, B, C'
- Snapshot version of X remains comprised of blocks A, B, C
- Moves active data to new consistent state

© 2010 NetApp, Inc. All rights reserved.

## HOW SNAPSHOT TECHNOLOGY WORKS (CONT.)



## Snapshot and WAFL

- Snapshot technology is a core feature of the storage system's WAFL (Write Anywhere File Layout) file system
- Every file in WAFL has at least one inode that organizes its data
- Each volume can contain up to 255 Snapshot copies

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT AND WAFL



## Inodes

- An inode is a data structure that is used to represent file system objects such as files and directories
- An inode is 192 bytes that describe a file's attributes and includes the following:
  - Type of file (regular file, directory, link, and so on)
  - Size
  - Owner, group, permissions
  - Pointer to xinode (ACLs)
  - Complete file data if the file is 64 bytes or less
  - Pointers to data blocks

© 2010 NetApp, Inc. All rights reserved.

## INODES

WAFL inodes are similar to Berkeley FFS (Fast File System) inodes. Veritas™ and Microsoft® file systems are based upon the Berkeley FFS, which force writes to pre-allocated locations. The primary difference is in the way WAFL writes contiguous data and metadata blocks to the next available block instead of predefined locations.

The most important metadata file is the inode file, which contains the inodes that describe all other files in the file system. The inode that describes the inode file itself is called the root inode. The root inode is in a fixed disk location.



## Snapshot Copies and Inodes

- Snapshot copies are a copy of the root inode of a volume
- The inodes of a Snapshot copy are read-only
- When the Snapshot inode is created:
  - It points to exactly the same disk blocks as the root inode
  - Brand new Snapshot copies consume only the space for the inode itself

© 2010 NetApp, Inc. All rights reserved.

### SNAPSHOT COPIES AND INODES

A Snapshot copy is a frozen, read-only image of a traditional volume, a FlexVol® volume, or an aggregate that reflects the state of the file system at the time the Snapshot copy was created. Snapshot copies are your first line of defense for backing up and restoring data. You can configure the Snapshot copy schedule.



## Managing Inodes

- To verify the amount of inodes:  
`df -i`
- To increase the maximum:  
`maxfiles`

© 2010 NetApp, Inc. All rights reserved.

## MANAGING INODES

### DF -I

The `df -i` command displays the amount of inodes in a volume. For more information about this command, see the manual pages.

### MAXFILES

The `maxfiles` command increases the number of inodes designated in a volume. For more information about this command, see the manual pages. **NOTE:** Do not use this command unless under the supervision of NetApp Global Support. There are certain increments that have to be used as well as memory implications.





## Creating Snapshot Copies

© 2010 NetApp, Inc. All rights reserved.

### CREATING SNAPSHOT COPIES



## Taking a Snapshot Copy

- Administrators can take Snapshot copies of:
  - Aggregates
    - Aggregate default for Snapshot reserve is 5% of aggregate
    - Restoring an aggregate Snapshot copy restores all volumes within that aggregate
  - Volumes
    - Volume default for Snapshot reserve is 20% of volume
    - Administrators can restore the entire volume or one or more files
- To change the amount of Snapshot reserve:  
`snap reserve [ -A | -V ] [volume_name] [percent]`

© 2010 NetApp, Inc. All rights reserved.

## TAKING A SNAPSHOT COPY

### VOLUMES

Snapshot copies for traditional and flexible volumes are stored in special subdirectories that can be made accessible to Windows® and UNIX® clients so that users can access and recover their own files without assistance. The maximum number of Snapshot copies per volume is 255.

### AGGREGATES


In an aggregate, 5% of space is reserved for Snapshot copies. In normal, day-to-day operations, aggregate Snapshot copies are not actively managed by a system administrator. For example, Data ONTAP automatically creates Snapshot copies of aggregates to support commands related to the SnapMirror software, which provides volume-level mirroring.

**NOTE:** Even if the Snapshot reserve is 0%, you can still create Snapshot copies. If there is no Snapshot reserve, Snapshot copies take their blocks from the active file system.



# Snapshot Reserve

- Aggregates
  - Each aggregate has 5% allocated for Snapshot reserve
- Flexible volumes
  - Each volume has 20% allocated for Snapshot reserve by default; the remainder is used for client data
  - For volumes used in a SAN configuration, NetApp recommends 0% reserve
- Snapshot reserve
  - The amount of space allocated for Snapshot reserve is adjustable; to use this space for data (which is not recommended) you must manually override the allocation used for Snapshot copies



© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT RESERVE



## CLI: Snapshot Creation

- Snapshot copies can be:
  - Scheduled
  - Manual
- To manually create Snapshot copies, use either:

```
system> snap create [-A aggrname | -V volname]
 [snapshotname]
```
- To rename Snapshot copies:

```
system> snap rename [-A aggrname | -V volname]
 [oldfilename] [newfilename]
```

© 2010 NetApp, Inc. All rights reserved.

## CLI: SNAPSHOT CREATION

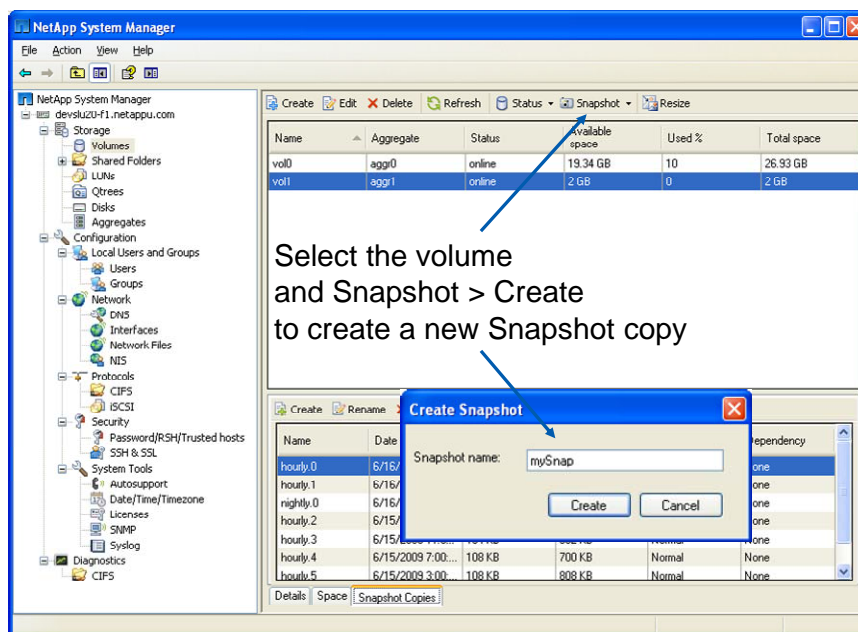
In the snap command, option A is used for aggregates and option V is used for volumes. If neither A nor V is specified, volume is the default.

The following table lists the commands used to create and manage Snapshot copies. If you omit the volume name from any of these commands, the command will apply to the root volume.

| EXAMPLE                                                         | RESULT                                                                                                                                                     |
|-----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>snap create engineering test</code>                       | Creates the Snapshot copy, test, in the engineering volume.                                                                                                |
| <code>snap list engineering</code>                              | Lists all available Snapshot copies in the engineering volume.                                                                                             |
| <code>snap delete engineering test</code>                       | Deletes the Snapshot copy test in the engineering volume.                                                                                                  |
| <code>snap delete -a vol2</code>                                | Deletes all Snapshot copies in vol2.                                                                                                                       |
| <code>snap rename engineering nightly.0<br/>firstnight.0</code> | Renames the Snapshot copy from nightly.0 to firstnight.0 in the engineering volume.                                                                        |
| <code>snap reserve vol2 25</code>                               | Changes the Snapshot reserve to 25 % on vol2.                                                                                                              |
| <code>snap sched vol2 0 2 6 @ 8, 12,<br/>16, 20</code>          | Sets the automatic schedule on vol2 to save the following weekly Snapshot copies: 0 weekly, 2 nightly, and 6 hourly at 8 a.m., 12 p.m., 4 p.m., and 8 p.m. |



# System Manager: Snapshot Copies



© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: SNAPSHOT COPIES



## System Manager: Snapshot Copies (Cont.)

The newly created Snapshot copy

| Name | Aggregate | Status | Available space | Used % | Total space |
|------|-----------|--------|-----------------|--------|-------------|
| vol0 | aggr0     | online | 19.34 GB        | 10     | 26.93 GB    |
| vol1 | aggr1     | online | 2 GB            | 0      | 2 GB        |

| Name     | Date Time         | Total Size | Cumulative Total Size | Status | Dependency |
|----------|-------------------|------------|-----------------------|--------|------------|
| mySnap   | 6/16/2009 4:05... | 52 KB      | 52 KB                 | Normal | None       |
| hourly.0 | 6/16/2009 4:00... | 132 KB     | 184 KB                | Normal | None       |
| hourly.1 | 6/16/2009 12:0... | 116 KB     | 300 KB                | Normal | None       |
| hourly.2 | 6/16/2009 8:00... | 152 KB     | 452 KB                | Normal | None       |
| hourly.3 | 6/16/2009 12:0... | 124 KB     | 576 KB                | Normal | None       |
| hourly.4 | 6/15/2009 8:00... | 108 KB     | 684 KB                | Normal | None       |
| hourly.5 | 6/15/2009 11:0... | 104 KB     | 788 KB                | Normal | None       |

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: SNAPSHOT COPIES (CONT.)



# System Manager: LUN Snapshot Creation

The screenshot displays the NetApp System Manager interface. The 'LUN Management' tab is active, showing a table with columns: Name, Container Path, Status, Size, and Type. A single LUN named 'lun1' is listed with a status of 'Online' and a size of '1.01 GB'. Below this, the 'LUN 'lun1' Snapshot Copies' window is open, showing a table with columns: Name and Access Time. A single snapshot copy named 'backup' is listed with an access time of 'Oct 15 09:00'. A third window, also titled 'LUN 'lun1' Snapshot Copies', is overlaid on the previous one, showing the same table. A blue arrow points from the text 'Ensure that LUN Snapshot copies are consistent' to the 'LUN 'lun1' Snapshot Copies' window. Another blue arrow points from the text 'NOTE: This is taking a snapshot copy of the volume containing the LUN' to the 'LUN 'lun1' Snapshot Copies' window. The bottom of the screenshot shows the copyright notice: '© 2010 NetApp, Inc. All rights reserved.'

Ensure that LUN Snapshot copies are consistent

NOTE: This is taking a snapshot copy of the volume containing the LUN

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: LUN SNAPSHOT CREATION



## Scheduling Snapshot Copies

© 2010 NetApp, Inc. All rights reserved.

### SCHEDULING SNAPSHOT COPIES





## Snapshot Schedule

- To print current schedule for all volumes:  
`system> snap sched`
- To print schedule per volume:  
`system> snap sched [volume_name]`
- To change current schedule per volume:  
`system> snap sched  
[volume_name[weeks[days[hours[@list]]]]]`
- Example:  
`system> snap sched vol2 1 6 5@4,8,12,16,20`
  - The Snapshot schedule above keeps the following Snapshot copies for vol2:
    - One weekly Snapshot copies
    - Six nightly Snapshot copies
    - Five hourly Snapshot copies taken at 4:00 a.m., 8:00 a.m., 12:00 p.m., 4:00 p.m., and 8:00 p.m.

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT SCHEDULE

The `snap sched` command sets a schedule to automatically create Snapshot copies and specifies how many of each type are stored. When the limit is reached, the oldest Snapshot copy for each interval is deleted and replaced by a new Snapshot copy.

Snapshot copies are like a *picture* of a volume. The only difference between a weekly Snapshot copy and a nightly or hourly copy is the time at which the Snapshot copy was created and any data that was changed between the Snapshot copies.



## Scheduling Snapshot Copies

- Default schedule:

```
system> snap sched vol0
```

```
Volume vol0: 0 2 6@8, 12, 16, 20
```

- Once nightly, Monday through Saturday, at midnight (12:00 a.m.)
- Four hourly at 8:00 a.m., 12:00 p.m., 4:00 p.m., and 8:00 p.m.
- Retains:

- Zero weekly
- Two most recent nightly
- Six most recent hourly

**NOTE:** If you change the root volume's Snapshot schedule, all new volumes will adopt the altered schedule by default

- First in, first out:

- Oldest nightly Snapshot copy
- Oldest hourly Snapshot copy

- Disable automatic Snapshot copies:

```
system> vol options volname nosnap [on|off]
```

© 2010 NetApp, Inc. All rights reserved.

## SCHEDULING SNAPSHOT COPIES



# System Manager: Snapshot Copies

**To configure volumes**

Select the volume and Snapshot > Configure to manage Snapshot technology on the volume

| Name | Aggregate | Status | Available space | Used % | Total space |
|------|-----------|--------|-----------------|--------|-------------|
| vol0 | aggr0     | online | 19.34 GB        | 10     | 26.93 GB    |
| vol1 | aggr1     | online | 2 GB            | 0      | 2 GB        |

Configuration details for vol1:

|                 |          |                     |                         |
|-----------------|----------|---------------------|-------------------------|
| Name:           | vol1     | Guarantee:          | volume                  |
| Status:         | online   | Maximum files:      | 62,244                  |
| Type:           | Flexible | Current files:      | 97                      |
| Root:           | No       | Character encoding: | undefined 0 (undefined) |
| Clone parent:   | NA       | Create unicode:     | Yes                     |
| Clone children: | NA       | Convert unicode:    | Yes                     |

Details | Space | Snapshot Copies

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: SNAPSHOT COPIES



## System Manager: Snapshot Copies (Cont.)

**Configure Volume Snapshots-vol1**

Snapshot reserve(%):   
Specify the percentage of volume space reserved for snapshot copies.

☒ Make snapshot directory (.snapshot) visible  
Specifies whether the .snapshot directory is visible on this volume at the client mount points.

☒ Enable scheduled snapshots

Schedule  
Select the number of scheduled snapshots to keep  
Weekly  Daily  Hourly

Hourly snapshot schedules:

Time zone: Central Standard Time

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: SNAPSHOT COPIES (CONT.)



## System Manager: Snapshot Copies (Cont.)

The screenshot shows the NetApp System Manager web interface. On the left is a navigation tree with categories like Storage, Configuration, Network, Protocols, Security, and Diagnostics. The main pane displays a table of volumes and a detailed view of the 'vol1' volume.

| Name | Aggregate | Status | Available space | Used % | Total space |
|------|-----------|--------|-----------------|--------|-------------|
| vol0 | aggr0     | online | 19.34 GB        | 10     | 26.93 GB    |
| vol1 | aggr1     | online | 2 GB            | 0      | 2 GB        |

The 'vol1' volume details are shown below the table:

- Volume:**
  - Total space: 2.4 GB
  - Snapshot reserve: 409.53 MB
- Available:**
  - Space: 2 GB
  - Snapshot reserve: 408.64 MB
  - Total: 2.4 GB
- Used:**
  - Data space: 1.11 MB
  - Snapshot copies space: 916 KB
  - Total: 2.01 MB

A callout box labeled "Select Space" points to the "Space" tab in the bottom navigation bar of the volume details section.

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: SNAPSHOT COPIES (CONT.)



## Restoring Snapshot Copies

© 2010 NetApp, Inc. All rights reserved.

### RESTORING SNAPSHOT COPIES



## Recovering Data

- When recovering data, you have two options:
  - Copy the data from a Snapshot copy
  - Use SnapRestore (see Module 18)
- To copy data from a Snapshot copy:
  - Locate the Snapshot copy
  - Recover the copy from `.snapshot` directory
    - To overwrite, copy to the original location
    - For a new, writeable version, copy to a new location

© 2010 NetApp, Inc. All rights reserved.

## RECOVERING DATA

### Using Snapshot copies to Recover Data

To recover data, you can:

- Restore a file from a Snapshot copy
- Use SnapRestore (license required)

To restore a file from a Snapshot copy:

Locate the Snapshot copy that contains the correct version of the file.

Restore the file from the `.snapshot` directory.

- To overwrite existing data, copy to the original location.
- To restore a writeable version, copy to a new location.



## Snapshot Visibility to Clients

- Make the `.snapshot` directory invisible to clients, and turn off access to the `.snapshot` directory:  
`vol options volname nosnapdir [on|off]`
- Make the `~snapshot` directory visible to CIFS clients:  
`options cifs.show_snapshot [on|off]`
- Make the `.snapshot` directory visible to NFS clients:  
`options nfs.hide_snapshot [on|off]`

**NOTE:** Default values are in bold

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT VISIBILITY TO CLIENTS

The following table lists the options available for controlling the creation of Snapshot copies and access to those copies and Snapshot directories on a volume:

Disable automatic Snapshot copies. Setting the `nosnap` option to on disables automatic Snapshot creation. You can still create Snapshot copies manually at any time.

Make the `.snapshot` directory invisible to clients and turn off access to the `.snapshot` directory. Setting the `nosnapdir` option to on disables access to the Snapshot directory that is present at client mountpoints and the root of CIFS directories, and makes the Snapshot directories invisible. (NFS uses `.snapshot` for directories, while CIFS uses `~snapshot`.) By default, the `nosnapdir` option is off (directories are visible).

Make the `~snapshot` directory visible to CIFS clients by completing the following steps:

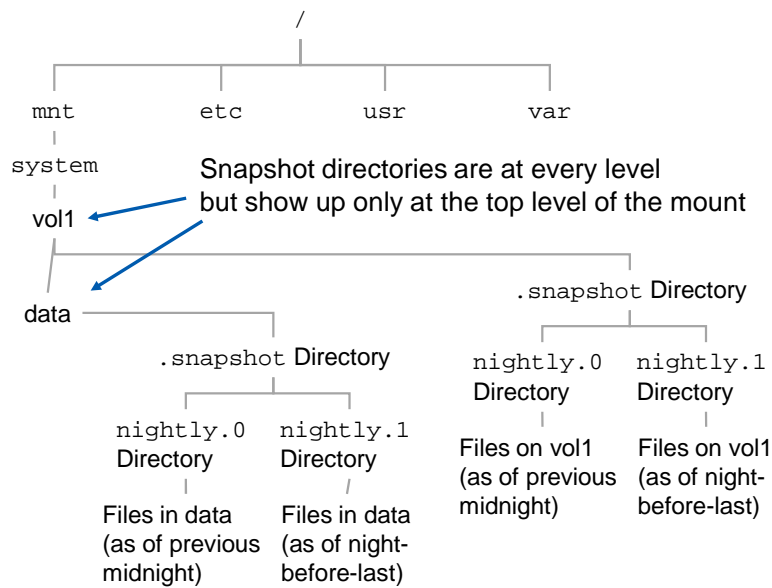
- 1. Turn the `cifs.show_snapshot` option on.
- 2. Turn the `nosnapdir` option off for each volume that you want directories to be visible.

**NOTE:** You must also ensure that Show Hidden Files and Folders is enabled on your Windows system.





# The .snapshot Directory



© 2010 NetApp, Inc. All rights reserved.

## THE .SNAPSHOT DIRECTORY

The .snapshot directory is at the root of a volume.

In the figure above, the directory structure is shown for an NFS client mounting vol1 to the mountpoint /mnt/system.



## Snapshot View from a UNIX/Linux Client

```
pwd
/system/vol0/.snapshot

ls -l
total 240
drwxrwxrwx 9 root other 2288 Jan 29 16:19 hourly.0
drwxrwxrwx 9 root other 3288 Jan 29 15:19 hourly.1
drwxrwxrwx 9 root other 4288 Jan 29 14:19 hourly.2
drwxrwxrwx 9 root other 5288 Jan 29 13:19 hourly.3
drwxrwxrwx 9 root other 6288 Jan 29 12:19 hourly.4
drwxrwxrwx 9 root other 7288 Jan 29 11:19 hourly.5
drwxrwxrwx 9 root other 12288 Jan 28 16:19 nightly.0
drwxrwxrwx 9 root other 22288 Jan 27 16:19 nightly.1
drwxrwxrwx 9 root other 32288 Jan 22 16:19 weekly.1
drwxrwxrwx 9 root other 42288 Jan 15 16:19 weekly.2
```

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT VIEW FROM A UNIX/LINUX CLIENT

### Snapshot Directories

Every volume in your file system contains a special Snapshot subdirectory that allows you to access earlier versions of the file system to recover lost or damaged files.

### Viewing Snapshot Copies from a UNIX Client

The Snapshot subdirectory appears to NFS clients as `.snapshot`. The `.snapshot` directories are usually hidden and are not displayed in directory listings.

To view a `.snapshot` directory, complete the following steps:

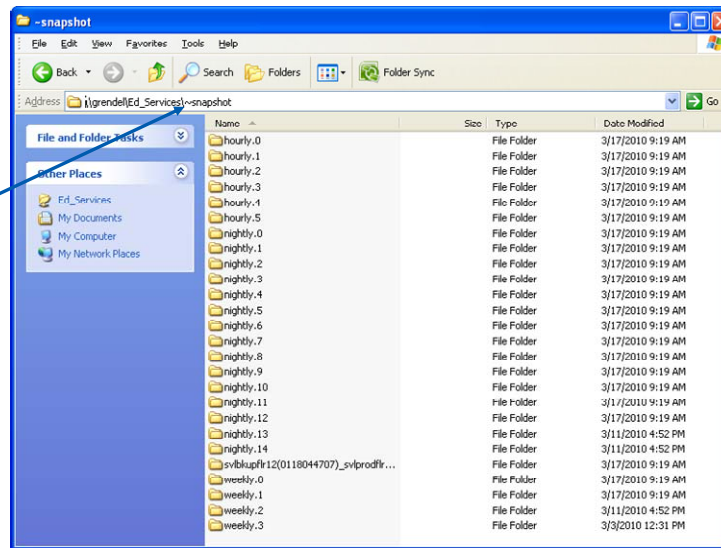
1. On the storage appliance, log in as root and ensure that the `nosnapdir` option is set to off.
2. To view hidden directories, from the NFS mountpoint, enter the `ls` command with the `-a` (all) option.

When listing the client Snapshot directories, the date/timestamp is usually the same for all directories. To find the actual date/time of each Snapshot copy, use the `snap list` command on the storage system.



## Snapshot View from a Windows Client

Snapshot copies are visible to Windows clients that have File Manager configured to display “hidden files”



© 2010 NetApp, Inc. All rights reserved.

### SNAPSHOT VIEW FROM A WINDOWS CLIENT

Snapshot directories are hidden on Windows clients. To view them, you must first configure the File Manager to display hidden files, then navigate to the root of the CIFS share and find the directory folder.

The subdirectory for Snapshot copies appears to CIFS clients as ~snapshot. Files displayed here are those files created automatically for specified intervals. Manually created Snapshot copies would also be listed here.

#### Restoring a File

To restore a file from the ~snapshot directory, rename or move the original file, then copy the file from the ~snapshot directory to the original directory.



## FlexClone

© 2010 NetApp, Inc. All rights reserved.

## FLEXCLONE



## FlexClone Volume Clones

- Enables multiple, instant dataset clones with no storage overhead
- Provides dramatic improvement for application test and development environments
- Renders competitive methods archaic

© 2010 NetApp, Inc. All rights reserved.

### FLEXCLONE VOLUME CLONES

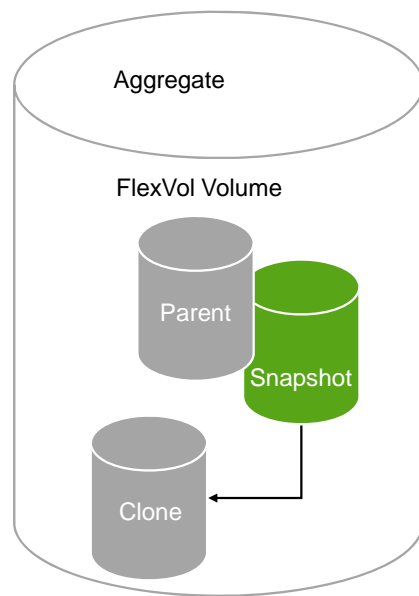
FlexClone volume clones provide an efficient way to copy data for:

- Manipulation
- Projection operations
- Upgrade testing

Data ONTAP allows you to create a volume duplicate with the original volume and clone volume sharing the same disk space for storing unchanged data.



## How Volume Cloning Works



Volume cloning:

- Starts with a volume
- Takes a Snapshot copy of the volume
- Creates a clone (a new volume based on the Snapshot copy)

Result:

Independent volume copies are efficiently stored

© 2010 NetApp, Inc. All rights reserved.

## HOW VOLUME CLONING WORKS

FlexClone volumes are managed similarly to regular FlexVol volumes, with a few key differences. The following is a list of important facts about FlexClone volumes:

- FlexClone volumes are a point-in-time, writable copy of the parent volume. Changes made to the parent volume after the FlexClone volume is created are not reflected in the FlexClone volume.
- You can only clone FlexVol volumes. To create a copy of a traditional volume, you must use the `vol copy` command, which creates a distinct copy with its own storage.
- FlexClone volumes are fully functional volumes managed just like the parent volume using the `vol` command.
- FlexClone volumes always exist in the same aggregate as parent volumes.
- FlexClone volumes can be cloned.
- FlexClone volumes and parent volumes share the same disk space for common data. This means that creating a FlexClone volume is instantaneous and requires no additional disk space (until changes are made to the clone or parent).
- A FlexClone volume is created with the same space guarantee as the parent.
- You can sever the connection between the parent and the clone. This is called *splitting* the FlexClone volume. Splitting removes all restrictions on the parent volume and causes the FlexClone volume to use its own storage.

**IMPORTANT:** Splitting a FlexClone volume from its parent volume deletes all existing Snapshot copies of the FlexClone volume and disables the creation of new Snapshot copies while the splitting operation is in progress.

- Quotas applied to a parent volume are not automatically applied to the clone.

When a FlexClone volume is created, existing LUNs in the parent volume are also present in the FlexClone volume, but are unmapped and offline.



# System Manager: FlexClone

Select the volume and Snapshot > Clone

| Name | Aggregate | Status | Available space | Used % | Total space |
|------|-----------|--------|-----------------|--------|-------------|
| vol0 | agg01     | online | 19.76 GB        | 8      | 26.93 GB    |
| vol1 | agg1      | online | 1.6 GB          | 0      | 2 GB        |

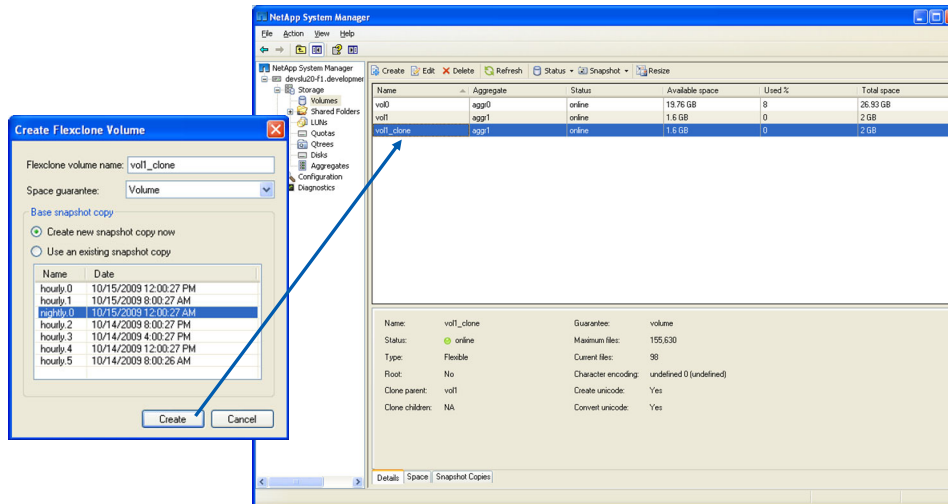
NOTE: FlexClone must be licensed first

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: FLEXCLONE



## System Manager: FlexClone (Cont.)



© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: FLEXCLONE (CONT.)





## Flexible Volume Clone Syntax

- Use the `vol clone create` command to create a flexible volume clone

– Syntax: `vol clone create volname`  
          `[-s none | file | volume]`  
          `-b parent_volname [parent_snapshot]`

- The following is an example of a CLI entry used to create a flexible volume clone:


```
system> vol clone create clone1 -b flexvol1 flex1
system> vol status clone1
Volume State Status Options
clone1 online raid_dp, flex
guarantee=volume(disabled)
Clone, backed by volume 'flex1' snapshot
clone_clone1.1'
Containing aggregate: 'aggr1'
```

© 2010 NetApp, Inc. All rights reserved.

## FLEXIBLE VOLUME CLONE SYNTAX



## Splitting Volumes



- Split volumes when most of the data on a volume is not shared
- Replicate shared blocks in the background

### Result

New, permanent volume is created for forking project data

© 2010 NetApp, Inc. All rights reserved.

## SPLITTING VOLUMES

Splitting a FlexClone volume from its parent removes any space optimizations currently employed by the FlexClone volume. After the split, both the FlexClone volume and the parent volume require the full space allocation specified by their space guarantees. After the split, the FlexClone volume becomes a normal FlexVol volume.

When splitting clones, keep the following in mind:

- When you split a FlexClone volume from its parent, all existing Snapshot copies of the FlexClone volume are deleted.
- During the split operation, no new Snapshot copies of the FlexClone volume can be created.
- Because the clone-splitting operation is a copy operation that could take some time to complete, Data ONTAP provides the `vol clone split stop` and `vol clone split status` commands to stop clone-splitting or check the status of a clone-splitting operation.
- The clone-splitting operation executes in the background and does not interfere with data access to either the parent or the clone volume.
- If you take the FlexClone volume offline while clone-splitting is in progress, the operation is suspended. When you bring the FlexClone volume back online, the splitting operation resumes.

After a FlexClone volume and its parent volume have been split, they cannot be rejoined.



## vol clone split Command

- To start a clone split:  
`system> vol clone split start volname`
- To stop a clone split:  
`system> vol clone split stop`
- To check the status of a clone split:  
`system> vol clone split status [volname]`
- To estimate the time of completion:  
`system> vol clone split estimate [volname]`

© 2010 NetApp, Inc. All rights reserved.

## VOL CLONE SPLIT COMMAND

### HOW TO VIEW THE RESULTS OF A CLONE SPLIT COMMAND

Example:

```
vol clone split status:
```

```
vol clone split start clone1
```

```
Tue Oct 12 23:49:43 GMT [wafl.scan.start:info]: Starting volume clone split on volume clone1.
```

```
Clone volume 'clone1' will be split from its parent.
```

```
Monitor system log or use 'vol clone split status' for progress.
```

```
vol clone split status
```

```
Volume 'clone1', 117193 of 364077 inodes processed (32%)
```

```
18578 blocks scanned. 18472 blocks updated.
```



## Space Usage

© 2010 NetApp, Inc. All rights reserved.

## SPACE USAGE



## Using the CLI to Monitor Space Used

- To monitor space used for Snapshot copies, use:
  - From the command-line interface
  - NetApp System Manager
- To determine how much space you will get back:
  - `snap list`
  - `snap reclaimable`
  - `snap delta`
- To delete:
  - A particular Snapshot copy:  

```
system> snap delete [-A | -V]
 [aggrname|volname] [snapshotname]
```
  - All Snapshot copies:  

```
system> snap delete [-A | -V] -a
 [aggrname|volname]
```

© 2010 NetApp, Inc. All rights reserved.

## USING THE CLI TO MONITOR SPACE USED



## The snap list Command

```
system> snap list
```

```
Volume vol0
```

```
working...
```

| %used     | %total  | date         | name     |
|-----------|---------|--------------|----------|
| 0% (0%)   | 0% (0%) | Apr 20 12:00 | hourly.0 |
| 17% (20%) | 1% (1%) | Apr 20 10:00 | hourly.1 |
| 33% (20%) | 2% (1%) | Apr 20 08:00 | hourly.2 |

| %Used                                                                                                                                                                                                                  | %Total                                                                                                                                                                                                      | Date                                                                                                                                                                                       | Name                                                                                                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The % used column shows the relationship between accumulated Snapshot copies and the total disk space consumed by the active file system; Values in parentheses show the contribution of this individual Snapshot copy | The % total column shows the relationship between accumulated Snapshot copies and the total disk space consumed by the volume; Values in parentheses show the contribution of this individual Snapshot copy | The date column shows the date and time the Snapshot copy was taken; Time is indicated on the 24-hour clock, and in this example reflects the hours set in the automatic Snapshot schedule | Scheduled Snapshot copies are automatically renumbered as new ones are taken so that the most recent is always "0"; This also ensures that the file with the highest number is always the oldest |

© 2010 NetApp, Inc. All rights reserved.

## THE SNAP LIST COMMAND

The `snap list` command displays a single line of information for each Snapshot copy in a volume. In the Snapshot List Example in the figure above, a list of Snapshot copies is displayed for the engineering volume. The following is a description of each column in the list:

- **%used**—Shows the relationship between accumulated Snapshot copies and the total disk space consumed by the active file system. Values in parentheses show the contribution of this individual Snapshot copy.
- **%total**—Shows the relationship between accumulated Snapshot copies in the total disk space consumed by the volume. Values in parentheses show the contribution of this individual Snapshot copy.
- **date**—Shows the date and time the Snapshot copy was taken. Time is indicated on the 24-hour clock, and in this example, reflects the hours set in the automatic Snapshot copy schedule.
- **name**—Lists the names of each of the saved Snapshot copies. Scheduled Snapshot copies are automatically renumbered as new ones are created so that the most recent copy is always .0. This also ensures that the file with the highest number (in this case, hourly.2) is always the oldest Snapshot copy.



## snap delta and snap reclaimable

### ■ snap delta (provides the rate of change)

```
system> snap delta vol0
Volume vol0
working...
From Snapshot To kB changed Time Rate(kB/hour)

nightly.0 AFS 46932 0d 23:00 3911.000
nightly.1 nightly.0 16952 1d 00:00 4237.705
nightly.2 nightly.1 16952 1d 00:00 4237.705
```

### ■ snap reclaimable

```
system> snap reclaimable vol0 hourly.0 nightly.0
Processing (Press Ctrl-C to exit)
snap reclaimable: Approximately 47108 Kbytes would
be freed.
```

© 2010 NetApp, Inc. All rights reserved.

## SNAP DELTA AND SNAP RECLAIMABLE

`snap delta [ vol_name [ snap ] [ snap ] ]` Displays the rate of change of data between Snapshot copies. When used without any arguments it displays the rate of change of data between Snapshot copies for all volumes in the system, or all aggregates in the case of `snap delta -A`. If a volume is specified, the rate of change of data is displayed for that particular volume. The query can be made more specific by specifying the beginning and ending Snapshot copies to display the rate of change between them for a specific volume. If no ending Snapshot is listed, the rate of change of data between the beginning Snapshot copy and the active file system is displayed.

The rate of change information is displayed in two tables. In the first table each row displays the differences between two successive Snapshot copies. The first row displays the differences between the youngest Snapshot in the volume and the active file system. Each following row displays the differences between the next older Snapshot copy and the previous Snapshot copy, stepping through all of the Snapshot copies in the volume until the information for the oldest Snapshot copy is displayed. Each row displays the names of the two Snapshot copies being compared, the amount of data that changed between them, how long the first Snapshot copy listed has been in existence, and how fast the data changed between the two Snapshot copies.

The second table shows the summarized rate of change for the volume between the oldest Snapshot copy and the active file system.



## Snapshot Automatic Delete

- Snapshot automatic delete determines when (if) Snapshot copies will be automatically deleted

- Set at volume level:

```
snap autodelete vol [on|off|show|reset]
```

- If autodelete is enabled, then options:

```
snap autodelete vol option val
```

### Options Value

|                   |                                       |
|-------------------|---------------------------------------|
| commitment        | try, disrupt                          |
| trigger           | volume, snap_reserve, space_reserve   |
| target_free_space | 1-100                                 |
| delete_order      | oldest_first, newest_first            |
| defer_delete      | scheduled, user_created, prefix, none |
| prefix            | <string>                              |

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT AUTOMATIC DELETE





## **snap autodelete: commitment**

What Snapshot copies can `autodelete` remove?

- The user can protect certain kinds of Snapshot copies from deletion
- The commitment option defines:
  - `try`  
Deletes Snapshot copies that are not being used by any data mover, recovery, or clones (NOT LOCKED)
  - `disrupt`  
Deletes Snapshot copies locked by applications that move data (such as SnapMirror), dump data, and restore data (mirror and dumps are aborted)

© 2010 NetApp, Inc. All rights reserved.

## **SNAP AUTODELETE: COMMITMENT**



## **snap autodelete: trigger**

When does snap autodelete occur?

When the “trigger” criteria is nearly full:

- volume

The volume is nearly full (98%)

- snap\_reserve

The reserve is nearly full

- space\_reserve

The space reserved is nearly full (useful for volumes with fractional\_reserve < 100)

© 2010 NetApp, Inc. All rights reserved.

### **SNAP AUTODELETE: TRIGGER**



## **snap autodelete: target\_free\_space**

When does `snap autodelete` stop?

- When the free space in the trigger criteria reaches a user-specified percentage, `snap autodelete` stops
  - This percentage is controlled by the value of `target_free_space`
  - The default percentage is 20%

© 2010 NetApp, Inc. All rights reserved.

### **SNAP AUTODELETE: TARGET\_FREE\_SPACE**



## `snap autodelete: order`

In what order are Snapshot copies deleted?

- The `delete_order` option defines the age order. If the value is set to:
  - `oldest_first`  
Delete oldest Snapshot copies first
  - `newest_first`  
Delete newest Snapshot copies first

© 2010 NetApp, Inc. All rights reserved.

## SNAP AUTODELETE: ORDER



## snap autodelete: order (Cont.)

Snapshot copies are deleted in the following order:

- The `defer_delete` option defines the order for deletion
- If the value is set to:
  - `scheduled`  
Delete the scheduled Snapshot copies last (identified by the scheduled Snapshot naming convention)
  - `user_created`  
Delete the administrator-created Snapshot copies last
  - `prefix`  
Delete the Snapshot copies with names matching the prefix string last

© 2010 NetApp, Inc. All rights reserved.

## SNAP AUTODELETE: ORDER (CONT.)



## **snap autodelete: prefix**

In what order are Snapshot copies deleted?

- The `prefix` option value pair is only considered when `defer_delete` is set to `prefix`
- Otherwise, it is ignored

© 2010 NetApp, Inc. All rights reserved.

### **SNAP AUTODELETE: PREFIX**



# System Manager: Space Monitoring

The space consumed in the volume

| Name | Aggregate | Status | Available space | Used % | Total space |
|------|-----------|--------|-----------------|--------|-------------|
| vol0 | aggr0     | online | 19.34 GB        | 10     | 26.93 GB    |
| vol1 | aggr1     | online | 2 GB            | 0      | 2 GB        |

| Name      | Date Time         | Total Size | Cumulative Total Size | Status | Dependency |
|-----------|-------------------|------------|-----------------------|--------|------------|
| mySnap    | 6/16/2009 4:05... | 52 KB      | 52 KB                 | Normal | None       |
| hourly.0  | 6/16/2009 4:00... | 132 KB     | 184 KB                | Normal | None       |
| hourly.1  | 6/16/2009 12:0... | 116 KB     | 300 KB                | Normal | None       |
| hourly.2  | 6/16/2009 8:00... | 152 KB     | 452 KB                | Normal | None       |
| nightly.0 | 6/16/2009 12:0... | 124 KB     | 576 KB                | Normal | None       |
| hourly.3  | 6/15/2009 8:00... | 108 KB     | 684 KB                | Normal | None       |
| hourly.4  | 6/15/2009 11:0... | 104 KB     | 788 KB                | Normal | None       |

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: SPACE MONITORING



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY





## Module Summary

In this module, you should have learned to:

- Describe the function of Snapshot copies
- Explain the benefits of Snapshot copies
- Identify and execute Snapshot commands
- Create and delete Snapshot copies
- Configure and modify Snapshot options
- Explain the importance of the `.snapshot` directory
- Describe how disk space is allocated by a Snapshot copy for volumes and aggregates
- Schedule Snapshot copies
- Configure and manage the Snapshot reserve

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 17: Snapshot Copies  
Estimated Time: 30 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What is a Snapshot copy?
- What are some of the NetApp products that are based on Snapshot technology?
- What are some of the Snapshot commands?
- What is the Snapshot schedule syntax?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# SnapRestore

Module 18  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## SNAPRESTORE



## Module Objectives

By the end of this module, you should be able to:

- Explain how SnapRestore® technology works with Snapshot™ copies
- Describe what SnapRestore reverts
- Revert a volume or a file using SnapRestore
- Explain how SnapRestore works with SnapMirror®
- Describe the effects of SnapRestore on backup operations

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## SnapRestore

### Back up to Local Storage

- Recover single file, volume, or aggregate

- SnapRestore reverts a file system back to any specified Snapshot copy or restores single file from a Snapshot copy
  - Fast online restores of files, volumes, and aggregate
  - Multiple recovery points
  - Easy recovery process based on a single command input
  - Requires the `snaprestore` license code
- Use SnapRestore to recover from data corruption or to revert a file system

© 2010 NetApp, Inc. All rights reserved.

## SNAPRESTORE

SnapRestore enables you to quickly revert a local volume or a file on a storage system to the state it was in when a particular Snapshot copy was taken. In most cases, reverting a file or volume is much faster than restoring files from tape or copying files from a Snapshot copy to the active file system.

You use SnapRestore to recover from data corruption. If a primary storage system application corrupts data files in a volume, you can revert the volume or specified files in the volume to a Snapshot copy that was taken before the data corruption. You can also use SnapRestore if you are testing a volume or file and want to restore that volume or file to pretest conditions.

You must purchase and install the `snaprestore` license code to enable and use the SnapRestore service.




## SnapRestore of a Volume or Aggregate

© 2010 NetApp, Inc. All rights reserved.

### SNAPRESTORE OF A VOLUME OR AGGREGATE



## How SnapRestore Works



- The active file system, or AFS, points at each of its 4k data blocks


© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPRESTORE WORKS





## How SnapRestore Works (Cont.)




- The AFS is captured in a first Snapshot copy named hourly.0
- hourly.0 points at the same blocks as the AFS

© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPRESTORE WORKS (CONT.)



## How SnapRestore Works (Cont.)



- As time goes by, some new files are created and other files are modified
- New and modified data is written to new 4k blocks


© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPRESTORE WORKS (CONT.)

As time goes, new files are created and others are modified. New or modified data is written to new 4k blocks. The Snapshot copy hourly.0 still points to the disk blocks where the files existed before they were modified.



## How SnapRestore Works (Cont.)



- The AFS continues to be modified while a new Snapshot copy is taken
- The original hourly.0 becomes hourly.1
- Some data blocks have new pointers from the AFS and from hourly.0 and hourly.1

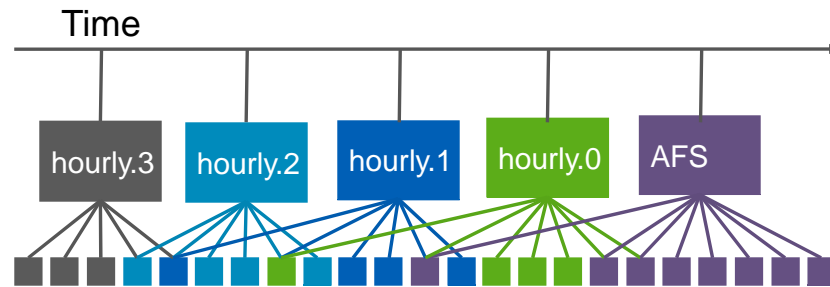
© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPRESTORE WORKS (CONT.)

The AFS continues to be modified while a new Snapshot copy, hourly.0 is taken. Some data blocks now have pointers from the AFS and from the Snapshot copies hourly.1 (the previous page's hourly.0) and hourly.0.



## How SnapRestore Works (Cont.)



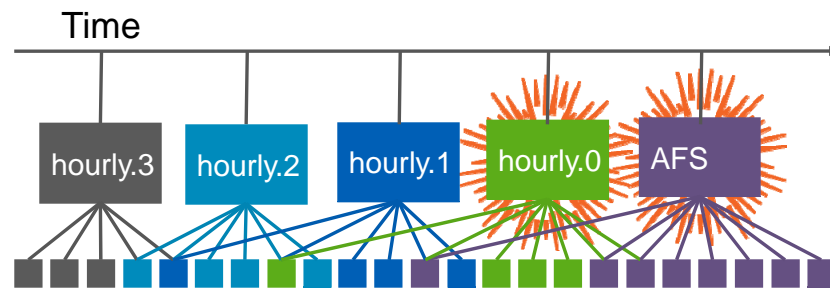
■ And so on...

© 2010 NetApp, Inc. All rights reserved.

## HOW SNAPRESTORE WORKS (CONT.)



## How SnapRestore Works (Cont.)



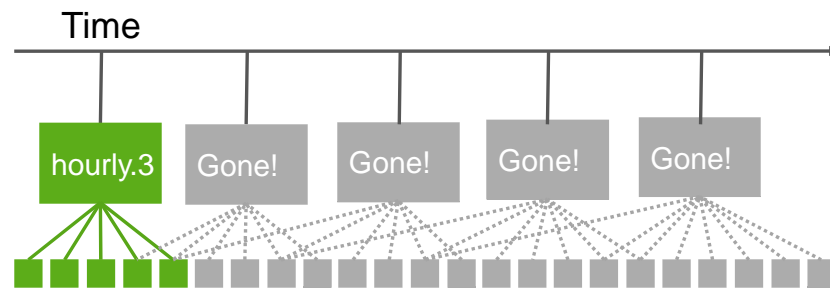
- A virus has attacked your data
  - The blocks that make the AFS are inconsistent
  - hourly.0 may point to inconsistent blocks as well
- You decide to revert to hourly.3 using SnapRestore technology

© 2010 NetApp, Inc. All rights reserved.

## HOW SNAPRESTORE WORKS (CONT.)



## How SnapRestore Works (Cont.)




- All data that existed before the reversion is gone
- Everything that was created after the former hourly.3 Snapshot copy ceases to exist
- Hourly.3 is renamed to what it was called back then: hourly.0

© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPRESTORE WORKS (CONT.)



## How SnapRestore Works (Cont.)



- All blocks included in hourly.0 (previously hourly.3) become the AFS

© 2010 NetApp, Inc. All rights reserved.

### HOW SNAPRESTORE WORKS (CONT.)



## Reverting a Volume

1. Verify that the aggregate is online and writable.
2. List the Snapshot copies in the volume:
3. Notify network users.
4. Initiate the restore:

```
system> snap list vol_name
```

```
system> snap restore -t vol -s snapshotname
/vol/volname
```

- Reverting a root volume:
  - Restores configuration files, including the registry
  - Will initiate a system reboot

© 2010 NetApp, Inc. All rights reserved.

## REVERTING A VOLUME





## Reverting a Aggregate

1. Verify that the volume is online and writable.
2. List the Snapshot copies in the volume:
3. Notify network users.
4. Initiate the restore:

```
system> snap list -A
```

```
system> snap restore -A -s snapshotname
 aggrname
```

### ■ NOTE:

- Reverting the aggregate to a Snapshot copy will affect all volumes within the aggregate
- Reverting an aggregate with the root volume requires a reboot

© 2010 NetApp, Inc. All rights reserved.

## REVERTING A AGGREGATE



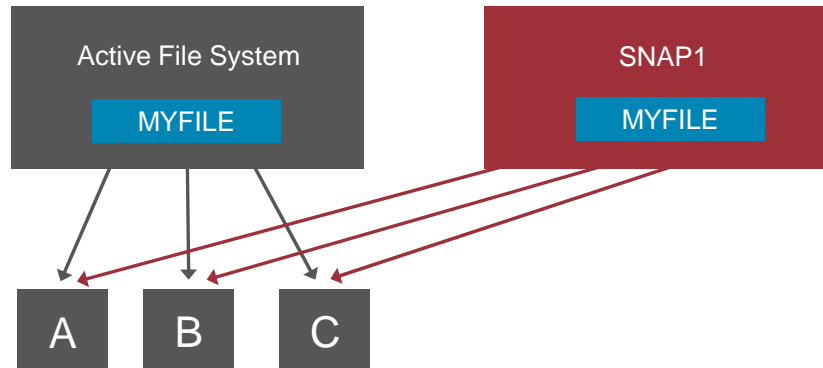
## SnapRestore of a Single File

© 2010 NetApp, Inc. All rights reserved.

### SNAPRESTORE OF A SINGLE FILE



## Use SnapRestore with a Single File



- MyFILE is made of disk blocks A, B, and C
- The AFS is captured in Snapshot copy SNAP1
- SNAP1 points to blocks A, B, and C and does not use additional disk space

© 2010 NetApp, Inc. All rights reserved.

### USE SNAPRESTORE WITH A SINGLE FILE

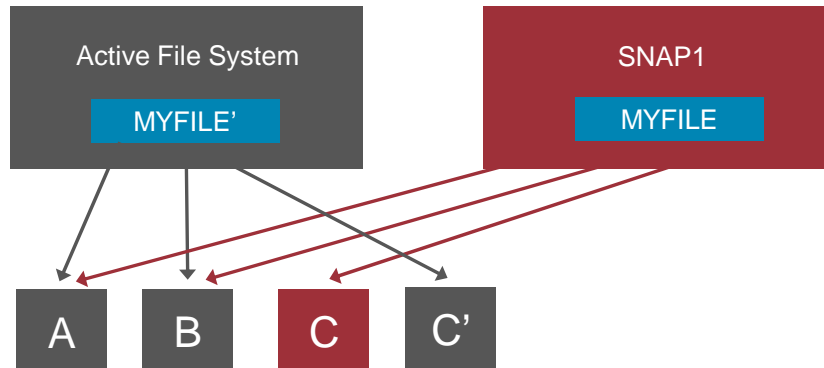
Data ONTAP® preserves pointers to all the disk blocks currently in use at the time the Snapshot copy is created.

In this illustration, the AFS contains a file named MYFILE made of three blocks: A, B, and C.

Based on a schedule or in response to the `snap create` command, the AFS is captured in a Snapshot copy named SNAP1. This Snapshot copy uses almost no additional disk space because the version of MYFILE within this Snapshot copy are the same blocks as the blocks for MYFILE in the AFS.



## Use SnapRestore with a Single File (Cont.)



- C block is modified; WAFL® writes the change in new disk block C'; MYFILE is now made of disk blocks A, B, and C'
- Snapshot copy SNAP1 still points to disk blocks A, B, and C

© 2010 NetApp, Inc. All rights reserved.

### USE SNAPRESTORE WITH A SINGLE FILE (CONT.)


When a file is changed, the Snapshot copy still points to the disk blocks where the file existed before it was modified, and changes are written to new disk blocks. Snapshot copies begin to consume extra space only when corresponding files in the AFS are changed or deleted.

In the illustration, a client modifies some data in MYFILE, causing the contents of block C to change. The WAFL (Write Anywhere File Layout) file system uses a copy-on-write policy that writes the modified block to a new location on disk, creating block C'. The AFS version of MYFILE is now composed of disk blocks A, B, and C', whereas the Snapshot copy SNAP1 still points to blocks A, B, and C.

In addition to the disk space used by block C' in the modified AFS, disk space used by the original blocks A, B, and C is still reserved in the Snapshot copy SNAP1. Deleting MYFILE in the AFS won't free-up disk blocks A, B, and C because the Snapshot copy SNAP1 still points to those blocks.



## Use SnapRestore with a Single File (Cont.)



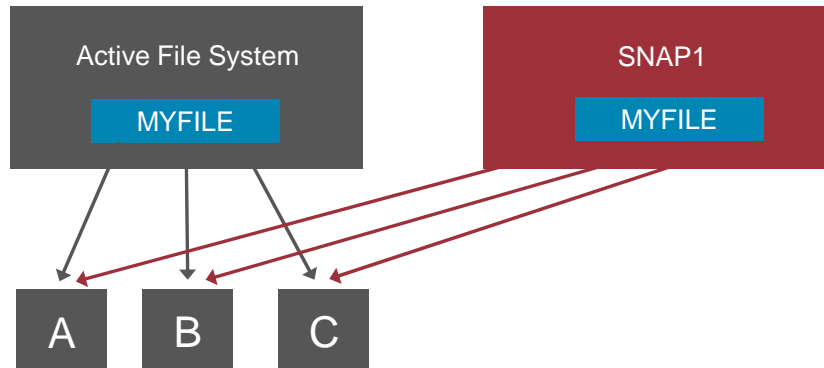
- During testing, let us say you want to revert back to MYFILE
- You could either:
  - Revert back by copying the data from the Snapshot directory
  - Use SnapRestore technology

© 2010 NetApp, Inc. All rights reserved.

## USE SNAPRESTORE WITH A SINGLE FILE (CONT.)



## Use SnapRestore with a Single File



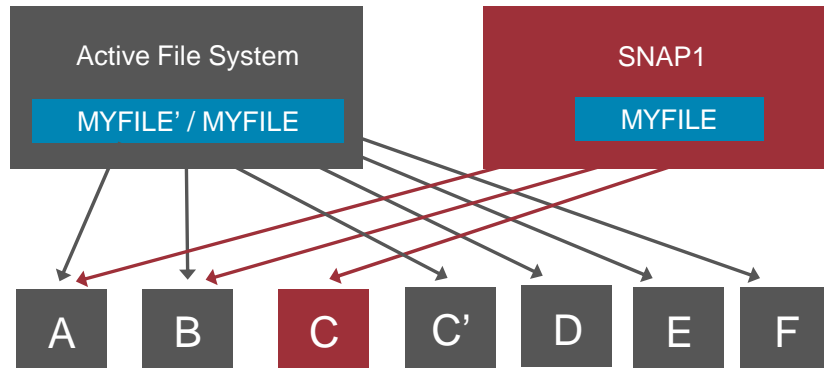
- With SnapRestore (to the original):
  - The original file (A, B, and C) is reverted back and is then managed by the AFS
  - The change block (C') is removed

© 2010 NetApp, Inc. All rights reserved.

## USE SNAPRESTORE WITH A SINGLE FILE



## Use SnapRestore with a Single File (Cont.)



- With SnapRestore (to different location than original):
  - The original file (A, B, and C) is copied to a new location on the disk (D, E, and F)
  - This new file can be placed in the same location as the original file or in an alternative location

© 2010 NetApp, Inc. All rights reserved.

## USE SNAPRESTORE WITH A SINGLE FILE (CONT.)



## Reverting a File

1. Verify that the volume is online and writable.
2. List the Snapshot copies in the volume:  

```
system> snap list volname
```
3. Notify network users.
4. Initiate the reversion:  

```
system> snap restore -t file -s snapshotname
path_and_file
```
5. Use the `-r` option to revert a file to a different location:  

```
system> snap restore -t file -s snapshotname
-r new_path_&_file old_path_&_file
```

© 2010 NetApp, Inc. All rights reserved.

## REVERTING A FILE

Follow these steps to revert a single file.

1. Verify that the volume is online and writable.
2. List the Snapshot copies in the volume.
  - ```
system> snap list /vol/vol_name
```
3. Notify network users that you are going to revert a file.
4. If you know the name of the Snapshot copy, initiate the restore using the following command:
 - ```
system> snap restore -t file -s snapshot_name path_and_file_name
```

    - `-t file` indicates that a file SnapRestore is to be performed.
    - `path_and_file_name` is the complete path to the name of the file to be reverted.

Data ONTAP displays a warning message and prompts you to confirm your decision to revert the file. Press **Y** to confirm that you want to revert the file. If you do not want to proceed enter Ctrl-C.

If the file already exists in the AFS, it will be overwritten with the version in the Snapshot copy.





## SnapRestore Versus Copying

- Using SnapRestore to revert a single file provides advantages over copying a single file when the file is large such as a database
  - Copying requires double the storage and time
  - Reverting saves time and reinstates the data
  - NetApp® recommends using SnapRestore technology over alternative technologies to ensure reliability

© 2010 NetApp, Inc. All rights reserved.

### SNAPRESTORE VERSUS COPYING

When restoring large quantities of data, it takes a long time to either copy files from a Snapshot copy or restore them from tape. In this case, using SnapRestore technology is the preferred method for recovering data because it saves time.



## SnapRestore Considerations

© 2010 NetApp, Inc. All rights reserved.

### SNAPRESTORE CONSIDERATIONS



## SnapRestore Rules

- You cannot undo a SnapRestore reversion
- Snapshot copy deletions are permanent
  - You cannot revert a volume to recover a deleted Snapshot copy
- While SnapRestore is in progress, Data ONTAP cannot delete and create Snapshot copies
  - Scheduled Snapshot copies will be suspended for the duration of the restore
  - A `dump` or other command attempts that depend on creating a Snapshot copy will fail

© 2010 NetApp, Inc. All rights reserved.

## SNAPRESTORE RULES



## SnapRestore and Backup Operations

- After a reversion, incremental backup and restore operations can no longer rely on the AFS timestamps
- Recommendations:
  - After the reversion, perform a level-0 backup
  - When restoring from tape, use only backups created after the volume reversion

© 2010 NetApp, Inc. All rights reserved.

### SNAPRESTORE AND BACKUP OPERATIONS



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Explain how SnapRestore technology works with Snapshot copies
- Describe what SnapRestore reverts
- Revert a volume or a file using SnapRestore
- Explain how SnapRestore works with SnapMirror
- Describe the effects of SnapRestore on backup operations

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 18: SnapRestore  
Estimated Time: 20 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What doesn't SnapRestore revert?
- Which option would you use in a `snaprestore` command to perform a file restore?
- Snapshot copy deletions are \_\_\_\_\_.

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING





Go further, faster®

# SnapVault

Module 19  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## SNAPVAULT



## Module Objectives

By the end of this module, you should be able to:

- Describe SnapVault® components and benefits
- Configure SnapVault on primary and secondary systems
- Administer SnapVault on primary and secondary systems
- Describe the application-consistent backup feature available in Data ONTAP® 8.0 7-Mode
- Restore data from secondary to primary systems

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## SnapVault Overview

© 2010 NetApp, Inc. All rights reserved.


### SNAPVAULT OVERVIEW



# SnapVault

Back up to Remote Storage  
- Recover qtrees, directories, or files

- Back up multiple storage systems to a central secondary system
- Minimize media consumption and system overhead through incremental backup
- Allow users to browse backed-up files online to perform restoration upon request



© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT

SnapVault is a disk-based storage backup feature of Data ONTAP. SnapVault enables data stored on multiple storage systems to be backed up to a central, secondary storage system quickly and efficiently as read-only Snapshot™ copies.

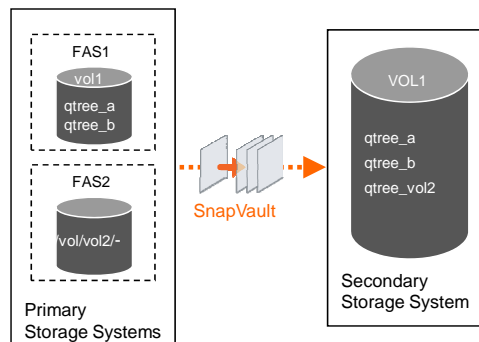
In the event of data loss or corruption on a storage system, backed-up data can be restored from the SnapVault secondary with less downtime and uncertainty than is associated with conventional tape backup and restore operations.

Additionally, users who wish to perform a restore of their own data may do so without the intervention of a system administrator. The SnapVault secondary may be configured with NFS exports and CIFS shares to let users copy the file from the Snapshot copy to the correct location.



## Theory of Operation

- The qtree is the basic unit of SnapVault backup and restore
- You can back up a primary qtree, non-qtree data, and volume data to a qtree on the secondary system
- If necessary, data is restored from the secondary qtrees back to their associated primary qtrees



© 2010 NetApp, Inc. All rights reserved.

## THEORY OF OPERATION

On storage systems running Data ONTAP, the qtree is the basic unit of SnapVault backup and restore. SnapVault backs up specified qtrees on the primary system to associated qtrees on the SnapVault secondary system. If data needs to be restored to the primary system, SnapVault transfers the specified versions of the qtrees back to their associated primary qtrees.

The non-qtree part of a primary system volume can be replicated to a SnapVault secondary qtree. Non-qtree data is any data on a storage system that is not contained in a qtree. The backed-up data can be restored to a qtree on the primary system, but cannot be restored as non-qtree data.

You can also back up a primary volume to a qtree on the secondary system. Any qtrees in the primary volume become directories in the secondary qtree. SnapVault cannot restore the data back to a volume. When restoring data, what was a source volume is restored as a qtree.

**NOTE** that volume-to-qtree backups are not supported for volumes containing Data ONTAP LUNs.



## Initial Transfer and Backup

- The initial transfer establishes the relationship between the primary and the secondary qtrees
- Incremental backup
  - The primary creates scheduled SnapVault Snapshot copies of the volume containing the qtrees to be backed up
  - The secondary carries out scheduled updates
    - Changed blocks are retrieved from the Snapshot copy of each primary qtree
    - When transfers complete, the secondary takes a Snapshot copy of its own volume

© 2010 NetApp, Inc. All rights reserved.

## INITIAL TRANSFER AND BACKUP

### INITIAL TRANSFER

In response to the `snapvault start` command, the secondary system requests initial transfers of qtrees specified for backup from the primary system volume to the secondary system volume. These transfers establish SnapVault relationships between the primary and secondary qtrees. To initialize qtrees, you do not need to create the qtrees on the secondary; the qtrees are created when the baseline transfers are started.

### INCREMENTAL BACKUP

In response to the `snapvault snap sched` command-line input, the primary system creates scheduled SnapVault Snapshot copies of the volume containing the qtrees to be backed up.

In response to the `snapvault snap sched -x` command-line input, the secondary system carries out scheduled update transfers and Snapshot copies creation.

For each secondary qtree, SnapVault retrieves, from the Snapshot data of each corresponding primary qtree, the incremental changes to the primary qtrees made since the last data transfer. Only the changed data blocks are sent to the secondary.

When the transfer is completed, the secondary takes a Snapshot copy of its own volume. Note that SnapVault does not transfer Snapshot copies; it only transfers selected data from within Snapshot copies.



## SnapVault Configuration

© 2010 NetApp, Inc. All rights reserved.

### SNAPVAULT CONFIGURATION



## Prerequisites

- A separate license for the primary, `sv_ontap_pri`, and for the secondary, `sv_ontap_sec`, is required
  - In Data ONTAP 7.3 and later, you can install both licenses on the same storage system
- You can increase the number of concurrent transfers by installing the `nearstore_option` license
- TCP port 10566 must be open on both sides and 10000 as well for central management (optional)

© 2010 NetApp, Inc. All rights reserved.

## PREREQUISITES

You must purchase and install a separate SnapVault license for each primary (`sv_ontap_pri`) and secondary (`sv_ontap_sec`) storage system.

SnapVault evaluation licenses are available upon request on the NOW™ (NetApp on the Web) site: [now.netapp.com/eservice/evallicense](http://now.netapp.com/eservice/evallicense)

In Data ONTAP 7.3 and later, you can install both the `sv_ontap_pri` and the `sv_ontap_sec` licenses on the same storage system. This system is then able to send and receive SnapVault backups, whether from other appliances or locally within itself.

### NOTE:

You cannot mix primary and secondary qtrees in the same volume, as this is unsupported and causes undesirable effects.

You cannot license a SnapVault secondary and a SnapVault primary on the same node of an active-active configured system.

Optionally, you can increase the number of possible concurrent streams on FAS2040, FAS3040, FAS3070, FAS3100 and FAS6000 storage systems by installing the `nearstore_option` license. This license should not be installed on these storage systems if they are intended to handle primary application workloads.

Port 10566 must be open in both directions for SnapVault backup and restore operations.

If NDMP is in use for control management, then port 10000 must be open on the primary and the secondary systems.





## Configuration Process

1. Identify the primary and secondary systems and the backup requirements
2. Configure the SnapVault primary
3. Configure the SnapVault secondary
4. Perform the initial baseline transfer

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURATION PROCESS



# 1. Hardware and Backup Requirements

- Identify the primary and secondary storage systems


– Console prompts:

- Primary: `pri>`
- Secondary: `sec>`

- Identify the qtrees and volumes that you are going to back up

– For example:

- Qtree: `/vol/vol1/users`
- Volume: `/vol/oracle`



© 2010 NetApp, Inc. All rights reserved.

## 1. HARDWARE AND BACKUP REQUIREMENTS



## 2. Configuration of Primary

- License SnapVault:

```
pri> license add xxxxxxxx
```

- Enable SnapVault:

```
pri> options snapvault.enable on
```

- Allow secondary to access primary:

```
pri> options snapvault.access host=sec
```

© 2010 NetApp, Inc. All rights reserved.

## 2. CONFIGURATION OF PRIMARY



## 2. Configuration of Primary (Cont.)

- SnapVault uses its own Snapshot copies so unless necessary, disable the normal Snapshot schedule

```
pri> snap sched vol1 0 0 0
```

```
pri> snap sched oracle 0 0 0
```

- Set up SnapVault Snapshot schedule

```
pri> snapvault snap sched vol1 sv_hourly 22@0-22
```

- This schedule is for the home directories volume vol1

- Creates hourly Snapshot copies, except 11:00 p.m.
- Keeps nearly a full day of hourly copies

```
pri> snapvault snap sched vol1 sv_daily 7@23
```

- This schedule is for the home directories volume vol1

- Creates daily Snapshot copy at 11:00 p.m.
- Retains a week's worth of daily copies

© 2010 NetApp, Inc. All rights reserved.

## 2. CONFIGURATION OF PRIMARY (CONT.)

Turn off the normal Snapshot schedules, which will be replaced by SnapVault Snapshot schedules.

```
pri> snap sched vol1 0 0 0
```

```
pri> snap sched oracle 0 0 0
```

Set up schedules for the home directory hourly Snapshot copies.

```
pri> snapvault snap sched vol1 sv_hourly 22@0-22
```

This schedule takes a Snapshot copy every hour, except for 11:00 p.m. It keeps nearly a full day of hourly copies, and combined with the daily or weekly backups at 11:00 p.m., ensures that copies from the most recent 23 hours are always available.

Set up schedules for the home directory daily Snapshot copies.

```
pri> snapvault snap sched vol1 sv_daily 7@23
```

This schedule takes a Snapshot copy once each night at 11:00 p.m. and retains the seven most recent copies.

The schedules created in Step 3 and Step 4 give 22 hourly and 7 daily Snapshot copies on the source to recover from before needing to access any copies on the secondary. This enables more rapid restores. However, it is not necessary to retain a large number of copies on the primary; higher retention levels are configured on the secondary.



### 3. Configuration of Secondary

- License SnapVault:

```
sec> license add xxxxxxxx
```

- Enable SnapVault:

```
sec> options snapvault.enable on
```

- Allow primary to access secondary:

```
sec> options snapvault.access host=pri
```

© 2010 NetApp, Inc. All rights reserved.

### 3. CONFIGURATION OF SECONDARY



### 3. Configuration of Secondary (Cont.)

- Create SnapVault destination FlexVol® volume:

```
sec> aggr create sv_flex 10
sec> vol create vault sv_flex 100g
```

- Optional: Set Snapshot reserve to zero

```
sec> snap reserve vault 0
```

- Turn off normal Snapshot copies schedule

```
sec> snap sched vault 0 0 0
```

© 2010 NetApp, Inc. All rights reserved.

### 3. CONFIGURATION OF SECONDARY (CONT.)

Create a FlexVol volume for use as a SnapVault destination.

```
sec> aggr create sv_flex 10
sec> vol create vault sv_flex 100g
```

The size of the volume should be determined by how much data you need to store and other site-specific requirements, such as the number of Snapshot copies to retain and the rate of change for the data on the primary FAS system.

Depending on site requirements, you may want to create several different SnapVault destination volumes. You may find it easiest to use different destination volumes for datasets with different schedules and Snapshot copy retention needs.

Optional: Set the Snapshot reserve to zero on the SnapVault destination volume.

```
sec> snap reserve vault 0
```

Due to the nature of backups using SnapVault, a destination volume that has been in use for a significant amount of time often has four or five times as many blocks allocated to Snapshot copies as it does to the active file system. Because this is the reverse of a normal production environment, many users find that it is easier to keep track of available disk space on the SnapVault secondary if SnapReserve is effectively turned off.

Turn off the normal Snapshot schedules, which will be replaced by SnapVault Snapshot schedules.

```
sec> snap sched vault 0 0 0
```



### 3. Configuration of Secondary (Cont.)

- Set up SnapVault Snapshot schedule

```
sec> snapvault snap sched -x vault sv_hourly
4@0-22
```

- This schedule checks all primary qtrees once per hour for a new Snapshot copy called sv\_hourly.0
  - Maintains 4 most recent copies

© 2010 NetApp, Inc. All rights reserved.

### 3. CONFIGURATION OF SECONDARY (CONT.)

Set up schedules for the hourly backups.

```
sec> snapvault snap sched -x vault sv_hourly 4@0-22
```

This schedule checks all primary qtrees backed up to the vault volume once per hour for a new Snapshot copy called sv\_hourly.0. If it finds such a copy, it updates the SnapVault qtrees with new data from the primary and then takes a Snapshot copy on the destination volume, called sv\_hourly.0.

**NOTE** that you are keeping only the four most recent hourly Snapshot copies on the SnapVault secondary. A user who wants to recover from a backup made within the past day has 23 backups to choose from on the primary FAS system and has no need to restore from the SnapVault secondary. Keeping four hourly Snapshot copies on the secondary merely ensures that you have at least the most recent four backups in the event of a major problem affecting the primary system.

**NOTE:** If you don't use the -x option, the secondary does not contact the primary and transfer the Snapshot copy. A Snapshot copy of the destination volume is merely created.



### 3. Configuration of Secondary (Cont.)

- Set up SnapVault Snapshot schedule (Cont.)

```
sec> snapvault snap sched -x vault sv_daily
12@23@sun-fri
```

- This schedule checks all primary qtrees each day at 11:00 p.m. for a new Snapshot copy called sv\_daily.0
  - Retains 12 most recent copies

© 2010 NetApp, Inc. All rights reserved.

### 3. CONFIGURATION OF SECONDARY (CONT.)

Set up schedules for the daily backups.

```
sec> snapvault snap sched -x vault sv_daily 12@23@sun-fri
```

This schedule checks all primary qtrees backed up to the vault volume once each day at 11:00 p.m. (except on Saturdays) for a new Snapshot copy called sv\_daily.0. If it finds such a copy, it updates the SnapVault qtrees with new data from the primary and then takes a Snapshot copy on the destination volume, called sv\_daily.0.

In this example, you maintain the most recent 12 daily backups, which, combined with the most recent 2 weekly backups (see next page).





### 3. Configuration of Secondary (Cont.)

- Set up SnapVault Snapshot schedule (Cont.)

```
sec> snapvault snap sched vault sv_weekly
13@23@sat
```

- This schedule creates a Snapshot copy of vault volume at 11:00 p.m. each Saturday for a new Snapshot copy called sv\_weekly.0
  - Retains 13 most recent copies

© 2010 NetApp, Inc. All rights reserved.

### 3. CONFIGURATION OF SECONDARY (CONT.)

Set up schedules for the weekly backups.

```
sec> snapvault snap sched vault sv_weekly 13@23@sat
```

This schedule creates a Snapshot copy of the vault volume at 11:00 p.m. each Saturday for a new Snapshot copy called sv\_weekly.0. There is no need to create the weekly schedule on the primary. Because you have all the data on the secondary for this Snapshot copy, you will simply create and retain the weekly copies on the secondary only.

In this example, you maintain the most recent 13 weekly backups, for a full 3 months of online backups.



## 4. Perform the Initial Baseline Transfer

- Initiate baseline for /vol/vol1/users on primary

```
sec> snapvault start -S pri:/vol/vol1/users
/vol/vault/pri_users
```

- Initiate baseline for /vol/oracle on primary

```
sec> snapvault start -S pri:/vol/oracle/-
/vol/vault/oracle
```

© 2010 NetApp, Inc. All rights reserved.

## 4. PERFORM THE INITIAL BASELINE TRANSFER

At this point, you have configured schedules on both the primary and secondary systems, and SnapVault is enabled and running. However, SnapVault does not yet know which qtrees to back up, or where to store them on the secondary. Snapshot copies will be taken on the primary, but no data will be transferred to the secondary.

To provide SnapVault with this information, use the SnapVault start command on the secondary:

```
sec> snapvault start -S pri:/vol/vol1/users
/vol/vault/pri_users

sec> snapvault start -S pri:/vol/oracle/- /vol/vault/oracle
```

If you later create another qtree called otherusers in the vol1 volume on the primary, it can be completely configured for backups with a single command:

```
sec> snapvault start -S
pri:/vol/vol1/otherusers/vol/vault/pri_otherusers
```

No additional steps are needed because the Snapshot schedules are already configured on both the primary and secondary for that volume.



## SnapVault Administration

© 2010 NetApp, Inc. All rights reserved.

### SNAPVAULT ADMINISTRATION



## SnapVault Administration

- After the initial relationship and schedule has been created between the primary and secondary, the storage administrator can:
  - Check the status of the transfer
  - Display a list of vaulted Snapshot copies
  - Manually update individual datasets
  - Create a Snapshot copy of the vault
  - Verify information using SnapVault logs
  - Configure Data ONTAP deduplication with SnapVault

© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT ADMINISTRATION



## Transfer Status

- To check the status of the most recent transfer:

```
sec> snapvault status
```

- To get detailed information on recent transfer:

```
sec> snapvault status -l
```

- To list all Snapshot copies scheduled:

```
pri> snapvault status -s
```

```
sec> snapvault status -s
```

- To list the configuration of all secondary qtrees:

```
sec> snapvault status -c
```

© 2010 NetApp, Inc. All rights reserved.

## TRANSFER STATUS



## Listing Snapshot Copies

- List Snapshot copies for source on the primary:

```
pri> snap list -q voll
Volume voll working...
qtree contents date source

sv_hourly.0 (Jan 22 15:00)
users Replica Jan 22 15:00 pri:/vol/voll/user
```

- List Snapshot copies for qtrees on the secondary:

```
sec> snap list -q vault
or
sec> snap list -o /vol/vault/pri_users
sec> snap list -o /vol/vault/oracle
```

© 2010 NetApp, Inc. All rights reserved.

## LISTING SNAPSHOT COPIES

Use the `snapvault status` command either from the primary or the secondary system to check the status of a data transfer, and to see how recently a qtree has been updated.

```
snapvault status [option] hostname:/vol/vol_name/qtree_name
```

Options can be one or more of the following:

- c lists all the secondary system qtrees, their corresponding primary system qtrees, maximum speed of scheduled transfers, and maximum number of times SnapVault attempts to start a scheduled transfer before skipping that transfer. This option can be run only from the secondary system.
- l displays the long format of the output, which contains more detailed information.
- s lists all the Snapshot copies scheduled on the primary or secondary storage system. Information includes volume, Snapshot copy base name, current status, and Snapshot copy schedule.

For details on the `snapvault status` command output, refer to the *Data ONTAP Data Protection Online Backup and Recovery Guide*.



## SnapVault Qtree Snapshot Copies

### ■ Example output of `snap list -o`:

```
sec> snap list -o /vol/vault/oracle
Qtree /vol/vault/oracle
working...
date source name

Nov 18 18:55 pri:/vol/oracle sv_hourly.0
Nov 18 18:55 pri:/vol/oracle sec(0007462703)_vault-base.0
Nov 18 18:50 pri:/vol/oracle sv_nightly.0
```

© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT QTREE SNAPSHOT COPIES



## Manually Updating the Vault

- After data on the primary has updated, administrators can initiate a manual update of the secondary vault by:

```
sec> snapvault update /vol/vol1/users
```

Primary location  
↙

- NOTE: Normally, we would have to wait for a scheduled Snapshot to occur before the data would be backed to the vault

- Verify:

```
sec> snap list -o /vol/vault/pri_users
```

© 2010 NetApp, Inc. All rights reserved.

## MANUALLY UPDATING THE VAULT

Update information in the primary source location of /vol/vol1/users. Normally, we would have to wait for a scheduled Snapshot to occur before data would be backed to the vault. We can force an unscheduled update by issuing the following command on the secondary:

```
sec> snapvault update /vol/vol1/users
```

SnapVault updates the qtree on the secondary storage system with the data from a new Snapshot copy of the qtree it creates on the primary storage system.

```
sec> snap list -o /vol/vault/pri_users
```





## Manual Snapshot Copies

- Like any volume managed by Data ONTAP, the storage administrator may create a Snapshot copy of the vault volume

```
sec> snapvault snap create voll sv_nightly
```

© 2010 NetApp, Inc. All rights reserved.

### MANUAL SNAPSHOT COPIES

Because we just carried out a manual update of a secondary qtree, we might want to immediately incorporate that update into the retained Snapshot copies on the secondary storage system.

```
sec> snapvault snap create voll sv_nightly
```

SnapVault creates a new Snapshot copy and, based on the specified Snapshot copy basename, numbers it just as if that Snapshot copy had been created by the SnapVault schedule process. SnapVault names the new Snapshot copy `sv_nightly.0`, renames the older Snapshot copies, and deletes the oldest `sv_nightly` Snapshot copy.

The `snapvault snap create` command does not update the data in the secondary storage system qtree from the data in the primary storage system prior to creating the new Snapshot copy.



## Log Files

- SnapVault logs are stored on the root volume, in the `/etc/log/snapmirror` file

- For example, on the primary

```
pri Wed Jan 22 08:55:52 PDT pri:/vol/vol1/users
sec:/vol/vault/pri_usersRequest(192.168.30.6)
pri Wed Jan 22 08:55:53 PDT pri:/vol/vol1/users
sec:/vol/vault/pri_usersStart
pri Wed Jan 22 08:57:31 PDT pri:/vol/vol1/users
sec:/vol/vault/pri_usersEnd (36900 KB)
```

- For example, on the secondary

```
sec Wed Jan 22 08:56:28 PDT pri:/vol/vol1/users
sec:/vol/vault/pri_usersRequest (Initialize)
sec Wed Jan 22 08:56:30 PDT pri:/vol/vol1/users
sec:/vol/vault/pri_usersStart
sec Wed Jan 22 08:58:08 PDT pri:/vol/vol1/users
sec:/vol/vault/pri_usersEnd (36084 KB)
```

© 2010 NetApp, Inc. All rights reserved.

## LOG FILES

The SnapVault logs record whether the transfer finished successfully or failed. If there is a problem with the updates, it is useful to look at the log file to see what has happened since the last successful update. The logs include the start and end of each transfer, along with the amount of data transferred.

The SnapVault logs' information is stored on the primary and secondary storage systems root volume, in the `/etc/log/snapmirror` file.



## Application-Consistent Backup

© 2010 NetApp, Inc. All rights reserved.

### APPLICATION-CONSISTENT BACKUP



## Named Snapshot Feature for SnapVault

- Data ONTAP 7.3.1 and earlier could not back up data from a user-specified Snapshot copy on the destination
  - Backups used the last created Snapshot copy
- Data ONTAP 8.0 7-Mode SnapVault allows the administrator to choose the Snapshot copy to transfer

© 2010 NetApp, Inc. All rights reserved.

### NAMED SNAPSHOT FEATURE FOR SNAPVAULT

This feature allows customers to back up data using SnapVault from any arbitrary Snapshot copy at the disaster recovery site.

In Data ONTAP 7.3.1 and earlier releases, SnapVault could not back up data from a specified Snapshot copy that was residing on the volume SnapMirror® destination volume. SnapVault would only transfer data from the latest Snapshot copy that was created by volume SnapMirror.

In Data ONTAP 8.0, SnapVault can back up data from any arbitrary Snapshot copy (either a user-specified or scheduled Snapshot copy) from the volume SnapMirror destination. The SnapVault backup from the disaster recovery site continues to be the same as the SnapVault backup from primary storage system to secondary, with following restrictions:

For a SnapVault scheduled update from the disaster recovery site, administrators need to set up the SnapVault primary schedule at the volume SnapMirror source.

In the case of a SnapVault update from a named Snapshot (that is, `snapvault update -s <snapname>`), the administrator needs to make sure that the named Snapshot copy exists at the volume SnapMirror source. To prevent any Snapshot copy from getting deleted by Data ONTAP applications, use the new `snapvault preserve` command.

Data ONTAP 8.0 only supports SnapVault in 7-Mode.



# SnapVault Preservations

## ■ Creating Preseervations:

```
pri> snapvault snap preserve
```

```
usage: snapvault snap preserve volname snapname [tagname]
```

```
pri> snapvault snap preserve oracle snap1 oracle
```


```
pri> snap list oracle
```

```
Volume oracle
```

```
working...
```

| %/used   | %/total | date         | name       |
|----------|---------|--------------|------------|
| -----    | -----   | ---          | -----      |
| 29%(29%) | 0%( 0%) | Jun 24 08:07 | snap1(acs) |

This is not the tag name;  
stands for Application-Consistent Snapshot



© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT PRESERVATIONS



## SnapVault Preservations (Cont.)

### ■ Viewing Preservations:

```
pri> snapvault snap preservations oracle
snapid snapname

1 snap1
pri> snapvault snap preservations primary snap1
type tagname

cli oracle
```

### ■ Removing Preservations:

```
pri> snapvault snap unpreserve
usage: snapvault snap unpreserve volname snapname
 {[tagname] | [-all]}
pri> snapvault snap unpreserve oracle snap1 oracle
```

© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT PRESERVATIONS (CONT.)



## Restore Data from Secondary to Primary

© 2010 NetApp, Inc. All rights reserved.

### RESTORE DATA FROM SECONDARY TO PRIMARY



## Single-File Restore

- To restore a single file from a secondary to a primary you can:
  - Copy the file from a Snapshot copy using CIFS or NFS; or
  - Issue the Data ONTAP `ndmcopy` command from the primary system; or
  - Use NetApp Protection Manager

© 2010 NetApp, Inc. All rights reserved.

### SINGLE-FILE RESTORE

Users who wish to perform a restore of their own data may do so without the intervention of a system administrator. The SnapVault secondary may be configured with NFS exports and CIFS shares to let users copy the file from the Snapshot copy to the correct location.

NOTE: SnapVault backups transfer all of the file permissions and access control lists held by the original data; if users were not authorized to access a file on the original file system, they will not be authorized to access the backup copies of that file. This allows self-service restores to be performed safely.

To restore a single file, you can also use the Data ONTAP `ndmcopy` command or the NetApp Protection Manager software (if deployed).

For information on the `ndmcopy` command, refer to the *Data ONTAP na\_ndmcopy* manual page.





## Qtree or Volume Qtree

- In the event of data loss on the primary, restoring data involves two steps:
  1. Use the `snapvault restore -S` command to restore a backed-up qtree from its last saved update
  2. Use the `snapvault start -r` command to restart the SnapVault backup relationship or release it with the `snapvault release` command
- Administrators can restore the data to an existing qtree on the primary system using:
  - Baseline restore
  - Incremental restore
    - Restores data in a nondisruptive way for datasets containing LUNs

© 2010 NetApp, Inc. All rights reserved.

## QTREE OR VOLUME QTREE

You use the `snapvault restore` command to restore a backed-up qtree saved to the secondary system.

Starting with Data ONTAP 7.3 and later, you can restore the data to an existing qtree on the primary storage system using a baseline restore or incremental restore.

```
snapvault restore [option] -s snapname -S
 sec_system:/vol/volname/sec_qtree
 pri_system:/vol/volname/pri_qtree
```

Options can be one or more of the following:

- The `-f` option forces the command to proceed without first asking for confirmation from the user.
- The `-k` option sets the maximum transfer rate in kilobytes per second.
- The `-r` option attempts an incremental restore. The incremental restore can be used to revert the changes made to a primary system qtree since any backed-up version on the secondary system.
- The `-s` option specifies that the restore operation must be from the Snapshot *snapname* on the secondary system.
- The `-w` option causes the command not to return after the baseline transfer starts. Instead, it waits until the transfer completes (or fails). At that time, it prints the completion status and then returns.

Starting with Data ONTAP 7.3 and later, the SCSI connectivity of applications to all LUNs within the qtree being restored will be maintained throughout the restore process in order to make the restore operation **nondisruptive** to applications. However, I/O operations will not be allowed during the restore operation. Only baseline restores and incremental restores can be nondisruptive.



## Baseline Restore of a Qtree

### ■ To restore a qtree:

```
pri> snapvault restore -S sec:/vol/vault/oracle
 /vol/oracle
```

Restore from  
secondary

Restore to  
primary

Restore will overwrite existing data in  
/vol/oracle.

Are you sure you want to continue (yes/no)? **Yes**

Transfer started.

Monitor progress with 'snapvault status' or the  
snapmirror log.

- We can replay the database logs (stored on a separate volume) to recover the complete database

© 2010 NetApp, Inc. All rights reserved.

## BASELINE RESTORE OF A QTREE

In this scenario, we have users home directories and an Oracle® database backup to SnapVault. Let us assume the database has become corrupt. The database logs are stored properly on another volume and what we want to do is to restore the database from the vault and then replay the database log files to get it back up to the point of corruption. We intend to restore a primary qtree to the exact qtree location on the primary storage system from which we backed it up, therefore we need to delete the existing qtree first from the primary storage system. We will delete the qtree by way of the regular CIFS or NFS routine.

We will now restore the qtree from the secondary to the primary by issuing the following command:

```
pri> snapvault restore -S sec:/vol/vault/oracle /vol/oracle
```

We can now replay the database logs to recover the complete database.

Use the -f flag to override the confirmation prompt and directly proceed with the restore.



## Incremental Restore

- Transfers only incremental changes from the secondary qtree to the specified primary qtree

```
pri> snapvault restore -r -S
 sec:/vol/vault/pri_users /vol/vol1/users
Restore will overwrite existing data in
/vol/vol1/users
Are you sure you want to continue (yes/no) ? Yes
Transfer started.
Monitor progress with 'snapvault status' or the
snapmirror log.
```

- If the incremental restore fails due to a lack of common Snapshot copies, attempt an in-place baseline restore

© 2010 NetApp, Inc. All rights reserved.

## INCREMENTAL RESTORE

SnapVault incremental restore is based on qtree SnapMirror resync-style Snapshot copy negotiation.

The primary qtree is resynced to the specified Snapshot copy on the secondary.

The resync rolls back the primary qtree to the specified Snapshot copy and an incremental restore transfer is initiated from the specified Snapshot copy.

The restore operation transfers only incremental changes from the secondary qtree to the specified primary qtree.

You use the new `-r` option to perform a SnapVault incremental restore.

Example:

```
pri> snapvault restore -r -S sec:/vol/vault/pri_users /vol/vol1/users
Restore will overwrite existing data in /vol/volname/pri_qtree
Are you sure you want to continue (yes/no)? Yes
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

When you want to restore over an existing primary qtree, it is recommended that you first attempt an incremental restore. If the incremental restore fails due to lack of common Snapshot copies, then attempt an in-place baseline restore. This is because the incremental restore is a more efficient restore.



## Nondisruptive Restore

- SCSI connectivity to the LUNs is maintained throughout the in-place baseline and incremental restores
  - The LUNs' attributes in the primary qtree are reserved in a temporary staging area
- When the restore completed, the LUNs' attributes are applied to the restored LUNs

```
pri> lun show staging
/vol/san_voll/Staging_19e45590-8948-11dc-
bb15-00a09802437a_19999999999999999999/LUN1
100m (104857600) (r/w, online, mapped)
```

© 2010 NetApp, Inc. All rights reserved.

## NONDISRUPTIVE RESTORE

SCSI connectivity to the LUNs is maintained throughout the in-place and incremental restores by way of the following process:

The primary qtree is made read-only. The LUNs' attributes in the primary qtree are reserved in a temporary staging area. LUN maps are updated. External SCSI requests are processed using the information stored in the staging area. Hosts see the same LUNs with the same drive letters at all times.

To display the LUNs in the preserved staging qtrees, use the command `lun show staging` command.

Example:

Original LUN location: `/vol/san_voll/qtree1/LUN1`

Staging qtree naming convention:

`/vol/<vol_name>/Staging_<Volume_UUID>_<Transaction_ID>`

```
pri> lun show staging
```

```
/vol/san_voll/Staging_19e45590-8948-11dc-bb15-
00a09802437a_19999999999999999999/LUN1 100m (104857600) (r/w,online,mapped)
```

When the restore has completed, the LUNs' attributes that are stored in the staging area are applied to the restored LUNs. The primary qtree is 'broken' to be write-enabled, and I/O operations are resumed to the restored LUNs.



## LUN Clone Backup

- With Data ONTAP 7.3 and later, SnapVault is able to back up LUN clones in optimized mode using SnapDrive® for Windows®
  - LUN clones are transferred as clones
  - Space savings with the parent LUN is preserved
  - After the SnapVault relationship is handed off to SnapDrive for Windows, update transfers must not be run from the command-line interface
  - The backing Snapshot copy is locked on the secondary after the backup Snapshot copy is transferred

© 2010 NetApp, Inc. All rights reserved.

### LUN CLONE BACKUP

When integrated with Data ONTAP, SnapDrive® for Windows® or SnapDrive for Windows with Microsoft® Volume Shadow Copy Service creates two Snapshot copies upon LUN backup:

A backing Snapshot copy containing the LUN to be cloned

A backup Snapshot copy containing both the LUN and the clone

In versions of Data ONTAP earlier than 7.3, SnapVault backs up a LUN clone as a new LUN during the initial baseline transfer. Therefore, the LUN clone and its backing LUN get replicated as two separate LUNs on the secondary.

With Data ONTAP 7.3 or later, SnapVault is able to back up LUN clones in optimized mode using SnapDrive for Windows. The LUN clones are transferred as clones and space savings with the parent LUN is preserved. The SnapVault initial baseline transfer is performed at the command-line interface but after the SnapVault relationship is handed off to SnapDrive for Windows, transfers must not be run from the command-line interface.

On the secondary, the backing Snapshot copy is locked after the backup Snapshot copy is transferred.

#### LIMITATIONS:

In optimized mode, the primary qtree must not contain LUN clones.

The transfer fails if the backing Snapshot copy is missing on the secondary.

A SnapVault restore will also fail if the backing Snapshot copy is missing on the primary.

Finally, in optimized mode, cascades from volume SnapMirror destinations are not supported.



## Restarting or Releasing SnapVault

- Remove a residual SnapVault Snapshot copy

```
pri> snap list oracle
```

Snapshot copy  
to delete

```
pri> snap delete oraclepri(1990911275)_oracle-sec.2
```

- To restart the SnapVault backup relationship:

```
sec> snapvault start -r -s pri:/vol/oracle
sec:/vol/vault/oracle
```

- To remove the SnapVault backup relationship:

```
sec> snapvault release /vol/vault/oracle
pri:/vol/oracle
```

© 2010 NetApp, Inc. All rights reserved.

## RESTARTING OR RELEASING SNAPVAULT

When you use the `snapvault restore` command to restore a primary qtree, SnapVault places a residual SnapVault Snapshot copy on the volume of the restored primary qtree. This Snapshot copy is not automatically deleted. If you have configured this volume to retain the maximum 255 Snapshot copies allowed by Data ONTAP, you must manually delete this residual Snapshot copy, or else no new Snapshot copies can be created.

We now will remove a residual SnapVault Snapshot copy to insure a proper functioning SnapVault relationship.

```
pri> snap list oracle
```

Find the residual Snapshot. It is will be distinguished by the following syntax: *Primaryhost (nvram-id)\_primaryvolume\_restoredqtree-dst.2* For example: `pri (1990911275)_oracle_tree-sec.2`

Remove this qtree with the following command:

```
pri> snap delete oracle pri(1990911275)_oracle-sec.2
```



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe SnapVault components and benefits
- Configure SnapVault on primary and secondary systems
- Administer SnapVault on primary and secondary systems
- Describe the application-consistent backup feature available in Data ONTAP 8.0 7-Mode
- Restore data from secondary to primary systems

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY





Go further, faster®

## Exercise

Module 19: SnapVault  
Estimated Time: 30 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What is the basic unit for a SnapVault backup and restore?
- True or false? In Data ONTAP 7.3 and later, you can install both the `sv_ontap_pri` and the `sv_ontap_sec` licenses on the same storage system.
- Which `snapvault` command and option would you use to perform an incremental restore?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# Open Systems SnapVault

Module 20  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## OPEN SYSTEMS SNAPVAULT



## Module Objectives

By the end of this module, you should be able to:

- Describe how Open Systems SnapVault® integrates with Data ONTAP® SnapVault
- List Open Systems SnapVault advanced features
- Configure and administer Open Systems SnapVault
- Perform Open Systems SnapVault backup and restore operations
- Troubleshoot and resolve Open Systems SnapVault transfer failures

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Open Systems SnapVault Overview

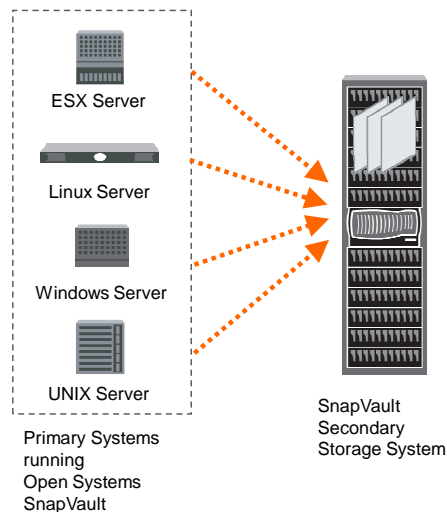
© 2010 NetApp, Inc. All rights reserved.

### OPEN SYSTEMS SNAPVAULT OVERVIEW



## Open Systems SnapVault

- The primary is a Microsoft® Windows® or a UNIX®-based open system
- The Open Systems SnapVault agent is installed on the primary and enables the system to back up its data to the secondary
- The secondary backs up the primary data using the Data ONTAP SnapVault technology



© 2010 NetApp, Inc. All rights reserved.

### OPEN SYSTEMS SNAPVAULT

Open Systems SnapVault is a disk-to-disk data protection solution that takes advantage of the NetApp SnapVault technology to protect data residing on the following platforms:

Microsoft Windows

Red Hat® Enterprise Linux

Novell® SUSE® Linux Enterprise Server

Sun™ Solaris

IBM® AIX®

HP HP-UX®

VMware® ESX Server

For a complete list of currently supported versions of these platforms, refer to the *Open Systems SnapVault 2.6 Installation and Administration Guide* on the NOW (NetApp on the Web) site.


The Open Systems SnapVault agent software module is installed on the primary and enables the system to back up its data to the secondary.

The secondary system is a NetApp storage appliance, such as a NearStore system. The secondary backs up the primary data using the Data ONTAP SnapVault technology.



## Theory of Operation

- The directory is the basic unit of Open Systems SnapVault backup
- Incremental update is based on file modification time
  - If block-level incremental is enabled, only changed blocks are transferred
- Snapshot copy creation for archiving takes place on the SnapVault secondary



© 2010 NetApp, Inc. All rights reserved.

## THEORY OF OPERATION

The directory is the basic unit of Open Systems SnapVault backup. Each directory, file, or drive you want to back up from a primary is backed up to its own qtree on the SnapVault secondary storage system.

For the first backup, you perform an **initial** baseline transfer of the identified directory or file. This transfer establishes the SnapVault relationship between the Open Systems SnapVault primary directory or file and its mapped SnapVault secondary qtree.

Subsequent transfers can either be initiated manually or configured for automatic scheduled updates. You specify the **schedules** on the secondary using the Data ONTAP `snapvault` command or using an NDMP-based management application such as Protection Manager or a third-party NDMP-based supported application.

Upon an **incremental** update request, the Open Systems SnapVault primary determines whether the directory or file has changed since the last successful transfer by examining the file modification time. If block-level incremental (BLI) is enabled (for Windows platforms only), Open Systems SnapVault determines which blocks' data changed and sends only those modified blocks to the secondary. After the secondary qtree is updated, SnapVault creates a Snapshot copy of the volume for archiving.

If a directory or file data needs to be restored to the primary, SnapVault retrieves the data from one specified Snapshot copy and transfers the data back to the primary system that requests it.



## Central Management

- Open Systems SnapVault backup schedules, retention policies, backup control, and monitoring can be centrally managed by the following NDMP-based applications:
  - NetApp Protection Manager
  - BakBone® NetVault®
  - Syncsort® Backup Express
  - CommVault® Galaxy®

© 2010 NetApp, Inc. All rights reserved.

## CENTRAL MANAGEMENT

Open Systems SnapVault can be managed from a variety of management applications. These applications use the NDMP protocol, TCP port 10000, to communicate with the Open Systems SnapVault clients and the storage systems over a TCP/IP network. Backup schedules, retention policies, backup control, and monitoring are centrally configured on these applications.

The applications that you can use to manage Open Systems SnapVault are as follows:

NetApp Protection Manager

BakBone, NetVault

Syncsort Backup Express

CommVault Galaxy

NOTE: The NetApp Host Agent is required to manage Open Systems SnapVault using Protection Manager and it is packaged with the Open Systems SnapVault software module.


For more information on central management using the NetApp Protection Manager, refer to the *Protection Manager* module of this course.





## Integrated Backup and DR Solution

- SnapVault based backup copies are replicated through SnapMirror® to a disaster recovery, or DR, site
- If the SnapVault secondary becomes unusable, you can redirect the subsequent Open Systems SnapVault transfers to the tertiary storage system



© 2010 NetApp, Inc. All rights reserved.

### INTEGRATED BACKUP AND DR SOLUTION

The Open Systems SnapVault primary systems have the same advantages as any other primary systems in a data protection and disaster recovery scenario.

You can protect the SnapVault secondary system from disasters by using the SnapMirror feature. The configuration involves setting up SnapMirror relationships from the volumes on the SnapVault secondary system to volumes on a remote (tertiary) NetApp storage system, as shown in the illustration.

If the SnapVault secondary becomes unusable because of a disaster, you can manually redirect the subsequent Open Systems SnapVault transfers to the tertiary system. Effectively, the tertiary system becomes the new SnapVault secondary, and the Open Systems SnapVault transfers continue using the most recent Snapshot copy common to both the primary and tertiary storage systems.



## OSSV Features

© 2010 NetApp, Inc. All rights reserved.

## OSSV FEATURES



## Backing Up Open Files

- Open File Manager (OFM)
  - Backs up open files on W2K platforms
  - Requires the `sv_windows_ofm_pri` license
  - Supports one active Snapshot copy per drive
- Volume Shadow Copy Service (VSS)
  - Backs up open files on W2K3 platforms
- OFM and VSS
  - Are integrated with the Open Systems SnapVault agent
  - Are tunable using the Open Systems SnapVault Configurator
- To disable open files backup on the secondary:  

```
snapvault modify -o back_up_open_files=off
```

© 2010 NetApp, Inc. All rights reserved.

## BACKING UP OPEN FILES

Open File Manager (OFM) allows Windows Server 2000 files that are open and in use to be backed up with only a very short disruption to users or their current applications. OFM is automatically installed at the same time the Open Systems SnapVault agent for Windows 2000 Server is installed. OFM is then enabled as soon as the Open Systems SnapVault agent installation completes, the primary system is rebooted, and the OFM component, `sv_windows_ofm_pri` is licensed on the secondary.

When an update is invoked, a Snapshot copy of the drive to be backed up is taken by OFM. To do so, OFM produces a temporary drive letter, and the Open Systems SnapVault agent backs up the files from that temporary drive's Snapshot copy. OFM supports one active Snapshot copy per disk volume letter. If multiple drives are backed up simultaneously, backups will proceed, but the drives will be backed up sequentially.

Microsoft Windows Server 2003 provides a native Snapshot copy mechanism as part of the Volume Shadow Copy Service (VSS). VSS Snapshot copy functionality is integrated with the Open Systems SnapVault agent as a standard feature. Unlike the OFM implementation, it is possible to have multiple copies of a shadow copy per drive.

OFM and VSS tunable parameters and drive exclusions can be configured using the Open Systems SnapVault Configurator. OFM and VSS settings should be changed only for troubleshooting. Open file backup can be disabled from the secondary using the `snapvault modify -o back_up_open_files=off` command.



## Snapshot Copies Management

- Open Systems SnapVault has the ability to retain Snapshot copies
- There are two possible configurations for common Snapshot management:
  - `MaxCPRestartWaitTime`
    - Set the maximum waiting time a Snapshot copy is retained after a transfer failure; default is 10 minutes
  - `FailCPRestartOnNewSnapshot`
    - If a Snapshot copy is not available during restart
      - When set to `FALSE`, the transfer to continues using a new Snapshot copy
      - When set to `TRUE` the transfer is aborted
- Use the `svsetstanza` command or edit the `snapvault.cfg` file to change these values

© 2010 NetApp, Inc. All rights reserved.

## SNAPSHOT COPIES MANAGEMENT

Common Snapshot copies management ensures that the same Snapshot copy is used for backup.

In releases earlier than Open Systems SnapVault 2.5, whenever the transfer of files failed, the Snapshot copies were deleted and a new Snapshot copy was created during the transfer restart.

Open Systems SnapVault 2.5 has the ability to retain old Snapshot copies and to use these copies subsequently during transfer restarts. There are two possible configurations for common Snapshot copies management:

`MaxCPRestartWaitTime` is the maximum waiting time (default is 10 minutes) a Snapshot copy is retained after the transfer failure. If the transfer of files restarts after the maximum waiting time, the Snapshot copy is lost and needs to be created again.

`FailCPRestartOnNewSnapshot` is the corresponding Snapshot copy is not available during restart of a transfer due to a system restart, elapsed time, or Open Systems SnapVault restart, either you allow the transfer to continue using a new Snapshot copy or abort the transfer. When the value is set to `TRUE`, the transfer is aborted.

When the value is set to `FALSE`, a new Snapshot copy is created and the transfer continues.

Use the `svsetstanza` command or edit the `snapvault.cfg` file on the Open Systems SnapVault primary system to set the `FailCPRestartOnNewSnapshot` value.



## Checkpoints for Restart

- Allows a failed transfer to resume from the last valid recorded checkpoint
  - Not used to resynchronize a broken relationship or to resume operations after a restore
- Supports checkpoints at block levels inside files
  - A failed transfer can restart even from the middle of a file
- Configure checkpoint interval
  - Default 300 seconds (5 minutes)
  - Minimum 60 seconds

© 2010 NetApp, Inc. All rights reserved.

## CHECKPOINTS FOR RESTART

When an Open Systems SnapVault backup process fails, checkpoint restart support allows the backup transfer (baseline or update) to be resumed from the last valid recorded checkpoint. Checkpoints are recorded by the Open Systems SnapVault primary when certain predetermined conditions or periodic intervals are met. The Open Systems SnapVault primary records the checkpoints and sends them to the SnapVault secondary system.

Checkpoint restarts are not used to resynchronize an Open Systems SnapVault relationship after a restore or to re-establish an Open Systems SnapVault broken relationship.

In releases earlier than Open Systems SnapVault 2.5, checkpoints were taken either at five-minute intervals or at the end of a file. This resulted in sending all the file data again, if there was a transfer failure for a large file.

Open Systems SnapVault 2.5 and later releases support the following improvements in the checkpoint mechanism:

**Enabling checkpoints at block levels inside files:** This enhancement is useful when the dataset contains large files (greater than 100 MB). Checkpoints are allowed inside files; therefore you can restart the transfer even from the middle of a file.

**Configuring checkpoint interval:** You can configure the checkpoint interval by editing the `snapvault.cfg` file on the Open Systems SnapVault primary system. The default value is set to 300 seconds (5 minutes), and the minimum value is 60 seconds.



## Block-Level Incremental

- Block-level incremental (BLI) backup recognizes that a file has changed based on timestamp and checksum algorithm
- BLI backup levels
  - OFF: No checksums computation
  - HIGH: Computes checksums on initial transfers and incremental updates (default)
  - LOW: Computes checksums on first and subsequent incremental updates
- Name-based BLI
  - Addresses known issues on file modification

© 2010 NetApp, Inc. All rights reserved.

### BLOCK-LEVEL INCREMENTAL

A block-level incremental (BLI) backup recognizes that a file has changed based on a **timestamp and checksum** algorithm. It also determines exactly which blocks in the file have changed, and then backs up only those blocks to the SnapVault secondary system.

Checksum database files are stored in the Open Systems SnapVault agent internal database directory. Each Open Systems SnapVault relationship has its own checksum file directory. Disk storage requirements to maintain the computed checksum values on the primary are expected to be approximately 2% of the baseline backup size.

BLI can operate at three levels: OFF, HIGH, or LOW. Those values are configurable in the Open Systems SnapVault Configurator.

OFF: No checksums are calculated. Full files are transferred once identified as being changed files

HIGH (default): Always computes checksums (initial and incremental transfers)

LOW: Computes checksums on first incremental and subsequent transfers

Applications such as Microsoft Word, Microsoft Excel®, and Microsoft PowerPoint® (referred to as name-based applications) modify files by inserting new data blocks in the file and rewriting all subsequent data blocks in the file to new positions in the file. As the modified file is considered new, a backup of all the rewritten blocks and a recalculation of checksum would be required.

Open Systems SnapVault agents work around this issue by recognizing files by name in addition to identifying the file by the file-system location. You can disable BLI backups for certain name-based applications with the Open Systems SnapVault Configurator.



## Exclude Lists and Encrypted Files

- Exclude lists
  - Excludes specified files, directories, sub-directories and paths from backups
- Support for Encrypted File System (EFS) files backup and restore
  - BLI does not support EFS files backup
  - Ensure that there is sufficient free space in the target volume to restore EFS files

© 2010 NetApp, Inc. All rights reserved.

### EXCLUDE LISTS AND ENCRYPTED FILES

Backup exclusion lists are used by the Open Systems SnapVault agent to exclude specified files, directories, sub-directories, and entire paths from backups.

Open Systems SnapVault supports two types of exclusion lists:

File exclusion list entries consist of single-path elements. A file or directory is excluded if the file name or any path element matches a file exclusion entry in the list. The file exclusion list is in the *install\_dir/etc/file-exclude.txt* file.

Path exclusion list entries consist of complete file system paths to either a directory or a file. The path exclusion list is in the *install\_dir/etc/path-exclude.txt* file.

Refer to the *Open Systems SnapVault Installation and Administration Guide* for details on supported characters and wildcards for those file entries.

Open Systems SnapVault is capable of backing up and restoring Encrypted File System, or EFS files, on Windows platforms. However, you cannot use block-level incremental backup to back up EFS files. Any time an EFS file is modified; Open Systems SnapVault backs up the entire EFS file.

The following are the requirements for backing up and restoring EFS files:

You cannot use a block-level incremental backup to back up EFS files. Any time an EFS file is modified, Open Systems SnapVault backs up the entire EFS file.

There must be a sufficient amount of free space in the target Windows volume. EFS-encrypted file restore requires the Open Systems SnapVault agent to create a temporary file that is equal to or greater than the size of the EFS-encrypted file that is being replaced.



## Open Systems SnapVault Database

- Allows restoration of the Open Systems SnapVault database without the need to initiate a baseline transfer
- On the Open Systems SnapVault primary, the database resides in `install_dir\snapvault\db\QsmDatabase`
  - By default the Open Systems SnapVault database is backed up upon each transfer
- On the secondary, database backup is stored under the qtree root as `OSSV_DATABASE_BACKUP`
- Backup levels: BLI (default), DB only, and NONE
- If the file system was modified between backup and restore, you must resync the relationship

© 2010 NetApp, Inc. All rights reserved.

### OPEN SYSTEMS SNAPVAULT DATABASE

The Open Systems SnapVault database consists of a set of files that contain information about the Open Systems SnapVault relationship between a primary and a secondary system:

- History file
- BLI checksums file (if BLI is enabled)
- Checkpoint file (if a backup transfer had failed with a checkpoint)

If the Open Systems SnapVault database becomes corrupt or gets out-of-sync with the secondary, data transfers fail and you will be forced to initiate a baseline transfer. However, if you have a valid backup of the database, you can restore it and continue with subsequent transfers.

By default, backup of the Open Systems SnapVault database occurs automatically every time data is transferred. A compressed file of the database is created and transferred to the secondary. The file is named `Open Systems SnapVault_DATABASE_BACKUP` and stored in the root of the destination qtree. There are three backup levels of the Open Systems SnapVault database:

- BLI (default): Backs up the history file and its corresponding BLI checksum file
- DB only: Backs up only the history file
- NONE: Disables the automatic database backup

To restore the database, issue the `snapvault restore` command from the Open Systems SnapVault primary, including the file name `Open Systems SnapVault_DATABASE_BACKUP` in the path. After the database file is restored, Open Systems SnapVault decompresses it and places the files where Open Systems SnapVault database files are located for the relationship. Data transfers can be performed from this point onward. However, if the file system was modified between the backup and the restore operations, you will have to resynchronize the relationship.





## System State and Event Log Backup

- Open Systems SnapVault supports W2K, W2K3, and W2K8 system state backup and restore
  - Allows to restore a system to an earlier state
  - Subsequent backups use BLI when enabled
- System state backup can include Windows Event Log
  - Used for troubleshooting problems or capacity planning
  - Application, Security, and System event log backup are configurable options
  - Incremental updates of Windows Event Log files are not supported

© 2010 NetApp, Inc. All rights reserved.

### SYSTEM STATE AND EVENT LOG BACKUP

You can back up and restore various components of the Windows 2000 Server and Windows Server 2003 system state. These components may include:

- Registry
- COM+ Class Registration database
- System files and settings, including the boot files
- Certificate Services database
- IIS Metadirectory
- System files that are under Windows file protection
- Performance counters
- System state data on domain controllers (Active Directory and SYSVOL data)

You can add backups of Windows system state data to existing Open Systems SnapVault backup schedules and use the backups to restore a system to an earlier state. You can also use an Open Systems SnapVault system state data backup in conjunction with complete file system backups as part of a disaster-recovery plan.

With Open Systems SnapVault support for Windows Event Log, you can maintain the records of all the events that occur in the system. It is necessary to record the events to help you carry out tasks such as troubleshooting problems or capacity planning.

You can enable or disable support for Windows Event Log as part of the system state backup using the Open Systems SnapVault Configurator. Application, Security, and System event log backup are configurable options.

Incremental updates of Event Log files are not supported. Every update does a full backup of event logs.



## Open Systems SnapVault Restore

- Methods to restore a directory or a file:
  - Copy file or directory using CIFS or NFS
  - Use the `snapvault restore` command
  - Use NetApp Operations Manager
- Restoring an entire primary system
  - There must be an operating system on the disk
- Restoring files to a primary system from tape
  - First restore the data from the tape to the secondary and then restore from that secondary to the Open Systems SnapVault primary

© 2010 NetApp, Inc. All rights reserved.

## OPEN SYSTEMS SNAPVAULT RESTORE

In the event of data loss or corruption on the Open Systems SnapVault primary system, the administrator can restore a directory or a file, the entire system, or files from a tape device.

### RESTORING A DIRECTORY OR A FILE

You can copy files from the secondary storage system to the Open Systems SnapVault primary system using NFS or CIFS. Some Windows and UNIX attributes are not preserved using this method, such as Windows sparse files, Windows EFS data, and UNIX access control lists.

You can also use the `snapvault restore` command from the Open Systems SnapVault primary command-line interface.

```
snapvault restore -s sv_snapshot -S sec_hostname:/vol/sec_vol/sec_qtree/file.doc
prim_hostname:dirpath
```

You can also use the NetApp Operations Manager restore wizard or supported NDMP management software.

### RESTORING AN ENTIRE PRIMARY SYSTEM

You can restore an entire primary system from a SnapVault secondary, but there must be an operating system on the Open Systems SnapVault primary's disk.

- 1.Reinstall the operating system on the primary.
- 2.Reformat the file system just as the original file system was formatted.
- 3.Install the Open Systems SnapVault agent.
- 4.(Optional) Restore the Windows system state data of the Open Systems SnapVault primary if you backed it up.
- 5.Restore the backed-up directories using the `snapvault restore` command.

## RESTORING TO A PRIMARY SYSTEM FROM TAPE

The process of restoring from tape to a primary system involves first restoring the data from the tape to a secondary and then restoring from that secondary to the Open Systems SnapVault primary.

1. Mount the tape that has the files that need to be restored on the secondary.
2. Use the Data ONTAP `restore` command to restore from the tape to the SnapVault secondary.
3. Restore the files from the secondary to the NFS or CIFS primary using the `snapvault restore` command.



## Restore Limitations

- When restoring using the `snapvault restore` command, restoring files to a directory that is part of a SnapVault relationship is not allowed
  - You have to release the SnapVault backup relationship from the primary before performing the restore to the same location
  - Or you can restore the files to a new location
- NTFS compressed files and directories lose compressed attributes when restored

© 2010 NetApp, Inc. All rights reserved.

## RESTORE LIMITATIONS

### RESTORING TO A DIRECTORY THAT IS PART OF A SNAPVAULT RELATIONSHIP

If you attempt to restore files to a directory involved in a SnapVault relationship using the `snapvault restore` command, the attempt fails with the message:

```
Invalid Qtree/Snapshot requested
```

```
Directory in wrong phase
```

```
Error performing restore. Check snapvault log.
```

The following information is logged in the SnapVault log files:

```
Invalid Qtree/Snapshot requested
```

```
Directory in wrong phase
```

You can restore to such a directory by first releasing the SnapVault backup relationship using the `snapvault release` command from the primary:

```
snapvault release pathname sec_hostname:/vol/sec_vol/sec_qtree
```

Alternatively, you can restore the files to a new location on the primary system.

### RESTORING NTFS-COMPRESSED FILES AND DIRECTORIES

Open Systems SnapVault does not restore the compressed attribute on a Windows file if both the following conditions are true:

The file has a compressed attribute set.

The file has a sparse or compressed alternate data stream attached.

In the case of compressed directories, the compressed attributes are lost when restored, regardless of the alternate data streams.



## Resynchronization

- Resync a broken relationship
  - Use the `snapvault start -r` command
  - Does not support resync of restored subdirectories and single files
- To resync after `snapvault restore`:
  1. Select the “Enable restart/resync on restore” option on the primary before performing the restore.
  2. Use the `snapvault start -r` command to resynchronize the relationship.

© 2010 NetApp, Inc. All rights reserved.

## RESYNCHRONIZATION

When the Open Systems SnapVault primary and the secondary systems become unsynchronized, the relationship needs to be resynchronized to be able to continue incremental backups.

An Open Systems SnapVault relationship can get out of sync when:

An older version of the Open Systems SnapVault database is restored to the Open Systems SnapVault primary system.

Data is restored using the `snapvault restore` command.

The state of the destination qtree in an Open Systems SnapVault relationship is changed to read-writable, because such an operation breaks the Open Systems SnapVault relationship.

Open Systems SnapVault resynchronization is supported from Data ONTAP 7.2 and Open Systems SnapVault 2.2 and later releases.

### RESYNCHRONIZING A BROKEN RELATIONSHIP

To resynchronize a broken Open Systems SnapVault relationship, on the SnapVault secondary, enter the following command:

```
Secondary> snapvault start -r -S
prim_hostname:dirpathsec_hostname:/vol/sec_vol/sec_qtree
```

NOTE: Open Systems SnapVault does not support resynchronizing restored subdirectories and single files.

### RESYNCHRONIZING AFTER A SNAPVAULT RESTORE

After data is restored using the `snapvault restore` command, if you want to perform subsequent incremental backups from the restored location to the same qtree on the secondary storage system, you must resynchronize the relationship.

Before restoring, enable the *Enable restart/resync on restore* check box in the Open Systems SnapVault Configurator. By default this option is not selected.

If you selected this option before you performed the `snapvault restore`, you can resynchronize the SnapVault backup relationship with the `snapvault start -r` command after the restore has completed.

If you did not select this option before performing the SnapVault restore, you will have to initialize the relationship after the restore using the `snapvault start` command from the SnapVault secondary.



## Open Systems SnapVault Deployment

© 2010 NetApp, Inc. All rights reserved.

### OPEN SYSTEM SNAPVAULT DEPLOYMENT





## Installing Open Systems SnapVault

- Install the Open Systems SnapVault agent on Windows or UNIX-based systems
  - Download from the NOW site
- Automatic post-installation check
  - `svinstallcheck`
- Unattended installation
  - Supported on all platforms on which Open Systems SnapVault is supported
  - Use `svconfigpackager` to create install image
- Open TCP ports 10000 and 10566 before install

© 2010 NetApp, Inc. All rights reserved.

### INSTALLING OPEN SYSTEMS SNAPVAULT

The Open Systems SnapVault agent can be installed from the CD-ROM or from the *Download Software* page on the NOW site. Refer to the *Open Systems SnapVault Release Notes* and the *Installation and Administration Guide* for details on the installation process for Windows and UNIX-based systems.

The Open Systems SnapVault agent installation will store the Open Systems SnapVault database, a set of executables, the logs file, and the Exclude Lists files on the primary. On a Windows primary, the default installation path is the C:\Program Files directory. On UNIX systems, the default installation location is the /usr directory. During the installation process, you will be prompted to enter the NDMP account to connect to the primary system, the NDMP listening port, and the host name or IP address of the secondary system. At the end of the installation routine, the `svinstallcheck` utility will verify successful installation and make sure that the services are running properly.

### UNATTENDED INSTALLATION

Unattended installation enables you to install or upgrade Open Systems SnapVault software over a large number of Open Systems platforms with minimal user intervention. To perform an unattended installation of Open Systems SnapVault, an installation script and other supporting files are required. A utility called `svconfigpackager` is available in the Open Systems SnapVault software. When run on an Open Systems SnapVault primary system, the utility saves the current configuration settings to a file. In addition, this utility can create an installation script that, in conjunction with the configuration settings file and other files, can be used to perform unattended installations or upgrades.

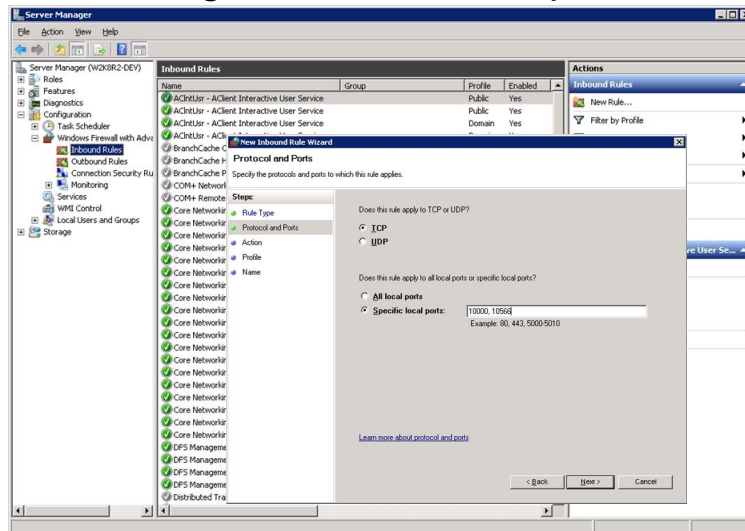
Ensure that the TCP ports 10000 (for central management using NDMP) and 10566 (QSMSEVER) are open before Open Systems SnapVault is installed. For NetApp Host Agent, HTTP port 4092 and HTTPS port 4093 must be open as well.



## W2k8R2: Firewall Configuration

- Windows Server 2008 R2 has a built-in firewall that must be configured to allow SnapVault traffic

Create an inbound rule and add the appropriate ports



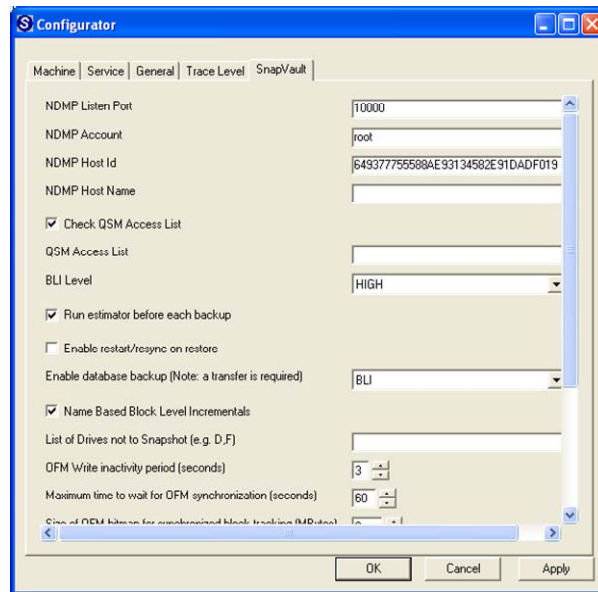
© 2010 NetApp, Inc. All rights reserved.

## W2K8R2: FIREWALL CONFIGURATION



## Open Systems SnapVault Primary

- Use the Open Systems SnapVault Configurator GUI to verify and modify Open Systems SnapVault parameters
- The `svsetstanza` command is an alternative to the Open Systems SnapVault Configurator utility



© 2010 NetApp, Inc. All rights reserved.

### OPEN SYSTEMS SNAPVAULT PRIMARY

You use the Open Systems SnapVault Configurator utility to verify and to modify the Open Systems SnapVault configuration parameters.

The Configurator utility is automatically installed during the Open Systems SnapVault agent installation and can be invoked from the *Windows > Start > Programs* menu or from the command line. The `svconfigurator.exe` program is located in the *Install\_dir/bin* directory.

From the Configurator, you can perform the following actions:

The **Machine tab** displays information about the version of Open Systems SnapVault agent software, and primary system machine information such as the IP address and OS version.

The **Service tab** allows you to stop and start the Open Systems SnapVault services.

The **General tab** enables you to generate debugging files and to modify the default db, tmp, and trace directories locations.

The **Trace Level tab** is used to choose which process to take traces of and the trace level. The Trace Level tab is used in conjunction with the General tab.

The **SnapVault tab** allows you to specify the SnapVault secondary hostname or IP address and to modify some important parameters such as the BLI level, OFM, VSS, and NDMP parameters.

Alternatively, you can use the `svsetstanza` utility from the primary for configuration purposes instead of using the Configurator. The `svsetstanza` command is located in the *install\_dir/util* directory.



## Licensing Open Systems SnapVault

- The Open Systems SnapVault licenses are installed on the SnapVault secondary to which Open Systems are backed up
- Open Systems SnapVault licenses
  - Not based on serial numbers but on platform and on the number of hosts managed
  - Evaluation licenses available on NOW site
- No limit on number of Open Systems SnapVault nodes

© 2010 NetApp, Inc. All rights reserved.

### LICENSING OPEN SYSTEMS SNAPVAULT

The Open Systems SnapVault licenses are installed on the secondary storage system to which they are backed up. You will need to install the following licenses on the SnapVault secondary storage system:

The SnapVault secondary `sv_ontap_sec` license.

The Open Systems SnapVault primary licenses for the platforms you want to back up to the secondary storage system. The following is a list of the primary licenses:

- `sv_windows_ofm_pri` (to use OFM for Windows 2000 Server)
- `sv_windows_pri` (for Windows systems)
- `sv_unix_pri` (for UNIX systems)
- `sv_linux_pri` (for Linux systems)
- `sv_vmware_pri` (for VMware ESX Server)

Licensing of the Open Systems SnapVault primary is based on the platform. Stored on the secondary, these licenses are independent and do not tie in to the serial number of the secondary system.

Evaluation licenses are available on the NOW site.



## Configuring SnapVault Secondary

- SnapVault secondary configuration steps are identical to the ones used in a SnapVault storage systems deployment:
  1. Install the required licenses
  2. Enable the NDMP service
  3. Specify Open Systems SnapVault primary systems to back up
  4. Initialize the baseline transfer
  5. Schedule updates

© 2010 NetApp, Inc. All rights reserved.

### CONFIGURING SNAPVAULT SECONDARY

The SnapVault secondary configuration steps are identical to the ones used to configure SnapVault in a NetApp storage systems environment.

1. Install the required licenses:

```
sec> license add <license_code>
```

2. Ensure that the NDMP service is enabled:

```
sec> options ndmpd.enable on
```

3. Specify the names or IP addresses of the Open Systems SnapVault primary systems to back up and restore:

```
sec> options snapvault.access host=prim_host1, prim_host2
```

4. For each Open Systems SnapVault platform directory or file to be backed up to the SnapVault secondary, execute an initial baseline transfer:

```
sec> snapvault start -S prim_hostname:dirpath
 sec_hostname:/vol/sec_vol/sec_qtree
```

Schedule incremental updates.

For example:

```
sec> snapvault snap sched -x vol1 sv_hourly 11@mon-fri@7-18
```

The `-x` parameter causes SnapVault to copy new or modified files from the Open Systems SnapVault system directory or file to their associated qtrees on the secondary. After all the secondary qtrees on the specified volume have been updated, the SnapVault secondary then creates a Snapshot copy of this volume for archiving.



## Monitor Transfers

- Use the `snapvault status` command to monitor the transfer progress, status and lag

```
C:\> snapvault status -l c:\data
Source: W2K3-Client:c:\data
Destination: system:/vol/voll/backup-windows
Status: Transferring
State: Source
Lag: -
Mirror Timestamp: -
Base Snapshot: -
Current Transfer Type: -
Contents: -
Last Transfer Type: -
Last Transfer From: -
Bytes transferred so far: 18 MB
Transfer Time Elapsed: 00:00:21
Total files to transfer: 12
Total files transferred: 7
Current File Size: 8323815
Current File Progress: 2097152
Current File Name: c:\data\file1
Transfer Error ID: -
Transfer Error Message: -
```

© 2010 NetApp, Inc. All rights reserved.

## MONITOR TRANSFERS

Use the `snapvault status` command to verify whether a transfer is in progress, the state of the destination, and how long ago the last successful transfer took place.

If the `-l` option is given, the output displays more detailed information such as the total files to transfer, how many were transferred, which file is currently transferring, and the transfer error ID, if any error occurred.

For details on the `snapvault status` command output and options, refer to the *Data ONTAP Data Protection Online Backup and Recovery Guide*.



## Open Systems SnapVault Commands

- The SnapVault commands available from the Open Systems SnapVault primary command-line interface are identical to the ones supported on a SnapVault storage primary system

```
C:\Program Files\netapp\snapvault\bin> snapvault
```

```
The following commands are available; for more
information type "snapvault help <command>"
```

```
abort destinations help release
restore status service diag
```

- The `service` and `diag` commands are specific to Open Systems SnapVault primary systems

© 2010 NetApp, Inc. All rights reserved.

## OPEN SYSTEMS SNAPVAULT COMMANDS

The SnapVault commands available from the Open Systems SnapVault primary command-line interface are identical to the ones supported on a SnapVault storage primary system. The SnapVault commands are stored in the `Install_dir/bin` directory. Refer to the SnapVault module of this course for details on the following SnapVault commands syntax:

Use the `status` command to monitor the status of an Open Systems SnapVault backup or restore.

Issue the `destinations` command to list all the known destinations for this Open Systems SnapVault primary system.

The `abort` command allows you to stop a restore transfer from the secondary to the primary system.

The `release` command allows you to release relationships that have been stopped on the secondary.

In the event of data loss or corruption on the primary system, use the `restore` command to restore the affected file, directory, drive, or even the entire system.

Use the `service` command to manage the Open Systems SnapVault service.

Issue the `diag` command to display details of a given error number.

NOTE: The `service` and `diag` commands are specific to Open Systems SnapVault primary systems.



## Log Files

- Open Systems SnapVault primary log files
  - Windows primary  
`C:\Program Files\netapp\snapvault\etc`
  - UNIX-based primary  
`/usr/snapvault`
  - A new file is created daily at midnight
    - File is named `snapvault.yyyymmdd`
    - No extension for the current file
- SnapVault secondary log files
  - Root volume: `/etc/log/snapmirror`

© 2010 NetApp, Inc. All rights reserved.

## LOG FILES

You can find the operational status and problem reports of the Open Systems SnapVault primary system in the log file named *snapvault*.

On the Open Systems SnapVault primary, the log files are found:

- On a Windows primary, in the `C:\Program Files\netapp\snapvault\etc` directory
- On a UNIX-based primary, in the `/usr/snapvault` directory

A new file is created daily at midnight or as soon after midnight as the first subsequent activity on the system takes place; the existing file is not archived until a new one is created.

The current file has no extension. The archived files have the `.yyymmdd` extension, where *yyyy* is the year, *mm* is the month, and *dd* is the date when the file was created.

**NOTE** that the log file is created during the initial baseline transfer, not by the install process. Therefore, if you search for the log file after completing the installation, but before you back up data for the first time, you will not find the log file.

On the secondary storage system, the SnapVault log files are stored in the `/etc/log/snapmirror` file in the root volume.





## Best Practices and Troubleshooting

© 2010 NetApp, Inc. All rights reserved.

### BEST PRACTICES AND TROUBLESHOOTING



## Considerations

- Open Systems SnapVault does not back up NFS mountpoints and CIFS shares
- Open Systems SnapVault built-in databases require free space
  - Use the Free Space Estimator utility
- OFM requires at least 15% disk free space in the drive being backed up
- Supports maximum 16 concurrent transfers
- Large number of small files may impact performance
- Application database must be dismounted before Open Systems SnapVault backups are initiated

© 2010 NetApp, Inc. All rights reserved.

## CONSIDERATIONS

Open Systems SnapVault does not back up remote NFS or CIFS file systems that have been mounted on or mapped to UNIX or Windows primary systems.

The Open Systems SnapVault database space requirements depend on the BLI level, the number of files, average file size, and number of directories to back up. Use the Free Space Estimator utility to ensure that there is sufficient space on the Open Systems SnapVault install drive to be able to perform an incremental update. Note that if OFM (Windows 2000 Server) is used, the file systems being backed up must have at least 15% disk space free in the drive being backed up. If free space is not available, disable OFM for those drives.

Open Systems SnapVault supports a maximum of **16** simultaneous transfers from a primary system. You should plan your backup schedules such that 16 or fewer transfers occur at the same time.

When using Open Systems SnapVault, there is always some overhead to be transferred for files in the Open Systems SnapVault relationship that have been modified. The Open Systems SnapVault primary will send one 4-KB header for every file and directory that exists in the relationship. In addition, for files and directories that are larger than 2 MB, an additional 4-KB header is transferred for every 2 MB.

A large number of small files may degrade performance and also result in a large amount of overhead data sent over the network. If a large number of files are not likely to be modified, consider changing the BLI level to LOW.

Open Systems SnapVault is not integrated with any database backup APIs. The database files need to be dismounted prior to using Open Systems SnapVault to back them up. If users want to use the “hot backup mode” method, they will need to script it and test it themselves and ensure that the procedure works reliably in their environment.



## Error Messages

- Collect error messages logged on the Open Systems SnapVault primary and the SnapVault secondary:  
`snapvault status -l`  
`/etc/log/snapmirror`  
`C:\Install_dir\snapvault\etc` (Windows example)
- On the primary, use the `snapvault diag <err_num>` to display details for a given error number
- Error messages, causes, and actions are listed in the *Open Systems SnapVault Installation and Administration Guide*

© 2010 NetApp, Inc. All rights reserved.

## ERROR MESSAGES

You can diagnose and troubleshoot Open Systems SnapVault issues by collecting and analyzing the error messages logged on the Open Systems SnapVault primary (*Install\_dir/snapvault*) and on the SnapVault secondary (*/etc/log/snapmirror*) systems.

On the Open Systems SnapVault primary command-line interface, use the `snapvaultdiag <err_num>` command to display details for `snapvault status -l` a given error code. The error code number is displayed in the command output.

Example:

```
$ snapvaultdiag 3016
```

Type: Error

Message: A network error has occurred.

Possible cause: The network socket was closed unexpectedly or the transfer was aborted by the user.

Possible action to take: Verify network connectivity between the Open Systems SnapVault primary system and the secondary storage system.

Error codes range:

- 1000-1999: Information messages
- 2000-2999: Warning messages
- 3000-3999: Errors causing an abort of the transfer

**NOTE** that not all the values in the range are used. A complete list of error codes, strings, causes and action is available in the *Open Systems SnapVault Installation and Administration Guide*.



## Data Collection

- OSSVINFO tool
  - Collects a complete set of information, commands outputs, and optionally ChangeLog and trace files
  - Packaged with the Open Systems SnapVault agent:
    - OSSVINFO.exe for W2K and W2K3
    - OSSVINFO.pl for Linux and UNIX-based
- Generate debug information
  - Traces are compressed into a .bin file

© 2010 NetApp, Inc. All rights reserved.

## DATA COLLECTION

When contacting NetApp Technical Support, you can either provide a set of files and command outputs, or run the Open Systems SnapVault INFO tool on the affected Open Systems SnapVault primary to collect a complete set of information from the primary and the secondary systems. The Open Systems SnapVault INFO tool is packaged with Open Systems SnapVault, and two versions are available:

- Open Systems SnapVault INFO.exe for Windows 2000 Server and Windows Server 2003 Open Systems SnapVault primary systems
- Open Systems SnapVault INFO.pl for Linux and UNIX-based Open Systems SnapVault primary systems

The Open Systems SnapVault INFO tool writes the data to the output directory as a text file in a specific format. Also, it collects the ChangeLog and trace files to this output directory if either the -q (for Windows) or -all (for all platforms) option is given, as illustrated in this example:

```
Open Systems SnapVaultINFO.exe [-s secondary] [-l username:password] [-q
qtreeid][-all] Output_Dir
```

The ChangeLog file captures duplicate inode issues and metadata corruption at run time. The ChangeLog capture is enabled by default in snapvault.cfg file.

NetApp Technical Support may ask you to enable Open Systems SnapVault debugging to troubleshoot the issue. The traces collected are automatically compressed into a .bin extension file. When generating debug files, ensure the following:

- Disable the generation of these files after you have sent a batch to Technical Support.
- Delete the debug files from the system after you have sent them to Technical Support to minimize the impact on performance.



## Baseline Transfer Assistance

- Logical replication, or LREP, is used to initialize SnapVault or Open Systems SnapVault data transfers over low-bandwidth connections using a portable device
- Two utilities
  - lrep\_reader
  - lrep\_writer
- Supports Open Systems SnapVault data restore, data compression and encryption
- LREP 2.0 is packaged with Open Systems SnapVault 2.5 and later releases and is also downloadable from the NOW site

© 2010 NetApp, Inc. All rights reserved.

### BASELINE TRANSFER ASSISTANCE

LREP is a tool used to initialize Open Systems SnapVault data transfers using a portable device and without traversing the network.

You can then ship the portable device to the location where the SnapVault secondary is, and move the data from the portable device to the secondary. After the relationship between the primary and secondary systems is established and modified, incremental transfers can occur directly. You can also restore data from the secondary storage system to the primary system using LREP.

The LREP tool consists of two utilities:

- **lrep\_reader**: Used at the remote office to write data from the Open Systems SnapVault primary system to the portable device.
- **lrep\_writer**: Used at the location of the SnapVault secondary storage system to write data from the portable device to the secondary storage system.

LREP enables the **compression** of data using a zlib library. Compression of LREP data is done in the memory before the data is written to the disk.

The Advanced Encryption Standard algorithm is used to encrypt and decrypt the LREP data.

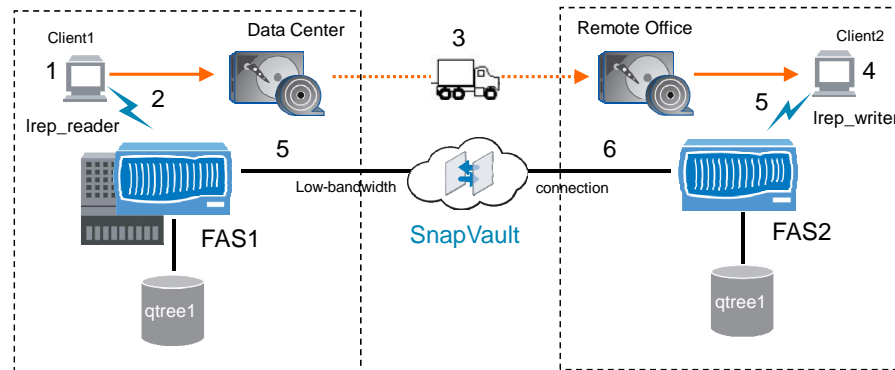
LREP 2.0 is packaged with Open Systems SnapVault 2.5 and later. You can also download the LREP tool from the NOW site. The tool can be run on Windows 2000/2003, Linux, and UNIX-based platforms.

For details on how to back up and restore Open Systems SnapVault data using LREP, refer to *Logical Replication (LREP) Tool 2.0 User Guide* on the NOW site.



## Baseline Transfer Assistance (Cont.)

- Open Systems SnapVault and SnapVault primaries are supported
- LREP is downloadable from the NOW site



© 2010 NetApp, Inc. All rights reserved.

## BASELINE TRANSFER ASSISTANCE (CONT.)

The LREP tool can also be used to perform SnapVault initial transfers using a portable device (tape or external drive) and without traversing the network. This process is often referred to as seeding the baseline transfer. LREP applies only to SnapVault. The LREP tool set is downloadable from the NOW site.

The illustration sets up a SnapVault baseline transfer from FAS1 (source system at the data center) to FAS2 (destination system at the remote office) using the LREP utilities.

1. Install the `lrep_reader` utility on Client1 at the data center.
2. Run the `lrep_reader` utility from Client1 to `snapmirror initialize` from the source to the portable media.
3. The portable media is then transferred to the remote office.
4. The `lrep_writer` utility is installed on Client2 at the remote office.
5. Run `lrep_writer` utility from Client2 to reconstruct the replication streams by reading from files, which were created by `lrep_reader`.
6. From the destination, run the `snapmirror initialize` command to read the logical replication stream data from the `lrep_writer`. The SnapVault initial transfer completes. Edit the `snapmirror.conf` file to reflect the correct source storage system.

## SUPPORTED PLATFORMS

The LREP tool is supported on the following operating systems:

- Windows 2000, Windows Server 2003, and Windows Storage Server 2003 on x86 and x86-64/EM64T platforms
- Red Hat Enterprise Linux 4.0 for x86 and x86-64/EM64T
- Red Hat Enterprise Linux 5.0 for x86 and x86-64/EM64T

- SuSE® Linux Enterprise Server 9 for x86 and x86-64/EM64T
- SuSE Linux Enterprise Server 10 for x86 and x86-64/EM64T
- Solaris™ 9 and 10 on UltraSPARC systems
- AIX 5L™ 5.1, 5.2, and 5.3 on IBM® PowerPC® and IBM POWER™ processor-based systems
- HP-UX® 11.23 and HP-UX 11.31 on PA-RISC® based systems.



## Documents and References

- *Open Systems SnapVault Installation and Administration Guide*
- *Open Systems SnapVault Release Notes*
- *Enabling Rapid Recovery with SnapVault*  
<http://media.netapp.com/documents/tr-3252.pdf>
- *Open Systems SnapVault Best Practices Guide*  
<http://media.netapp.com/documents/tr-3466.pdf>

© 2010 NetApp, Inc. All rights reserved.

## DOCUMENTS AND REFERENCES

For more information, see the NOW site.





## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe how Open Systems SnapVault integrates with Data ONTAP SnapVault
- List Open Systems SnapVault advanced features
- Configure and administer Open Systems SnapVault
- Perform Open Systems SnapVault backup and restore operations
- Troubleshoot and resolve Open Systems SnapVault transfer failures

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 20: Open Systems  
SnapVault  
Estimated Time: 30 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What is the basic unit of an Open Systems SnapVault backup?
- True or false? A separate license is required to be able to back up open files on Windows Server platforms.
- Which command and utility would you use to configure or modify Open Systems SnapVault parameters?

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# High-Availability

Module 21

Accelerated NCDA Boot Camp Data  
ONTAP 8.0 7-Mode



## HIGH-AVAILABILITY



## Module Objectives

By the end of this module, you should be able to:

- Define the high-availability controller configuration
- Describe the three modes of high-availability operation with a high-availability pair
- Analyze the effect on client protocols during failover and giveback operations


© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Loss of Controller

- Loss of a controller may be overcome by configuring a high-availability pair




© 2010 NetApp, Inc. All rights reserved.

## LOSS OF CONTROLLER



## High-Availability Controller Configuration



Connected to its own disk shelves  
Connected to the other controller's disk shelves  
Add multipathing to the disk shelves  
Storage controllers are connected to each other  
If a storage controller fails, the surviving partner serves  
the data of failed controller

© 2010 NetApp, Inc. All rights reserved.

### HIGH-AVAILABILITY CONTROLLER CONFIGURATION


A high-availability configuration is two storage systems (nodes) whose controllers are connected to each other either directly or through switches.

The nodes are connected to each other through a cluster adapter or NVRAM adapter, which allows one node to serve data to the disks on its failed partner node. Each node continually monitors its partner, mirroring data for the partner's NVRAM.





## High-Availability Features



- High-availability controller configuration provides:
  - Fault tolerance
  - Nondisruptive software upgrades
  - Nondisruptive hardware maintenance

© 2010 NetApp, Inc. All rights reserved.

## HIGH-AVAILABILITY FEATURES

High-availability configurations provide fault tolerance and the ability to perform nondisruptive upgrades and maintenance.

High-availability configurations provide the following benefits:

- **Fault tolerance**—When one node fails or becomes impaired, a takeover occurs and the partner node continues to serve the data of the failed node.
- **Nondisruptive software upgrades**—When you halt one node and allow takeover, the partner node continues to serve data for the halted node, allowing you to upgrade the halted node. For more information about nondisruptive software upgrades, see the *Data ONTAP® Upgrade Guide*.
- **Nondisruptive hardware maintenance**—When you halt one node and allow takeover, the partner node continues to serve data for the halted node, allowing you to *replace or repair hardware* on the halted node.



## Requirements for High Availability

- Architecture compatibility
- Storage capacity
- Disk and disk shelf compatibility
- Cluster interconnect adapters and cables installed
- Nodes attached to the same networks
- Same software licensed and enabled

© 2010 NetApp, Inc. All rights reserved.

### REQUIREMENTS FOR HIGH AVAILABILITY

The number of disks in a standard high-availability configuration must not exceed the maximum configuration capacity. In addition, the total amount of storage attached to each node must not exceed the capacity of a single node.

To determine your maximum configuration capacity, see the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml).

**NOTE:** When a failover occurs, the takeover node temporarily serves data from all the storage in the high-availability configuration. When the single-node capacity limit is less than the total high-availability configuration capacity limit, the total disk space in a cluster can be greater than the single-node capacity limit. It is acceptable for the takeover node to temporarily serve more than the single-node capacity would normally allow, as long as it does not own more than the single-node capacity.

### DISKS AND DISK-SHELF COMPATIBILITY

Both Fibre Channel (FC) and SATA storage is supported in standard high-availability configurations, as long as the two storage types are not mixed on the same loop.

If needed, a node can have only FC storage and the partner node can have only SATA storage.

Cluster interconnect adapters and cables must be installed.

Nodes must be attached to the same network and the network interface cards must be configured correctly.

System features such as CIFS, NFS, or SyncMirror® software must be licensed and enabled on both nodes.



## Partner Communication

- In a high-availability controller configuration, partners communicate through the interconnect with a heartbeat
  - System state is written to disk in a “Mailbox”
  - Data not committed to disk is written to the local and partner nonvolatile RAM (NVRAM)

© 2010 NetApp, Inc. All rights reserved.

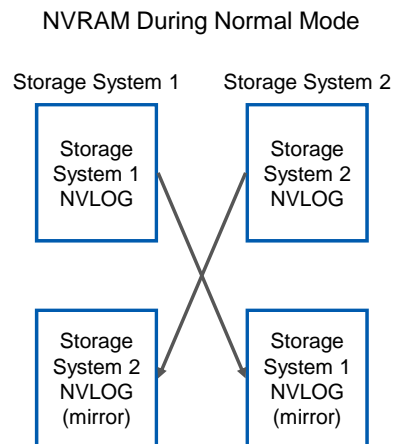
### PARTNER COMMUNICATION

To ensure that both nodes in a high-availability controller configuration maintain the correct and current status of the partner node, heartbeat information and node status are stored on each node in the mailbox disks. The mailbox disks are a redundant set of disks used in coordinating takeover or giveback operations. If one node stops functioning, the surviving partner node uses the information on the mailbox disks to perform takeover processing, which creates a virtual storage system. In the event of an interconnect failure, the mailbox heartbeat information prevents an unnecessary failover from occurring. Moreover, if cluster configuration information that is stored on the mailbox disks is out of sync during boot, the high-availability controller nodes automatically resolve the situation. The FAS system failover process is extremely robust, preventing split-brain issues from occurring.



## High-Availability Controllers and NVRAM

- Each node reserves half of the total NVRAM for the partner's data
- During takeover, the surviving partner performs the down system's reads and writes using the mirror NVLOG



© 2010 NetApp, Inc. All rights reserved.

### HIGH-AVAILABILITY CONTROLLERS AND NVRAM

Data ONTAP uses the WAFL® file system to manage data processing and NVRAM to guarantee data consistency before committing writes to disks. If the storage controller experiences a power failure, the most current data is protected by the NVRAM, and file system integrity is maintained.

In the high-availability controller environment, each node reserves half of the total NVRAM size for the partner node's data to ensure that exactly the same data exists in NVRAM on both storage controllers. Therefore, only half of the NVRAM in the high-availability controller is dedicated to the local node. If failover occurs, when the surviving node takes over the failed node, all WAFL checkpoints stored in NVRAM are flushed to disk. The surviving node then combines the split NVRAM.

### HOW THE INTERCONNECT WORKS

The interconnect adapters are a critical component in the high-availability controller configuration. Data ONTAP uses these adapters to transfer system data between the partner nodes, which maintain data synchronization in the NVRAM on both controllers. Other critical information is also exchanged through the interconnect adapters, including the heartbeat signal, system time, and details about temporary disk unavailability due to pending disk-firmware updates.



## Configuring High Availability

- License the high-availability service called cf:  
system and system2> `license add xxxxxx`
- Reboot:  
system and system2> `reboot`
- Enable the service on one of the two systems:  
system or system2> `cf enable`
- Check the status:  
system or system2> `cf status`
- To check the partner:  
system or system2> `cf partner`

© 2010 NetApp, Inc. All rights reserved.

### CONFIGURING HIGH AVAILABILITY

To add the license, enter the following command on both node consoles for each required license:

```
license add xxxxxx
```

where xxxxx is the license code you received for the feature

To reboot both nodes, enter the following command:

```
reboot
```

To enable the license, enter the following command on the local node console:

```
cf enable
```

To verify that controller failover is enabled, enter the following command on each node console:

```
cf status
```



## Setting Matching Node Options

1. Analyze the values of the options for each nodes.
2. Verify that the options settings are the same.
3. Correct any mismatched options.

- The following table lists parameters that must be the same for both nodes:

| Parameter             | Setting for        |
|-----------------------|--------------------|
| Date                  | date, rdate        |
| NDMP                  | NDMP (on or off)   |
| Published route table | route              |
| Route                 | routed (on or off) |
| Time zone             | timezone           |

© 2010 NetApp, Inc. All rights reserved.

### SETTING MATCHING NODE OPTIONS


Because some Data ONTAP options need to be the same on both the local and partner node, you need to check these options with the options command on each node and change them as necessary.

#### STEPS

1. View and note the values of the options on the local and partner nodes, using the following command on each console:  
`options`  
The current option settings for the node are displayed on the console. Output similar to the following is displayed:  
`autosupport.doit TEST`  
`autosupport.enable on`
2. Verify that the options with comments in parentheses are set to the same value for both nodes. The comments are as follows:
  - Value might be overwritten in takeover
  - Same value required in local+partner
  - Same value in local+partner recommended
3. Correct any mismatched options using the following command:  
`options option_name option_value`



## Normal Operation




- Each storage controller handles its own storage requests

© 2010 NetApp, Inc. All rights reserved.

## NORMAL OPERATION



## Takeover Operation



- Surviving partner has two identities, with each identity able to access appropriate volumes and networks only
- You can access the failed node using console commands

© 2010 NetApp, Inc. All rights reserved.

### TAKEOVER OPERATION

When a takeover occurs, the functioning partner node takes over the functions and disk drives of the failed node by creating an emulated storage system that:

- Assumes the identity of the failed node
- Accesses the failed node's disks and serves its data to clients

The partner node maintains its own identity and its own primary functions, but also handles the added functionality of the failed node through the emulated node.





## Takeover Events

- Takeover occurs on the following events:
  - A node undergoes a software or system failure that leads to a panic
  - A node undergoes a system failure (for example, a loss of power) and cannot reboot
  - There is a mismatch between the disks that one node believes it owns and the disks that the other node believes it owns
  - One or more network interfaces configured to support failover becomes unavailable
  - A node cannot send heartbeat messages to its partner and no other mechanism is available
  - A node is halted: `halt`
  - A takeover is manually initiated

© 2010 NetApp, Inc. All rights reserved.

## TAKEOVER EVENTS



## Configurable Takeover Events

- To allow a high-availability storage system to take over if the partner node panics:

```
system> options cf.takeover.on_panic on
```

Default value

- To allow a high-availability storage system to take over if the partner node reboots:

```
system> options cf.takeover.on_reboot on
```

Default value, unless  
FC or iSCSI is licensed

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURABLE TAKEOVER EVENTS



## partner

- To access the failed storage controller:

```
system(takeover)> partner
```

```
system2/system>
```

Failed controller / Takeover controller >

- Execute commands as needed:

- **NOTE:** Some commands are unavailable in partner mode

- To toggle back to the prompt of the first system:

```
system2/system> partner
```

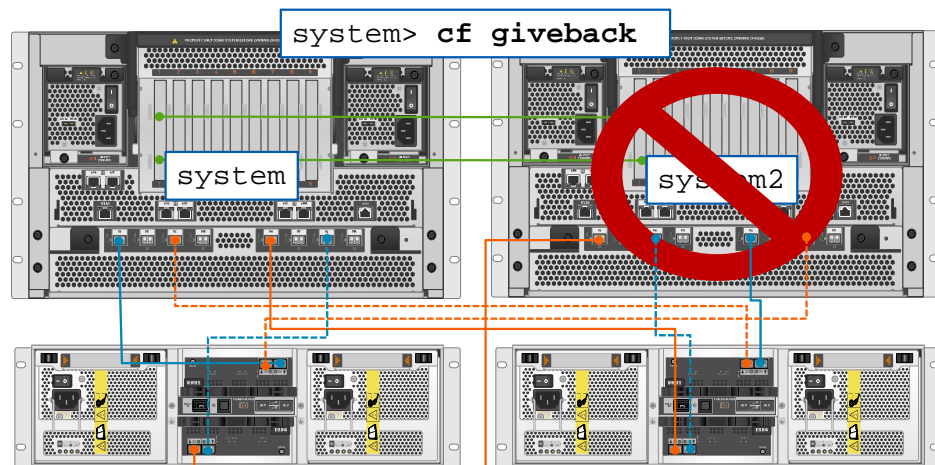
```
system(takeover)>
```

© 2010 NetApp, Inc. All rights reserved.

## PARTNER



## Giveback Operation



- `cf giveback` command terminates the emulated node
- The failed node resumes normal operation
- The high-availability configuration resumes normal operation

© 2010 NetApp, Inc. All rights reserved.

### GIVEBACK OPERATION

After a partner node is repaired and operating normally, you can use the `cf giveback` command to return operations to the partner.

When the failed node is functioning again, the following events can occur:

You initiate a `cf giveback` command that terminates the emulated node on the partner.

The failed node resumes normal operation, serving its own data.

The high-availability configuration resumes normal operation, with each node ready to take over for its partner if the partner fails.



## Automatic Giveback

- When a storage system that had given control over to its partner successfully boots up, the high-availability partner can perform an automatic giveback if configured

```
system> options cf.giveback.auto.enable on
```

Default value  
is off

- To adjust the giveback delay time for automatic giveback:

```
system> options cf.giveback.auto.delay.seconds 300
```

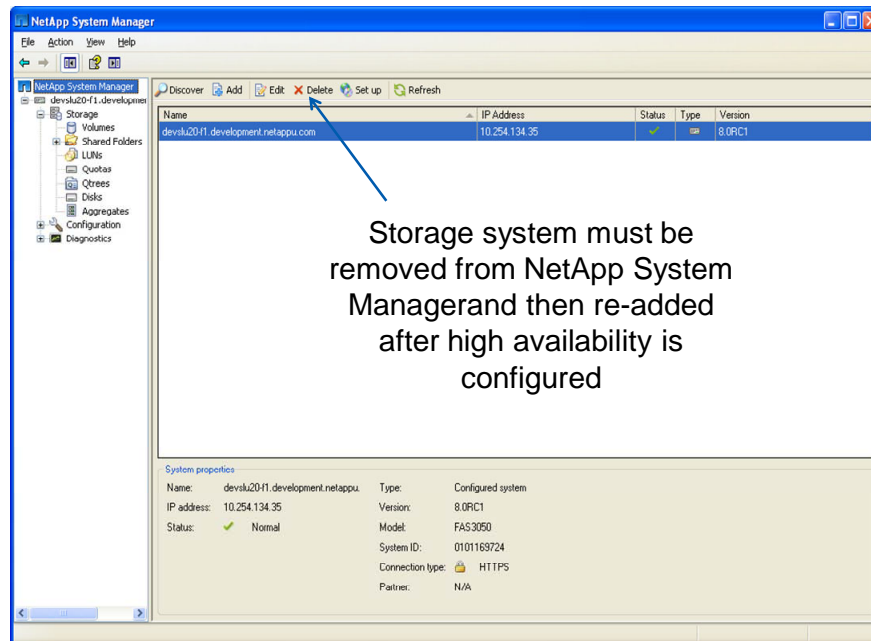
Default value

© 2010 NetApp, Inc. All rights reserved.

## AUTOMATIC GIVEBACK



## System Manager: High-Availability



© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY



## System Manager: High-Availability (Cont.)

- Currently, we cannot enable high-availability from NetApp System Manager so we will have to perform this operation from the command-line interface:
  - License controller failover (cf):

```
system> license add xxxxxx
system2> license add xxxxxx
```
  - Reboot:

```
system> reboot
system2> reboot
```
  - Enable controller failover:

```
system> cf enable
```
  - Check status:

```
system> cf status
```

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

The screenshot displays the NetApp System Manager web interface. At the top, a menu bar includes 'File', 'Action', 'View', and 'Help'. Below it, a toolbar contains icons for 'Discover', 'Add', 'Edit', 'Delete', 'Set up', and 'Refresh'. A table lists discovered storage systems:

| Name                                | IP Address    | Status | Type | Version |
|-------------------------------------|---------------|--------|------|---------|
| devslu20-f1.development.netappu.com | 10.254.134.35 | ✓      |      | 8.0x17  |
| devslu20-f2.development.netappu.com | 10.254.134.36 | ✓      |      | 8.0x17  |

An 'Add Storage System' dialog box is open, prompting the user to 'Enter the host name or IP address of the storage system you want to add.' The input field contains '10.254.134.35'. The dialog has 'Options...', 'Add System', and 'Cancel' buttons. A text box on the right states: 'Add one of the storage systems to NetApp System Manager and the partner is automatically identified'. Below the table, the 'System properties' for the selected system are shown:

**System properties**

|             |                                  |                  |                                 |
|-------------|----------------------------------|------------------|---------------------------------|
| Name:       | devslu20-f1.development.netappu. | Type:            | Configured active/active system |
| IP address: | 10.254.134.35                    | Version:         | 8.0x17                          |
| Status:     | ✓ Normal                         | Model:           | FAS3050                         |
|             |                                  | System ID:       | 0101169724                      |
|             |                                  | Connection type: | HTTPS                           |
|             |                                  | Partner:         | devslu20-f2                     |

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)





## System Manager: High-Availability (Cont.)

The screenshot displays the NetApp System Manager web interface. On the left, a tree view shows the hierarchy: 'devslu20-f2/devslu20-f1' is selected, with sub-items for 'devslu20-f2.development.netappu.com', 'devslu20-f1.development.netappu.com', and 'Active/Active Configuration'. An arrow points from the text 'The high-availability pair' to the 'Active/Active Configuration' item. Another arrow points from the text 'High-availability configuration problems' to the 'Notifications' pane. The 'Properties' pane shows details for the HA Pair: Controller 1 (devslu20-f2.development.netappu.com) and Controller 2 (devslu20-f1.development.netappu.com), both with status 'ac'. Other details include Model: FAS3050, Version: 8.0k17, Volumes: 4, Aggregates: 3, and Disks: 56 (Spare: 47). The 'Performance' pane contains four graphs: CPU (%), I/O (bytes/sec), Operations (ops/sec), and Latency (sec/op). The 'Notifications' pane lists two items: '(devslu20-f1) License mismatch' and '(devslu20-f2) License mismatch', both with status 'ac'. The 'Reminders' pane lists four items: '(devslu20-f2) Configure', '(devslu20-f2) Configure', '(devslu20-f1) Provide', and '(devslu20-f2) Provide', all with status 'ac'.

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

NetApp System Manager

devslu20-f2/development.netappu.com

Storage

Configuration

Local Users and Groups

Network

Protocols

Security

System Tools

Autosupport

Date/Time/Timezone

Licenses

syslog

Diagnostics

devslu20-f1/development.netappu.com

Active/Active Configuration

| Name                 | Type       | Key      | Expires on            |
|----------------------|------------|----------|-----------------------|
| cf                   | Evaluation | NWAAAGDG | 7/20/2009 12:14:07 AM |
| Windows Shares(CIFS) | Evaluation | JMWMPDQA | 7/20/2009 12:30:17 AM |
| ISCSI                | Evaluation | TRNVKWH  | 7/20/2009 12:30:52 AM |
| UNIX Exports(NFS)    | Evaluation | TPQIZBN  | 7/20/2009 12:29:38 AM |

Verify licenses match with the partner

License description

No additional information available for this license. This license exists in a release of Data ONTAP released after this version of NetApp System Manager. A future release of the NetApp System Manager will provide details on this license.

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

**Configure an IP address to take over the partner's IP and enable the interface**

| Name | Type      | IP Address    | Partner | Status   |
|------|-----------|---------------|---------|----------|
| e0a  | Ethernet  | 10.254.134.36 | NA      | Enabled  |
| e0b  | Ethernet  |               | NA      | Disabled |
|      | Ethernet  |               | NA      | Disabled |
|      | Ethernet  |               | NA      | Disabled |
|      | Localloop | 127.0.0.1     | NA      | Enabled  |
|      |           | 127.0.20.1    | NA      | Enabled  |

**Edit Network Interface**

General | Advanced

Name: e0b  
Type: Ethernet  
Status: Disabled

**HA failover mode**

☐ Shared  
On takeover, this interface will assume the address of its partner interface, in addition to its current address.

☐ Dedicated  
This interface has a fixed address and does not take over a partner address on takeover

☒ Standby  
Until takeover, this interface will be idle. On takeover, this interface will assume the address of its partner interface.

IP address:   
Subnet mask:   
Partner Interface: e0a [10.254.134.35]

OK Cancel Apply

e0b Partner: NA  
Ethernet Media Type: unknown (auto-negotiated)  
Disabled Flow Control: full  
NA Trusted: Yes  
NA WINS Used: Yes  
NA

Alias

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

Do the same task for the other partner; remember to enable the interface

| Name | Type      | IP Address    | Partner | Status   |
|------|-----------|---------------|---------|----------|
| e0a  | Ethernet  | 10.254.134.35 | NA      | Enabled  |
| e0b  | Ethernet  |               | NA      | Disabled |
|      | Ethernet  |               | NA      | Disabled |
|      | Ethernet  |               | NA      | Disabled |
|      | Localloop | 127.0.0.1     | NA      | Enabled  |
|      |           | 127.0.20.1    | NA      | Enabled  |

**Edit Network Interface**

General | Advanced

Name: e0b  
Type: Ethernet  
Status: Disabled

HA failover mode

☐ Shared  
On takeover, this interface will assume the address of its partner interface, in addition to its current address.

☐ Dedicated  
This interface has a fixed address and does not take over a partner address on takeover

☒ Standby  
Until takeover, this interface will be idle. On takeover, this interface will assume the address of its partner interface.

IP address:   
Subnet mask:   
Partner Interface: e0a (10.254.134.35)

OK Cancel Apply

e0b Partner: NA  
Ethernet Media Type: 1000t-fd (auto-negotiated)  
Disabled Flow Control: full  
NA Trusted: Yes  
NA WINS Used: Yes  
NA

Alias

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

The screenshot displays the NetApp System Manager web interface. On the left, a tree view shows the hierarchy: 'devslu20-f2/devslu20-f1' > 'devslu20-f2.development.netappu.com' > 'Active/Active Configuration'. An arrow points from the text 'Fixed the high-availability configuration problems' to the 'Active/Active Configuration' link. The main content area is divided into four panels: 'Properties', 'Performance', 'Notifications', and 'Reminders'. The 'Properties' panel shows the HA Pair as 'devslu20-f2/devslu20-f1', with Controller 1 and 2 both marked as 'devslu20-f2.development' and 'ac'. Other details include Model: FAS3050, Version: 8.0x17, Volumes: 4, Aggregates: 3, and Disks: 56 (Spare: 47). The 'Performance' panel contains four graphs: CPU (%), I/O (bytes/sec), Operations (ops/sec), and Latency (sec/op). The 'Notifications' panel states 'There are no notifications at this time'. The 'Reminders' panel lists four tasks: '(devslu20-f2) Configure', '(devslu20-f2) Configure', '(devslu20-f1) Provide', and '(devslu20-f2) Provide', each with a green arrow icon.

Fixed the high-availability configuration problems

NetApp System Manager

Properties

- HA Pair: ✓ devslu20-f2/devslu20-f1
- Controller 1: ✓ devslu20-f2.development → ac
- Controller 2: ✓ devslu20-f1.development → ac
- Model: FAS3050
- Version: 8.0x17
- Volumes: 4
- Aggregates: 3
- Disks: 56 (Spare: 47)

Performance

- CPU (%)
- I/O (bytes/sec)
- Operations (ops/sec)
- Latency (sec/op)

Notifications

There are no notifications at this time

Reminders

- (devslu20-f2) Configure →
- (devslu20-f2) Configure →
- (devslu20-f1) Provide →
- (devslu20-f2) Provide →

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

The screenshot displays the NetApp System Manager web interface. On the left, a tree view shows the hierarchy: **NetApp System Manager** > **devslu20-f2/devslu20-f1** > **Active/Active Configuration**. A blue box with the text "To configure high availability" has an arrow pointing to this menu item.

The main panel shows the **Active/Active Settings** and **Active/Active Status** for two systems: **devslu20-f2** and **devslu20-f1**.

**Active/Active Settings:**

- Active/Active configuration: ☒ Enabled
- Takeover on short bootstrap time: ☒ Enabled

**Active/Active Status:**

| System name | Active/Active state | Interconnect status | System state | Action       |
|-------------|---------------------|---------------------|--------------|--------------|
| devslu20-f2 | Normal              | Up                  | Up           | Takeover <-- |
| devslu20-f1 | Normal              | Up                  | Up           | --> Takeover |

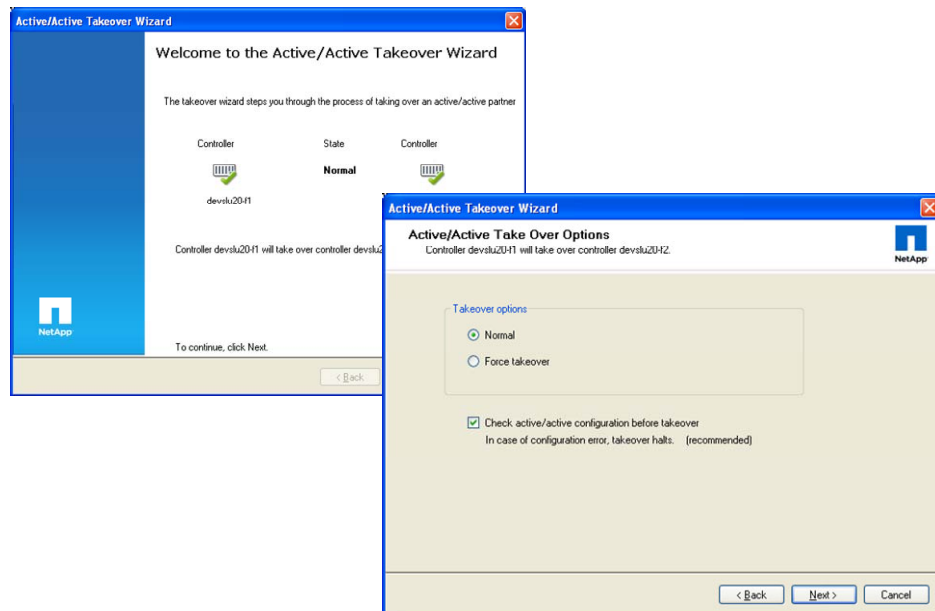
A blue box with the text "To perform a takeover" has an arrow pointing to the **Takeover <--** button for devslu20-f2 and another arrow pointing to the **--> Takeover** button for devslu20-f1. A third arrow points from the **--> Takeover** button to the right edge of the screenshot.

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)

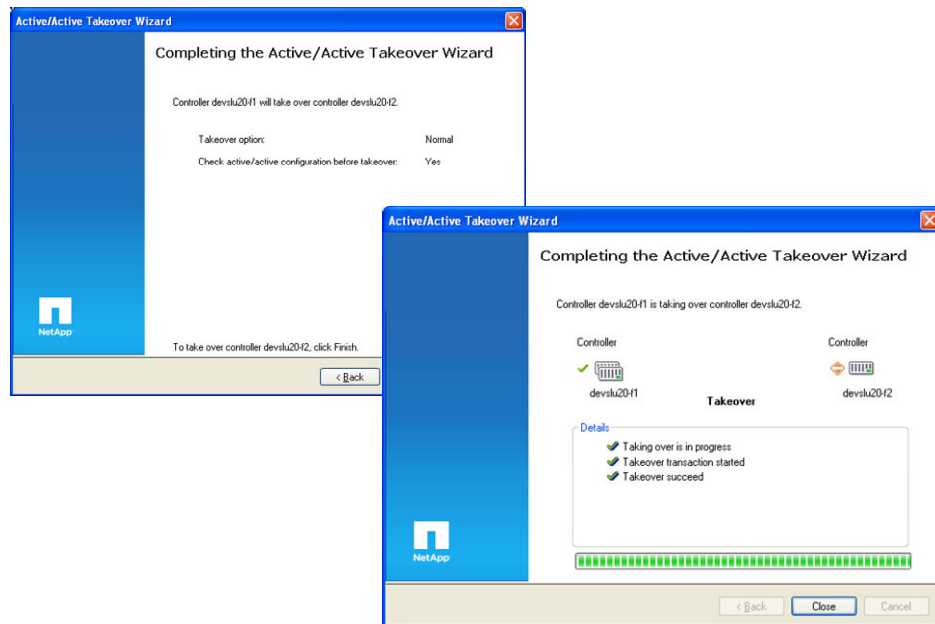


© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## System Manager: High-Availability (Cont.)



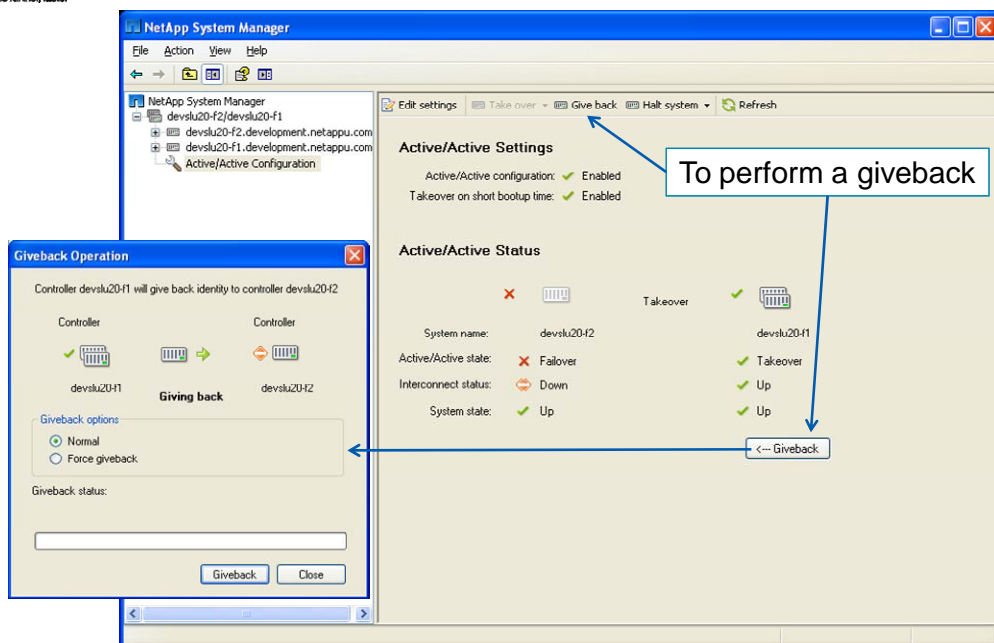
© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)





## NetApp System Manager: High-Availability Configuration (Cont.)

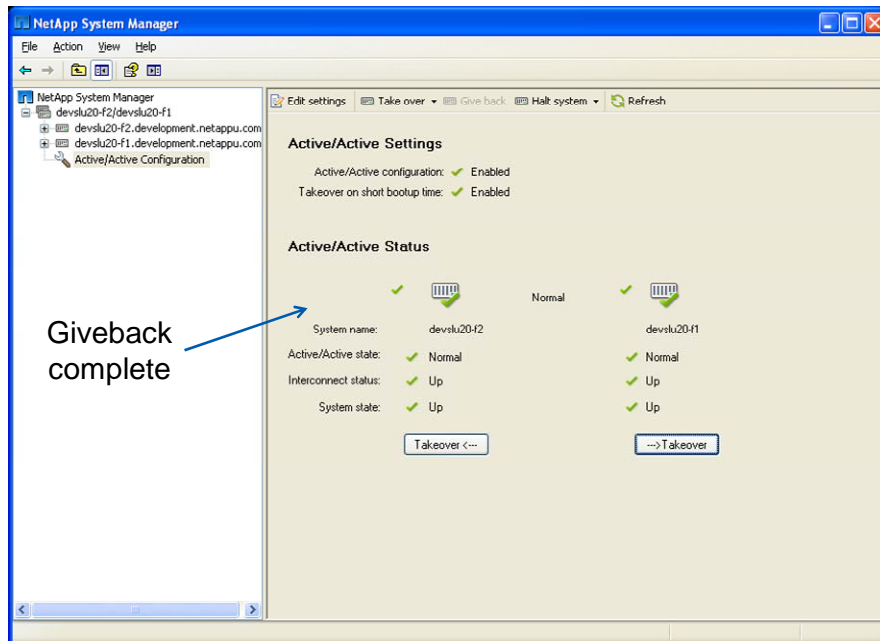


© 2010 NetApp, Inc. All rights reserved.

## NETAPP SYSTEM MANAGER: HIGH-AVAILABILITY CONFIGURATION (CONT.)



## System Manager: High-Availability (Cont.)



Giveback  
complete

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: HIGH-AVAILABILITY (CONT.)



## Failover Effects on Client Connections

- For clients or applications using NFS v3, NFS v4, HTTP, FC, and iSCSI protocols, I/O requests are suspended during the takeover/giveback period
- Connections can usually be resumed when the takeover or giveback process is complete
- For CIFS (SMB 1.0), sessions are lost
- Stateful clients and applications may—and usually do—attempt to re-establish the session

© 2010 NetApp, Inc. All rights reserved.

### FAILOVER EFFECTS ON CLIENT CONNECTIONS



## Negotiated Failover

- Data ONTAP allows failover to occur with failure of one or more network interfaces to ensure continual client interaction
- To enable negotiated failover (NFO): Off by default  

```
system> options
cf.takeover.on_network_interface_failure on
```
- To configure policy for marked NICs:  

```
system> options
cf.takeover.on_network_interface_failure.policy
<all_nics | any_nics>
```

 ← Set on each system in high-availability pair
- To mark a NIC to participate in NFO:  

```
system> ifconfig <interface> nfo
```
- To unmark NIC and prevent it from participating in NFO:  

```
system> ifconfig <interface> -nfo
```

© 2010 NetApp, Inc. All rights reserved.

## NEGOTIATED FAILOVER

To enable negotiated failover in the event of a failed network interface, you must explicitly enable the `cf.takeover.on_network_interface_failure` option, set the failover policy, and mark each interface that can trigger a negotiated failover (NFO).

**NOTE:** The `cf.takeover.on_network_interface_failure.policy` option must be set *manually* on each controller in a high-availability pair: `all_nics`= ALL interfaces marked for failover must fail before takeover will occur `any_nic`= ANY interface marked for failover will trigger a high-availability takeover. The use of the `cf.takeover.on_network_interface_failure` option is not the first line of defense against a network switch being a single point of failure. This option should only be considered when a single-mode vif or second-level vif cannot be used. Controller failover is disruptive to CIFS clients and can be disruptive to NFS clients using soft mounts. However, vif failover is completely nondisruptive and is therefore the preferred method. However, negotiated failover is used increasingly in a MultiStore® environment.



## Best Practices

- Test failover and giveback operations before placing high-availability controllers into production
- Monitor:
  - Performance of network
  - Performance of disks and storage shelves
  - CPU utilization of each controller to ensure it does not exceed 50%
- Enable AutoSupport

**NOTE:** For more information about high availability, please see the [High Availability](#) Web-based course

© 2010 NetApp, Inc. All rights reserved.

## BEST PRACTICES

General best practices require comprehensive testing of all mission-critical systems before introducing them into a production environment. High-availability controller testing should include takeover and giveback, or functional testing as well as performance evaluation. Extensive testing validates planning.

### **Monitor network connectivity and stability.**

Unstable networks not only affect total takeover and giveback times, they adversely affect all devices on the network in various ways. NetApp® storage controllers are typically connected to the network to serve data, so if the network is unstable, the first symptom is degradation of storage-controller performance and availability. Client service requests are retransmitted many times before reaching the storage controller, appearing to the client as slow responses from the storage controller. In a worst-case scenario, an unstable network can cause communication to time-out, and the storage controller appears to be unavailable.

During takeover and giveback operations in the high-availability controller environment, storage controllers attempt to connect to numerous types of servers on the network, including Windows® domain controllers, DNS, NIS, LDAP, and application servers. If these systems are unavailable or the network is unstable, the storage controller continues to retry establishing communications, which delays takeover or giveback times.



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Define the high-availability controller configuration
- Describe the three modes of high-availability operation with a high-availability pair
- Analyze the effect on client protocols during failover and giveback operations

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 21: High-Availability  
Estimated Time: 30 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.





## Check Your Understanding

- What are three modes of operation for a high-availability controller configuration?
- What is the purpose of using a high-availability controller configuration?
- What happens during a takeover?
- True or false?
  - Options must be set the same on both nodes.
  - The license must be set the same on both nodes.
  - Both nodes must have the same number of disks.
  - Both nodes must be part of the same domain.

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# MetroCluster

Module 22

Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## METROCLUSTER



## Module Objectives

By the end of this module, you should be able to:

- Describe a stretch MetroCluster environment
- List the basic steps to implement a stretch MetroCluster
- Describe a fabric-attached MetroCluster environment
- List the basic steps to implement a fabric-attached MetroCluster

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## MetroCluster Overview


© 2010 NetApp, Inc. All rights reserved.

### METROCLUSTER OVERVIEW



## MetroCluster

- MetroCluster is a cost-effective replication solution for combined high-availability and SyncMirror disaster recovery within a campus or metro area



### Configurations

- Stretch MetroCluster provides campus disaster recovery protection
  - Can stretch up to 500m
- Fabric MetroCluster provides Metropolitan disaster recovery protection
  - Can stretch up to 100km with FC switches
- V-Series MetroCluster

© 2010 NetApp, Inc. All rights reserved.


## METROCLUSTER

MetroCluster is a way to stretch a cluster beyond the 500 meter distance limitation. This is very valuable for sites that need a cluster on a campus or metropolitan area to allow for some localized failures as well as run as a cluster with failover integration.

This is very popular in industries and countries where a metropolitan separation is mandated for disaster recovery.



## MetroCluster: Overview



**NOTE:** This diagram displays the fabric-attached MetroCluster configuration; stretch MetroCluster available also

© 2010 NetApp, Inc. All rights reserved.


## METROCLUSTER: OVERVIEW

MetroCluster combines the reliability of a high-availability pair with the synchronous replication of SyncMirror over a distance.



## SyncMirror

- MetroCluster uses SyncMirror
  - Combines RAID 1 and RAID 4 / RAID-DP



© 2010 NetApp, Inc. All rights reserved.

## SYNCMIRROR



## Stretch MetroCluster Implementation

© 2010 NetApp, Inc. All rights reserved.


### STRETCH METROCLUSTER IMPLEMENTATION





## Loss of Building

- Loss of a entire building can be overcome by implementing stretch MetroCluster




© 2010 NetApp, Inc. All rights reserved.

## LOSS OF BUILDING



## Stretch MetroCluster



© 2010 NetApp, Inc. All rights reserved.

### STRETCH METROCLUSTER

To implement a stretch MetroCluster, the following hardware and software components must be available:

- High-availability controllers with the controller failover license, which provide automatic failover capability between sites
- SyncMirror® software, which provides an up-to-date copy of data at the remote site allowing data to be ready for access after failover without administrator intervention
- Controller Remote license, which provides a mechanism for administrators to declare a site disaster and initiate a site failover
- VI interconnects, which provide connectivity between the storage systems; the controller heartbeat is through InfiniBand connectivity

**NOTE:** Spindle capacity limitations for a stretch MetroCluster is based upon the limitation of the platform.



## Cable Types

| Fiber Type      | Data Rate | Max Distance (M) |
|-----------------|-----------|------------------|
| OM-2 (50/125UM) | 1 Gb/s    | 500              |
|                 | 2 Gb/s    | 300              |
|                 | 4 Gb/s    | 150              |
| OM-3 (50/125UM) | 1 Gb/s    | 860              |
|                 | 2 Gb/s    | 500              |
|                 | 4 Gb/s    | 270              |
| OM-3+           | 1 Gb/s    | 110              |
|                 | 2 Gb/s    | 750              |
|                 | 4 Gb/s    | 500              |

Interconnect and partner's plex 1 shelves require special cables to connect stretch MetroCluster over certain distances

© 2010 NetApp, Inc. All rights reserved.

### CABLE TYPES

Special cables are required for stretch MetroCluster at certain distances. Always refer to the NOW site for current support configurations. For a current description of supported cables, see: [http://now.netapp.com/NOW/knowledge/docs/san/guides/CFO\\_cable](http://now.netapp.com/NOW/knowledge/docs/san/guides/CFO_cable).



## Data ONTAP 8.0 7-Mode Configuration

- Storage System Platforms:
  - FAS3040/3070
  - FAS31XX (two single-controller chassis)
  - FAS60XX
- Disk Ownership Method:
  - Software only
- Interconnect hardware:
  - FC/VI adapter (FAS31XX only)
  - Copper/Fibre converters for interconnect (FAS30XX and FAS60XX)

See the MetroCluster Compatibility Matrix on the NOW site

© 2010 NetApp, Inc. All rights reserved.

### DATA ONTAP 8.0 7-MODE CONFIGURATION



## Configure Local Node

- Add the following licenses:
  - cf
  - syncmirror\_local
  - cf\_remote
- Cable local node as described previously
  - Set all used onboard FC adapters as initiators

```
system1> fcadmin config -d adapter
```

```
system1> fcadmin config -t initiator adapter
```
- Use software disk ownership to assign the disks to the correct pool
  - Remember:
    - Always assign all disks on the same loop to the same system and pool
    - Always assign all loops connected to the same adapter to the same pool
  - Verify with `storage show disk -p` command

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURE LOCAL NODE



## Software-Based Ownership

- Determined by storage system administrator
- To verify current ownership:

```
system1> disk show -v
```

| DISK  | OWNER             | POOL  | SERIAL NUMBER |
|-------|-------------------|-------|---------------|
| 0b.43 | Not Owned         | NONE  | 41229013      |
| ...   |                   |       |               |
| 0b.29 | system (84165672) | Pool0 | 41229011      |
| ...   |                   |       |               |

- To view all disks without an owner:

```
system1> disk show -n
```

| DISK  | OWNER     | POOL | SERIAL NUMBER |
|-------|-----------|------|---------------|
| 0b.43 | Not Owned | NONE | 41229013      |
| ...   |           |      |               |

© 2010 NetApp, Inc. All rights reserved.

## SOFTWARE-BASED OWNERSHIP



## Software-Based Ownership (Cont.)

- To assign disk ownership, use:

```
system> disk assign {disk_list|all|
 [-T storage_type] -n count|auto} [-p pool]...
```

- *disk\_list* is the Disk IDs of the unassigned disk
- T is either ATA, FCAL, LUN, SAS, or SATA
- p is either 0 or 1

- To assign a specific set of disks:

```
system1> disk assign 0b.43, 0b.41, 0b.39 -p 1
```

- To unassign disks:

```
system1> disk assign 0b.39 -s unowned -f
```

- s is used to specify the sysid to take ownership
- f is used to force assignment of previously assigned disks
- **NOTE:** Unassign only hot spare disks

Specify  
the Disk  
IDs that  
you wish  
to work  
with

© 2010 NetApp, Inc. All rights reserved.

## SOFTWARE-BASED OWNERSHIP (CONT.)



## Network Configuration

- Prior to Data ONTAP 7.3.2, the MetroCluster pair must be in the same subnet
- Data ONTAP 7.3.2 and later the MetroCluster pair may be in separate subnet
  - If required:
    - Set the `cf.takeover.use_mcrs_file` option to `on`
    - Edit the `/etc/mcrs` file with appropriate `ifconfig` commands
  - If configured:
    - system2 will use system1's `/etc/mcrs` file upon takeover instead of system1's `/etc/rc` file

© 2010 NetApp, Inc. All rights reserved.

## NETWORK CONFIGURATION





## Verify Configuration

- Run controller failover configuration checker
  - Cf-config-check.cgi
  - Downloaded from the NOW™ (NetApp® on the Web) site
  - Run from a host machine
- Compare options and parameters on each node
  - Adjust if necessary

© 2010 NetApp, Inc. All rights reserved.

## VERIFY CONFIGURATION



## Configure Remote Node

- Add licenses
  - cf
  - syncmirror\_local
  - cf\_remote
- Cable remote controller as described previously
  - Set all used onboard FC adapters as initiators

```
system2> fcadmin config -d adapter
```

```
system2> fcadmin config -t initiator adapter
```
- Set software disk ownership
- Configure network takeover if required
- Verify configuration
- Test failover and giveback
  - Local to remote
  - Remote to local

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURE REMOTE NODE



## Set Up Mirrors

- To create a new mirrored aggregate:

```
system> aggr create aggrname -m disk#
```

- To add a mirror to an existing aggregate:

```
system> aggr mirror aggrname
```

- For example:

```
system1> aggr mirror aggr0
```

```
system2> aggr mirror aggr0
```

Double the  
number of  
expected  
disks

NOTE: Root volumes  
must be mirrored

© 2010 NetApp, Inc. All rights reserved.

## SET UP MIRRORS



## Fabric-Attached MetroCluster Implementation


© 2010 NetApp, Inc. All rights reserved.

### FABRIC-ATTACHED METROCLUSTER IMPLEMENTATION



## Loss of Site

- Loss of a site can be overcome by implementing fabric-attached MetroCluster

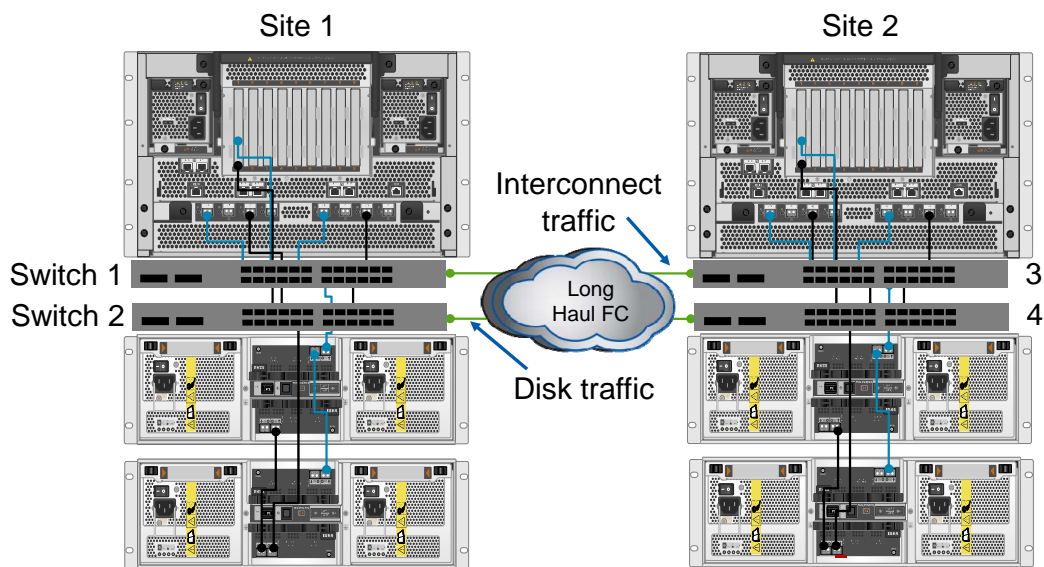


© 2010 NetApp, Inc. All rights reserved.

## LOSS OF SITE



## Fabric-Attached MetroCluster



Fabric-attached MetroCluster expands high-availability up to 100km

© 2010 NetApp, Inc. All rights reserved.

### FABRIC-ATTACHED METROCLUSTER

To implement a fabric-attached MetroCluster configuration, the following hardware and software components must be available:

- High-availability controllers with the controller failover license, which provide automatic failover capability between sites
- SyncMirror software, which provides an up-to-date copy of data at the remote site allowing data to be ready for access after failover without administrator intervention
- Controller Remote license, which provides a mechanism for administrators to declare a site disaster and initiate a site failover
- Fabric switches, which provide connectivity between the storage systems

**NOTE:** Spindle capacity limitations for a fabric-attached MetroCluster is the limitation of the platform or 672, whichever is lower.



## Supported Configurations

- Storage System Platforms:
  - FAS3040/3070
  - FAS31XX (two single-controller chassis)
  - FAS60XX
- Fabric-attached MetroCluster supports:
  - Brocade switches
    - Brocade 200E
    - Brocade 5000
    - Brocade 300
    - Brocade 5100
  - Brocade Fabric Operating System version 6.0.x or later
  - Brocade licenses:
    - Full-fabric license
    - Extended distance license (if over 10km)
    - Ports-on-demand licenses for additional ports if necessary


© 2010 NetApp, Inc. All rights reserved.

## SUPPORTED CONFIGURATIONS



## Virtual Channels

- Virtual channels (VC) separate traffic when the distance between switches is less than 10km



- Virtual Channel usage
  - # 2 - Interconnect and ISL
  - # 3 - FC ports and disk shelves
  - # 4 - FC ports and disk shelves
  - # 5 - FC ports and disk shelves

© 2010 NetApp, Inc. All rights reserved.

## VIRTUAL CHANNELS

Virtual channels for an 8-port switch and a 16-port switch only work at distances of less than 10 kilometers. If the ISL port is configured for a distance greater than 10 kilometers on these switches, all four virtual channels (2 through 5) collapse into one channel.





## Cable Types

| Fiber Type                   | Data Rate | Max Distance (M) |
|------------------------------|-----------|------------------|
| OM-2 (50/125UM)              | 1 Gb/s    | 500              |
|                              | 2 Gb/s    | 300              |
|                              | 4 Gb/s    | 150              |
| OM-3 (50/125UM)              | 1 Gb/s    | 860              |
|                              | 2 Gb/s    | 500              |
|                              | 4 Gb/s    | 270              |
| OM-3+                        | 1 Gb/s    | 110              |
|                              | 2 Gb/s    | 750              |
|                              | 4 Gb/s    | 500              |
| OS1 Single Mode<br>(9/125UM) | 2 Gb/s    | 10,000*          |
|                              | 4 Gb/s    | 10,000*          |

Special cables between the FC switches are required over certain distances

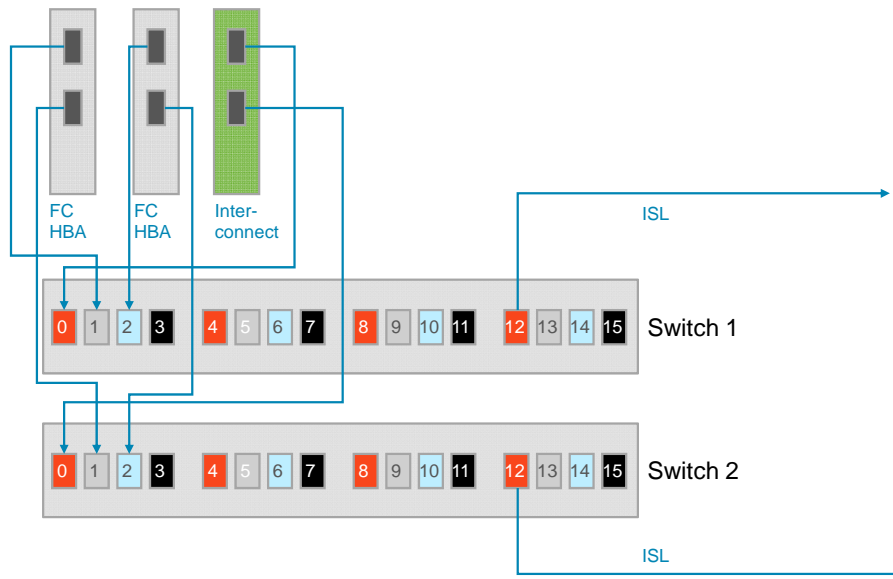
\*The maximum distance shown here is typically due to the standard 1310nm SFPs; Using 1550nm high-power SFPs, a distance of 70-100km can be achieved

© 2010 NetApp, Inc. All rights reserved.

## CABLE TYPES



## Cable Local Node (Controller)

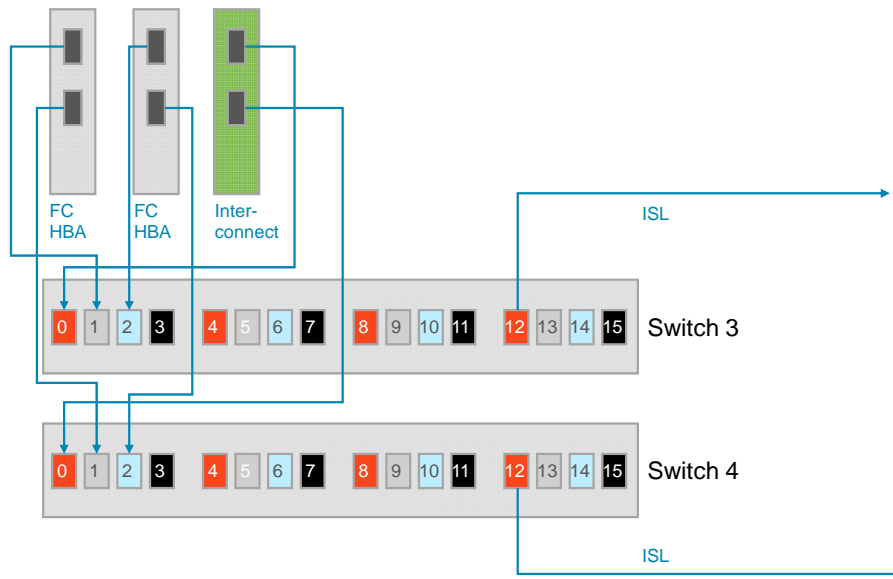


© 2010 NetApp, Inc. All rights reserved.

## CABLE LOCAL NODE (CONTROLLER)



## Cable Remote Node (Controller)

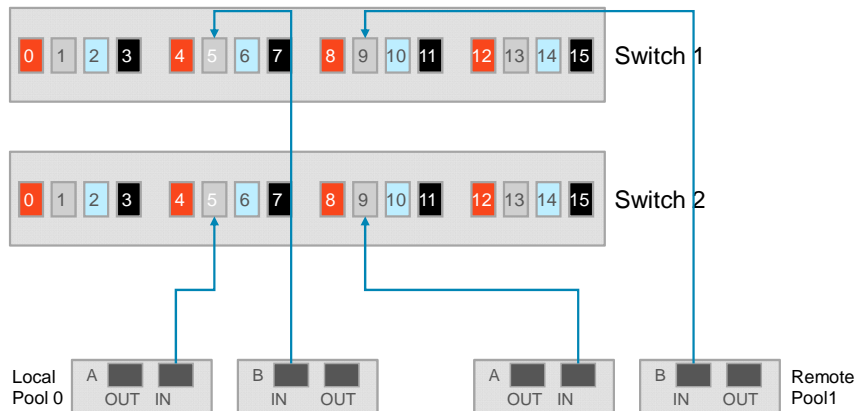


© 2010 NetApp, Inc. All rights reserved.

## CABLE REMOTE NODE (CONTROLLER)



## Cable Local Node (Shelves)



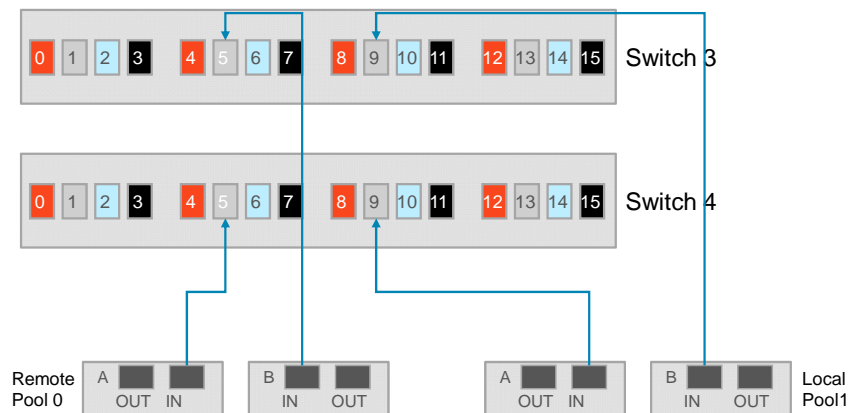
**Best Practice:** While there can be a maximum of two disk shelves per loop, performance can be maximized by installing each shelf on a different loop until all disk ports are used; then start adding the second shelves to each existing loop

© 2010 NetApp, Inc. All rights reserved.

## CABLE LOCAL NODE (SHELVES)



## Cable Remote Node (Shelves)



© 2010 NetApp, Inc. All rights reserved.

## CABLE REMOTE NODE (SHELVES)



## Configure the Switches

| Step | Action                                                                                                                                                                  | Command                                                                  |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| 1    | Log in to the switch                                                                                                                                                    |                                                                          |
| 2    | Check the licenses<br>Verify full-fabric license<br>Extended distance license (if over 10km)<br>Ports-on-demand licenses for additional ports are necessary             | <code>licenseshow</code>                                                 |
| 3    | If necessary, add needed licenses                                                                                                                                       | <code>licenseadd "license key"</code>                                    |
| 4    | Check the switch firmware<br>If necessary<br>a. Download the switch firmware from <a href="http://now.netapp">http://now.netapp</a> and update<br>b. Reboot the switch. | <code>version</code><br><br>{Exact steps omitted}<br><code>reboot</code> |

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURE THE SWITCHES



## Configure the Switches (Cont.)

| Step | Action                                                                                                                          | Command                                                                                                                                 |
|------|---------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Disable the switch:                                                                                                             | <code>switchdisable</code>                                                                                                              |
| 2    | Clear any preexisting configuration:                                                                                            | <code>cfgclear cfgdisable cfgsave</code>                                                                                                |
| 3    | Set switch to default settings:                                                                                                 | <code>configdefault</code>                                                                                                              |
| 4    | Set the switch parameters by entering the following command:<br>You should set only the following parameters: any unique number | <code>configure</code><br>Fabric parameters = <code>y</code><br>domain_id = <code>XXX</code><br>Disable device probing = <code>1</code> |
| 5    | Exit the configuration utility :                                                                                                | <code>ctrl-d</code>                                                                                                                     |
| 6    | Reboot the switch:                                                                                                              | <code>fastboot</code>                                                                                                                   |

© 2010 NetApp, Inc. All rights reserved.

### CONFIGURE THE SWITCHES (CONT.)

As a best practice, set the domain ID according to the switch number in the configuration. For example, at the local site, switches 1 and 2 would have domain IDs 001 and 002, and switches 3 and 4 at the remote site would be 003 and 004, respectively.



## Configure the Switches (Cont.)

| 1                                  | Log in to the switch and disable the switch:                                                                                                                                                                                                                                                                                                                                                                                                      | switchdisable                      |  |              |                  |    |          |     |          |    |          |    |           |                                          |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|--|--------------|------------------|----|----------|-----|----------|----|----------|----|-----------|------------------------------------------|
| 2                                  | Set all ports attached to disk loops to half duplex by entering the following command: Perform this command only for the ports to which disks are attached.                                                                                                                                                                                                                                                                                       | portcfglport <port#>,0,0,1         |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 3                                  | For Brocade 200E, 300, 5000, and 5100 only:<br>Set all ports attached to disk loops to Locked L-Port.                                                                                                                                                                                                                                                                                                                                             | portcfglport <port#>,1,0,1         |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 4                                  | Verify disk loop port is showing ON in the Locked Loop HD field                                                                                                                                                                                                                                                                                                                                                                                   | portcfglport                       |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 14                                 | Disable trunking on ISL Ports:                                                                                                                                                                                                                                                                                                                                                                                                                    | portcfgtrunkport <ISL_port#> 0     |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 15                                 | To configure the long-distance ISL port for an ISL length of up to 10 km, enter the command:<br><table border="1"><thead><tr><th colspan="2">Switch Setting Matrix for Distance</th></tr><tr><th>Port Setting</th><th>Maximum Distance</th></tr></thead><tbody><tr><td>LE</td><td>&lt;= 10 Km</td></tr><tr><td>L.5</td><td>&lt;= 25 Km</td></tr><tr><td>L1</td><td>&lt;= 50 Km</td></tr><tr><td>L2</td><td>&lt;= 100 Km</td></tr></tbody></table> | Switch Setting Matrix for Distance |  | Port Setting | Maximum Distance | LE | <= 10 Km | L.5 | <= 25 Km | L1 | <= 50 Km | L2 | <= 100 Km | portcfglongdistance<br><ISL_port#>, "LE" |
| Switch Setting Matrix for Distance |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                    |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| Port Setting                       | Maximum Distance                                                                                                                                                                                                                                                                                                                                                                                                                                  |                                    |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| LE                                 | <= 10 Km                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                    |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| L.5                                | <= 25 Km                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                    |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| L1                                 | <= 50 Km                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                    |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| L2                                 | <= 100 Km                                                                                                                                                                                                                                                                                                                                                                                                                                         |                                    |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 16                                 | Enable the switch:                                                                                                                                                                                                                                                                                                                                                                                                                                | switchenable                       |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 17                                 | Set the switch name                                                                                                                                                                                                                                                                                                                                                                                                                               | switchname Metro1_switch           |  |              |                  |    |          |     |          |    |          |    |           |                                          |
| 18                                 | Verify that the switch settings are correct                                                                                                                                                                                                                                                                                                                                                                                                       | configshow                         |  |              |                  |    |          |     |          |    |          |    |           |                                          |

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURE THE SWITCHES (CONT.)

**NOTE:** When setting the parameters for the ports for disk loops, go ahead and set the parameters for what might be unused disk loop ports now. Then later on when installing additional storage the switches will already be configured. For more information, please see the Fabric-attached MetroCluster Brocade Switch Configuration Guide found at the NOW site.





## Zoning

- NetApp recommends zoning to isolate FC traffic
  - Required if it is a greater than 336-disks configuration
- Create separate zones for FC-VI and the storage disks
- Switch 1 and 3 example:

```
switch> zonecreate "FCVI", "1,0; 3,0"
switch> zonecreate "STOR", "1,1; 1,2; 1,3; 1,5; 1,6;
1,7; 1,9; 1,10; 1,11; 1,13; 1,14; 1,15; 1,17; 1,18;
1,19; 1,21; 1,22; 1,23; 3,1; 3,2; 3,3; 3,5; 3,6; 3,7;
3,9; 3,10; 3,11; 3,13; 3,14; 3,15; 3,17; 3,18; 3,19;
3,21; 3,22; 3,23"
switch> cfgcreate "Zone_netapp", "FCVI; STOR"
switch> cfgenable "Zone_netapp"
switch> cfgsave "Zone_netapp"
switch> cfgshow
```

© 2010 NetApp, Inc. All rights reserved.

## ZONING


Example for switch 2 and 4:

```
switch> zonecreate "FCVI", "2,0; 4,0"
switch> zonecreate "STOR", "2,1; 2,2; 2,3; 2,5; 2,6; 2,7; 2,9; 2,10; 2,11; 2,13;
2,14; 2,15; 2,17; 2,18; 2,19; 2,21; 2,22; 2,23; 4,1; 4,2; 4,3; 4,5; 4,6; 4,7;
4,9; 4,10; 4,11; 4,13; 4,14; 4,15; 4,17; 4,18; 4,19; 4,21; 4,22; 4,23"
switch> cfgcreate "Zone_netapp", "FCVI; STOR"
switch> cfgenable "Zone_netapp"
switch> cfgsave "Zone_netapp"
switch> cfgshow
```



## Traffic Isolation Technique

- Allows flow control of interswitch traffic by creating a dedicated path for traffic flowing from a specific set of source ports
- Requires:
  - Data ONTAP 7.2.6.1 or later
  - Brocade 300, 5000, or 5100 switches
  - Fabric OS 6.0.0b or later



© 2010 NetApp, Inc. All rights reserved.

### TRAFFIC ISOLATION TECHNIQUE

Effective in the Brocade fabric operating system (FOS) version 6.0.0b, it is now possible to dedicate inter-switch links to certain traffic. The traffic isolation feature allows you to control the flow of interswitch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N\_Ports). For example, you might use traffic isolation for the following scenarios:

- To dedicate an ISL to high-priority cluster-interconnect traffic such as NVRAM mirroring traffic
- To isolate high-priority traffic from the disruptions that might be caused by N\_ports and E\_ports handling high-volume but low-priority traffic

Traffic isolation is implemented using a special zone, called a *traffic isolation zone (TI zone)*. A TI zone indicates the set of ports and ISLs to be used for a specific traffic flow. When a TI zone is activated, the fabric attempts to isolate all inter-switch traffic entering from a member of the zone to only those ISLs that have been included in the zone. The fabric also attempts to exclude traffic not in the TI zone from using ISLs within that TI zone.

Diagram on the slide shows a fabric with a TI zone consisting of N\_Ports “1,8” and “4,6” and E\_Ports “1,1”, “3,9”, “3,12”, and “4,7”. The dotted line indicates the dedicated path from domain 1 to domain 4. All traffic entering domain 1 from port 8 is routed through the ISL on port 1. Similarly, traffic entering domain 3 from port 9 is routed to the ISL on port 12, and traffic entering domain 4 from the ISL on port 7 is routed to the device through port 6. Traffic coming from other ports in domain 1 would not use port 1, but would use port 2 instead. Other traffic is excluded from the dedicated path as long as other equal-cost routes through the fabric exist. For example, if the ISL formed by E\_Ports “1,2” and “3,10” failed, all traffic between domains 1 and 3 would use the ISL formed by E\_Ports “1,1” and “3,9” even though that ISL is a dedicated path in a TI zone.



## Configure Local Node

- Add licenses
  - `cf`
  - `syncmirror_local`
  - `cf_remote`
- Cable remote controller as described previously
  - Set all used onboard FC adapters as initiators

```
system1> fcadmin config -d adapter
system1> fcadmin config -t initiator adapter
```
- Set software disk ownership
- Configure network takeover if required
- Verify configuration

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURE LOCAL NODE



## Verify Connectivity

- Use the following command:

```
system1> storage show disk -p
```

|                                   | PRIMARY  | PORT | SECONDARY | PORT | SHELF | BAY  |
|-----------------------------------|----------|------|-----------|------|-------|------|
|                                   | -----    | ---- | -----     | ---- | ----- | ---- |
| Local<br>Pool0                    | sw2:5.16 | A    | sw1:5.16  | B    | 1     | 0    |
|                                   | sw1:5.17 | B    | sw2:5.17  | A    | 1     | 1    |
|                                   | sw2:5.18 | A    | sw1:5.18  | B    | 1     | 2    |
|                                   | sw1:5.19 | B    | sw2:5.19  | A    | 1     | 3    |
| ...                               |          |      |           |      |       |      |
| switch_name:switch_port.device_id |          |      |           |      |       |      |
| Remote<br>Pool1                   | sw2:9.16 | A    | sw1:9.16  | B    | 1     | 0    |
|                                   | sw1:9.17 | B    | sw2:9.17  | A    | 1     | 1    |
|                                   | sw2:9.18 | A    | sw1:9.18  | B    | 1     | 2    |
|                                   | sw1:9.19 | B    | sw2:9.19  | A    | 1     | 3    |
| ...                               |          |      |           |      |       |      |

Verify dual paths

© 2010 NetApp, Inc. All rights reserved.

## VERIFY CONNECTIVITY



## Configure Remote Node

- Add licenses
  - `cf`
  - `syncmirror_local`
  - `cf_remote`
- Cable remote controller as described previously
  - Set all used onboard FC adapters as initiators

```
system2> fcadmin config -d adapter
system2> fcadmin config -t initiator adapter
```
- Set software disk ownership
- Configure network takeover if required
- Verify configuration
- Test failover and giveback
- Set up mirror aggregates

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURE REMOTE NODE




## Site Failover and Recovery

© 2010 NetApp, Inc. All rights reserved.

### SITE FAILOVER AND RECOVERY



## Host Failure



### Recovery - Automatic

- Host1 and Host2 clustered
- Host1 fails
- Host2 will access its data that Host1 was using from System1
- No difference than high-availability


© 2010 NetApp, Inc. All rights reserved.

## HOST FAILURE

If a configuration has host clustering, the application fails over to Host 2. This is the same as if it was a standard storage system cluster with the exception that Host 2 will be accessing its data from System1 as normal.



## Controller Failure



### Recovery - Automatic

- System1 fails
- Host1 will get its data from System2 but access the shelf in Data Center #1
- No difference than high-availability

© 2010 NetApp, Inc. All rights reserved.


## CONTROLLER FAILURE

This process is the same as a normal high-availability failover because storage system #2 can access the disks in both data centers.





## Controller and Disk Failure



### Dual Failure Recovery

- This is a site failure
- MetroCluster is required for business continuity
- One step failover with `cf forcetakeover -d` command

© 2010 NetApp, Inc. All rights reserved.

## CONTROLLER AND DISK FAILURE

If a controller and a disk shelf (or shelves) fails at one site:


No automated storage system takeover occurs. Storage system #2 will continue to run serving its LUNs/volumes.

The application may or may not fail over to host #2 depending on the clustering software being used. There is no access to data being served by storage system #1.

The manual `cf forcetakeover -d` command causes a storage system controller takeover to occur. Data then needs to be set online to allow operations to continue.



## Interconnect Failure



### Recovery

- No failover; mirroring disabled
- Both controller heads will continue to run serving their LUNs/volumes
- Re-syncing happens automatically after the interconnect is reestablished

© 2010 NetApp, Inc. All rights reserved.

## INTERCONNECT FAILURE

If there is an interconnect failure:

No automated storage system takeover occurs. Both storage system heads will continue to run serving their LUNs/volumes.

Both storage system heads are still running. This is not a problem because each head is “only” serving its LUNs/volumes.


Assuming that each host is accessing its data locally (that is, host #1 accessing LUNs/volumes served by storage system #1, host #2 accessing LUNs/volumes served by storage system #2) applications continue to run without problems.

Resyncing happens automatically after the interconnect is reestablished.

The host cluster will probably fail over because the MetroCluster pair cannot communicate.



## Interconnect and Host Failure



### Recovery

- Host1 and Host2 clustered
- To avoid “split brain”, perform a site failover
- One step failover with `cf forcetakeover -d` command
- Make sure to shut down the System1

© 2010 NetApp, Inc. All rights reserved.

## INTERCONNECT AND HOST FAILURE

If the interconnect and a host failure occurs:

No automated storage system take over occurs. Both storage system heads will continue to run serving their LUNs/volumes.


Both storage system heads are still running. This is not a problem because each head is “only” serving its LUNs/volumes.

The application will fail over to host #1 because the heartbeat is dead and only the storage system partner has access to the quorum/lock disk but it cannot access the data being served by storage system #2. There is no access to data being served by storage system #2.

The manual `cf forcetakeover -d` command causes storage system takeover to happen to allow operations to continue. The other storage system needs to be shut down to avoid storage system split-brain behavior. The LUNs /volumes then need to be set online to allow operations to continue.



## MetroCluster: Site Failure



One step failover :  
system2> **cf forcetakeover -d**

© 2010 NetApp, Inc. All rights reserved.

### METROCLUSTER: SITE FAILURE

There are several situations that could necessitate a site takeover. They include:

- Complete environmental failure (for example: air conditioning or power).
- Geographic disaster (for example: earthquake, fire, or flood).

During a site failure:

No automated storage system takeover occurs. Storage system #2 will continue to run serving its LUNs/volumes.

The application may or may not fail over to host #2 depending on the clustering software being used. There is no access to data being served by storage system #1.

The manual `cf forcetakeover -d` command causes a storage system controller takeover to occur. Data then needs to be set online to allow operations to continue.



## Site Failover and High-Availability Failover

- High-availability failover:
  - Automatic (if enabled)
  - Process relies on the surviving controller being able to read information (quorum) from the disks on the failed controller
  - If the quorum of disks is unavailable, automatic takeover doesn't occur
- Site failover:
  - Manual process by execute using `cf forcetakeover -d` command
  - Takeover occurs in spite of the lack of access to quorum disks

© 2010 NetApp, Inc. All rights reserved.

### SITE FAILOVER AND HIGH-AVAILABILITY FAILOVER



## Site Failover Consequences

- Breaks the mirrored relationship
  - Volumes have a new file system ID (FSID) to avoid conflict
    - LUNs have a new serial number & need to be brought online  
`system> lun online /vol/vol1/lun1`
    - Previous NFS mounts are stale and will need to be remounted
  - Data ONTAP 7.2.4 and later has a new `cf.takeover.change_fsid` option
    - Default off
    - If set on:
      - Allows LUNs to retain original serial number and brought online automatically
      - NFS mounts to be brought online automatically

© 2010 NetApp, Inc. All rights reserved.

## SITE FAILOVER CONSEQUENCES



## Split Brain Scenario

- If `cf forcetakeover -d` is used, then when the problem at the failed site is resolved, administrators must isolate the failed site to a split brain scenario
- Split Brain occurs when the failed site comes back up and doesn't know a takeover occurred
  - For example: System1 reads/writes data from System1's pool0 and System2 reads/writes data from System1's pool1
  - Possible data corruption
- To prevent a split brain:
  - Restrict access to previous failed site controller until proper site recovery

© 2010 NetApp, Inc. All rights reserved.

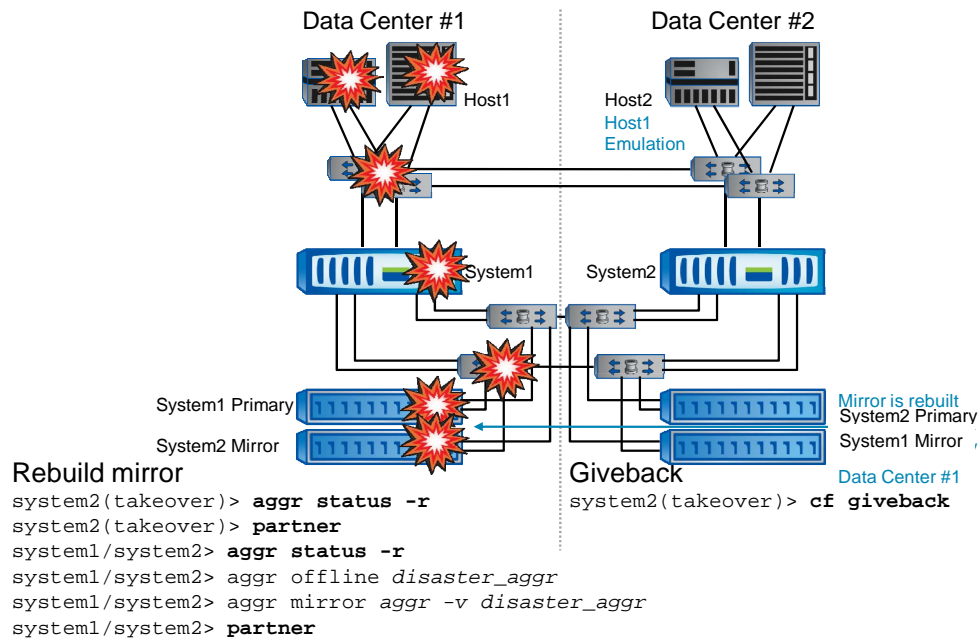
### SPLIT BRAIN SCENARIO

To prevent a split brain:

1. Turn off power to the previous failed node. Disk shelves should be left on.
2. Disconnect the interconnect and Fibre Channel adapter cables of the node at the surviving site.
3. Use network management procedures to enable the storage systems at the disaster site to be isolated from the external public network.
4. Use any application-specific method that either prevents the application from restarting at the disaster site or prevents the application clients from accessing the application servers at the disaster site.



## MetroCluster: Site Recovery



© 2010 NetApp, Inc. All rights reserved.

### METROCLUSTER: SITE RECOVERY

During a site recovery:

Storage system #1 is repaired.

The mirror is rebuilt. While the re-establishment of the aggregate mirrors is performed, there is no disruption to users since the storage systems are still in failover mode.

The storage system high-availability pair is returned to normal with a `cf giveback` command.

Host #1 and Host #2 may operate as normal.





## Module Summary

© 2010 NetApp, Inc. All rights reserved.

### MODULE SUMMARY

See Technical Report 3548 *Best Practice for MetroCluster Design and Implementation* for more information about stretch and fabric MetroCluster.



## Module Summary

In this module, you should have learned to:

- Describe a stretch MetroCluster environment
- List the basic steps to implement a stretch MetroCluster
- Describe a fabric-attached MetroCluster environment
- List the basic steps to implement a fabric-attached MetroCluster

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 22: MetroCluster  
Estimated Time: 0 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- MetroCluster may be configured in what three configurations?
- What three special licenses are required to implement a MetroCluster?
- True or false? With special cables and FC switches, Fabric-attached MetroCluster configurations can span up to 100km.

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# SnapMirror

Module 23

Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## SNAPMIRROR



## Module Objectives

By the end of this module, you should be able to:

- Explain the SnapMirror® Async, Sync, and Semi-Sync modes of operation
- Describe how volume SnapMirror and qtree SnapMirror replicate data
- Configure SnapMirror
- Perform advanced SnapMirror operations
- Explain SnapMirror performance impact

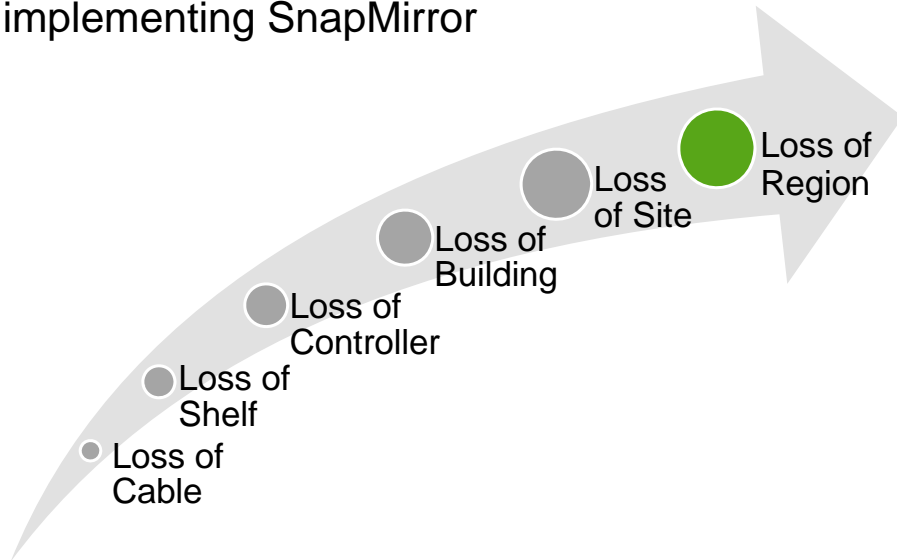
© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Loss of Region

- Loss of a region can be overcome by implementing SnapMirror




© 2010 NetApp, Inc. All rights reserved.

## LOSS OF REGION



# SnapMirror



- SnapMirror allows mirroring volumes or qtrees

© 2010 NetApp, Inc. All rights reserved.

## SNAPMIRROR





## Features and Benefits of SnapMirror

- A fast and flexible enterprise solution that addresses critical applications areas:
  - Data migration
  - Data replication
  - Remote access and load sharing
  - Disaster recovery
  - Remote tape archival

© 2010 NetApp, Inc. All rights reserved.

### FEATURES AND BENEFITS OF SNAPMIRROR

SnapMirror provides a fast and flexible enterprise solution for replicating data over local area, wide area, and Fibre Channel networks. SnapMirror addresses multiple application areas such as mission-critical data protection, and business continuance in case of a disaster.

Data **migration** from one storage system to another can be done without interrupting network service.

SnapMirror replication allows the **distribution** of large amounts of data to remote sites as a read-only replica. Remote data access provides fast access to data by local clients.

If critical data is replicated to a different location, in case of a **disaster** at the source site, the replica can be made available to clients across the network until the damage caused by the disaster is repaired.

Additionally, as the source data can be replicated at a time chosen by systems administrators, the solution **minimizes network utilization**.

SnapMirror is also used for **backup offloading**. SnapMirror technology attaches the off-site storage device to the SnapMirror destination system, offloading tape backup overhead from production servers.



## Replication Modes

- Async
  - Replicates Snapshot™ copies from a source volume or qtree to a destination volume or qtree
  - Incremental updates are based on schedules
- Sync
  - Replicates writes from a source volume to a secondary volume at the same time it is written to the source volume
- Semi-Sync
  - Minimizes source system performance impact

© 2010 NetApp, Inc. All rights reserved.

### REPLICATION MODES

The Data ONTAP SnapMirror feature enables an administrator to replicate data either asynchronously or synchronously.

The SnapMirror **Async** mode replicates Snapshot copies from a source volume or qtree to a destination volume or qtree. Incremental updates are based on a schedule or are performed manually using the `snapmirror update` command. Async mode works with both volume SnapMirror and qtree SnapMirror.

SnapMirror **Sync** mode replicates writes from a source volume to a destination volume at the same time it is written to the source volume. SnapMirror Sync is used in environments that have zero tolerance for data loss.

SnapMirror **Semi-Sync** provides a middle-ground solution that keeps the source and destination systems more closely synchronized than Async mode, but with less impact on performance.



## Volume and Qtree SnapMirror


© 2010 NetApp, Inc. All rights reserved.

### VOLUME AND QTREE SNAPMIRROR



## Volume SnapMirror

- Block-for-block replication
- Supports Async, Sync, and Semi-Sync
- Can be initialized using a tape device
- Support volume cascade in a series
- Source volume
  - Online/writable
- Destination volume
  - Online/read-only



© 2010 NetApp, Inc. All rights reserved.

### VOLUME SNAPMIRROR

Volume SnapMirror enables block-for-block replication. The entire volume, including its qtrees, and all the associated Snapshot copies, are replicated to the destination volume.

Volume SnapMirror can be initialized using a tape device and a volume's replica can be cascaded in a series.

Volume SnapMirror can be asynchronous, synchronous, or semi-synchronous.

The source volumes are visible, accessible, and writable by the clients.

The destination volumes are visible, accessible, and read-only and usually on a separate system to which the source volumes are replicated.



## Initial Transfer and Replication

- Initial baseline transfer
  - Create a non-root restricted destination volume
  - All data in all Snapshot copies on the source are transferred to the destination volume
  - Read-only destination volume brought online after initial transfer completed
- Replication
  - Scheduled process updates the mirror
  - Current Snapshot copy is compared with the previous Snapshot copy
  - Changes are synchronized from source to destination

© 2010 NetApp, Inc. All rights reserved.

### INITIAL TRANSFER AND REPLICATION

To initialize a volume, you first have to restrict the destination volume in which the replica will reside. During the baseline transfer, the source storage system takes a Snapshot copy of the volume. All data blocks referenced by this Snapshot copy, including volume metadata such as language translation settings, as well as all Snapshot copies of the volume, are transferred and written to the destination volume.

After the initialization completes, the source and destination file systems have one Snapshot copy in common. Updates occur from this point and are based on the schedule specified in a flat-text configuration file known as the `snapmirror.conf` file or by using the `snapmirror update` command.



## Requirements and Limitations

- Destination's Data ONTAP® version must be equal to or more recent than the source
- Like-to-like transfers only: flex-to-flex
- Destination volume capacity equal to or greater than source
- Quota cannot be enabled on destination volume
- TCP port range 10565–10569 must be open

© 2010 NetApp, Inc. All rights reserved.

## REQUIREMENTS AND LIMITATIONS

With volume SnapMirror, the destination must run a version of Data ONTAP that is equal to or more recent than the source. In addition, the source and destination must be on the same Data ONTAP release.

Volume SnapMirror replication can only occur with volumes of the same type: both traditional volumes or both flexible volumes.

Volume SnapMirror requires that the size of the destination volume be equal to or greater than the size of the source volume. Administrators can thin provision the destination so that it appears to be equal to or greater than the size of the source volume.

The source volume disk's checksum type (block or zone checksum) must be identical to the destination disk's checksum type.

**NOTE:** For traditional volume, disks' checksum type, size, and geometry must be identical. The destination volume has to contain the same number of disks and the same size disks as the source volumes, allowing more efficient deployment of resources. This limitation applies only to traditional volumes.


Quotas cannot be enabled on destination volume.

It is recommended that you allow a range of TCP ports from 10565 to 10569.



## Qtree SnapMirror

- Logical replication
- Independent of the type of volume
- Source volume and qtree are online/writable
- Destination volume is online/writable
- Destination qtree is read-only



© 2010 NetApp, Inc. All rights reserved.

### QTREE SNAPMIRROR

Qtree SnapMirror is a logical replication. All the files and directories in the source file system are created in the target destination qtree.

Qtree SnapMirror replication occurs between qtrees regardless of the type of the volume (traditional or flexible). You can mirror qtrees using SnapMirror from a traditional volume to a flexible volume and from a flexible volume to a traditional volume. Note that volume size and the disk geometry do not make any difference for qtree SnapMirror.

Qtrees from different sources can be replicated to a destination and qtree SnapMirror replication can occur between different releases of Data ONTAP.

With qtree SnapMirror, the source volume and qtree are online and writable. The destination qtree is read-only, while the destination volume remains writable and may contain replicated qtrees from multiple source volumes, and qtrees or nontree data not managed by SnapMirror.

**NOTE:** Unlike volume SnapMirror, qtree SnapMirror does not require that the size of the destination volume be equal to or greater than the size of the source qtree.



## Initial Transfer and Replication

- Initial baseline transfer
  - Destination qtree is created automatically upon first-time replication
- Replication
  - Scheduled process updates the qtree replica
  - Current Snapshot copy is compared with the previous Snapshot copy
  - Only changed blocks (identified through the inodes) are transferred to the destination; Snapshot copy is not transferred
  - When transfer completes, qtree SnapMirror creates a Snapshot copy associated with the qtree replica

© 2010 NetApp, Inc. All rights reserved.

### INITIAL TRANSFER AND REPLICATION

To initialize a qtree, you do not need to create a destination qtree; the qtree is automatically created when the baseline transfer is started. The baseline transfer is created when qtree SnapMirror creates a Snapshot copy of the source volume that contains the qtree to be replicated. This Snapshot copy contains all the source volume data, including both the data in the qtree to be replicated and the metadata.

After the initialization completes, the source and destination file systems have one Snapshot copy in common. Updates occur from this point and are based on the schedule specified in a flat-text configuration file known as the `snapmirror.conf` file or by using the `snapmirror update` command.

Qtree SnapMirror determines changed data by first looking through the inode file for inodes that have changed and then looking through the changed inodes of the replicated qtree for changed data blocks. Only new or changed blocks associated with the designated qtree are transferred to the destination. Qtree SnapMirror does not transfer the Snapshot copy from the source to the destination. When the transfer completes, qtree SnapMirror creates a Snapshot copy of the destination volume associated with the replicated qtree.





## Requirements and Limitations

- Supports Async mode only
- Destination volume must have 5% extra space
- Destination qtree cannot be `/etc`
- Cannot be initialized using a tape device
- Does not support cascading of mirrors
- Deep directory structure and large number of small files may impact performance

© 2010 NetApp, Inc. All rights reserved.

## REQUIREMENTS AND LIMITATIONS

Qtree SnapMirror is available in asynchronous mode only. The destination volume must contain 5% more free space than the source qtree. A destination qtree cannot be the `/etc` file. Qtree SnapMirror cannot be initialized using a tape device and does not support cascading of mirrors. Qtree SnapMirror performance is impacted by deep directory structure and large number (tens of millions) of small files replicated.

To determine changed data, qtree SnapMirror looks at the inode file and defines which inodes are in the qtree of interest and which inodes have changed. If the inode file is large, but the inodes of interest are few, qtree SnapMirror spends a lot of time going through the inode file to find very few changes.



## SnapMirror Deployment

© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR DEPLOYMENT



## Licensing SnapMirror

- One `snapmirror` license is required per source and destination storage system
  - Evaluation license keys are available upon request on the NOW™ (NetApp® on the Web) site
- SnapMirror Sync and Semi-Sync require the additional `snapmirror_sync` license
  - Special license keys are available in the *Data ONTAP Data Protection Online Backup and Recovery Guide*

© 2010 NetApp, Inc. All rights reserved.

### LICENSING SNAPMIRROR

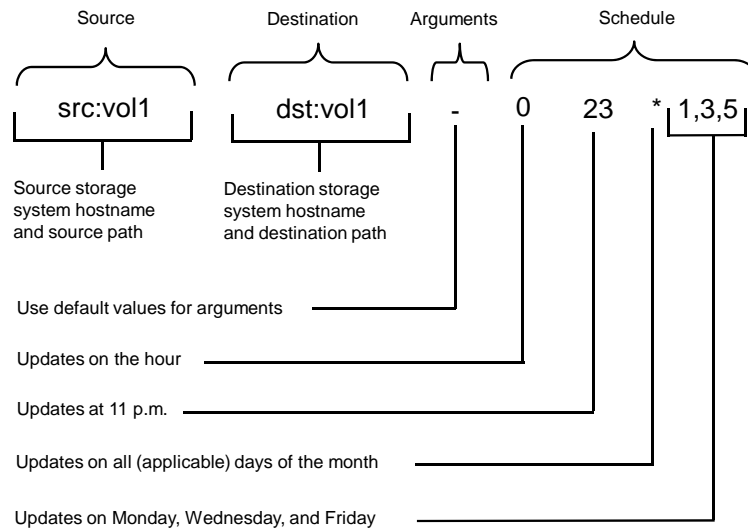
SnapMirror requires a `snapmirror` license on both the source and the destination storage systems. If the SnapMirror source and destination are on the same storage system, only one license is to be installed.

SnapMirror Sync and Semi-Sync require an additional `snapmirror_sync` free license available in the *Data ONTAP Data Protection Online Backup and Recovery Guide* on the NOW site.





## snapmirror.conf



© 2010 NetApp, Inc. All rights reserved.

## SNAPMIRROR.CONF

The `snapmirror.conf` configuration file entries define the relationship between the source and the destination, the mode of replication, and the arguments that control SnapMirror when replicating data. The syntax for entries in the `snapmirror.conf` file is as follows:

```
src_system: /vol/src_vol/[src_qtree]
dest_system: /vol/dest_vol[/dest_qtree] [arguments] [schedule]
```

The **arguments** field lets you define the transfer speed and the restart mode. In this field you can also enable checksum, set the synchronicity level and the visibility interval. A dash (-) indicates that all arguments' default values apply. The **schedule** consists of four space-separated fields in order: *minute*, *hour*, *day\_of\_month*, and *day\_of\_week*.

All possible values can be applied with an asterisk (\*). A single dash (-) means “never” and prevents this schedule entry from executing.

**NOTE:** SnapMirror updates can be scheduled to occur as frequently as every minute.

For more information on the `snapmirror.conf` file entries, refer to the *Data ONTAP Data Protection Online Backup and Recovery Guide*.



## snapmirror.conf File Examples

Example entries from `snapmirror.conf` file that resides on the destination

- Qtree SnapMirror

```
src:/vol/vol1/q1 dst:/vol/vol1/q1 - 15 * * *
```

- Volume SnapMirror

```
src:vol2 dst:vol2 kbs=2000 10 8,20 * *
```

© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR.CONF FILE EXAMPLES

```
src:/vol/vol1/q1 dst:/vol/vol1/q1 - 15 * * *
```

The source qtree `q1` is replicated to the destination qtree `q1` every hour on the hour 15 minutes past the hour, every day of the week, and every day of the month.

```
src:vol2 dst:vol2 kbs=2000 10 8,20 * *
```

The source volume `vol2` is replicated to the destination volume `vol2` at 10 minutes past 8:00 a.m. and 8:00 p.m. every day of the month and every day of the week. In other words, the source volume `vol2` is replicated at 8:10 a.m. and 8:10 p.m. Data ONTAP can use a maximum of 2,000 kilobytes per second to transfer data.



## Monitoring Transfer

- Use the `snapmirror status` command to monitor transfer progress, check relationship status, and control the lag

```
dst> snapmirror status
Snapmirror is on.
Source Destination State Lag Status
src:vol1 dst:vol1 Snapmirrored 00:05:30 Idle
src:/vol/vol2/q1 dst:/vol/vol2/q1 Snapmirrored 00:09:53 Quiescing
src:/vol/vol2/q2 dst:/vol/vol2/q2 Snapmirrored 00:15:20 (Transferring
 12288 MB done)
```

© 2010 NetApp, Inc. All rights reserved.

## MONITORING TRANSFER

Use the `snapmirror status` command either from the source or the destination storage system to monitor the transfer progress, check relationships state and status, and control the transfer lag (age of backup).

Command syntax:

```
snapmirror status [options] [system:][path]...
```

The value for *options* can be `-l` or `-q`

The `-l` option displays the long format of the output.

The `-q` option displays which volumes or qtrees are quiesced or quiescing.

The `snapmirrorquiesce` command waits for all existing transfers to both volumes and qtrees to complete and blocks any further updates. If a qtree is not in a stable state (is in transition), the `snapmirrorquiesce` command forces it into a stable state. You can quiesce only volumes and qtrees that are online and that are SnapMirror destinations. You cannot quiesce a restricted or offline volume or a qtree in a restricted or offline volume.

For details on the `snapmirror status` command output, refer to the *Data ONTAP Data Protection Online Backup and Recovery Guide*.



## Listing Snapshot Copies

- Use the `snap list` command to monitor Snapshot copies deletion and creation
- DO NOT delete SnapMirror Snapshot copy

```
dst> snap list vol1
Volume vol1
DateName

Sep 13 21:00 hourly.0
Sep 13 20:45 dst(0050404361)_vol1.0 (snapmirror)

Destination system _____
Destination system ID _____
Destination volume _____
Transfer incremental number _____
```

© 2010 NetApp, Inc. All rights reserved.

### LISTING SNAPSHOT COPIES

Use the `snap list` command to list all Snapshot copies, including the SnapMirror Snapshot copies stored in the source and the destination volumes.

SnapMirror Snapshot copies are distinguished from system Snapshot copies by a more elaborate naming convention. The default name of a SnapMirror volume Snapshot copy is as follows:

`dest_system(sysid)_name.number`

`dest_system` is the host name of the destination storage system

`sysid` is the destination system ID number

`name` is the name of the destination volume

`number` is the number of successful transfers for the Snapshot copy, starting at 1. Data ONTAP increments this number for each transfer.

The `snap list` command displays the keyword ***snapmirror*** next to the necessary Snapshot copy.

**CAUTION:** Do not delete Snapshot copies that SnapMirror creates in the source volume. The most recent SnapMirror Snapshot copy is referred to as the newest common Snapshot copy. Incremental changes to the destination depend on this Snapshot copy. If SnapMirror cannot find the required Snapshot copy on the source, it cannot perform incremental changes to the destination. The affected relationship will have to be reinitialized.





## Log Files

### ■ SnapMirror logging

`options snapmirror.log.enable[on|off]`

- The option is on by default
- Log files `/etc/log/snapmirror.[0-5]` are saved in the root volume

### ■ Example:

```
dst Thu Sep 13 20:41:09 GMT src:vol1 dst:vol1 Request (Initialization)
dst Thu Sep 13 20:41:32 GMT src:vol1 dst:vol1 Abort(Destination not allowed)
dst Thu Sep 13 20:45:31 GMT src:vol1 dst:vol1 Request (Initialization)
dst Thu Sep 13 20:45:35 GMT src:vol1 dst:vol1 Start
dst Thu Sep 13 21:10:40 GMT src:vol1 dst:vol1 End (104857600 KB)
```

© 2010 NetApp, Inc. All rights reserved.

## LOG FILES

The SnapMirror logs record whether the transfer finished successfully or failed. If there is a problem with the updates, it is useful to look at the log file to see what has happened since the last successful update. The logs include the start and end of each transfer, along with the amount of data transferred.

Use the option `snapmirror.log.enable` to record SnapMirror data transfer logs. By default, the option is on.

`options snapmirror.log.enable [on|off]`

Log files are stored on the source and the destination storage system's root volume, in the `/etc/logs/snapmirror` directory.

A new log file is generated every week as `snapmirror.0`. Older log files are renamed `snapmirror.[1-5]` and the oldest log file is deleted.



## NearStore Personality

- Converts the destination storage system to a NearStore® system
- Increases the number of concurrent transfers on those destination systems
- Requires the `nearstore_option` license on the secondary and Data ONTAP 7.1 or later

© 2010 NetApp, Inc. All rights reserved.

### NEARSTORE PERSONALITY

NearStore® Personality allows you to utilize FAS systems as secondary systems. This feature requires the `nearstore_option` license.

When enabled, the `nearstore_option` license increases the number of possible concurrent destination qtree SnapMirror and SnapVault® replications by optimizing the transfer resources required for those replications. This license should not be installed on these storage systems if they intend to handle primary application workloads.

For more information on transfer resources required for NearStore replications, refer to the latest Data ONTAP *Data Protection Online Backup and Recovery Guide* on the NOW site.



## Concurrent Transfers

- Each storage system model supports a maximum number of simultaneous replication operations
- In Data ONTAP 7.3 and later, stream counts are increased for certain platforms

| SOURCE      | Platform          | Volume SnapMirror | Qtree SnapMirror |
|-------------|-------------------|-------------------|------------------|
|             | FAS6080           | 150               | 128              |
|             | FAS6080_nearstore | 150               | 512              |
| DESTINATION | Platform          | Volume SnapMirror | Qtree SnapMirror |
|             | FAS6080           | 150               | 128              |
|             | FAS6080_nearstore | 300               | 512              |

© 2010 NetApp, Inc. All rights reserved.

## CONCURRENT TRANSFERS

Each storage system model supports a maximum number of simultaneous replication operations.

In Data ONTAP 7.3, the volume SnapMirror and qtree SnapMirror maximum stream counts are increased for certain platforms. This enhancement allows customers to accommodate the use of large numbers of flexible volumes and multiple concurrent transfers.

The new concurrent stream counts apply only to the `snapmirror initialize` and the `snapmirror update` Data ONTAP commands.

Refer to the latest *Data ONTAP Online Data Protection Backup and Recovery Guide* for new supported stream counts.



## Managing Transfers

- In Data ONTAP 7.3 and later, use the following options to manage volume SnapMirror transfers:

`replication.volume.reserved_transfers <n>`

- Guarantees that specified number of volume SnapMirror source/destination transfers always start
- Default is 0

`replication.volume.transfer_limits [current|previous]`

- Allows reversion to stream counts from versions of Data ONTAP earlier than 7.3
- Default is `current`

© 2010 NetApp, Inc. All rights reserved.

## MANAGING TRANSFERS

In Data ONTAP 7.3, you can specify the number of volume SnapMirror transfers for which you want resources by using the following Data ONTAP option:

`options replication.volume.reserved_transfers <n>`

`<n>` is the number of volume SnapMirror transfers for which you want resources reserved.

The default value is 0.

The reason for reserving transfers is that reserved resources will not be available for other replication types like qtrees SnapMirror or SnapVault transfers.

The stream count setting was increased in Data ONTAP 7.3. You can revert back to the previous stream count setting by using the following option:

`options replication.volume.transfer_limits [current | previous]`

If the value is set to `current`, the maximum for the current release will be used. If the value is set to `previous`, the maximum for Data ONTAP 7.2.0 will be used. The default value is `current`.



## Throttling Network

- Per transfer
  - Use the `kbs` argument in `snapmirror.conf`
- Dynamic throttle
  - Allows changing the throttle for a SnapMirror relationship while the transfer is active

```
snapmirror throttle <n> dst_hostname:dst_path
```
- System-wide throttle
  - Limits the total bandwidth for all transfers

```
options replication.throttle.enable on
options replication.throttle.incoming.max_kbs
options replication.throttle.outgoing.max_kbs
```

© 2010 NetApp, Inc. All rights reserved.

## THROTTLING NETWORK

Throttle network usage can be configured on a per transfer basis, using the `kbs` argument in the `snapmirror.conf`.

**Dynamic throttle** allows you to change the throttle value for a SnapMirror relationship while the transfer is active. This feature is available from Data ONTAP 7.1 and later.

```
snapmirror throttle <n> dst_hostname:dst_path
```

`<n>` is the new throttle value in kilobytes per second

**System-wide throttling** is available from Data ONTAP 7.2 and later and limits the total bandwidth used by all transfers at any time (SnapMirror and SnapVault transfers).

There are three options.

Enable or disable system-wide throttling on all systems: `replication.throttle.enable [on|off]`

Set maximum bandwidth for all incoming transfers: `replication.throttle.incoming.max_kbs <value>`

Set maximum bandwidth for all outgoing transfers: `replication.throttle.outgoing.max_kbs <value>`

The default value is unlimited, which means there is no limit on total bandwidth used. Valid transfer rate values are 1 to 125,000 kilobytes per second.



## Space Guarantee

- In Data ONTAP 7.3 and later, space is guaranteed for SnapMirror destination flexible volumes
  - Preallocates space in the aggregate for the volume
  - Volume guarantee is enabled by default
  - Guarantee is maintained even when the volume is made offline
- To enable or disable space guarantee on an existing SnapMirror destination volume:  
`vol options vol_name guarantee [volume|none]`

© 2010 NetApp, Inc. All rights reserved.

## SPACE GUARANTEE

In Data ONTAP 7.3 and later, space is guaranteed for SnapMirror destination flexible volumes. This new feature preallocates space in the aggregate for the volume. However, it is still possible for transfers to fail when the aggregate is full. When you create a flexible volume, by default its *volume* guarantee is enabled.

Support for space guarantee on the SnapMirror destination volume allows for maintenance of this guarantee after the initial baseline transfer is completed. The destination volume is set to an internal guarantee type called Replica (RAID label). The guarantee is also maintained when the volume is offline. When the SnapMirror destination storage system is upgraded to Data ONTAP 7.3 or later, you can enable or disable the space guarantee on the existing SnapMirror destination volumes using the `vol options` command. Note that file guarantee is not supported on SnapMirror destination volumes.

```
dst> vol options vol_name guarantee [volume | none]
```

An *upgraded\_replica* volume option is added to display if a volume is an upgraded SnapMirror destination as a part of the `vol status` command.

```
dst> vol status
Volume State Status Options
dst_vol online raid_dp, flex snapmirrored=on,
 snapmirroredupgraded_replica,
 read-only fs_size_fixed=on,
 guarantee=volume(disabled)

vol0 online raid4 root
```



## Synchronous SnapMirror

© 2010 NetApp, Inc. All rights reserved.

### SYNCHRONOUS SNAPMIRROR



## Sync Mode

- Volume SnapMirror Async is the base of SnapMirror Sync
  - Replicates Snapshot copies from a source volume to a secondary volume at the same time it is written to the source
- To configure Sync mode, replace the schedule field by `sync` in the `snapmirror.conf` file
  - `src:vol1 dst:vol1 - sync`
  - **NOTE:** If you edit the `snapmirror.conf` file while in `sync` mode, the relationship will drop to `async` and then will attempt to re-establish `sync` mode

© 2010 NetApp, Inc. All rights reserved.

## SYNC MODE

SnapMirror in synchronous mode is a mode of replication that sends updates from the source to the destination as they occur, rather than according to a predetermined schedule. This guarantees that data written on the source system is protected on the destination even if the entire source system fails.

Volume SnapMirror Async is the base of SnapMirror Sync. The first step involved in SnapMirror sync mode replication is a one-time baseline transfer of the source volume. When the baseline transfer is completed, SnapMirror transitions through a series of states, becoming more and more synchronous until the relationship gets in synchronous mode.

To configure SnapMirror Sync mode, the schedule field is replaced by `sync` in the `snapmirror.conf` configuration file.

Example:

```
SystemA:vol1 SystemB:vol1 - sync
```

Note that when changes are made to the `snapmirror.conf` file for a SnapMirror Sync entry, the SnapMirror relationship will go out of sync momentarily and then attempt to return to synchronous state.





## CP and NVLOG Forwarding

- SnapMirror Sync forwards consistency point (CP) to the destination to keep consistency between the source and the destination volumes
- SnapMirror Sync also forwards the NVRAM requests to the destination
  - Stored in NVLOG files  
`/etc/sync_snapmirror_nvlog/<dstfsid>.log[0|1]`
  - NVLOG files are replayed only in disaster recovery scenario

© 2010 NetApp, Inc. All rights reserved.

### CP AND NVLOG FORWARDING

**SnapMirror Sync forwards all consistency point (CP)** writes to the destination to keep consistency between the source and the destination volumes.

A CP is generally taken when the NVRAM is half-full, when the timer generates a CP (10 seconds), and when a Snapshot copy is created although other events can cause a CP.

Before Data ONTAP 7.2.2, the source CP will not complete until the destination had completed its CP. With Data ONTAP 7.2.2 and later, CPs are not synchronized but are forwarded to the destination and data is saved in memory.

**SnapMirror Sync forwards the NVRAM logs** to the destination to ensure that NVRAM operations on the source are replicated to the destination for replay in case of a disaster on the source system.

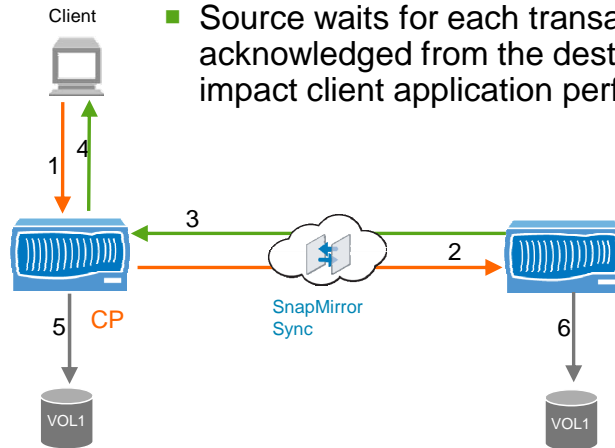
NVLOG data is treated as a stream of writes to a pair of special files named the NVLOG files  
`/etc/sync_snapmirror_nvlog/<dstfsid>.log[0|1]`

Before Data ONTAP 7.2.2, NVLOG files are written in the root volume of the destination system. With Data ONTAP 7.2.2 and later, NVLOG files are written in the parent aggregate of the destination volume.



## Theory of Operation

- Client writes are acknowledged after the writes have been logged to the NVRAM on the source and to the NVLOG files on the destination
- Source waits for each transaction to be acknowledged from the destination; it may impact client application performance



© 2010 NetApp, Inc. All rights reserved.

## THEORY OF OPERATION

Client writes are acknowledged after the writes have been logged to the NVRAM on the source and to the NVLOG files on the destination. Because the source waits for each transaction to be acknowledged from the destination before moving forward, it may impact client application performance.

The source system receives a write request from a client. The request is journaled in the system's NVRAM and recorded in cache memory.

The request and the NVLOG metadata are forwarded to the SnapMirror destination system where they are also journaled in NVRAM and cache memory.

The destination responds to the source system.

Data ONTAP acknowledges the write to the client system, and the application that requested the write is free to continue processing.

When a consistency point is triggered, Data ONTAP uses the transaction data in cache memory to build a list of data block changes that need to be written to disk.

This list of data blocks is sent to the destination, which initiates its own write to disk and returns an acknowledgment to the WAFL® software on the source system.



## Semi-Sync Mode

- Prior to Data ONTAP 7.3, Semi-Sync mode provided different synchronicity levels to control synchronicity versus performance
  - Use the `outstanding` argument in the SnapMirror configuration file to set the synchronicity level
  - `src:vol1 dst:vol1 outstanding=5s sync`

`outstanding={x ops | x ms | x s}`

| Value                   | Description                                                                                                           |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------|
| No value (default)      | SnapMirror operates in a fully synchronous manner                                                                     |
| <code>x ops</code>      | Allows <code>x</code> number of outstanding write operations before forcing the clients to wait for an acknowledgment |
| <code>x s   x ms</code> | Defines the amount of time (seconds or milliseconds) a destination waits before sending a write acknowledgment        |

© 2010 NetApp, Inc. All rights reserved.

## SEMI-SYNC MODE

Versions of Data ONTAP earlier than 7.3, allowed for Semi-Sync mode, which provided different synchronicity levels to control synchronicity versus performance.

The field in the configuration file that controls the synchronicity level is the `outstanding` argument. This variable allows you to modify the amount of time or the number of operations a destination waits before sending a write acknowledgment to the source.

`outstanding={x ops | x ms | x s}`

Default is no value: SnapMirror operates in a fully synchronous manner.

The ops suffix allows `x` number of outstanding write operations before forcing the clients to wait for an acknowledgment.

s or ms defines the amount of time (seconds or milliseconds) a destination waits before sending a write acknowledgment.

When the outstanding value is **less than 10 seconds**, the source forwards the NVLOG like it would in sync mode, but it does not wait for the acknowledgment from the destination system. This provides performance improvement for the client writing to the source volume. However, there is a possibility of data loss during that interval should the source crash prior to forwarding the NVLOG.

When the outstanding value is set to **10 seconds or more**, only the CP streams are forwarded to the destination; NVLOG data is not forwarded. Eliminating NVLOG forwarding reduces the load on the storage systems.



## Semi-Sync Mode (Cont.)

- In Data ONTAP 7.3 and later, the `outstanding` argument is deprecated
- To configure Semi-Sync mode, replace the `schedule` field by `semi-sync`  
`src:vol1 dst:vol1 - semi-sync`
- Only the CP streams are forwarded to the destination; NVLOG data is not forwarded
  - Saves processing power and bandwidth



© 2010 NetApp, Inc. All rights reserved.

## SEMI-SYNC MODE (CONT.)

In Data ONTAP 7.3, the `outstanding` argument is deprecated. To configure Semi-Sync mode, replace the `schedule` field by `Semi-Sync` in the `snapmirror.conf` file.

Example:

```
src:vol1 dst:vol1 - semi-sync
```

Only the CP streams are forwarded to the destination. NVLOG data is not forwarded. This mode minimizes performance impact on client applications, reduces storage systems' processing power, and saves bandwidth.



## Deployment Examples

Compare these entries from the destination `snapmirror.conf` file

- **Qtree SnapMirror**

```
src:/vol/vol1/q1 dst:/vol/vol1/q1 - 15 * * *
```

- **Volume SnapMirror**

```
src:vol2 dst:vol2 kbs=2000 10 8,20 * *
```

- **Sync**

```
src:vol3 dst:vol3 - sync
```

- **Semi-Sync**

```
src:vol4 dst:vol4 - semi-sync
```

© 2010 NetApp, Inc. All rights reserved.

## DEPLOYMENT EXAMPLES

### Example 1: Qtree SnapMirror

```
src:/vol/vol1/q1 dst:/vol/vol1/q1 - 15 * * *
```

The source qtree q1 is replicated to the destination qtree q1 every 15 minutes, every day of the week and every day of the month.

### Example 2: Volume SnapMirror

```
src:vol2 dst:vol2 kbs=2000 10 8,20 * *
```

The source volume vol2 is replicated to the destination volume vol2 at 10 minutes past 8:00 a.m. and 8:00 p.m. every day. In other words, the source volume vol2 is replicated daily at 8:10 a.m. and 8:10 p.m. Data ONTAP can use a maximum of 2,000 kilobytes per second to transfer data for this relation.

### Example 3: SnapMirror Sync

```
src:vol3 dst:vol3 - sync
```

The source volume vol3 is synchronously replicated to the destination volume vol3.

### Example 4: SnapMirror Semi-Sync

```
src:vol4 dst:vol4 - semi-sync
```

The source volume vol4 is replicated to the destination volume vol4 in Semi-Sync mode. Only the CP-Sync mechanism is used. NVLOG data is not forwarded to the destination.



## Performance Tuning

- Do not change the `visibility_interval` default value (3 minutes)
  - Snapshot copies are taken more often
- Before Data ONTAP 7.2.2, NVLOG files are written in the root volume
  - Impacts on how quickly NVLOG may be committed to disk
  - For traditional volume, ensure that the root volume spans enough disks
- After Data ONTAP 7.2.2, NVLOG files are written to destination volume's parent aggregate
- `crc32` checksums are more CPU-intensive than TCP checksums

© 2010 NetApp, Inc. All rights reserved.

## PERFORMANCE TUNING

In Sync mode, changes are shown on the destination only after the source takes a Snapshot copy of the source volume (every three minutes by default) and then transfers it to the destination. To control the view of the data on the destination, you use the `visibility_interval` argument in the `snapmirror.conf` file. If `visibility_interval` is set too low, the source system will be kept very busy creating Snapshot copies and this can impact performance. Changing the `visibility_interval` default value of three minutes is not recommended.

In versions of Data ONTAP earlier than 7.2.2, NVLOG data is written in the root volume. This may impact on how quickly NVLOG data may be committed to disk. For traditional volumes, ensure that the root volume spans enough disks. Note that with Data ONTAP 7.2.2 and later, the NVLOG files are written in the destination volume's parent aggregate.

The checksums algorithm is used to protect SnapMirror transmitted data. Cyclic redundancy check checksums, also known as `crc32c`, are computed by the CPU on the destination storage system and may have undesired effects on performance. TCP checksums are computed directly on the network interface card, or NIC, and are less CPU-intensive; therefore the TCP checksums computation is the recommended method.



## Requirements and Limitations

- Supports only volume replication
- Supports bidirectional sync relationships from Data ONTAP 7.2.1 and later
- Storage system platforms must be identical and run the same major Data ONTAP release
- One source cannot have sync relationships to multiple destinations
- Cascading sync relationships are not supported
- Source and destination cannot be on the same storage system
- A high-availability configuration cannot have SnapMirror Sync relationships from one half of the configuration to the other half of the configuration

© 2010 NetApp, Inc. All rights reserved.

## REQUIREMENTS AND LIMITATIONS

SnapMirror Sync can only be used on volumes, not qtrees.

SnapMirror Sync follows the same volume type matrix as volume SnapMirror. Replication must be from traditional volume to traditional volume or flexible volume to flexible volume.

Bidirectional SnapMirror Sync is supported on storage systems using Data ONTAP 7.2.1 and later.

Replications are allowed only between identical storage system platforms running the same major Data ONTAP release.

One source system cannot have SnapMirror Sync relationships to multiple destination systems.

Cascading sync relationships is not supported.

The source and destination of the SnapMirror Sync relationship cannot be on the same storage system, such as: SystemA:src\_vol->SystemA:dst\_vol

A high availability configuration cannot have SnapMirror Sync relationships from one half of the configuration to the other half of the configuration.



## Advanced Features

© 2010 NetApp, Inc. All rights reserved.

### ADVANCED FEATURES





## SnapMirror over Multiple Paths

- Allows Fibre Channel and/or Ethernet as transport
- Supports Async and Sync modes
- One or two paths allowed
  - Two FC NIC adapters per storage system
  - Two Ethernet NIC adapters per storage system
  - One of each per storage system
- Multiplexing: both paths are used at the same time for load balancing
- Failover: first path specified is active, the second path is in standby and becomes active if the first path fails

© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR OVER MULTIPLE PATHS

SnapMirror supports up to two paths for a particular SnapMirror relationship. The paths can be Ethernet, Fibre Channel, or a combination of Ethernet and Fibre Channel.

Multiple paths are supported by SnapMirror Async and Sync replication modes.

The two paths can be used in one of two modes:


**Multiplexing mode:** SnapMirror uses both paths at the same time, essentially load balancing the transfers. If one path fails, the transfers occur on the remaining path. After the failed path is repaired, the transfers resume using both paths.

**Failover mode:** SnapMirror uses the first specified path as the desired path and uses the second specified path only after the first path fails.



## Configuring Multiple Paths

- Add a connection name line in the `snapmirror.conf` file
  - Define connection mode and network interfaces
- Edit the schedule entry to reflect the new connection name as the source system



```
/etc/snapmirror.conf
FAS1_conf = multi (FAS1-e0a,FAS2-e0a) (FAS1-e0b,FAS2-e0b)
FAS1_conf:vol1 FAS2:vol1 - sync
```

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURING MULTIPLE PATHS

To implement multiple paths between the source and destination storage system, edit the `snapmirror.conf` file to add a connection name line that defines the mode of the connection and what the two connections are. Then, edit the schedule entry to reflect the new connection name as the source system.

In this following illustration, the source volume `vol1` on the storage system FAS1 is synchronously replicated to the destination volume `vol1` on the storage system FAS2. Two gigabit Ethernet paths are configured and replication occurs using both connections in multiplexing mode as specified in the `snapmirror.conf` file.

```
FAS1_conf = multi (FAS1-e0a,FAS2-e0a) (FAS1-e0b,FAS2-e0b)
```

The first entry defines the connection name (`FAS1_conf`), the mode of the connection (`multi`) and what the two connections are (`FAS1-e0a` connected to `FAS2-e0a` and `FAS1-e0b` connected to `FAS2-e0b`).


```
FAS1_conf:vol1 FAS2:vol1 - sync
```

The second entry defines the sync relationship for volume `vol1`.



## SnapMirror over Fibre Channel

- Source and destination must have NetApp X1024 or X1124 FC HBA (IP over FC)
- At least one FC switch between the source and the destination
  - Switches must be from the same vendor
  - Supported vendors: Cisco® or Brocade®
- SnapMirror traffic must use dedicated zones



© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR OVER FIBRE CHANNEL

SnapMirror over Fibre Channel enables you to use the SnapMirror Async and Sync features over a Fibre Channel SAN environment.

SnapMirror over FC requires at least one Fibre Channel switch in the data path. The supported switch vendors are Cisco or Brocade. To comply with SnapMirror over Fibre Channel certification, use only switches from one vendor in the SnapMirror data path. Supported switches and firmware versions are specified in *Requirements for SnapMirror Over Fibre Channel Transport (Asynchronous, Synchronous, and Semi-synchronous modes)* on the NOW site.

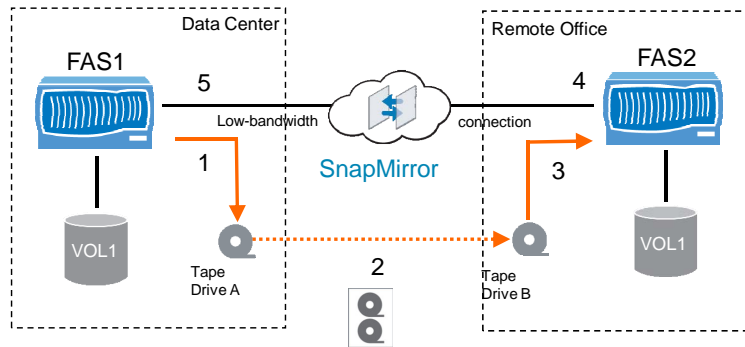
The storage system must be configured with NetApp X1024 FC HBA (two 2-GB ports) or X1124 (two 4-GB ports) for IP over Fibre Channel functionality. The adapters translate the SnapMirror IP packets to and from FC frames and supports the multiple path configurations.

To begin the SnapMirror over FC configuration process, first configure the FC NIC adapters and connect the systems to the Fibre Channel switches. Optionally, you can set up multiple SnapMirror traffic zones. SnapMirror traffic must be configured within dedicated zones. After the switches have been set up, configure SnapMirror and baseline the initial transfers. For details on SnapMirror over Fibre Channel configuration steps, refer to the latest *Data ONTAP Data Protection Online Backup and Recovery Guide* on the NOW site.



## SnapMirror to Tape

- Accommodates the initial transfer over low-bandwidth connections using a physically transported tape
- Incremental updates are performed over the network
- SnapMirror to tape supports volume replication only



© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR TO TAPE

SnapMirror to tape is a deployment that supports SnapMirror replication over low-bandwidth connections by accommodating the initial transfer between the source and the destination systems using a physically transported tape. When baseline transfer has been carried out with the local tape device, incremental SnapMirror updates can be performed over the network. The SnapMirror-to-tape function is available for volume replication only.

On the source system, use the `smtape backup` command to copy all volume Snapshot copies, including the base Snapshot copy, to tape. If more than one backup tape is necessary, use the `smtape continue` command to continue the copying on a new tape.

Physically transport the backup tapes from the source system to the destination system.

On the destination system, use the `vol create` and `vol restrict` commands to set up a SnapMirror target volume.

Use the `smtape restore` command to copy the initial SnapMirror tape to the destination system. If the backup spans more than one tape, use the `smtape continue` command to continue the restore on a subsequent tape.

Use the `snapmirror update` command to trigger an incremental update from the source to the destination system over the low-bandwidth connection, or edit the `snapmirror.conf` file to set up an incremental update schedule from the source to the destination volume.

Finally, use the `snapmirror release` command to eliminate the source-to-tape relationship and associated Snapshot copy.



## Converting a Replica to a Writable

- To convert a replica to a writable file system, break the SnapMirror relationship

```
dst> snapmirror break dst_vol
```

- To resume the replication operations, resynchronize the broken off relationship

```
src> snapmirror resync dst_hostname:dst_vol
```

- To make the break permanent, release the relationship from its direct source

```
src> snapmirror release src_vol
dst_hostname:dst_vol
```

© 2010 NetApp, Inc. All rights reserved.

## CONVERTING A REPLICA TO A WRITABLE

You might want to convert a read-only replica to a writable qtree or volume to migrate data to a new location or in case of a disaster, when the source becomes unavailable and you wish to redirect a CIFS or NFS client's access to the destination.

To convert a replica to a read/write volume or qtree, use the `snapmirror break` command:

```
snapmirror break dst_vol
```

Note that to convert a qtree replica to be writable, you must first quiesce the destination qtree:

```
snapmirror quiesce /vol/dst_vol/dst_qtree
snapmirror break /vol/dst_vol/dst_qtree
```

After breaking a SnapMirror relationship, to resume incremental updates, use the `snapmirror resync` command. To avoid losing data, always resync from the storage system that has the *less* up-to-date file system: `snapmirror resync dst_hostname:dst_vol`

SnapMirror identifies the newest common Snapshot copy created for the last successful update. This Snapshot copy will be used as the basis for resynchronization.

To make the break permanent, release the volume or the qtree relationship from its immediate source:

```
snapmirror release.
```




## Reestablishing the Relationship

- To resume the replication operations, resynchronize the broken-off relationship

NOTE: Changes stored on the destination during the break will be lost upon resynchronization

```
dst> snapmirror resync src_hostname:src_vol
```



© 2010 NetApp, Inc. All rights reserved.

## REESTABLISHING THE BROKEN RELATIONSHIP

You can use the `snapmirror resync` command to restore or redefine a SnapMirror source or destination relationship that was broken with the `snapmirror break` command.

**Applied to the original destination**—the `snapmirror resync` command will put a volume or qtree back into a SnapMirror relationship and resynchronize its contents with the source without repeating the initial transfer.

**Applied to the source volume**—the `snapmirror resync` command can turn the source volume into a copy of the original destination volume. In this way, the roles of source and destination can be reversed.




## Reestablishing the Relationship (Cont.)

- Resynchronizing from the original source reverses the roles of the source volume and destination volume

```
src> snapmirror resync dst_hostname:dst_vol
```

SnapMirror Source  
is now  
the new destination

SnapMirror Destination  
is now  
the new source



© 2010 NetApp, Inc. All rights reserved.

## REESTABLISHING THE RELATIONSHIP (CONT.)

You can use the `snapmirror resync` command to restore or redefine a SnapMirror source or destination relationship that was broken with the `snapmirror break` command.

**Applied to the original destination**—the `snapmirror resync` command will put a volume or qtree back into a SnapMirror relationship and resynchronize its contents with the source without repeating the initial transfer.

**Applied to the source volume**—the `snapmirror resync` command can turn the source volume into a copy of the original destination volume. In this way, the roles of source and destination can be reversed.

## CONSIDERATIONS

You might want to resynchronize a source and a destination volume or qtree when:

You are changing the current source to a different volume or qtree.

You make a destination volume writable for application testing, and then want to make it a SnapMirror destination again.

You need to recover from a disaster that disabled the source.


You want to reverse the functions of the source and the destination.



## Making the Break Permanent

- To make the break permanent, release the relationship from its direct source

```
src> snapmirror release src_vol
dst_hostname:dst_vol
```



© 2010 NetApp, Inc. All rights reserved.

## REESTABLISHING THE BREAK PERMANENT

To permanently end a SnapMirror relationship between a source and destination pair of volumes or qtrees, you need to use different commands on the source and destination storage systems.

### Considerations

**Source system**—Use the `snapmirror release` command. Releasing a source from a destination volume or qtree allows the source to delete its base Snapshot copy for the SnapMirror relationship.

**Destination system**—Use the `snapmirror break` command. After breaking the relationship, you need to scrub the destination with additional steps. Unless these extra steps are taken, the Snapshot copies associated with the broken relationship remain stored on the destination system, and a `snapmirror status` command will continue to list the former destination object as a current destination object.





## Migrating SnapMirror Volumes

- Use the `snapmirror migrate` command to migrate data between volumes that are in a SnapMirror relationship
- SnapMirror migration:
  - Performs a SnapMirror incremental transfer to the destination volume
  - Stops NFS and CIFS services to the source volume
  - Migrates NFS file handles to the destination volume
  - Makes the source volume restricted
  - Makes the destination volume read-write

© 2010 NetApp, Inc. All rights reserved.

### MIGRATING SNAPMIRROR VOLUMES

SnapMirror can migrate data between volumes and redirect NFS clients to the new volume without rebooting the storage system or remounting the volume on NFS clients. The migration must be run on two volumes that are currently the source volume and destination volume in a SnapMirror relationship.

You use the `snapmirror migrate` command on the storage system, which holds the source volume.

```
snapmirror migrate src_hostname:src_volume
 dst_hostname:dst_volume
```

The SnapMirror migration process does the following:

Performs a SnapMirror incremental transfer to the destination volume

Stops NFS and CIFS services to the source volume

Migrates NFS file handles to the destination volume

Makes the source volume restricted

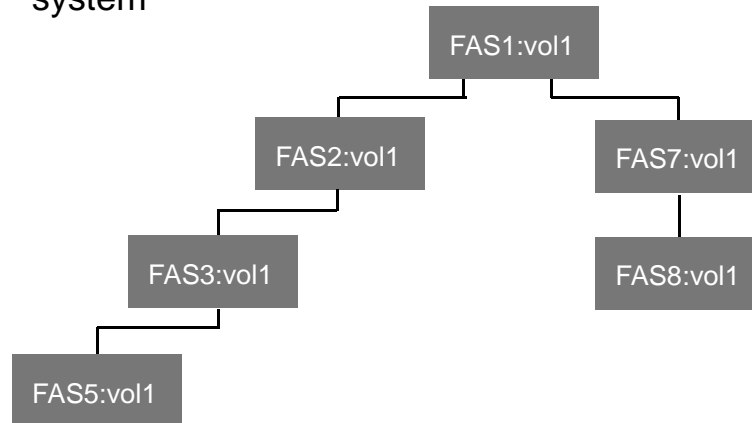
Makes the destination volume read-write

Note that SnapMirror does not transfer IP addresses, license keys, or quota information. You must remount on the NFS clients. SnapMirror does not migrate CIFS clients. You must reestablish CIFS client sessions after migrating data to the destination volume.



## Cascading SnapMirror Volumes

- SnapMirror creates and retains the Snapshot copies on the original source volume
- The SnapMirror Snapshot copies are cascaded down the line to replicate the volumes on each destination system



© 2010 NetApp, Inc. All rights reserved.

### CASCADING SNAPMIRROR VOLUMES

Instead of propagating data from one central master site to many destinations, which would require expensive network connections and excessive CPU time, you can propagate data from one volume to another volume and from that one to the next, in a series.

In a volumes cascade, SnapMirror creates and retains the Snapshot copies on the original source volume. The SnapMirror Snapshot copies are cascaded down the line to be able to replicate the volumes on each destination system.

In this illustration, volume vol1 on storage system FAS1 is replicated to seven storage systems. To set up cascading volumes on each storage system as shown in the diagram, the snapmirror.conf file would look like this:






















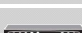





```
FAS1:vol1 FAS7:vol1 - 15 * * 1,2,3,4,5
FAS1:vol1 FAS2:vol1 - 15 * * 1,2,3,4,5
FAS7:vol1 FAS8:vol1 - 25 * * 1,2,3,4,5
FAS2:vol1 FAS3:vol1 - 35 * * 1,2,3,4,5
FAS2:vol1 FAS4:vol1 - 35 * * 1,2,3,4,5
FAS3:vol1 FAS5:vol1 - 45 * * 1,2,3,4,5
FAS3:vol1 FAS6:vol1 - 45 * * 1,2,3,4,5
```

To remove a destination from the cascade, use the `snapmirror release` command from the immediate source. SnapMirror will delete the Snapshot copies associated with that destination.



## Cascade Support Matrix

- SnapMirror Sync is allowed only on the source
- Qtree SnapMirror cannot cascade more than one hop

| Configuration of SnapMirror Relationships                                                                                                                                                                                                                                     | Support |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|
|  → SM Sync →  → VSM →      | ✓       |
|  → SM Sync →  → SM Sync →  | ✗       |
|  → SM Sync →  → QSM →      | ✗       |
|  → VSM →  → VSM →          | ✓       |
|  → VSM →  → SM Sync →      | ✗       |
|  → VSM →  → QSM →          | ✓       |
|  → QSM →  → VSM →          | ✓       |
|  → QSM →  → SM Sync →      | ✗       |
|  → QSM →  → QSM →          | ✗       |

© 2010 NetApp, Inc. All rights reserved.

## CASCADE SUPPORT MATRIX

Not all cascading configurations are supported. The following limitations apply to both traditional and flexible volumes as of Data ONTAP 7.2.2 and later.

This table can be reduced to two simple rules:

1. Synchronous SnapMirror is allowed only on the source.
2. Qtree SnapMirror cannot cascade more than one hop.



## SnapMirror Interactions

© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR INTERACTIONS



## SnapRestore and SnapMirror

- You can revert a source volume but not a destination volume
- You can revert to any Snapshot copy, including SnapMirror Snapshot copies
- DO NOT select a Snapshot copy taken before a SnapMirror Snapshot copy
  - Incremental update will fail
  - The relationship must be reinitialized

© 2010 NetApp, Inc. All rights reserved.

## SNAPRESTORE AND SNAPMIRROR



## SnapVault and SnapMirror

- The SnapVault and SnapMirror bundle provides a consolidated data protection and disaster recovery solution
  - Fast and space-efficient disk-based backup
  - Recover from system disaster by making any online backup copy writable
  - Rapid disaster recovery from hourly, nightly, or weekly backup
  - Offsite disaster recovery by replicating backups to remote sites

© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT AND SNAPMIRROR



## SnapVault Versus SnapMirror

|                   | SnapMirror                                  | SnapVault                                                        |
|-------------------|---------------------------------------------|------------------------------------------------------------------|
| Deployment        | Data migration and replication              | Data archiving                                                   |
| Disaster Recovery | Integrated failover capability              | SnapMirror and SnapVault bundle is required for failover         |
| Cascading         | Supports cascading volume in a series       | Cascading SnapVault qtrees replica is not supported              |
| Snapshot Copy     | Does not create Snapshot copy for archiving | Provides Snapshot copy scheduling and retention on the secondary |
| Update Frequency  | Up to per-minute updates                    | Up to per-hour updates                                           |

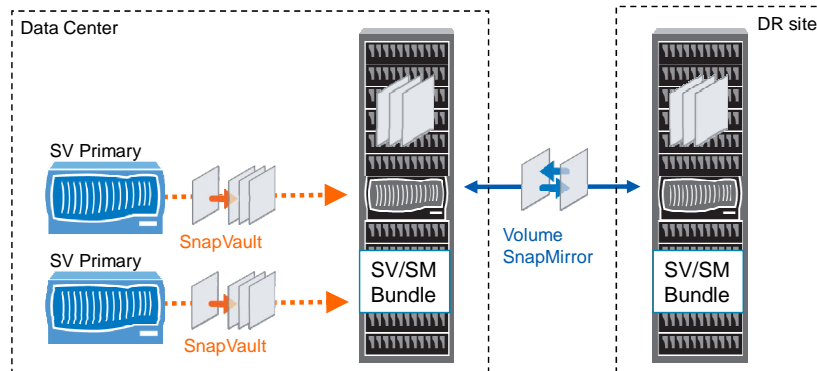
© 2010 NetApp, Inc. All rights reserved.

## SNAPVAULT VERSUS SNAPMIRROR



## Protecting SnapVault with SnapMirror

- This solution consists of replicating SnapVault based backup copies through SnapMirror to a disaster-recovery site to provide:
  - Backup and standby service for SnapVault
  - Backup and restore protection for SnapVault



© 2010 NetApp, Inc. All rights reserved.

## PROTECTING SNAPVAULT WITH SNAPMIRROR





## Considerations

- If a new SnapVault update is triggered before the volume SnapMirror transfer has completed, then the ongoing transfer will abort
- Transfers fail when the storage system reaches the maximum simultaneous transfers supported
- In both cases, ensure that SnapMirror and SnapVault schedules do not overlap and do not stretch

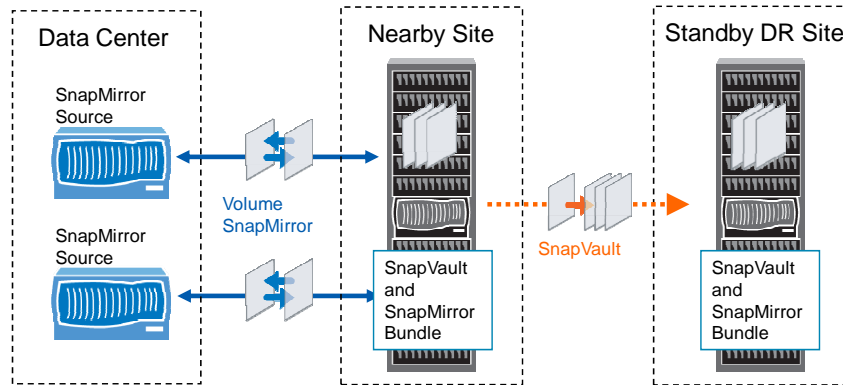
© 2010 NetApp, Inc. All rights reserved.

## CONSIDERATIONS



## SnapMirror and SnapVault for DR

- In this solution, data at the production center is replicated at the volume level to a SnapMirror destination system
- Then the SnapMirror replicas are protected at the qtree level to a tertiary appliance at the DR site using SnapVault



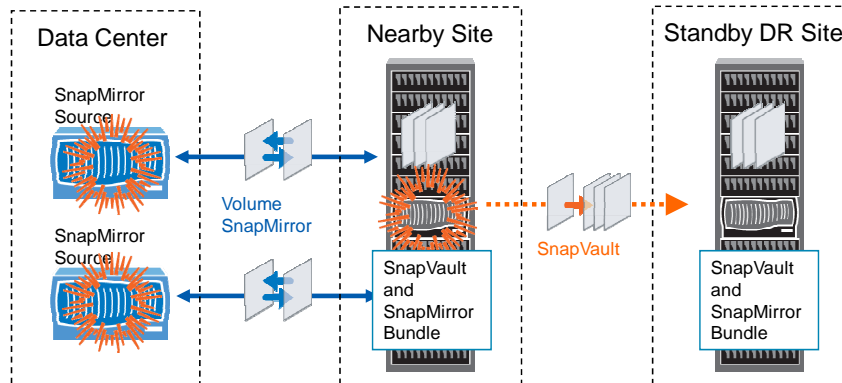
© 2010 NetApp, Inc. All rights reserved.

## SNAPMIRROR AND SNAPVAULT FOR DR



## Recover from Site and Regional Disasters

- In the event of a site disaster, you can fail over SnapMirror to the nearby destination system and fail back without requiring a complete data transfer
- In the event of a regional disaster, you can fail over SnapVault to the tertiary device and then resume vault operations after recovery



© 2010 NetApp, Inc. All rights reserved.

## RECOVER FROM SITE AND REGIONAL DISASTERS



## SnapMirror Performance

© 2010 NetApp, Inc. All rights reserved.

### SNAPMIRROR PERFORMANCE



## Volume SnapMirror Performance

- With SnapMirror Async, performance tends to be centered on:
  - Update frequency
  - Having enough network bandwidth
  - Storage system utilization
- Volume SnapMirror Async performance is particularly affected by:
  - Volume size and rate of data changed
  - Storage system utilization

© 2010 NetApp, Inc. All rights reserved.

### VOLUME SNAPMIRROR PERFORMANCE

Volume SnapMirror performance is centered on the update frequency, the network bandwidth, and the storage system utilization. Volume SnapMirror Async performance is particularly affected by the volume size, the rate of data changed, and the disk geometry for traditional volumes.

#### DISK GEOMETRY

For versions of Data ONTAP earlier than 7.0 and traditional volumes, it is recommended that the source and destination volumes contain disks of the same size, and be organized in the same RAID group configuration to gain optimal performance. For flexible volumes, disk geometry matching is no longer a consideration.

#### SNAPSHOT COPY CREATION AND UPDATE FREQUENCY

SnapMirror creates a Snapshot copy before every update and deletes a Snapshot copy at the end. On heavily loaded storage systems, Snapshot copy creation time can stretch out and restricts the frequency of SnapMirror updates. Stretched SnapMirror schedules result in SnapMirror creating many Snapshot copies on the source storage system at the same time, which can impact client access. For this reason staggered SnapMirror schedules are recommended to avoid system blockages.


#### VOLUME SIZE AND CHANGED BLOCKS

To perform an incremental update, the block map in the new Snapshot copy is compared to the block map in the baseline Snapshot copy. The time required to determine the block changes depends on the volume size. With Data ONTAP 7.0 and later, you can use the `snap delta` command to determine the rate of data change between Snapshot copies on a volume.



## Qtree SnapMirror Performance

- Qtree SnapMirror performance is impacted by:
  - Directory structure
  - Large number (tens of millions) of small files
  - Transfer size



© 2010 NetApp, Inc. All rights reserved.

## QTREE SNAPMIRROR PERFORMANCE

Qtree SnapMirror performance is impacted by deep directory structure and large numbers, such as tens of millions, of small files replicated.

### DIRECTORY STRUCTURES AND LARGE NUMBERS OF SMALL FILES

To determine changed data, qtree SnapMirror looks at the inode file and defines which inodes are in the qtree of interest and which inodes have changed. If the inode file is large, but the inodes of interest are few, qtree SnapMirror spends a lot of time going through the inode file to find very few changes. Disk I/Os used to access the data become small and inefficient.

### TRANSFER SIZE

When a qtree SnapMirror update is transferring, the `snapmirror status -l` command shows how many kilobytes have been transferred so far; the value may be greater than the expected delta (changes expected). This overhead is due to metadata transfer, for example: 4-KB header, file creation, deletion, ACLs, and so on.

When the update has completed, you can use the Data ONTAP `df` command for the destination volume to verify that the expected change size is correct.



## Concurrent Transfer Limitation

- Updates fail when the system exceeds the maximum simultaneous replication operations it supports
  - Each transfer beyond the limit will reattempt to run once per minute
- To optimize:
  - Stagger update schedules
  - For qtree SnapMirror, if there are too many qtrees per destination volume, rebaseline those qtrees to another volume

© 2010 NetApp, Inc. All rights reserved.

### CONCURRENT TRANSFER LIMITATION

The transfer fails when the system reaches the maximum number of simultaneous replication operations. Each transfer beyond the limit will reattempt to run once per minute.

To optimize SnapMirror deployment, it is recommended that the schedules be staggered. For qtree SnapMirror, if there are too many qtrees per destination volume, the solution is to re-baseline those qtrees to another volume.



## CPU Utilization

- SnapMirror consumes available CPU cycles on a storage system
  - 100% CPU utilization does not mean that performance is degraded
  - SnapMirror may have some impact, but in the majority of cases, it is NOT very significant
- Monitor storage system CPU using
  - Operations Manager Performance Advisor
  - Data ONTAP `sysstat` command

© 2010 NetApp, Inc. All rights reserved.

### CPU UTILIZATION

SnapMirror consumes available CPU cycles on a storage system.

When the source storage system shows that the CPU utilization is up to 100%, it does not mean that the system performance or the SnapMirror throughput is degraded.

SnapMirror may have some impact, but in the majority of cases, it is not very significant.

You can monitor storage system CPU using Operations Manager Performance Advisor or the Data ONTAP `sysstat` command.





## System Activities

- On heavily loaded systems, SnapMirror competes with other processes and may impact response times
- To address this problem consider these alternatives:
  - Use FlexShare® software
  - Schedule SnapMirror updates at times when NFS or CIFS traffic is low
  - Reduce update frequency
  - Upgrade to a more powerful NetApp controller

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM ACTIVITIES

On heavily loaded systems, SnapMirror competes with other processes and may impact response times.

To address this problem you can set the system priority to High or Very High on dedicated storage systems for SnapMirror replication using FlexShare® software.

You can also schedule SnapMirror updates at times when NFS or CIFS traffic is low and reduce the frequency of updates.

Finally, consider upgrading to a more powerful NetApp controller when the system resources become the system blockage.



## Network Distance and Bandwidth

- Network distance causes write latency
- To address network issues:
  - Limit the bandwidth using network throttle features
  - Utilize a dedicated network for SnapMirror
  - Use multipath for load balancing/failover
  - Look for typical network problems
    - For example, duplex mismatches

© 2010 NetApp, Inc. All rights reserved.

### NETWORK DISTANCE AND BANDWIDTH

When deploying SnapMirror, you have to consider the round-trip travel time of a packet from the source to the destination storage system, because network distance causes write latency. The round trip has a latency of approximately 2 milliseconds if the source and the destination storage systems are 100 miles apart.

Networking issues impacting SnapMirror performance can be addressed by limiting the bandwidth using the system-wide or per-transfer network throttle features.

Networking issues can also be addressed by using a dedicated path for SnapMirror transfers or using multiple paths for load balancing and failover.

If the network still does not perform up to expectations, look for typical network problems. For example, duplex mismatches can cause networks to be very slow.



## Documents and References

- *Data ONTAP Data Protection Online Backup and Recovery Guide*
- *SnapMirror Async Overview and Best Practices Guide*
  - <http://media.netapp.com/documents/tr-3446.pdf>
- *SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations*
  - <http://media.netapp.com/documents/tr-3326.pdf>

© 2010 NetApp, Inc. All rights reserved.

## DOCUMENTS AND REFERENCES

You can obtain additional information about SnapVault and related technologies from the following:

### MANUAL

*Data ONTAP Data Protection Online Backup and Recovery Guide*

### TECHNICAL REPORTS

TR-3446: *SnapMirror Async Overview and Best Practices Guide*

TR-3326: *SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations*



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Explain the SnapMirror Async, Sync, and Semi-Sync modes of operation
- Describe how volume SnapMirror and qtree SnapMirror replicate data
- Configure SnapMirror
- Perform advanced SnapMirror operations
- Explain SnapMirror performance impact

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 23: SnapMirror  
Estimated Time: 45 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What are the main differences between volume and qtree SnapMirror?
- True or false? When the `visibility_interval` is reached it causes the source system to create a Snapshot copy.
- True or false? Ethernet and FC interfaces may be logically combined in multi-mode or failover mode.

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING



Go further, faster®

# Performance

Module 24  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## PERFORMANCE





## Module Objectives

By the end of this module, you should be able to:

- Use Data ONTAP® tools to identify networking, disk I/O, FC loop saturation, and CPU bottlenecks using `systat`, `stats`, and `perfstat`
- Discuss how increasing utilization can affect performance
- Use the `reallocate` command to maintain performance
- Use recommended techniques to optimize Data ONTAP configuration for SAN and NAS

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



**sysstat**

© 2010 NetApp, Inc. All rights reserved.

**SYSSTAT**



## Write Performance: sysstat Command

```
system> sysstat -c 10 -s 5
```

| CPU | NFS | CIFS | HTTP | Net kB/s |     | Disk kB/s |       | Tape kB/s |       | Cache age |
|-----|-----|------|------|----------|-----|-----------|-------|-----------|-------|-----------|
|     |     |      |      | in       | out | read      | write | read      | write |           |
| 2%  | 0   | 0    | 0    | 0        | 0   | 9         | 23    | 0         | 0     | >60       |
| 0%  | 0   | 0    | 0    | 0        | 0   | 0         | 0     | 0         | 0     | >60       |
| 5%  | 0   | 0    | 0    | 0        | 0   | 21        | 27    | 0         | 0     | >60       |
| 1%  | 0   | 0    | 0    | 0        | 0   | 0         | 0     | 0         | 0     | >60       |
| 5%  | 0   | 0    | 0    | 0        | 0   | 20        | 28    | 0         | 0     | >60       |
| 1%  | 0   | 0    | 0    | 0        | 0   | 0         | 0     | 0         | 0     | >60       |
| 4%  | 0   | 0    | 0    | 0        | 0   | 21        | 26    | 0         | 0     | >60       |
| 1%  | 0   | 0    | 0    | 0        | 0   | 0         | 0     | 0         | 0     | >60       |
| 5%  | 0   | 0    | 0    | 0        | 0   | 22        | 27    | 0         | 0     | >60       |
| 0%  | 0   | 0    | 0    | 0        | 0   | 0         | 0     | 0         | 0     | >60       |

```
--
Summary Statistics (10 samples 5.0 secs/sample)
CPU NFS CIFS HTTP Net kB/s Disk kB/s Tape kB/s Cache
in out read write read write age
Min
0% 0 0 0 0 0 0 0 0 0 0 >60
Avg
2% 0 0 0 0 0 9 13 0 0 >60
Max
5% 0 0 0 0 0 22 28 0 0 >60
```

© 2010 NetApp, Inc. All rights reserved.

## WRITE PERFORMANCE: SYSSTAT COMMAND

The best command for viewing system utilization is `sysstat [interval]`, where `interval` is the incremental interval in seconds (the default is every 15 seconds). The `sysstat` command is a little like a speedometer for your storage system, allowing you to view real-time activity per second.

The statistics displayed by the `sysstat` command should help you answer questions such as:

- Is the system usage steady or does it fluctuate?
- Is the CPU percentage high without corresponding input/output activity?

## INTERPRETING SYSSTAT RESULTS

The following is a description of `sysstat` command results:

- CPU—Displays an average of the busiest CPUs. **NOTE:** The `sysstat -m` command displays statistics for each CPU in a multiprocessor system.
- NFS—Number of NFS operations per second
- CIFS—Number of CIFS operations per second
- HTTP—Number of HTTP operations per second
- Net KBps in and out—The kilobytes per second of data requested from the network as a read or write
  - This is the network traffic displayed in KBps, which tells you how much network traffic the storage appliance is handling, how constant that traffic is, and if the system is exceeding its network traffic limitations.
- Disk KBps reads and writes—Shows disk activity
  - Disk reads occur if data is not cached. Disk writes should occur ideally every 10 seconds.
- Cache age—Displays the age, in minutes, of the oldest read-only blocks in the buffer cache (not information relevant to diagnosing performance).



**stats**


© 2010 NetApp, Inc. All rights reserved.

## STATS



## Performance Counters

- Counters are organized in an object-instance-counter hierarchy
  - Counters are collected from the Counter Manager
- The `stats` command allows users to look at any object-instance and the corresponding counter
  - Support for preset files

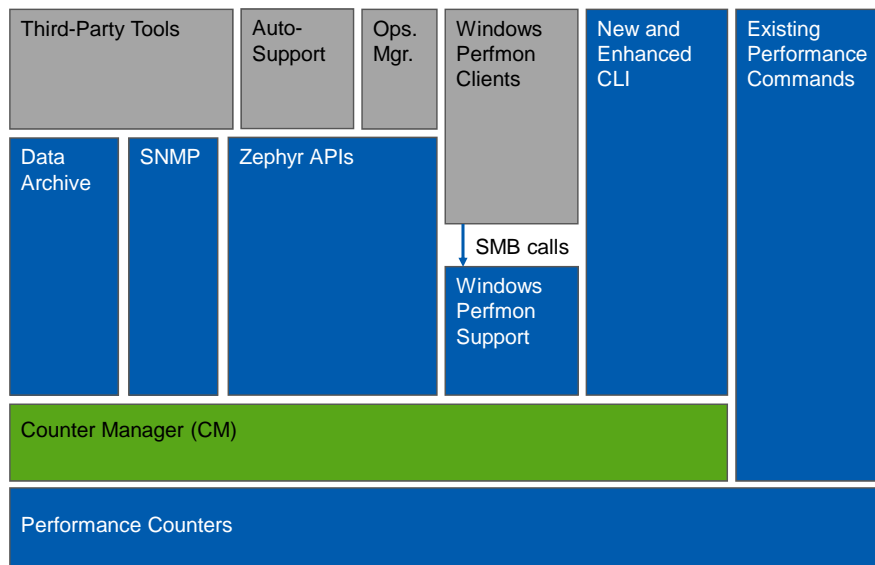


© 2010 NetApp, Inc. All rights reserved.

## PERFORMANCE COUNTERS



## Counter Manager (Review)



© 2010 NetApp, Inc. All rights reserved.

### COUNTER MANAGER (REVIEW)

Counter Manager is a thin layer built into the Data ONTAP architecture that provides a single view of Data ONTAP performance counters and a standard performance API set for all clients. Clients include ZAPI, AutoSupport, Windows® perfmon, SNMP, and the command-line interface.

#### Motivation

Counter Manager was engineered into the architecture of Data ONTAP 6.5 to create a complete set of performance metrics that can supply you with statistics for analysis of configuration mistakes.

Counter Manager provides an infrastructure to:

- Improve customer and internal performance monitoring
- Provide simple performance problem diagnosis
- Enhance existing sizing processes
- Provide capacity planning capabilities

For a complete list of the performance counters available in Data ONTAP, look in the Operations Manager documentation for Performance Objects and Counters.



## stats: Command Syntax

- The `stats` command is a way to collect or view statistical data on a storage appliance.
- The `stats` command may be run in one of three ways:
  - Single, in which current counter values are displayed:
    - `stats show`
  - Repeating, in which counter values are displayed multiple times at a fixed interval:
    - `stats show -i 1`
  - Period, in which counters are gathered over a single period of time and then displayed:
    - `stats start then stats stop`

© 2010 NetApp, Inc. All rights reserved.

### STATS: COMMAND SYNTAX

`stats` Command

`stats list objects`

`stats list instances [ object_name ]`

`stats list counters [-p preset][ object_name ]`

`stats explain counters [ object_name][ counter_name ]`

`stats show [ -n num ][ -i interval ][ -o path ] [ -I identifier ][ -d delimiter ][ -p preset ][ -r | -c ][ object_def { object_def } ]`

`stats start [-p preset][-I identifier][{object_def}]`

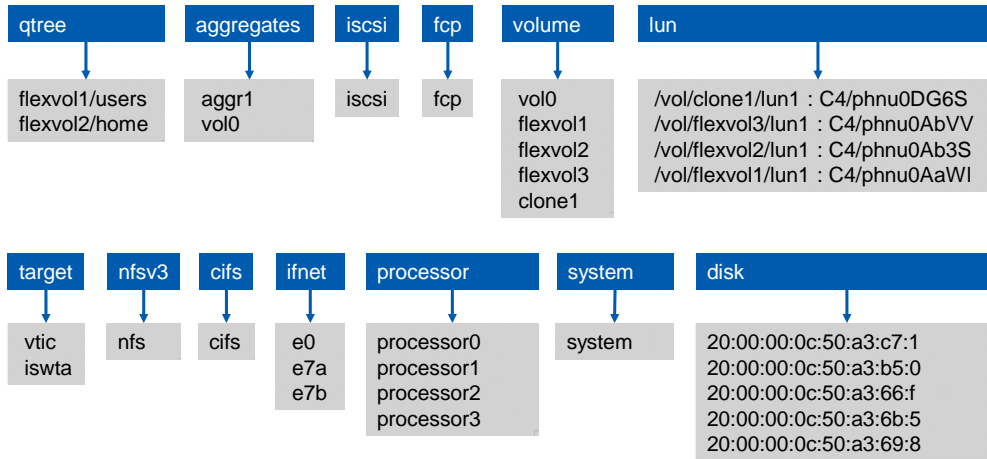
`stats stop [-p preset][-I identifier][-r] [-c] [-o path_name]`



## stats: Command Example

```
system> stats list objects
```

```
system> stats list instances
```



© 2010 NetApp, Inc. All rights reserved.

## STATS: COMMAND EXAMPLE





# stats Command

```
system> stats list counters qtree
```

```
system> stats explain counters
qtree nfs_ops
```

qtree

flexvol1/users  
flexvol2/home

nfs\_ops  
cifs\_ops

Counters for object name: qtree  
Name: nfs\_ops  
Description: Number of NFS  
operations per second to the qtree  
Properties: rate  
Unit: per\_sec

```
system> stats list counters volume
```

```
system> stats explain counters volume
write_ops
```

volume

vol0  
flexvol1  
clone1

total\_ops  
avg\_latency  
read\_ops  
read\_data  
read\_latency  
write\_ops  
write\_data  
write\_latency  
other\_ops  
other\_latency

Counters for object name:  
volume  
Name: write\_ops  
Description: Number of writes  
per second to the volume  
Properties: rate  
Unit: per\_sec

© 2010 NetApp, Inc. All rights reserved.

## STATS COMMAND

A single "\*" means all counters in all instances of all objects.



## stats Command (Cont.)

disk

20:00:00:0c:50:a3:c7:1  
20:00:00:0c:50:a3:b5:0  
20:00:00:0c:50:a3:66:f  
20:00:00:0c:50:a3:6b:5  
20:00:00:0c:50:a3:69:8

total\_transfers  
user\_reads  
user\_writes  
cp\_reads  
guaranteed\_reads  
guaranteed\_writes  
user\_read\_chain  
user\_write\_chain  
cp\_read\_chain  
guaranteed\_read\_chain  
guaranteed\_write\_chain  
user\_read\_blocks  
user\_write\_blocks  
cp\_read\_blocks  
guaranteed\_read\_blocks  
guaranteed\_write\_blocks  
user\_read\_latency  
user\_write\_latency  
cp\_read\_latency  
guaranteed\_read\_latency  
guaranteed\_write\_latency  
disk\_busy

```
Select Telnet 10.32.91.125

system> stats show disk:.*
disk:20:00:00:0c:50:a3:c7:11:total_transfers:0/s
disk:20:00:00:0c:50:a3:c7:11:user_reads:0/s
disk:20:00:00:0c:50:a3:c7:11:user_writes:0/s
disk:20:00:00:0c:50:a3:c7:11:cp_reads:0/s
disk:20:00:00:0c:50:a3:c7:11:guaranteed_reads:0/s
disk:20:00:00:0c:50:a3:c7:11:guaranteed_writes:0/s
disk:20:00:00:0c:50:a3:c7:11:user_read_chain:0
...
disk:20:00:00:0c:50:a3:67:5a:user_read_chain:0
disk:20:00:00:0c:50:a3:67:5a:user_write_chain:0
```

In the sample above, we are listing stats for all the disks

```
Select Telnet 10.32.91.125

system>stats show disk:20:00:00:0c:50:a3::6b::58:disk_busy
disk:20:00:00:0c:50:a3:6b:58:disk_busy:0%
system>
```

Note: The disk instance name contains colons, therefore it must de-referenced by using the colon twice

In the sample above, we are listing a specific counter for a disk instance

© 2010 NetApp, Inc. All rights reserved.

## STATS COMMAND (CONT.)

**NOTE:** The command `storage show disk -a` will show the worldwide name for a specific disk in your system.



## Preset sysstat.xml File

system> stats show -p sysstat -i 1

| CPU | NFS | CIFS | HTTP | Net in | Net out | Disk rea | Disk wri |
|-----|-----|------|------|--------|---------|----------|----------|
| %   | /s  | /s   | /s   | KB/s   | KB/s    | KB/s     | KB/s     |
| 0   | 0   | 0    | 0    | 0      | 0       | 0        | 0        |
| 1   | 0   | 0    | 0    | 0      | 1       | 48       | 268      |
| 0   | 0   | 0    | 0    | 1      | 0       | 0        | 0        |
| 2   | 0   | 34   | 0    | 924    | 23      | 0        | 0        |

```
#cat /etc/stats/preset/sysstat.xml
<?xml version="1.0" ?>
<!-- This preset is similar to the traditional
'sys-stat' command, using column
output -->
<preset orientation="column"
 print_instance_names="false"
 catenate_instances="true" >
 <object name="SYSTEM">
 <counter name="cpu_busy">
 <width>4</width>
 <title>CPU</title>
 </counter>
 <counter name="nfs_ops">
 <width>6</width>
 <title>NFS</title>
 </counter>
 <counter name="cifs_ops">
 <width>6</width>
 <title>CIFS</title>
 </counter>
 ...
 </object>
</preset>
#
```

You can create customized XML files to display only the statistics that are important to you

© 2010 NetApp, Inc. All rights reserved.

## PRESET SYSSTAT.XML FILE

The `stats` command supports preset configurations that contain commonly used combinations of statistics and formats. The preset to be used is specified with the `-p` command-line argument. For example:

```
stats show -p sysstat
```

Each preset is stored in a file, the `/etc/stats/preset` directory of the root volume. This directory contains a few template files that may be customized.



**perfstat**

© 2010 NetApp, Inc. All rights reserved.

## PERFSTAT



## perfstat

- Data collection script
  - Collects statistics
  - Good for collecting time-sequence data
- Captures both storage controller and host-side data
- Always use the latest perfstat version available from the NOW™ (NetApp® on the Web) site
  - Windows
  - UNIX® (shell)

© 2010 NetApp, Inc. All rights reserved.

### PERFSTAT

The `perfstat` script is available for download from the NOW (NetApp on the Web) site. It is a simple script that administrators can run from a client. The script runs several storage system commands and client commands to collect data.

#### PERFSTAT TOOL CAPTURES JUST ABOUT EVERYTHING

`perfstat` for UNIX is a simple Bourne shell script that captures performance and configuration statistics. Output from `perfstat` is typically captured in an output file for later analysis. `perfstat` is capable of capturing information from the host(s) and NetApp storage systems simultaneously. Currently, `perfstat` supports the following operating system platforms: Solaris™, HP-UX®, OSF1, Linux®, AIX®, FreeBSD.

`perfstat` is typically run as the root user from the host or as a user with root-level permissions.



## perfstat Examples

- Use with a workload running in the background to monitor performance:  
`perfstat -f system -t 10 > perfstat.out`
- To gather information from multiple storage systems, use:  
`perfstat -f system1, system2 -t 10 > perfstat.out`
- Alternative technique  
Step1: `perfstat -b -f system -t 10 > perfstat.out`  
Step 2: Run workload of interest  
Step 3: `perfstat -e -f system >> perfstat.out`  
End statistic collection

Begin  
collecting  
statistics

© 2010 NetApp, Inc. All rights reserved.

### PERFSTAT EXAMPLES

Use with a workload running in the background to monitor performance:

```
perfstat -f storagesystemname -t 10 > perfstat.out
```

Send perfstat.out to NetApp Technical Support for analysis.

To gather information from multiple storage systems, use:

```
perfstat -f storagesystemname1, storagesystemname2 -t 10 > perfstat.out
```

An alternative technique:

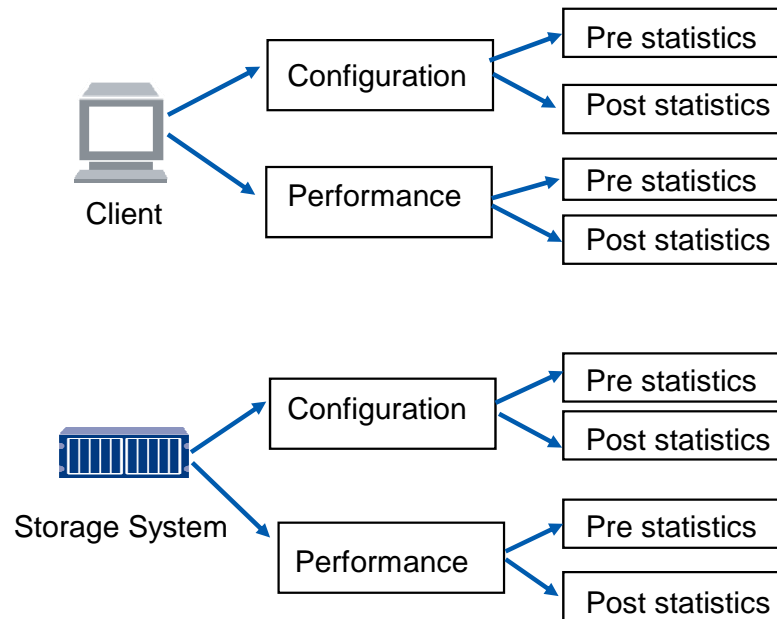
```
perfstat -b -f storagesystemname -t 10 > perfstat.out
```

Run a workload of interest:

```
perfstat -e -f storagesystemname >> perfstat.out
```



## perfstat Output



© 2010 NetApp, Inc. All rights reserved.

### PERFSTAT OUTPUT

`perfstat` output is divided into storage system output called storage system-side or host-side output. This output is further divided into the iterations given at the command line. You can view the output of the commands prior to the iteration (PRESTATS) and after the iteration (POSTSTATS).

Using utilities like PerfViewer can help you find all the information in the output of `perfstat`.



## Data Growth Management

© 2010 NetApp, Inc. All rights reserved.


## DATA GROWTH MANAGEMENT





## Data Layout

- Ongoing disk utilization can affect performance
- There are two main types of data layout:



© 2010 NetApp, Inc. All rights reserved.

## DATA LAYOUT

Free space locality can affect disk performance when the free space available on an aggregate is scattered unevenly across a RAID group or RAID groups. Free space locality issues occur over time as old data is deleted and the data blocks associated with that data become available. It is part of the natural aging process for a volume. Administrators can take measures to prevent any performance degradation.

File space locality issues occur when the data blocks associated with a file are not stored together in contiguous areas of the disks. When the WAFL® file system writes files, it attempts to write using the most efficient data layout. When the disks are very full, WAFL cannot optimize the writes for efficiency.



## Free Space Usage

- As aggregates and volumes age:
  - There will be many small updates (writes)
  - Freed data blocks may be overwritten
- Nearly full aggregates:
  - Tend to use most of the space available
  - Will not have large areas of free space
- Adding disks after an aggregate is nearly full means:
  - WAFL will need to store writes to the new disks

© 2010 NetApp, Inc. All rights reserved.

### FREE SPACE USAGE

Space locality issues are caused when an aggregate becomes close to full and the file system is aging. As older data is removed, the Snapshot™ copies associated with that data expire, resulting in free blocks. However, the newly freed blocks are often not in contiguous areas of the disk. The result is space locality issues.

Adding disks in small increments (in other words, less than a RAID group at a time) means that writes to the new disks will be in partial stripes. WAFL will try to increase efficiency by writing to the new disks.



## Symptoms of Space Locality Issues

- Can be diagnosed with `perfstat` output:
  - Back-to-back consistency points (CPs) and poor chain lengths
  - Excessive CP reads (for example, 1-to-1 ratio between CP read operations and writes)
  - Skewed chain length histogram (certain disks alone have good chain lengths)
  - Partial stripe histogram in `perfstat` output; look for disk statistics
- Use the optimization value in the output of `reallocate status` (`reallocate measure /vol/voll`)

```
3942.08 stripes written 3923.73 partial stripes
 18.35 full stripes 9434.67 blocks written
11083.34 blocks read 1217.20 1 blocks per stripe size 7
 1103.87 2 blocks per stripe size 7 868.77 3 blocks per stripe size 7
 464.72 4 blocks per stripe size 7 198.94 5 blocks per stripe size 7
 70.24 6 blocks per stripe size 7 18.35 7 blocks per stripe size 7
```

© 2010 NetApp, Inc. All rights reserved.

## SYMPTOMS OF SPACE LOCALITY ISSUES

To calculate the number of blocks written for a particular RAID group, multiply how many times a stripe was written to each RAID group by the number of blocks that were written to in that stripe. For instance, in the example above:

One disk was written to in the stripe of seven = **1217.20 stripes (1.24 stripes \* 7 disk = 8520.4 blocks)**

Two disks were written to in the stripe of seven = **1103.87 stripes (0.59 stripes \* 7 disks = 7727.09 blocks)**

Three disks were written to in the stripe of seven = **868.77 stripes (1.87 stripes \* 7 disks = 6081.39 blocks)**

**NOTE:** This output has been simplified for the sake of discussion. Large RAID groups improve performance. The RAID group size of seven has been used only for an educational sample.

Things to remember:

- This output is for data disks only.
- Data is calculated from per-second averages.
- Data is only calculated for the time when `statit` was running.



## Space Locality Issue Prevention

- Steps to avoid spatial locality issues:
  - Maintain 25% free space in aggregates with database loads
  - Maintain 15% free space in aggregates with file-based loads
  - Add at least four disks to an aggregate at a time
- Data ONTAP 7.0 and later has additional tools to mitigate space and file space locality issues
  - Use flexible volumes
  - Use the reallocation command

© 2010 NetApp, Inc. All rights reserved.

### SPACE LOCALITY ISSUE PREVENTION

Steps to avoid space locality issues:

- Maintain 25% free space in aggregates for database loads
- Maintain 15% free space in aggregates for file-based loads
- Add at least four disks to an aggregate at a time

Data ONTAP 7.0 has additional tools to mitigate space (and file) spatial locality issues on flexible volumes.

**NOTE:** NetApp best practice is to add disks equal to the number of disks already in the RAID group.



## Solutions for Space Locality Issues

- Use `reallocate` to reduce file spatial locality issues
  - The `reallocate` command can be run at the aggregate, FlexVol and file level
- Reallocating a file or a volume
  - Can increase space consumed by the file when the file has data blocks in common with a Snapshot copy
  - Requires a substantial amount of free space (at least 25%) when processing a volume
- With Data ONTAP 7.0, WAFL supports continuous file reallocation using schedules
  - Better UI and scheduling, as well as the ability to check, optimize, and recheck built into the design

© 2010 NetApp, Inc. All rights reserved.

## SOLUTIONS FOR SPACE LOCALITY ISSUES

The reallocation family of commands manages the allocation, or layout, of large files and LUNs on a storage system. Additionally, all files in a volume may be reallocated and the block layout of aggregates may be optimized. Using the `reallocate` command, layout measurement and optimization (reallocation) can be automated.

The allocation management process consists of three main steps:

**Measure the current layout.** If the optimization is less than a threshold value, then take no action. This step is optional.

**Perform reallocation.** When performing aggregate reallocation only Step two currently applies. This is split into two phases:

- Perform block reallocation of the aggregate.
- Fix the flexible volume information within the aggregate.

**Measure the layout again.** If the optimization is above the threshold value, repeat Step two and Step three as necessary.

These steps, together with scheduling reallocation comprise a reallocation job.



## Reallocation Scheduling

- `reallocate` jobs are scheduled to run periodically
- Three ways to schedule:
  1. Interval (default)
    - Default = 1 day between scans
      - Next job is scheduled 24 hours after completion
  2. Specific schedule
    - Cron-type scheduling
      - Use “\*” for all, “-” for range, “,” for list
      - Do not use “\*” for minutes
  3. No schedule (once only)
    - `reallocate start -f <path>`
    - `waf1 scan reallocate <path>`

© 2010 NetApp, Inc. All rights reserved.

## REALLOCATION SCHEDULING

Before scheduling a job, remember to start the job. Starting the job creates a default schedule that you then modify when you use the `reallocate schedule` command.



## Reallocation Scheduling (Cont.)

| Minute           | Hour | Day of Month | Day of Week | Result                                           |
|------------------|------|--------------|-------------|--------------------------------------------------|
| 0                | 23   | *            | 6           | Scan at 11:00 p.m. every Saturday                |
| 0                | 10   | 15           | *           | Scan at 10:00 a.m. on the fifteenth of the month |
| 0,10,20,30,40,50 | *    | *            | *           | Scan every 10 minutes                            |
| 0                | 23   | *            | 1-5         | Scan Monday through Friday at 11:00 p.m.         |
| 0                | 21   | *            | 0,6         | Scan Saturday and Sunday at 9:00 p.m.            |

© 2010 NetApp, Inc. All rights reserved.

## REALLOCATION SCHEDULING (CONT.)

### Samples

#### Example 1:

```
reallocate start /vol/flex1/abc
```

Explanation: Starts a reallocate job on /vol/flex1/abc, setting the interval to one day by default.

#### Example 2:

```
reallocate schedule -s "0 23 * 6" /vol/flex1/abc
```

Explanation: Sets the scheduled reallocate job on /vol/flex1/abc to start at 11:00 p.m. every Saturday.

#### Example 3:

```
reallocate schedule -d /vol/flex1/abc
```

Explanation: Reverts the reallocate job for /vol/flex1/abc to the default interval of one day.

#### Example 4:

```
reallocate stop /vol/flex1/abc
```

Explanation: Stops all reallocate jobs on /vol/flex1/abc.

#### Example 5:

```
reallocate start -f /vol/flex1
```

Explanation: Starts a full reallocate on the volume named flex1.



## reallocate Command

```
system> reallocate start /vol/flex1/abc
Fri Nov 5 20:17:11 GMT [wafl.scan.start:info]: Starting WAFL
layout measurement on volume flex1.
Reallocation scan will be started on '/vol/flex1/abc'.
Monitor the system log for results.
system> reallocate status -v
Reallocation scans are on
/vol/flex1/abc:
 State: Idle
 Flags: maybe_realloc,repeat
 Threshold: 4
 Schedule: n/a
 Interval: 1 day
 Optimization: 1
```

By default, a reallocation on a file or LUN will have an interval of 1 day and a threshold of 4

© 2010 NetApp, Inc. All rights reserved.

## REALLOCATE COMMAND





## reallocate Command (Cont.)

```
system> reallocate measure /vol/flex1
Tue Sep 9 14:52:24 PDT [wafl.scan.start:info]: Starting WAFL
layout measurement on volume vol1.
Tue Sep 9 14:52:24 PDT [wafl.reallocate.check.value:info]:
Allocation measurement check on '/vol/flex1' is 1.
system> reallocate start -f /vol/flex1
Fri Nov 5 21:07:27 GMT [wafl.scan.start:info]: Starting file
reallocating on volume flex1.
Reallocation scan will be started on '/vol/flex1'.
Monitor the system log for results.
system> reallocate status -v
Reallocation scans are on
/vol/flex1:
 State: Reallocating: Inode 596, block 0 of 1168
 Flags: doing_force,whole_vol
 Threshold: 4
 Schedule: n/a
 Interval: n/a
 Optimization: n/a
```

© 2010 NetApp, Inc. All rights reserved.

## REALLOCATE COMMAND (CONT.)



## reallocate Command (Cont.)

```
system> reallocate start -A aggr1
Reallocation scan will be started on 'aggr1'.
Monitor the system log for results.
Tue Sep 9 14:46:23 PDT [wafl.scan.start:info]: Starting block
reallocation on aggregate aggr1.
Tue Sep 9 14:47:54 PDT [wafl.scan.br.realloc.done:info]: Block
reallocation scan on aggregate aggr1 is complete.
system> reallocate start -p /vol/flex1
Reallocation scan will be started on '/vol/flex1'.
Monitor the system log for results.
Tue Sep 9 14:58:34 PDT [wafl.scan.start:info]: Starting WAFL
layout measurement on volume flex1.
```

© 2010 NetApp, Inc. All rights reserved.

## REALLOCATE COMMAND (CONT.)

### REALLOCATE START -A AGGR1

**Perform reallocation on the aggregate aggr1.** Aggregate-level reallocation optimizes the location of physical blocks in the aggregate, improving contiguous free space in the aggregate. Aggregate Snapshot copies should be deleted prior to running aggregate reallocation. Blocks in an aggregate Snapshot copy will not be reallocated.

Volumes in an aggregate on which aggregate reallocation has started but has not successfully completed will have the 'active\_redirect' status. Read performance to such volumes may be degraded until aggregate reallocation is successfully completed. Volumes in an aggregate that has previously undergone aggregate reallocation will have the 'redirect' status.

Do not use -A after growing an aggregate if you wish to optimize the layout of existing data; instead use `reallocate start -f /vol/<volname>` for each volume in the aggregate.

### REALLOCATE START -P FILENAME

A physical reallocation (using the -p option of the `reallocate start` command) reallocates user data on the physical blocks in the aggregate, while preserving the logical block locations within a flexible volume. You can perform physical reallocation with flexible volumes or with files and LUNs within flexible volumes.

Physical reallocation might reduce the extra storage requirements in a flexible volume when reallocation is run on a volume with Snapshot copies. It might also reduce the amount of data that needs to be transmitted by SnapMirror® on its next update after reallocation is performed on a SnapMirror source volume. Physical reallocation is not supported on flexible volumes or on files and LUNs within flexible volumes that are in an aggregate created by a version of Data ONTAP earlier than Data ONTAP 7.2.



## Deduplication


© 2010 NetApp, Inc. All rights reserved.

## DEDUPLICATION



# Deduplication

## NetApp Deduplication



- NetApp® deduplication
  - 20:1 or greater for backup
- Integrated with Data ONTAP
  - General-purpose volume deduplication
  - Identifies and removes redundant data blocks
- Application agnostic
  - Primary storage
  - Backup data
  - Archival data
- Service
  - Runs as a background process and is transparent to any client

© 2010 NetApp, Inc. All rights reserved.

## DEDUPLICATION

Deduplication can be thought of as the process of “unduplicating” data. The term deduplication was first coined by database administrators many years ago as a way of describing the process of removing duplicate records after two databases had been merged.

In the context of disk storage, deduplication refers to any algorithm that searches for duplicate data objects (for example, blocks, chunks, files) and discards those duplicates. When duplicate data is detected, it is not retained, but instead a “data pointer” is modified so that the storage system references an exact copy of the data object already stored on disk. This deduplication feature works well with datasets that have lots of duplicated data (for example, full backups).

When configured, NetApp deduplication runs as a background process that is transparent to any client accessing data from a storage system. This feature allows a reduction of storage costs by reducing the actual amount stored over time. For example, if a 100 GB full backup is taken on the first night and then there is 5 GB of data change during the next day, the second nightly backup will only need to store the 5 GB changed data. This amounts to a 95% spatial reduction on the second backup. A full backup can yield more than 90% spatial reduction with incremental backups averaging about 30% of the time. With non-backup scenarios, such as with virtual machine images, gains of up to 40% space savings may be realized. To estimate your own savings, please visit our deduplication calculator at <http://www.dedupecalc.com>.



## Deduplication in Action

presentation.ppt



Original file  
20 blocks

presentation.ppt




Identical file  
20 blocks


presentation.ppt




Edited file  
10 blocks added

 = Identical blocks

With NetApp deduplication  
30 total blocks



Without NetApp deduplication  
70 total blocks



© 2010 NetApp, Inc. All rights reserved.

### DEDUPLICATION IN ACTION


In this example, one user creates a PowerPoint® presentation (presentation.ppt) that is 20 blocks in size. This presentation is then copied to another location by another user. Finally, a third user copies the presentation to a third location and edits the file, adding 10 blocks.

When we store these files on a storage system with deduplication configured, the original presentation file will be saved, while the second copy, because it is identical to the original, merely references the original file's location on the storage system. The third location of the presentation file is not completely duplicated. Because the third user edited the file, the edits are saved to the storage system while referencing all unedited blocks back to the original file.

With NetApp deduplication, there are 30 blocks being used to store a total of 70 blocks worth of data. This is a 58% space savings.



# NetApp Deduplication: Internals



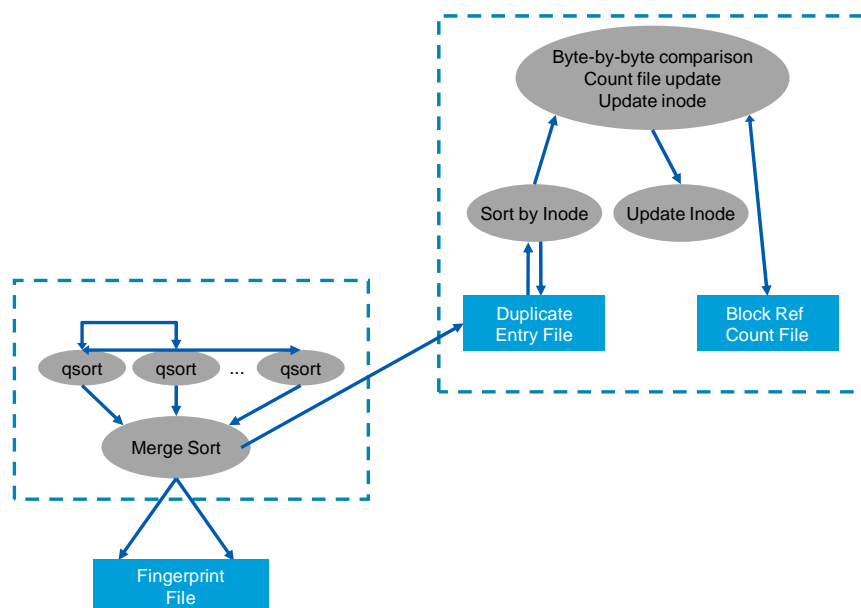
© 2010 NetApp, Inc. All rights reserved.

## NETAPP DEDUPLICATION: INTERNALS

Typically, when deduplication is enabled on a volume, data already exists on it. NetApp deduplication must then scan the existing blocks in the flexible volume and create a fingerprint file. This is accomplished by an administrator running the `sis start -s` command. During this phase, a gatherer process will identify all existing files, generate fingerprints and place them in a gatherer file. A fingerprint is a combination of a calculated value and the block location (that is [fingerprint value, block location]). The results are a 32-bytes-per-fingerprint record or 0.8% overhead. The gatherer will then pass this information to a Fingerprint Manager which performs a sorting of fingerprints using quick sort and merge sort techniques (“qsort” and merge sort in the figure). New fingerprints are then written to the fingerprint file. Over time, the fingerprint file might have a number of stale entries due to the files being deleted or moved to another volume. After 20% of the entries become stale, a stale remover phase occurs to purge the fingerprint file of outdated records.



## NetApp Deduplication: Internals (Cont.)




© 2010 NetApp, Inc. All rights reserved.

### NETAPP DEDUPLICATION: INTERNALS (CONT.)

Duplicates are identified during the merge sort process. Identified duplicate records are sorted by inode and then duplicate blocks are eliminated by the block sharing engine one after the other in the order of the inode number. Fingerprints are used to find potential duplicate blocks, but data comparison is always done before duplicates are eliminated. After the block has been identified as a true duplicate, indirect blocks are updated by pointing to the already existing data block. The reference count metadata is incremented. The duplicate block, having no inode or indirect blocking to it (that is refcount value of '0'), is considered free by WAFL.



## NetApp Deduplication: Internals (Cont.)



Fingerprint  
File

© 2010 NetApp, Inc. All rights reserved.

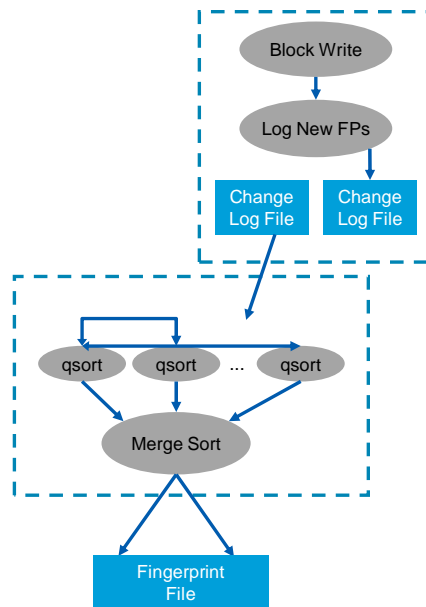
### NETAPP DEDUPLICATION: INTERNALS (CONT.)

When new write requests come in to the storage system, a new fingerprint is calculated and is written to a change log in the flexible volume metadata.





## NetApp Deduplication: Internals (Cont.)



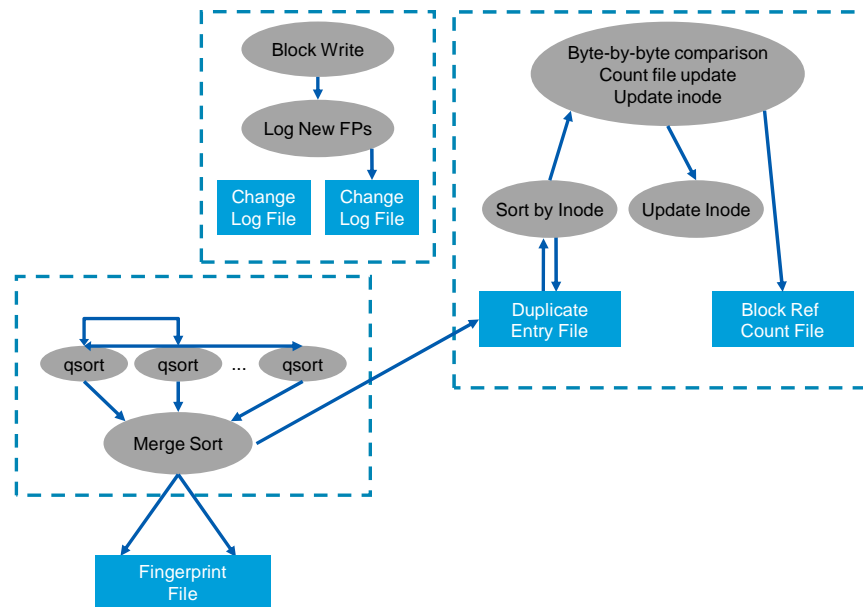
© 2010 NetApp, Inc. All rights reserved.

### NETAPP DEDUPLICATION: INTERNALS (CONT.)

The change log is then sorted by the Fingerprint Manager and the new fingers are merged into the fingerprint file. While the first change log is being processed, all new data written to the storage system is fingerprinted and its fingerprint is written to a second change log file.



## NetApp Deduplication: Internals (Cont.)




© 2010 NetApp, Inc. All rights reserved.

### NETAPP DEDUPLICATION: INTERNALS (CONT.)

Duplicates are then identified, sorted by inode, and then after a byte-by-byte comparison to verify the blocks are truly duplicate, indirect blocks are updated pointing to the already existing data block. The reference count metadata is updated. The duplicate block, having no inode or indirect blocking to it (that is refcount value of '0'), is considered free by WAFL.



## NetApp Deduplication: Internals (Cont.)



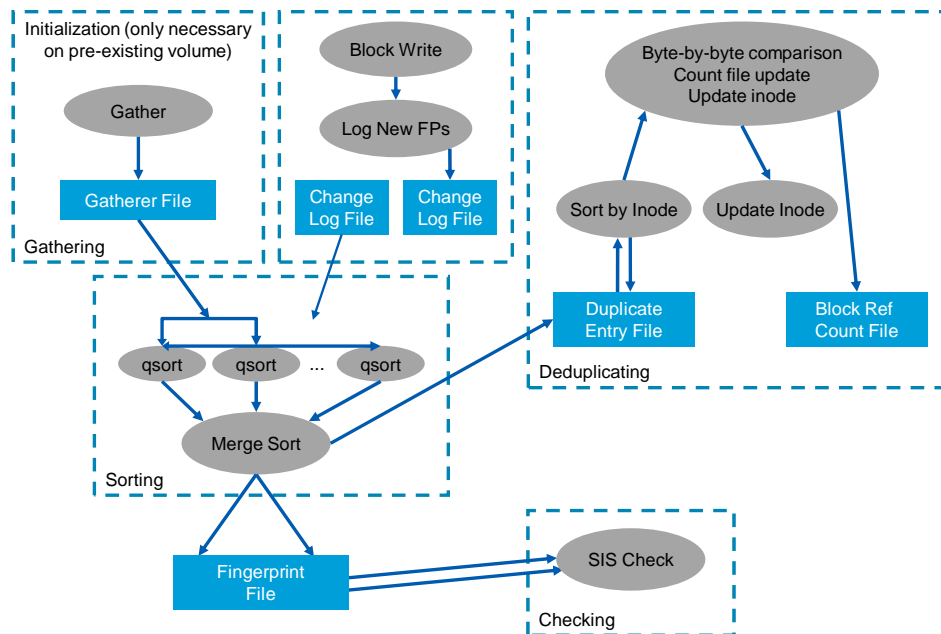
© 2010 NetApp, Inc. All rights reserved.

### NETAPP DEDUPLICATION: INTERNALS (CONT.)

For maintenance, a storage administrator may run the `sis check` command, which verifies the integrity of the fingerprint file. This is automatically triggered by the deduplication operation when 20% of fingerprint entries become stale.



## NetApp Deduplication: Stages



### NETAPP DEDUPLICATION: STAGES

As we have seen, NetApp deduplication eliminates duplicated data through sharing across files. This can be summarized in three back stages of gathering or initialization, sorting, and deduplicating files. Additionally, there is a checking stage that verifies the integrity of the fingerprint file.



## Configuration Overview

- License it:  
`system> license add <license>`
- Turn it on:  
`system> sis on <vol>`
- Deduplicates existing data:  
`system> sis start -s <vol>`
- Schedule when to deduplicate or run manually:  
`system> sis config [-s schedule] <vol>`  
`system> sis start <vol>`
- For maintenance:  
`system> sis status [-l] <vol>`  
`system> sis check`
- View the space savings:  
`system> df -s <vol>`

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURATION OVERVIEW

To configure NetApp deduplication, you must first license it on the storage system. Use the `license add` command to perform this task. Next, you must turn it on with the `sis on` command to the volume you wish to deduplicate.

If data already exists on the storage system's volume that you wish to deduplicate, you need to run the `sis start -s` command on the volume. This will scan the file system to collect fingerprints to each data block and will sort fingerprints to identify duplicate blocks. Each fingerprint entry maps a fingerprint value to the location of a disk block: [fingerprint value, block location]. Such data structure enables us to query blocks based on block contents.

The system can then be configured to run the deduplicate process at a particular time with the `sis config` command. The storage administrator may then run the `sis start` command to process fingerprints present in the change log, which are recorded while writing data to disk. During this step, new duplicate blocks will be eliminated and a list of new fingerprints will be added to the database. This can be done manually by running the `sis start` command or it may be automatically triggered by a scheduled deduplication process.

The storage administrator may view `sis status` to verify the status of the deduplication operation and use `df -s` to view the amount of space savings.

**NOTE:** When files are removed, the fingerprints are not automatically purged. Stale fingerprints are purged after a certain threshold is reached or a `sis check` command is run explicitly on the volume.



## Configuring Deduplication

```
system> sis on /vol/vol1
```

SIS for "/vol/vol1" is enabled.

Already existing data could be processed by running  
"sis start -s /vol/vol1".

```
system> sis start -s /vol/vol1
```

The file system will be scanned to process existing  
data in /vol/vol1.

This operation may initialize related existing  
metafiles.

Are you sure you want to proceed with scan (y/n)? y

Fri Nov 10 11:42:58 EST [waf1.scan.start:info]:

Starting SIS volume scan on volume vol1.

The SIS operation for "/vol/vol1" is started.

© 2010 NetApp, Inc. All rights reserved.

## CONFIGURING DEDUPLICATION

Here is an example of turning on deduplication on a volume named **vol1**.

Next, the storage administrator scans the volume to identify current space savings and add the existing data's fingerprint records to the fingerprint database by using the `sis start -s` command.



## Configuring Deduplication (Cont.)

```
system> sis status /vol/vol1
```

| Path      | State   | Status | Progress      |
|-----------|---------|--------|---------------|
| /vol/vol1 | Enabled | Active | 12 GB Scanned |

...

```
system> sis status /vol/vol1
```

| Path      | State   | Status | Progress          |
|-----------|---------|--------|-------------------|
| /vol/vol1 | Enabled | Idle   | Idle for 00:01:26 |

```
system> df -s /vol/vol1
```

| Filesystem | used     | saved   | %saved |
|------------|----------|---------|--------|
| /vol/vol1  | 20568268 | 3768732 | 15%    |

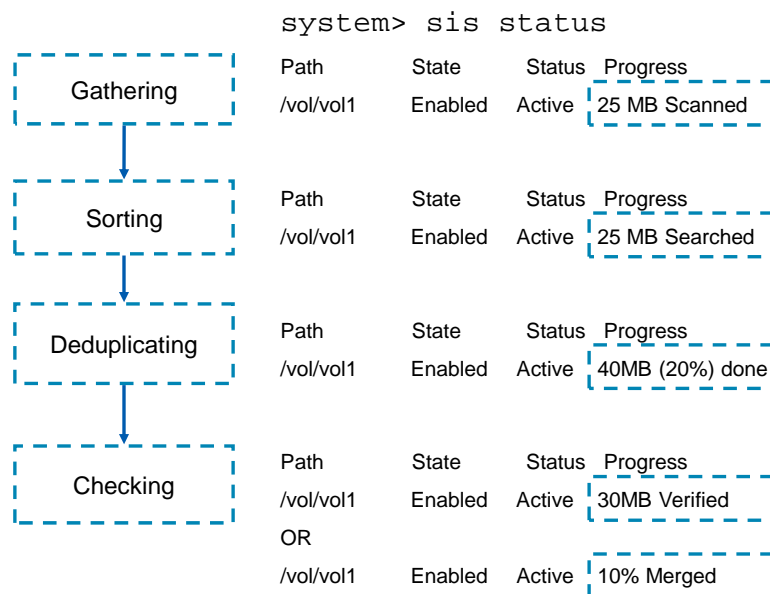
© 2010 NetApp, Inc. All rights reserved.

### CONFIGURING DEDUPLICATION (CONT.)

Here, the storage administrator uses the `sis status` command to confirm the initialization scan progress that was started with the `sis start -s` command. When the system is idle, the amount of time the process has been idle appears with the `sis status` command. Finally, a storage administrator can verify the amount of savings by using the `df -s` command.



## sis status Progress and Stages



© 2010 NetApp, Inc. All rights reserved.

### SIS STATUS PROGRESS AND STAGES

The `sis status` command will display different messages depending on stage of the deduplication process that is occurring on the storage system. This slide shows the four basic stages and the associated progress message.





## Scheduling Deduplication

### ■ Default schedule:

```
system> sis on /vol/vol1
system> sis status
/vol/vol1 sun-sat@0
```

### ■ To configure a schedule:

```
system> sis config -s - /vol/vol1
system> sis config -s 23@sun-fri /vol/vol1
system> sis config -s auto /vol/vol1
system> sis config -s sat@6 /vol/vol1
```

© 2010 NetApp, Inc. All rights reserved.

## SCHEDULING DEDUPLICATION

By default, deduplication occurs at midnight every day. This schedule may be configured by using the `sis config` command.

The schedule (-s) parameter can be specified in one of four ways:

1. If "-" is specified, there won't be a scheduled deduplication operation on the flexible volume.
2. The hours list can be given and then the day list, separated by the '@' sign.
3. If "auto" is specified, then deduplication will run on the flexible volume whenever there are 20% new fingerprints in the change log.
4. The days list can be given and then the hour list, separated by the '@' sign.



## Other Commands

- `vol status` command
  - `SIS` keyword will be listed in the output for deduplication volumes

```
system> vol status
Volume State Status Options
Vol0 online raid_dp, flex root
Vol1 online raid_dp, flex sis
```

© 2010 NetApp, Inc. All rights reserved.

## OTHER COMMANDS

The `vol status` command can be used to confirm whether a volume is a deduplication volume. The `sis` keyword will appear with in the status column if the volume is a deduplication volume.



# System Manager: Deduplication

**Start the dedupe**

| Name | Aggregate | Status | Available space | Used % | Total space |
|------|-----------|--------|-----------------|--------|-------------|
| vol0 | aggr0     | online | 19.75 GB        | 8      | 26.93 GB    |
| vol1 | aggr1     | online | 1.6 GB          | 0      | 2 GB        |

**Volume**

**Deduplication**

☒ Enable deduplication on volume

☐ Run dedupe manually  
No dedupe schedule will be set on the volume. Dedupe can be run manually.

☐ Auto  
Automatically start dedupe updates based on amount of new data in the volume.

☒ Attach custom schedules

**Schedule the days**

☒ Sunday ☒ Monday ☒ Tuesday ☒ Wednesday  
☒ Thursday ☒ Friday ☒ Saturday

**Schedule the hours**

Hours: 12:00 AM, 1:00 AM, 2:00 AM, 3:00 AM, 4:00 AM, 5:00 AM, 6:00 AM, 7:00 AM

Hour Ranges: Add, Delete

Start Time, End Time, Frequency (Hours)

OK, Cancel, Apply

© 2010 NetApp, Inc. All rights reserved.

## SYSTEM MANAGER: DEDUPLICATION



## Deduplication and SnapVault

- Integrated with SnapVault
  - Starts automatically when both SnapVault and deduplication are licensed and enabled
  - Starts after a SnapVault transfer completes
  - Does not use a schedule

© 2010 NetApp, Inc. All rights reserved.

### DEDUPLICATION AND SNAPVAULT

Deduplication and SnapVault work together in the following ways:

Deduplication starts automatically after a SnapVault transfer completes.

**NOTE:** After every SnapVault transfer, deduplication checks the log of changes. The deduplication of blocks is initiated only if the number of changed blocks is greater than 20 percent of the blocks in the volume.

Deduplication of a SnapVault volume cannot be run based on a schedule. Deduplication internally synchronizes with the SnapVault schedule. Therefore, no external synchronization is required between deduplication runs and the schedule of SnapVault transfers. However, deduplication can be run manually.

The maximum number of deduplication operations allowed to run concurrently on a storage system is eight. This includes the deduplication operations linked to SnapVault volumes and those that are not linked to SnapVault volumes.



## Monitoring Deduplication

### ■ Monitor deduplication process progress:

```
sec> sis status
```

| Path      | State   | Status | Progress          |
|-----------|---------|--------|-------------------|
| /vol/vol1 | Enabled | Idle   | Idle for 10:45:23 |
| /vol/vol2 | Enabled | Active | 25 GB Scanned     |
| /vol/vol3 | Enabled | Active | 25 MB Searched    |
| /vol/vol4 | Enabled | Active | 40 MB (20%) Done  |

### ■ Measure space savings:

```
sec> snapvault status -b
```

Snapvault secondary is ON.

| Volume | actual  | used    | saved  | %saved | ratio  |
|--------|---------|---------|--------|--------|--------|
| -----  | -----   | ----    | -----  | -----  | -----  |
| vol1   | 478GB   | 226GB   | 251GB  | 53%    | 2.11:1 |
| vol2   | 956GB   | 237GB   | 719GB  | 75%    | 4.02:1 |
| vol3   | 21478MB | 16623MB | 4854MB | 23%    | 1.29:1 |
| vol3   | 62GB    | 20GB    | 42GB   | 68%    | 3.12:1 |

© 2010 NetApp, Inc. All rights reserved.

## MONITORING DEDUPLICATION

You can use the Data ONTAP `sis status` command to monitor the deduplication process progress.

In the above example:

Volume **vol1** is *Idle*. Deduplication on the volume finished 10:45:23 ago.

Volume **vol2** is *Active*. Deduplication is doing the whole volume scanning. So far, it has scanned 25 GB of data.

Volume **vol3** is *Active*. The operation is searching for duplicated data. There is 25 MB of data already searched.

Volume **vol4** is also *Active*. The operation has saved 40 MB of data. This is 20% of the total duplicate data found in the searching stage.

You can find out how much space you have saved with deduplication by using the Data ONTAP `snapvault status -b` command. Do not use the `df -s` command.

Descriptions of the `snapvault status -b` command output fields:

`actual` - Total data sent from Veritas™ NetBackup™ and “stored”.

`used` - Total disk space consumed on the volume.

`saved` - Total storage savings on the volume.

`%saved` - Percentage of storage saved.

`ratio` - The effective compression ratio storage savings is providing.



## Optimizing Data ONTAP Configuration for NAS

© 2010 NetApp, Inc. All rights reserved.

### OPTIMIZING DATA ONTAP CONFIGURATION FOR NAS



## Optimizing NAS

- Techniques for optimizing Data ONTAP configuration for NAS:
  - Check network configuration settings
  - Use NAS best practices
  - Prevent spatial locality issues by scheduling `reallocate jobs`
  - Tune the client
  - Use FlexCache® software for volumes used mostly for reads

© 2010 NetApp, Inc. All rights reserved.

## OPTIMIZING NAS



## The `nfsstat` Command

- Too many NFS reply cache hits indicate a potential client or network problem
  - Use the `nfsstat -d` command to determine the threshold

NFS reply cache statistics:

TCP:

| In progress | Delay hits | Misses  | Idempotent | Nonidempotent |
|-------------|------------|---------|------------|---------------|
| 33          | 0          | 2627385 | 0          | 0             |

UDP:

| In progress | Delay hits | Misses   | Idempotent | Nonidempotent |
|-------------|------------|----------|------------|---------------|
| 574833      | 0          | 41387473 | 12838      | 7233          |

594,937 hits / 44,014,858 misses = 1.4%

© 2010 NetApp, Inc. All rights reserved.

### THE NFSSTAT COMMAND

A large number of NFS reply cache hits can indicate a potential client problem or a network problem.

Use the `nfsstat -d` command to determine if the threshold for NFS reply cache is too high. A threshold for this should be around 0.1%. To calculate the threshold:

Add “misses” for TCP and UDP.

Add all other numbers to determine hits.

Divide the hits by the misses to get a percentage.





## NFS Client-Side Tuning

- Use TCP over UDP
  - TCP is more reliable than UDP
  - Use the `nfsstat -c` command to view the percentage of TCP versus UDP calls over NFS

Server rpc:

TCP:

| calls           | badcalls | nullrecv | badlen | xdr call |
|-----------------|----------|----------|--------|----------|
| <b>20312393</b> | 121      | 0        | 0      | 121      |

UDP:

| calls           | badcalls | nullrecv | badlen | xdr call |
|-----------------|----------|----------|--------|----------|
| <b>86233228</b> | 4        | 0        | 0      | 4        |



~81% of NFS calls are UDP

© 2010 NetApp, Inc. All rights reserved.

## NFS CLIENT-SIDE TUNING

For NFS, the storage system supports both TCP and UDP. NetApp recommends TCP over UDP.



## NFS Client-Side Tuning (Cont.)

- Mount parameters can optimize performance depending on the application
- Set `rsize` and `wsizesize` to multiples of 4096
  - 4096, 8192, 32768
  - Gives NFS transfer sizes that are multiples of the disk block size of 4 KB
  - Reduces the number of partial payloads
- Use the highest possible `rsize` and `wsizesize` to reduce the number of NFS requests: 32768
  - Properly utilizes jumbo frames
  - Client sends as few NFS requests as possible
  - Certain switches have issues with certain sizes
  - Some clients have issues with certain sizes

© 2010 NetApp, Inc. All rights reserved.

### NFS CLIENT-SIDE TUNING (CONT.)

The `rsize` and `wsizesize` mount options determine how large a network read or write operation can be before the client breaks it into smaller operations. If jumbo frames are enabled, it is important to use an `rsize` and `wsizesize` that will effectively fill the jumbo frame.

Low `rsize` and `wsizesize` values can be appropriate if adverse network conditions prevent NFS from working with higher values. However, when you encounter poor performance because of network problems, switching to NFS over TCP is a better way to achieve good performance than using small read or write sizes over UDP. The client and server fragment large UDP datagrams, such as single read or write operations more than a kilobyte, into individual IP packets. RPC-over-UDP retransmits a whole RPC request if any part of it is lost on the network, whereas RPC-over-TCP efficiently recovers a few lost packets and reassembles the complete request at the receiving end.

Therefore using NFS over TCP, with 32K read and write sizes, usually provides good performance by allowing a single RPC to transmit or receive a large amount of data. With NFS over UDP, 32K read and write sizes may provide good performance, but often using NFS over UDP results in terrible performance if the network is at all congested. A good compromise value when using NFS over UDP is 8K or less. If you find even that does not work well, and you cannot improve network conditions, we recommend switching to NFS over TCP if possible.



## NFS Client-Side Tuning (Cont.)

- Turn off access time updates on the storage system (not the client)
  - Generally used by system administrators on local file systems
  - The storage system, controls a file's timestamp
  - Storage system command:

```
vol options no_atime_update on
```

© 2010 NetApp, Inc. All rights reserved.

## NFS CLIENT-SIDE TUNING (CONT.)

### The `no_atime` COMMAND

It is a common trick for system administrators to set the *noatime* mount option on local file systems to improve disk performance. Because NFS servers, not clients, control the values contained in a file's timestamps (access time, metadata change time, and data modify time) by default, this trick is not effective for NFS mounts. However, storage systems allow you to reduce the overhead caused by aggressive *atime* flushing if you set a volume's *no\_atime\_update* option. On a storage system console, type `help vol options` for details.



## NFS Client-Side Caching

- Increase RAM for the client
- Lengthen the attribute cache time-out (`actimeo`)
- Use the no-access-time (`noatime` or `noac`) option

© 2010 NetApp, Inc. All rights reserved.

### NFS CLIENT-SIDE CACHING

#### INCREASE RAM FOR THE CLIENT

The most effective way to improve client-side performance is to increase the client cache abilities by adding RAM to your clients. This should greatly reduce the cache turnover rate and should result in fewer read requests and faster client response time.

#### HOW CAN YOU DETERMINE IF THE CLIENT HAS THE RIGHT AMOUNT OF RAM?

Lengthen the attribute cache time-out (`actimeo`). Use the “no acknowledgement” option (`noac`).

Linux has a special mount option for file systems called `noatime` that can be added to each line that addresses one file system in the `/etc/fstab` file. If a file system has been mounted with this option, reading accesses to the file system will no longer result in an update to the `atime` information associated with the file like we have explained above. The importance of the `noatime` setting is that it eliminates the system’s need to make writes to the file system for files that are simply being read. Because writes can be somewhat expensive, this can result in measurable performance gains. Note that the write time information to a file will continue to be updated any time that the file is written to. We will set the `noatime` option to our `/chroot` file system.

#### USE THE NO ACCESS TIME (NOATIME AND NODIRTIME) OPTION

By default, the file system is mounted with normal access time (`atime`) recording. If `noatime` is specified, the file system will ignore access time updates on files, except when they coincide with updates to the `ctime` or `mtime` values. This option reduces disk activity on file system where access times are unimportant. Consult the manual pages for your version of UNIX for the right option. The option is `noatime` for Solaris and Linux. The option is `noac` (no attributes) for FreeBSD.



## Tuning the Client Operating System

- Operating system tuning
  - Install the latest kernel patches
  - Check the auto-negotiation settings for Ethernet
  - Increase maximum NFS threads, high and low watermarks, and stream settings



© 2010 NetApp, Inc. All rights reserved.

## TUNING THE CLIENT OPERATING SYSTEM



## Optimize CIFS Client Configuration

- Consider client hardware:
    - The faster the clients are, the better the overall performance
    - The larger the client memory, the better the performance
  - Increase window size setting
    - Add a registry value on the Windows client to increase window size to 64,240
    - On the storage system, set `cifs.tcp_window_size` to 64,240
- NOTE:** The maximum setting for Data ONTAP is 512K

© 2010 NetApp, Inc. All rights reserved.

## OPTIMIZE CIFS CLIENT CONFIGURATION

### HARDWARE AND OPERATING SYSTEM DEPENDENCIES

CIFS performance is sensitive to client performance (mostly due to opportunistic locking). The faster clients provide better overall performance. Clients with larger memory will also have better performance.

### WINDOW SIZE SETTING

Large window size basically increases the number of messages that can be in flight. The maximum window size that is supported on our storage system is 64,240. Increasing this on both the storage system and clients can dramatically improve performance for large transfers. You need to set the `cifs.tcp_window_size` option to 64240. The window size on the client is controlled by adding the registry value (dword):

```
\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpWindowSize
```

Then set this value to 64240 (0xFAF0 in hex).



## Optimize CIFS Client Configuration (Cont.)

- Set `cifs.max_mpx` appropriately
  - Tells how many operations per client can be pending on the storage system at the same time
  - Should never be below 50
  - Setting this automatically to a high number can potentially cause problems
  - Always contact NetApp Technical Support for more information
- Watch `cifs.neg.buf.size` if you are experiencing poor write performance:
  - See Microsoft article *Q223140*
  - See NetApp Knowledge Base article *ntapcs675*

© 2010 NetApp, Inc. All rights reserved.

## OPTIMIZE CIFS CLIENT CONFIGURATION (CONT.)

### `Cifs.max_mpx`

This tells the client how many operations can be pending on the storage system at the same time. The value should never be set below the NT default of 50. If `cifs stat` shows a maximum multiplex value greater than 32, you need to raise this number. Do not set this to any other values without an explicit conversation with NetApp Global Services. One way you can tell how much is enough is to run `perfmon` on the client that is consuming the most resources and look at the Redirector\current commands' statistics. The `Max_mpx` value should be set to a healthy margin above that number. It should be noted that just automatically switching the setting this high is not a good idea. High values consume a lot of resources in the clients. That is one of the reasons that this is a hidden option. Setting this to an unapproved value will cause errors on some clients.

### **CIFS.NEG.BUF.SIZE**

If you are having CIFS write performance issues, this is a configuration value you may wish to tune. For more information on this, refer to:

Microsoft: <http://support.microsoft.com/kb/q223140/>

NetApp: <http://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs675>



## Optimize CIFS Client Configuration (Cont.)

- Opportunistic locks (oplocks)
  - For clean networks, oplocks reduce network traffic by not doing regular updates to the server that a particular file is using
  - Data loss can happen for any application that has write-cached data under the following circumstances:
    - It has an exclusive oplock on the file
    - It is told to either break that oplock or close the file
    - During the process of flushing the write cache, the network or target system generates an error

© 2010 NetApp, Inc. All rights reserved.

### OPTIMIZE CIFS CLIENT CONFIGURATION (CONT.)

Opportunistic locks (oplocks) enable a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then read from or write to a file without regularly reminding the server that it needs access to the file in question. This improves performance by reducing network traffic.





## Optimize CIFS Client Configuration (Cont.)

### ■ Oplocks

- By default, this option is enabled
- You may want to turn oplocks off if:
  - The documentation for the database application you are using recommends that they be turned off
  - The CIFS clients are on an unreliable network
  - You are handling critical data; that is, you have a good network but you cannot afford even the slightest data loss

© 2010 NetApp, Inc. All rights reserved.

## OPTIMIZE CIFS CLIENT CONFIGURATION (CONT.)

You can turn CIFS oplocks off on individual clients. If you turn them off at the storage system, this will disable all oplocks to or from the storage system.



## Optimizing Data ONTAP Configuration for SAN

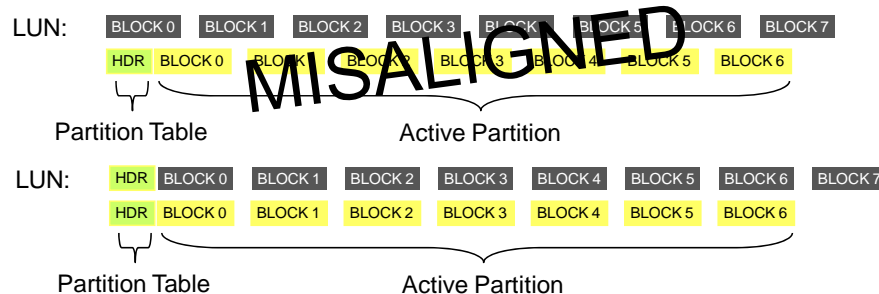
© 2010 NetApp, Inc. All rights reserved.

### OPTIMIZING DATA ONTAP CONFIGURATION FOR SAN



## Misaligned I/O and Partial Writes

- Recall:
  - To WAFL, a LUN is a set of 4-KB blocks
  - All I/O between WAFL and disks will be in multiples of 4 KB
- Consider:



© 2010 NetApp, Inc. All rights reserved.

### MISALIGNED I/O AND PARTIAL WRITES

Key point: each I/O to a block in the active partition (yellow) maps onto two partial blocks in the LUN (gray).

Consider this scenario: block 0 of the active partition is written. Sometime later block 1 is written. Sometime later block 0 is read. It will require two disk operations to fetch LUN blocks 0 and 1 because they will be on distinct places on the disk.

Note that the solution is to create the LUN properly in the first place. Properly created LUNs have a special place to store the partition table.

### WAYS TO PREVENT THE MISALIGNED I/O AND PARTIAL WRITES

Use SnapDrive® to create LUNs.

Create LUNs using the correct type (Windows, Linux, Solaris).



## Misaligned I/O and Partial Writes (Cont.)

- Impact of misaligned I/O on random reads:
  - Each read request would require two (or more) disk reads
  - Response time and throughput will be affected
- Impact of misaligned I/O on random writes:
  - Each write request will generate two partial writes
- Detect misaligned LUN using:
  - `stats show lun`
  - mbrscan tool available on the NOW site
- To correct a misaligned I/O
  - mbralign tool available on the NOW site

© 2010 NetApp, Inc. All rights reserved.

### MISALIGNED I/O AND PARTIAL WRITES (CONT.)

WAFL always moves data to and from disks in multiples of 4-KB chunks. If a LUN write does not start or end on what WAFL believes is a 4-KB boundary, a partial write is done. A partial write is a write that covers only part of one of the WAFL file system's 4-KB buffers. WAFL must read the old 4-KB chunk into memory, overlay the partial write, and then commit the merged result back to disk.

This process is slow and wasteful, and can also slow down the processing of consistency points (CPs).

WAFL will only allow a certain number of partially written buffers to exist in a CP. This prevents extending the CP and causing back-to-back CPs.



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Use Data ONTAP tools to identify networking, disk I/O, FC loop saturation, and CPU bottlenecks using `systat`, `stats`, and `perfstat`
- Discuss how increasing utilization can affect performance
- Use the `reallocate` command to maintain performance
- Use recommended techniques to optimize Data ONTAP configuration for SAN and NAS

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



Go further, faster®

## Exercise

Module 24: Performance  
Estimated Time: 0 minutes



### EXERCISE

Please refer to your Exercise Guide for more instruction.



## Check Your Understanding

- What command provides statistics from the Counter Manager?
- What directory holds `.xml` template files for customizing the output of the `stats` command?
- True or false? The `reallocate` command incurs significant overhead and can cause CPU bottlenecks.

© 2010 NetApp, Inc. All rights reserved.

## CHECK YOUR UNDERSTANDING





Go further, faster®

# Protection Manager Overview

Appendix A  
Accelerated NCDA Boot Camp  
Data ONTAP 8.0 7-Mode



## PROTECTION MANAGER OVERVIEW



## Module Objectives

By the end of this module, you should be able to:

- Describe the basic function and operation of Protection Manager

© 2010 NetApp, Inc. All rights reserved.

## MODULE OBJECTIVES



## Managing NetApp Data Protection with Protection Manager


© 2010 NetApp, Inc. All rights reserved.

### MANAGING NETAPP DATA PROTECTION WITH PROTECTION MANAGER



## Managing NetApp Protection Technology

NetApp protection solutions are based on Data ONTAP Snapshot technology



© 2010 NetApp, Inc. All rights reserved.

### MANAGING NETAPP PROTECTION TECHNOLOGY

NetApp data protection technology, based on Data ONTAP Snapshot technology, efficiently replicates data stored on NetApp storage systems, UNIX hosts, or Windows hosts using SnapVault, Open Systems SnapVault, or SnapMirror. The DataFabric Manager database engine, the Operations Manager application, and the Protection Manager administrative interface automate and simplify the management of heterogeneous data and application backup and recovery. DataFabric Manager can send and receive messages in the widely used SOAP (Simple Object Access Protocol) API for storage management in heterogeneous, traditional data center environments as well as in public, private, or hybrid ITaaS environments.

Protection Manager provides high level of assurance for data protection by proactively identifying unprotected data, checking for errors in configurations, diagnosing root cause of issues and suggesting corrective actions, and providing detailed status reports.

Protection Manager has a conformance engine that eliminates common setup errors ahead of time. Protection Manager helps you manage disaster recovery by providing DR verification for the failover and failback process.

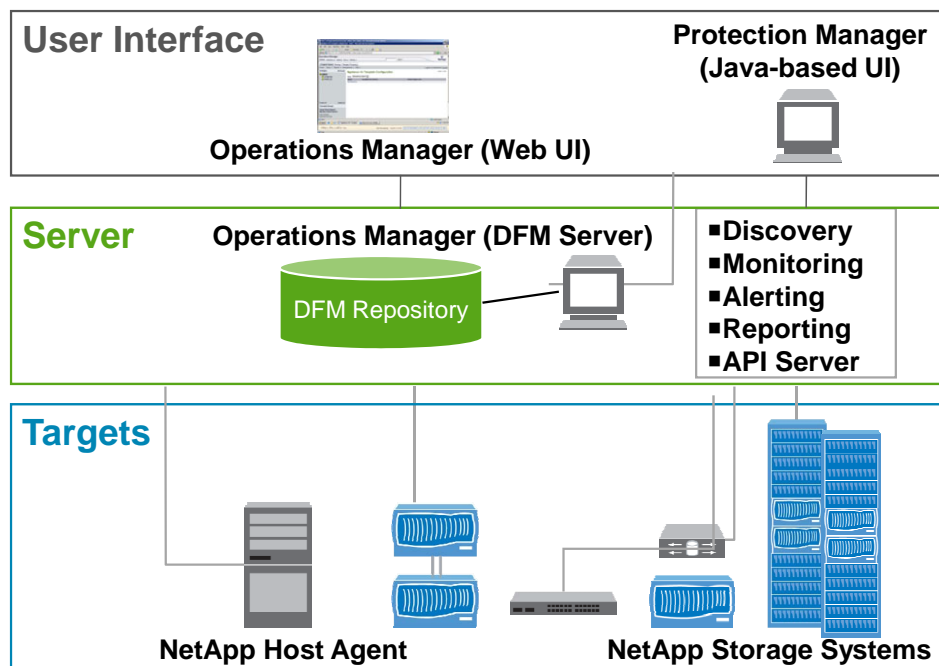
Download the NetApp Management Console (NMC) software, which contains Protection Manager, from Operations Manager and install the package on either a UNIX or Windows host.

Protection Manager uses the information stored in the Operations Manager database to be aware of primary and secondary storage systems, OSSV clients, and existing SnapVault and SnapMirror relationships.

Protection Manager operates with Snapshot technology, SnapVault, Open Systems SnapVault and SnapMirror to create new or import existing SnapVault or SnapMirror relationships and to automate replication from primary to secondary and tertiary relationships.



# Operations Manager Architecture



© 2010 NetApp, Inc. All rights reserved.

## OPERATIONS MANAGER ARCHITECTURE

Operations Manager is the user interface for a Web-based application called DataFabric Manager. DataFabric Manager discovers, monitors, and manages NetApp storage systems and can be implemented in groups of DataFabric Manager installation for larger or multiple data centers.

The DataFabric Manager server is divided into three major components:

- **User Interface** — Operations Manager. Operations Manager provides central database and agent coordination for Operations Manager, Provisioning Manager, Protection Manager, and Performance Advisor.
- **NetApp Management Console**— Protection Manager, Provisioning Manager, and Performance Advisor. The NetApp Management Console can be installed on any Windows or Linux system and is a Java-based client.
- **Targets**— the NetApp storage system hosts and open system hosts. The NetApp Host Agent is installed on a file server and is responsible for file-level reporting and tracking SAN components. Operations Manager does not require agents to monitor NetApp storage systems. However, NetApp host agent must be installed on target hosts for file-level reporting.

Protection Manager consolidates storage system aggregates into resource pools and can also manage thin provisioning. Protection Manager makes it easier for non-experts who may be called upon to manage backups because Protection Manager eliminates manual configuration, by using pre-designed policies for “configure once – apply many” management to protect new primary data. You can also preview details of set up steps.

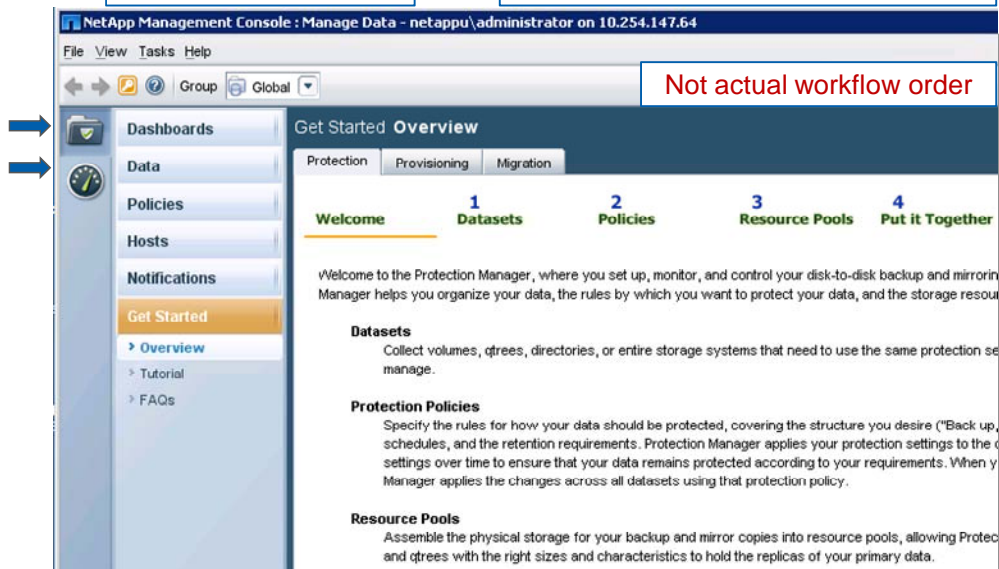
Protection Manager simplifies management of credentials. You can set up role-based access with Operations Manager and Protection Manager will inherit permissions. Protection Manager integrates the NetApp suite of products with other products in the market for management of ITaaS deployments as well as industry tape backup solutions. You can also use the Protection Manager command line interface so that you can implement scripting if necessary.



## Protection Manager Overview and Tutorials

The Get Started Overview provides flash tutorials

Overview defines each component of Protection Manager



### PROTECTION MANAGER OVERVIEW AND TUTORIALS

After downloading and installing the NetApp Management Console (NMC), from Operations Manager, you will launch the NMC on your desktop and log in using the same credentials as for Operations Manager.

Protection Manager will open; you will see buttons for Provisioning Manager and Provisioning Manager on the far left panel.

On the bottom left panel, click the Get Started button for tutorials and definitions of datasets, protection policies, and resource pools.


You will notice the “1, 2, 3, 4” on the Overview screen that is the first screen you will see. Note that the numbers do not denote workflow. Datasets are not necessarily the first component that you will add in your administration of Protection Manager. As you will see, there are some preliminary tasks performed to set up Protection Manager and then some routine tasks performed in the ongoing management of data protection.




## Protection Manager Components Defined



Hosts




Datasets



Data Resource Pools



Protection Policies  
and Schedules



Provisioning Manager  
Provisioning Policies



Conformance Checker

© 2010 NetApp, Inc. All rights reserved.

## PROTECTION MANAGER COMPONENTS DEFINED

Protection Manager has several components that work together to provide data protection.


- Datasets are collections of units of primary storage such as storage system volumes, qtrees, and directories stored on Windows or UNIX hosts. The units of primary storage are grouped together to be protected by the same protection policy and schedule.
- Hosts are NetApp or IBM storage systems, V-Series storage systems, vFilers, or open systems Windows or UNIX hosts. Hosts can be primary or secondary storage and are added to Protection Manager to be protected or to be used to provide storage for secondary copies of data.
- Data Resource Pools, (referred to as resource pools) are collections of secondary storage that are configured similarly and are assigned to datasets to provide secondary or tertiary storage.
- Protection policies are the rules established by the administrator for how data should be handled when the protection policy is attached to a dataset.
- Schedules are created in accordance with the disaster recovery plan (DRP) and assigned to protection policies.
- Conformance Checker is an algorithm of Protection Manager that compares configurations with Protection Manager rules, thresholds, alarms, and administrator policies to ensure that data is properly protected.
- Provisioning policies are a function of Provisioning Manager. Provisioning policies set rules for the selection of secondary storage systems. Provisioning policies, if created, provisioning policies are added to the criteria used by the Conformance Checker to ensure that the intent of the administrator is upheld.

The workflow entails setting up each component and then putting the components together to protect primary data that has been organized into datasets.




## Components are Created and Put Together


Add primary hosts and  
assign to Resource Pools



Add secondary hosts and  
assign to Datasets



Define Protection Policies,  
Schedules, and assign to Datasets



Define Provisioning and  
Protection Policies, then  
assign to Datasets



© 2010 NetApp, Inc. All rights reserved.

### COMPONENTS ARE CREATED AND PUT TOGETHER

To put it all together with the goal of protecting primary data, a logical beginning is to add primary and secondary hosts. Hosts are the physical resources assigned to resource pools or to datasets.

After hosts have been added to Protection Manager, the next logical step would be to define resource pools. The topic of how you will provide secondary storage comes up in the Add Datasets Wizard. Resource pools are assigned to datasets. (Note that resource pools are not mandatory; you can add secondary storage to datasets manually by selecting a storage system or an aggregate of a storage system.)

The next logical move would be to define protection policies with schedules that will live up to your disaster recovery plan. Protection policies are assigned to datasets to carry out the data protection plan.

Another option to consider as a setup step is to define provisioning policies to govern the selection of secondary storage and to set the configuration options of volumes being created to store mirror and backup copies. Provisioning policies are assigned to datasets.

The next segment explains the integration of Provisioning Manager to provide management tools for configuring storage devices.





## Protection Manager Workflow

### 1. Setup tasks that require planning and design:



- Adding Hosts (both primary and secondary)
- Adding Data Resource Pools (designate secondary storage)
- Designing protection policies with schedules
- Defining provisioning policies to select appropriate storage automatically

### 2. Administrative tasks performed routinely:



- Adding primary data to be protected in datasets
- Monitoring protected and unprotected data
- Restoring data

© 2010 NetApp, Inc. All rights reserved.

## PROTECTION MANAGER WORKFLOW

The setup tasks require planning and design. These tasks involve the disaster recovery plan of the organization, the inventory of primary data and available secondary storage, schedules that are based on RPO and RTO. The data protection setup tasks for Protection Manager include:

- Adding hosts, both primary and secondary, to be managed by Protection Manager
- Adding resource pools that ensure adequate secondary storage
- Designing protection policies that work with the disaster recovery plan
- Creating schedules that work with RPO
- Adding provisioning policies to govern the selection of secondary storage and configure volume settings

Administrative tasks performed routinely:

- Adding data to be protected in datasets
- Monitoring protected and unprotected data
- Restoring data

A setup task for using Protection Manager is to add the primary and secondary storage systems on the Hosts screen. You can add physical storage systems, vFilers (virtualized storage systems), or vFiler templates.

A vFiler template is a set of vFiler configuration settings, including the corresponding CIFS, DNS, NIS, and administrative host configuration settings, that you want to use as default settings for one or more vFiler units that you plan to add as hosts. You can configure as many vFiler templates as you need.

When adding a vFiler unit as a host, you can specify a vFiler template that provides the default configuration settings for that vFiler unit. In addition to the configuration settings provided by the vFiler template, you also must specify those values that are unique to the vFiler unit, such as name and IP address.

To collect data and manage physical or virtualized storage systems, Protection Manager needs credentials. You can set these credentials Protection Manager on the Hosts page. Select the host and use the Edit button or

the Diagnose button. You also need to enable NDMP and enter the storage system NDMP user name. Protection Manager will then configure the NDMP credentials automatically.

Administrators can also add the open system hosts that are to be protected. We will provide more details about protecting Windows, Linux, and UNIX systems when we cover Open Systems SnapVault later in this course.

We are referring to this task as a setup task because you will not have to repeatedly add hosts to Protection Manager. Once this step is completed for your data center, you will only add new hosts when you change or acquire new storage or open systems.

Another setup task not repeated often is creating resource pools. Once you organize the storage system hosts into resource pools for backups and mirror copies, Protection Manager can provision storage out of these resource pools, as needed. As part of planning for data protection you have already created aggregates of unused space on the storage system hosts you intend to assign to resource pools.

Each resource pool can be assigned to multiple datasets. Protection Manager performs calculations to ensure that there is adequate space to contain the mirror copies and backups designated in new and existing protection policies.

Once you create your resource pools, you can maintain adequate storage by adding disks to aggregates or adding more aggregates to a resource pool.

Another setup task is to set up protection policies. Data Protection provides pre-defined policies that perform common data protection sequences; such as, Back up, then mirror, Local backups only, Mirror and back up. Protection Manager provides end-to-end, policy-based management and seamless integration with Snapshots, SnapVault, SnapMirror, and Open Systems SnapVault. Protection Manager selects the appropriate technology to perform the action required in the protection policy.

The disaster recovery plan, expected RTO and RPO will determine what pre-defined protection policies will be selected. RTO and RPO will also determine what schedules will be assigned to the protection policy.

The data protection administrator can set up hosts, resource pools, and several protection policies with names that indicate the type of datasets each protection policy was intended for. After the preliminary tasks have been set up according to the organizations disaster recovery plan, even an administrator who is unfamiliar with data protection accomplish on-going management of data protection by working with the established hosts, resource pools, and protection policies.



## Module Summary

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY



## Module Summary

In this module, you should have learned to:

- Describe the basic function and operation of Protection Manager

© 2010 NetApp, Inc. All rights reserved.

## MODULE SUMMARY