



Welcome to NAS Solutions Design Concepts.

The AUDIO portion of this course is supplemental to the material and is not a replacement for the student notes accompanying this course.

EMC recommends downloading the Student Resource Guide from the Supporting Materials tab, and reading the notes in their entirety.

These materials may not be copied without EMC's written consent.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC is a registered trademark and Celerra is a trademark of EMC Corporation.

All other trademarks used herein are the property of their respective owners.

Course Objectives

Upon completion of this course, you will be able to:

- Describe the important technical data to gather regarding the use of NAS
- Describe how to collect technical data for a NAS environment
- Describe how to interpret and comprehend the gathered data
- State the parameters to set and tools used to control and manage the NAS environment
- Discuss the best practices to configure and deploy NAS

The objectives for this course are shown here. Please take a moment to read them.

Concepts and Terminology

Upon completion of this lesson, you will be able to:

- Describe NAS implementation options and terminology



The objectives for this lesson are shown here. Please take a moment to read them.

NAS Terminology Review

- Network Attached Storage (NAS)
- Storage Area Network (SAN)
- Control Station
- Data Mover
- File System

You should be familiar with the terminology listed here and on the next slide.

NAS: A method of file system access that uses a dedicated file server, and runs a multi-threaded, real-time OS which is optimized for file I/O that is addressable as a node on an IP network.

SAN: A dedicated network of storage devices which serve data at the block level and appear as physical devices. They most commonly use the Fibre Channel protocol.

Control Station: The management server in a Celerra which allows an administrator to install, configure, and manage Data Movers.

Data Mover: A file server in a Celerra.

File System: A method for the storage and organization of file, directories, navigation, manipulation, access and retrieval of data, based on an abstract (rather than physical) addressing scheme.

NAS Terminology Review (Cont.)

- Automatic Volume Manager (AVM)
- Quota
- Virtual Data Mover (VDM)
- SnapSure
- Celerra Replicator

AVM: System used by the Celerra to allow for easy and efficient creation of File Systems and their underlying storage containers (volumes).

Quota: Method of tracking, reporting and/or limiting the amount of storage consumed by a particular user, group or directory in a file system environment.

VDM: A virtual container that holds all the information necessary to support one or multiple CIFS servers. It acts as a discrete instance of a subset of the CIFS service, providing logical isolation between the VDM and the CIFS server(s) it contains.

SnapSure: The virtual point-in-time copy product for Celerra. It uses the combination of a pointer-based reserved volume and the original production file system to create and preserve a point-in-time view of that file system.

Replicator: A Celerra software product that replicates file systems between two Celerra Data Movers by leveraging SnapSure technology.

NAS New Terminology

- Internet SCSI (iSCSI)
- Celerra WORM (CWORM)
- Distributed File System (DFS)
- Nested Mount File System (NMFS)

The terminology listed on this slide and the next may be unfamiliar to you. It is important to acquaint yourself with these terms:

iSCSI (SCSI protocol over IP): Data transfer protocol which uses SCSI format data encapsulated within IP (Internet Protocol) frames for transmission over an Ethernet network.

CWORM: Technology that allows administrators to archive data to WORM storage on standard rewritable magnetic disks to create a permanent, unalterable set of files and directories in both UNIX and Windows clients.

DFS: Microsoft DFS is a client/server protocol that allows administrators to organize many resource shares into one distributed file system. By using DFS, selected shared folders can be viewed in the namespace, assigned names, and designed into a tree hierarchy in which the folders appear. Users can navigate through the namespace without needing to know the server names or the actual shared folders hosting the data.

NMFS: Allows the Celerra to present a number of different NFS file system exports to a client as a single mount point.

NAS New Terminology (Cont.)

- **Network Data Management Protocol (NDMP) Backup to Disk**
 - The Celerra can use the Network Data Management Protocol to provide backup service directly to a disk device on the Celerra, rather than tape
- **Virtual Tape Library Unit (VTLU)**
 - Allows Celerra to emulate a Tape Library Unit (TLU) for backup
- **iPlanet LDAP – Sun Java System Directory Server**
 - Allows the Celerra to authenticate with iPlanet LDAP services for UNIX/LINUX environments

NDMP: With DART version 5.4, the Celerra offers support for NDMP backup to disk via the Virtual Tape Library Unit (VTLU) feature.

VTLU: With DART version 5.4, the Celerra offers support for NDMP backup to disk using the virtual tape library unit (VTLU) feature. A VTLU emulates a physical TLU by sending the same information as a physical TLU to the backup software. Because the NDMP-compliant 3rd party backup software cannot distinguish between a virtual and physical TLU, the VTLU (once configured) can be managed like any physical TLU. A VTLU can be configured and managed using the command line interface (CLI) or Celerra Manager.

SUN JAVA System Directory Server is a distributed LDAP-based directory server that provides a central repository for storing and managing identity profiles, access privileges, application and network resource information. In a Celerra environment, this may be used to provide user account, group, hosts, and netgroups information. EMC recommends you continue to use DNS to get information about hostnames and their IP addresses.

EMC Global Education EMC²
where information lives™

EMC Celerra Platforms

Broadest Range of NAS Products

← **SIMPLE WEB-BASED MANAGEMENT** →

NS350 NS40* NS500 NS700	NS704	NS40G* NS500G NS700G	NS704G NS80G*	NSX
High availability	Advanced clustering	High availability	Advanced clustering	Advanced clustering
One or two Data Movers * X-Blade Data Movers	Four Data Movers	One or two Data Movers * X-Blade Data Movers	Four Data Movers * X-Blade Data Movers	Four to eight X-Blades
Upgradeable to gateway	Upgradeable to gateway	NAS gateway to SAN	NAS gateway to SAN	NAS gateway to SAN
Integrated CLARiiON	Integrated CLARiiON	CLARiiON, Symmetrix	CLARiiON, Symmetrix	CLARiiON, Symmetrix

© 2006 EMC Corporation. All rights reserved. NAS Solutions Design Concepts - 8

This slide displays the entire EMC NAS line of equipment. The NS Series has the following component features:

- Available as gateway and integrated
- Intended for small to medium sized use
- From one to four data movers
- One or two control stations
- CLARiiON or DMX storage
- NS500, NS500G, NS700, NS700G, and NS704G
- Sometimes has “S” at end to indicate single data mover, i.e. NS700GS
- Also can use last number to identify number of data movers, i.e. NS702G has two data movers, NS501G has one
- *Requires copper to fiber MIA adapters for storage connectivity*

EMC Global Education

EMC²
where information lives™

Integrated vs. Gateway NAS

- **Integrated NAS**
 - Dedicated Storage
 - Easy to deploy
 - Simple to manage
 - Gateway upgradeable
- **Gateway NAS**
 - Shared storage
 - Add NAS to SAN
 - High storage utilization
 - Consolidated management

The diagram illustrates two NAS architectures. On the left, 'Integrated NAS' is shown as a single server rack with two yellow cylinders labeled 'NAS' inside. On the right, 'Gateway NAS' is shown as a central cloud labeled 'SAN' connected to two server racks labeled 'NSxG' and two server racks labeled 'Storage'. The 'Storage' racks contain blue cylinders labeled 'SAN' and yellow cylinders labeled 'NAS'. Below the 'SAN' cloud are two server icons labeled 'Servers'.

© 2006 EMC Corporation. All rights reserved.

NAS Solutions Design Concepts - 9

Mercer Road
Global Education & Productivity

The Celerra NS series has two distinct versions: integrated and gateway.

Integrated NAS combines the NAS head and dedicated storage into a single packaged solution. Integrated NAS has a lowest acquisition and implementation cost, and simple management. An example model would be the NS500.

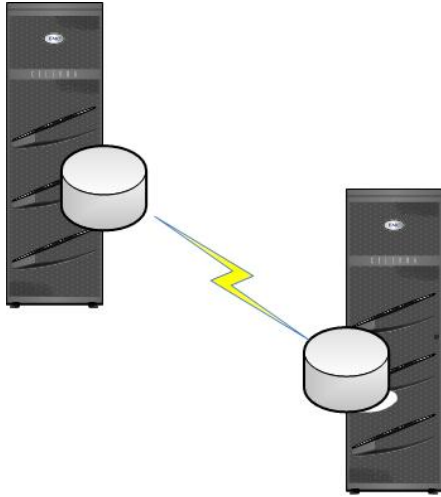
Gateway NAS is the term used for the NAS head-only, and accesses external storage. Gateway NAS is ideal for existing SANs and allows back-end storage to be pooled between NAS and open hosts. This improves the storage utilization and consolidates the management. An example model is the NS704G.

EMC Global Education

EMC²
where information lives™

Celerra Replication Technologies - 1

- Celerra Replicator
 - Asynchronous
 - File system, or Virtual Data Mover (VDM) level replication
 - Local or remote
- Celerra Replicator for iSCSI
 - Asynchronous
 - iSCSI LUN level
 - Local or remote
- Celerra with CLARiiON Mirrorview/S
 - Synchronous (Mirrorview/A not supported)
 - Block level replication
 - Leverages Mirrorview/S on CLARiiON storage
 - Limited to Mirrorview Synchronous distances
- Celerra with SRDF
 - Synchronous or Asynchronous
 - Block level replication
 - Leverages SRDF on DMX storage



© 2006 EMC Corporation. All rights reserved.

NAS Solutions Design Concepts - 10

Mercer Road
Global Education & Productivity

There are several replication technologies available with EMC NAS devices.

Two of the replication technologies are dictated by the backend storage arrays.

For a Symmetrix backend, you can use the Symmetrix Remote Data Facility feature to provide connectivity between Campus distance and extended distance Disaster Recovery facilities with either Synchronous or Asynchronous data replication.

When using a CLARiiON backend, the MirrorView volume replication methodology can be used for Synchronous data replication only because the MirrorView/A is not currently supported. This infers that the distance between sites is limited to synchronous distances only (see the CLARiiON training module for explanation of these limitations).

The other two replication methodologies, Celerra Replicator and Celerra Replicator for iSCSI, are network based and therefore Asynchronous in nature, but do not have theoretical distance limitations. However, network latency and the amount of data required to replicate places restrictions on distances if scheduled usage of the replicated data has constraints. For example, backups that need to be run and necessary data needs to be in place within a certain time window. If this schedule is not flexible, then adaptation to this solution may have to be considered.

Celerra Replication Technologies - 2

- **Celerra SnapSure**
 - Point in time logical copy
 - On Windows clients, integrates with Microsoft Volume Shadow copy Services (VSS)
 - Allows individual users to restore files
 - CLARiiON or DMX storage
- **Celerra Snapshot for iSCSI**
 - Point in time logical copy
 - On Windows clients, integrates with Microsoft Volume Shadow copy Services (VSS)
 - Transportable to other servers
 - Module available for iSCSI SnapShot integration with MS Exchange 2000


The replication methodologies previously mentioned are all built upon the Celerra SnapSure feature, except for the Symmetrix SRDF methodology.

Celerra SnapSure feature provides an instantaneous point in time view of the data to replicate. The way these products provide this is slightly different to each other. The Celerra SnapSure feature uses a discrete SavVol volume for data block storage prior to their modification in the Production File System. When using Celerra SnapSure and Celerra Replicator together on the same Celerra, (as they both make use of the same technology), it is important that sufficient storage is provided to accommodate both features as their SavVols are mutually exclusive of each other.

Celerra Snapshot for iSCSI works slightly different then Celerra SnapSure, by maintaining the point-in-time view of data within the same file system as the production volume. There are two ways this can be achieved. You can use sparse creation so no data blocks are allocated to the point-in-time view until modifications are made. Or, use dense creation which creates a block allocation equal in size to the production volume upon implementation for the point-in-time view. The choice of methodology must be carefully considered due to very different storage allocation requirements (see Celerra iSCSI training modules for further information on these methodologies).

Celerra Replication Technologies - 3

- Celerra TimeFinder/FS
 - Point in time physical copy
 - DMX storage only
- Celerra Time Finder/FS Near Copy or Far Copy
 - Point in time physical copy
 - Remote file system cloning
 - DMX storage only
- FS_COPY
 - Point in time physical copy
 - File system level
 - CLARiiON or DMX storage



Another data replication methodology is also backend governed. TimeFinder/FS is based upon the Symmetrix feature of TimeFinder. This allows for creating discrete, re-startable images of a volume and its file system. The Symetrix can perform volume replication to special volumes called Business Continuance Volumes and the Celerra tracks file system extends in the NAS_DB on the Control Station. This allows for the presentation of a BCV volume back to a Celerra and, by using a different Data Mover to the Production Data Mover, have a totally discrete copy of the Production File System for different business needs, such as testing, backup, recovery, etc.

TimeFinder/FS can be combined with SRDF for Campus distance data replication or extender distance replication Asynchronously.

Another file system based replication methodology is the FS_Copy feature, which creates a full point-in-time copy of a file system locally or over the network. This copy is independent of the backend storage array used. Over the network, copies of a file system need to be carefully planned due to network bandwidth utilization and latency issues.

Celerra iSCSI

- Traditionally, the Celerra supported file system level exports. iSCSI allows the Celerra to present block-level LUNs to a client over an Ethernet network
- **Benefits of Celerra iSCSI**
 - Block level iSCSI LUN support on the Celerra
 - Lower cost than traditional Fibre Channel SAN
 - Requires a supported iSCSI HBA and software on the client
 - Appears as a LUN on client
 - iSCSI LUNs are portable
 - Supports snapshots and replication

One of the new and most sought after technologies in the industry today is the utilization of SCSI block level commands over existing network infrastructures. One of the methodologies used to achieve this is the implementation of iSCSI support.

With EMC NAS devices it is now possible to implement traditional SCSI bound applications, MS Exchange, MS SQL, etc., over the network infrastructure by using iSCSI volumes presented to the host running the applications.

The benefits to this implementation are listed here. Please take a moment to review them.

Important Information About NAS

Upon completion of this lesson, you will be able to:

- Describe the technical data that you need to gather about the NAS environment



The objectives for this lesson are shown here. Please take a moment to read them.

NAS Consultative Analysis - 1

- Investigation of the Environment
 - Document and confirm all configurations
 - Servers: Number, OS, networking
 - Storage: Amount, utilization, number of files, size of files
 - Shares: Permissions, local users and groups
 - Network: Switches, bandwidth, utilization
 - Backups: Current capability
 - Management
 - Document all applications targeted for NAS, *AND* those that would NOT be well-suited for NAS
 - Network OS servers being used as storage devices
 - Block-level applications such as databases
 - This investigation can reveal consolidation opportunities or applications that may be better suited for SAN storage

When NAS solutions are designed, a thorough knowledge of the existing network and application environments is required. The design of effective NAS solutions requires a thorough knowledge of IP networking, an administrator's perspective on Windows and/or UNIX/LINUX network operating systems, an end-user's knowledge of applications, AND an expert knowledge of EMC's Celerra products.

Some aspects of environment analysis are shown here.

NAS Consultative Analysis - 2

- Interviews
 - IT Management
 - IT Administrators
 - Application users
 - End users

- Document any expectations
 - Capacity
 - Expandability
 - Performance
 - Availability
 - Security
 - Ease of Management

Additionally, many NAS expectations are discovered too late in the process. Early interviews of personnel involved and end-users help to understand the big picture. NAS tends to directly affect more people, such as users home directories. You need to know the experience with the current environment *and* future expectations of the new NAS solution.

The major categories of expectations are listed on the slide.

NAS Consultative Analysis - 3

- Why consider a NAS solution?
 - Save money by using their existing LAN connectivity
 - Lower cost per unit (megabyte, gigabyte)
 - Internal applications that utilize files, objects, and data
 - A need for object security through network protocols
 - A need for simultaneous Windows and UNIX client access to the same file systems
 - Server & Storage Consolidation
 - Reduction in management time and costs
 - Ease of replication of large amounts of data

Why should a NAS solution be considered? The slide suggests several reasons.

If there is existing LAN connectivity in place, it is natural to try to leverage it. If file system-level applications are already in use, NAS would make much more sense than a switch to block-level access. A NAS solution can also take advantage of existing network protocol-based object security. If there is a need for simultaneous Windows and UNIX client access to the same files, NAS is the only suitable technology. Finally, if the server and storage environments have grown in an uncontrolled fashion and become unmanageable, consolidation of servers and storage in an optimized fashion has great advantages.

NAS Consultative Analysis - 4

- What are the biggest challenges, both today and in the future?
- What are the most critical projects/deliverables?
- What are the major pain points?

Getting answers to all of these questions, or as many as is reasonably possible within the time available, can be the difference between success and failure in a NAS implementation. During the pre-solution phase, try to get this information so the most effective solution can be architected. After the solution is agreed upon, the implementation team must confirm all information that is handed over from the prior work, and augment that information to fill in any blanks.

- What are the biggest challenges, both today and in the future?
- What are the most critical projects/deliverables?
- What are the major pain points?

Knowledge of priorities allows prioritization of the features to include now and what to plan for in the next cycle.

NAS Consultative Analysis - 5

- How is the data protected?
- How is the data backed up?
- What are the Business Continuity (BC) requirements?
- How will a mixed client environment be managed?
- What service levels have to be met?
- What skills are there with SAN storage and NAS?
- Will checkpoints be used?
- Will replication be required?

Other questions to ask include:

- How is data protected?
 - Is there some style of data protection in place? Could any of the Celerra-integrated protection features (TimeFinder/FS, SnapSure, Replicator, SRDF) be effective in a NAS solution?
- How is data backed-up?
 - The Celerra has integrated NDMP backup support. If there is already a backup solution in place, it may be possible to leverage it.
- What are the Business Continuity (BC) requirements?
 - The Celerra has protection against planned and unplanned outages, including Data Mover failover, Fail Safe Network devices, and cooperation with industry-standard network connectivity protection schemes such as Etherchannel and LACP
- How is a mixed client environment managed?
 - The Celerra NAS architecture allows native-style client management in mixed protocol environments. For example, Windows and UNIX permission structures are preserved and respected. Competitors may take a different approach, which involves permission translation.

Other important questions are also listed on the slide.

NAS Consultative Analysis - 6

- Are Windows clients being used, which OS and service pack?
- Are UNIX clients being used, which NFS version?
- Is Mac used, which version and what authentication?
- How many servers are used and for what purposes?
 - Create a detailed list:
 - Data storage, utilization, capacity, number and size of files
 - File service
 - Applications service
 - Home directories
- What applications do these servers support?
 - Commercial or custom “home-grown” applications
 - Web services
 - Home Directories
 - FTP

Additional questions are listed on the slide.

- Are Windows clients used and, if so, at which OS and service pack level?
- Are UNIX clients used and, if so, at which NFS version?
 - Interoperability issues, such as client type, level and version, may seem to be minor issues, but taking them for granted can lead to serious compatibility issues. Besides the advantage of EMC’s native-style permission management, the Celerra offers support for many client features, such as CIFS auditing and SMB signing.
- How many servers are used and what are they being used for? Possible uses include:
 - Data storage
 - File service
 - Applications service
- Which applications do these servers support? For example:
 - Commercial/“Home-grown” applications
 - Intranet/Internet
 - Home Directories

Ideally, the information derived from these questions should be gathered in a structured way using the SVC Qualifier. It helps determine how much distributed storage could be consolidated onto a Celerra NAS. Detailed knowledge of application environments helps determine whether or not NAS is a good solution, and if it is, which specific features and capabilities of the Celerra architecture should be applied.

Gather Information About NAS

Upon completion of this lesson, you will be able to:

- Describe the tools and resources that allow you to gather NAS configuration, management, and performance information



The objectives for this lesson are shown here. Please take a moment to read them.

NAS Analysis and Configuration Tools

- Remote Office Assessment and Migration tool
- Server inventory
- Storage audit (available, usage, etc.)
- Third Party Storage
- SAN analysis
- Network topography and performance analysis
- NAS Support Matrix (NSM)
- NAS Code Release Notes
- SVC Celerra Qualifier
- RPQ

This slide lists some of the available tools to help analyze an environment. After the data gathering and analysis phase, some of these tools can be used to configure a potential NAS solution. The next series of slides goes into more detail about some of the tools listed here.

Remote Office Assessment and Migration ROAM

- An EMC developed solution for data migration and consolidation from remote offices to a central location (data center). Solution built around EMC OnCourse data migration software
 - SIG (Solutions Implementation Guide) provides a step-by-step method for data migration from remote locations
 - Installation
 - Assessment
 - Migration Planning
 - Migration
 - Cut over
 - ROAM Assistant tool used for migration planning

Another powerful tool available in the deployment of NAS is EMC's Remote Office Assessment and Migration tool (ROAM). This tool allows for current office/business data assessment and, if required, migration of the data. This is built around a specialized implementation of the OnCourse data migration software (see the ROAM training materials for further information).

Other NAS Analysis and Configuration Tools - 1

- **Server Inventory**
 - OS type and version
 - Interfaces, addresses, DNS names
 - Type, utilization, and size of file systems
 - Number of files and size of files
 - Local users and groups
 - File shares
- **Storage audit**
 - Current storage types: DAS, NAS, SAN
 - Versions (FLARE for CLARiiON and Microcode for Symmetrix)
 - Disk types (FC, SCSI, ATA)
 - How much room for expansion?
 - Available storage ports or FA boards

Server inventories are used to verify OS type and version, name, type, size of file systems, file shares and information about local users and groups.

A storage audit is conducted to verify type, code versions, disk types, and available storage ports or FA boards.

Other NAS Analysis and Configuration Tools - 2

- **Tools**

- **EMCGRAB**

- Gathers information for hosts
- Used with HEAT, identifies software and patch levels and configurations

- **HEAT**

- Web based tool that generates reports against EMCGRAB and EMCREPORT

- **Celerra Health Check**

- For existing Celerra environment

- **Pre Upgrade Health Check**

- For existing Celerra environment

- **Collect Support Materials**

- For existing Celerra environment

EMCGrab is a UNIX script that collects host information about HBA, patch levels, driver configuration etc. These “grabs” are available for all UNIX platforms (Solaris, AIX, LINUX etc.).

EMCReport is specific to Windows and performs the same data collections as EMCGrab.

EMCGrab and EMCReport must be installed and executed on each host connected to the array.

HEAT is a web based tool that generates reports against the EMCGrab and EMCReport output files. The generated reports are html files and provide a multitude of information. A HEAT report must be generated for each host connected to the array.

Other NAS Analysis & Configuration Tools - 3

- **Third Party Storage**
 - Are there any incumbent third party storage vendors?
 - Any re-training requirement to support an EMC NAS solution?
 - Will third party support be used?
- **SAN analysis**
 - Type of Fibre Channel Switches
 - Firmware version
 - Available ports
 - Speed of ports

Are there any incumbent third party storage vendors? For example, NetApp snapshots in the “.snapshot” directories must be identified before a data migration.

Is there a re-training requirement in order to support an EMC NAS solution?

Will third party support be utilized?

A SAN analysis should be conducted to verify the type of Fibre Channel switches in use and number of available ports.

Other NAS Analysis & Configuration Tools - 4

- Network topography
 - Create a network Visio diagram
- Network analysis
 - Number and type of Ethernet switches
 - Media type, copper or fiber
 - Number of available ports
 - Speed (gigabit, 10/100)
 - Types of trunking supported?
 - Ethernet Channel
 - LACP
 - Are there any performance issues or requirements?
 - What is the current bandwidth utilization? (Is this a fit for a NAS deployment or a SAN implementation?)

Part of any network environment analysis includes network topology. A Visio diagram can go a long way to provide an accurate understanding of the network infrastructure.

A network analysis should include, but not be limited to:

- Number and type of Ethernet switches
- Media type, copper or fiber
- Open ports
- Speed (gigabit, 10/100)
- Types of trunking supported?
 - Ethernet Channel
 - LACP
- Are there any performance issues or requirements?

NAS Support Matrix

- Formerly known as NAS Interoperability Matrix or “NIM”
- The primary reference tool for EMC NAS Configuration
- Applications supported
- File system protocol support
- Sizing and Capacity
- Disk layout templates
- Hardware Compatibility
- Backup software and tape hardware support
- Celerra Replicator and SnapSure configuration guidelines
- Located on PowerLink

The NAS Support Matrix, formerly known as the NAS Interoperability Matrix or “NIM”, is the primary reference tool for EMC NAS configuration. In this document, you find out which applications are supported, file system protocol support, sizing and capacity, disk layout templates, hardware compatibility, backup software and tape hardware support, and Celerra Replicator and SnapSure configuration guidelines. The latest version can be found on PowerLink.

NAS Code Release Notes

- Included with NAS Code documentation
- Last minute additions and corrections
- Bug fixes
- Outstanding issues
- Always check BEFORE upgrading NAS code

Another important resource for in when planning NAS deployments are the NAS Code release notes. These notes address key areas of concern when deploying NAS, some of which are listed above. Please take a moment to review the bullets on this slide.

SVC Celerra Qualifier - 1

- **Why do we require a Qualifier?**
 - Ensure requirements and environments are understood
 - Ensure correct solution is proposed
 - Resolve potential problems before they occur
 - Ensure proposal is supportable and meets EMC best practices
 - Approved qualifier is required for CCA process
 - Must be completed before sale is proposed
- **Where to get the qualifier and submit once completed?**
 - <http://gig.corp.emc.com/>

Why do we require a Qualifier?

- To ensure that requirements are fully understood, and the appropriate data is gathered by the account team.
- To ensure you are proposing the most correct and appropriate solution. This process helps verify not only that the solution proposed actually works, but addresses the needs in the best possible way. It also verifies compatibility with the existing environment. If a technical or application problem is found in the proposal, the SVC Celerra team that reviews the Celerra Qualifier can help propose a possible alternative.
- To help resolve potential problems before they occur.

SVC Celerra Qualifier - 2

- Mandatory pre-sales tool to validate the solution
- TC/SE/CSL (Technical Pre-Sales Account Member) owns the qualifier
- Typical time to complete and submit is 10-20 minutes, depending on configuration details
- NAS Support Matrix (NSM): on PowerLink and emc.com
- Process provides value-added features such as best practices, technical notes, and SVC tips for use in the field
- Efficient, quick, and effective

The SVC Qualifier is an interactive process for the mandatory pre-sales qualification of proposed solutions. It involves the completion of an interactive questionnaire, which is then submitted with supporting documentation (client meeting notes, network diagrams and solution presentation slides, etc).

The screenshot shows a software window titled "UserForm1" with the main heading "Preliminary NAS Environment Questionnaire 1/2". On the left side, there is a sidebar with the EMC logo and navigation tabs for "1 Custom", "2 SVC Rev", and "3 Propose". The main area contains several sections: "NAS Model" with radio buttons for CNS-14, NSx00G, and NSx00; "Backend Storage" with radio buttons for Symmetrix and Clarion; and "Involve SAN Connection" with radio buttons for Yes and No. Below these are 13 numbered questions with Yes/No radio button options. At the bottom, there are sections for "TS Service Offer" with five radio button options, and a table for "Proposed Hardware" and "Proposed Software". The table has columns for Backend Storage, NAS Head, SAN, DART Version, and Management Software. The footer of the window shows "© 2006 EMC Corporation. All rights reserved." and "NAS Solutions Design Concepts - 32".

The SVC Qualifier begins with a single page, the “Preliminary NAS Environment Questionnaire”, which contains a series of high level questions. The answers provided determine which detail pages will be presented. Only those detail pages that directly pertain to the environment described in the preliminary questionnaire must be completed.

RPQ

- Mechanism to provide support for a non-standard configuration
- Completed by TC and included with Qualifier
 - Custom or older applications
 - Hardware or software not listed in the NAS Support Matrix
 - Solution is reviewed to determine if it can be supported
 - If it can be supported, proposed solution can be sold
 - If it can NOT be supported, proposed solution should be re-evaluated
 - Also required to upgrade to significant NAS code releases within 90 days of release to general availability

The internal RPQ (Request for Price Quotation) process allows for Technical Architects to submit for approval configurations that are not found in the standard Support Matrix for specific individual customized deployments to mean customer demands. These are submitted to NAS engineering teams for review and approval or re-evaluation.

This process is also required for significant code upgrades within the 90 day grace period of the GA announcement.

Best Practices

Upon completion of this lesson, you will be able to:

- Describe best practices for optimizing NAS configurations



The objectives for this lesson are shown here. Please take a moment to read them.

Data Mover and X Blade Failover

- **Active/Standby Ratios**
 - **Standard Availability**
 - 7 Active to 1 Standby
 - **Moderate Availability**
 - 3 Active to 1 Standby
 - **High Availability**
 - 1 Active to 1 Standby

In environments where the cost of downtime is relatively low, and hardware costs are a major consideration, a 7 to 1 Active standby ratio will probably suffice.

In environments where availability is more important, but where there are still cost limitations, a 3 to 1 ratio is recommended.

In environments where downtime has a high cost and connectivity is critical, a 1 to 1 ratio is recommended.

Although it may be physically possible to configure a single Primary Data Mover to multiple Standby Data Movers, the failover process remains one to one. For example, if a Primary Data Mover has multiple Standby Data Movers configured and it fails, the first Standby assumes the Primary Data Mover personality. If the failed over Data Mover subsequently fails, none of the other configured Standby Data Movers will assume a new personality. This configuration is only useful if there are multiple Standbys configured, thereby allowing the Primaries to failover to ANY available Standbys when failure occurs.

CLARiiON Disk Layout

- Always use AVM profiles whenever possible
- Follow disk layout templates in the NAS Support Matrix
- When using a single DAE, do not stripe across two LUNs from the same RAID Group. Instead, try to stripe over as many spindles as possible
- When using multiple DAEs, avoid striping across LUNs of different RAID types and configurations
- Do not span file systems across multiple storage arrays

For optimal performance, stripe across different volumes. Striping across a single volume is possible, but it does not improve performance.

On an NSxxx with a single DAE, do not stripe a file system across two LUNs from the same RAID group. Instead, concatenate LUNs from a single RAID group together using Celerra, then create a stripe volume across that concatenated metavolume.

With a single DAE system, stripe file systems over as many spindles as possible, even if this means crossing RAID types or configurations.

With multiple DAE systems, avoid striping a file system across LUNs of different RAID types and configurations. Do not mix RAID1, 4+1 RAID5, and 8+1 RAID5 LUNs in a single file system. Do not mix LUNs composed of different sized spindles.

Other NAS Best Practices

- Enhanced security
- Enhanced connectivity protection
- Enhanced scalability and portability
- Ethernet trunking
 - Use for availability or aggregated bandwidth of switches support it
 - Ethernet Channel
 - Link Aggregation Control Protocol (LACP)
- Fail Safe Networking
 - High availability, even if switches do not support trunking
 - Fail Safe Networking does not trunk ports for throughput

For enhanced security, use SSH rather than telnet for remote Control Station connection. It is no longer possible to make a default connection to the Control Station using telnet. However, SSH is always available.

For enhanced connectivity protection, Fail Safe Networking (FSN) devices should be connected to redundant switches.

For enhanced scalability and portability, CIFS Servers should be configured on VDMs (Virtual Data Movers) whenever possible.

NAS Solution Migration - 1

- Data migration considerations
 - Service levels to maintain, types of data, CIFS cleanliness
- Migration methods
 - Celerra Data Migration Service (CDMS)
 - Windows migration tools
 - EMCopy
 - Robocopy
 - Secure copy
 - Drag and drop
 - Remote Office Assessment and Migration tools
 - UNIX migration tools
 - Cp
 - Rsync
 - Backup to tape and restore
 - RainFinity

When developing a NAS solution, migration of the existing data must be considered. Some considerations are the service level up-time requirements, data type to migrate (CIFS, NFS, or both), and cleanliness of the CIFS environment (if the Domain Administrator does not have permissions on all files, the migration may fail).

Migration methods and tools include:

- CDMS allows seamless migration of data using the CIFS or NFS protocol from source file servers to the Celerra Network Server with only limited interruption to normal business operations.
- Windows migration tools:
 - EMCopy duplicates a directory tree from one server to another, keeping the Windows security intact including ACLs with local groups' ACEs.
 - RoboCopy is one of Microsoft's Resource Kit's utilities. It is a very robust file copying utility with many features.
 - Secure Copy is a powerful and comprehensive data migration solution that automates the copying of data between Windows NT/2000/2003 servers without agents or scripts.
 - Drag and drop selects the files to copy and drags them to where they need to be.
- UNIX migration tools:
 - Cp, the UNIX shell command, can be used to copy all files in a directory tree, or all files in a directory and its subdirectories.
 - Rsync compares all files in a directory subtree and copies any changed files to a mirrored directory on the target machine.
- Backup to tape and restore.

NAS Solution Migration - 2

- **Replacing older Celerra hardware**
 - In place upgrade for CNS-14
 - CNS-14 with 8 or fewer data movers can be upgraded to an NSX without migration
 - More than 8 data movers require consolidation before upgrade or would not be eligible
 - For other models, migrating mixed protocol data in a Celerra to Celerra migration can be an issue
 - Celerra Replicator can be used, but has limitations
 - Code levels on each Celerra must be the same – could require an upgrade of the production Celerra
 - Limited to the number of allowed replications based on NAS code version
 - If the number of file systems to migrate exceeds the number of possible replications, other migration options must be considered or added

This slide lists points to consider when older Celerra Hardware is replaced with newer models. Please take a moment to review to gain insight into some caveats you may encounter.

NAS Solution Development

- Do you understand the technical and business needs?
- Have you prioritized those needs?
- Does the solution address the technical criteria *and* business requirements that drove them?

In order to develop a solution, you must understand many facets of the business, technical environment, current problems, and how EMC products might satisfy these requirements. It is imperative to understand business needs as well as technical needs. An assessment must then be made as to how well a proposed solution addresses those technical and business needs.

We have seen tools to help assess the specific environment, such as the SVC Celerra Qualifier. These tools provide a structured methodology for identifying and documenting the NAS opportunity. The use of this structured methodology, along with knowledge and experience with NAS, simplifies the decision making process to determine the most appropriate NAS solution.

NAS Solution Development Checklist

- Has the environment been documented?
- Has the solution been qualified?
 - Has the SVC been engaged early to provide input and alternatives?
 - Has the NAS Qualifier been submitted and approved?
 - Has an RPQ been approved if necessary?
- Does the solution address the short and long term goals?
- Has software been positioned?
- Has an appropriate migration strategy been considered?
- Has a project plan been developed?

It cannot be stated too often that in NAS, the more known about the specific environment, the more likely the implementation planning is successful. SVC Qualifier is an excellent repository for everything known about the technical environment. Regardless of the level of detail available, SVC Qualifier should be the first step.

The proposed solution should be submitted to the SVC for qualification as early in the process as possible.

The more aware you are of long term needs, and the more they are accounted for in the initial solution, future NAS opportunities and successes will follow. In most cases, the initial NAS solution proposed can open the door.

The value of EMC's suite of NAS software products cannot be overestimated. SnapSure, Celerra Replicator, Celerra Manager Advanced Edition, and the other available software products allow the configuration of a fully integrated software solution that addresses all or most of the identified needs.

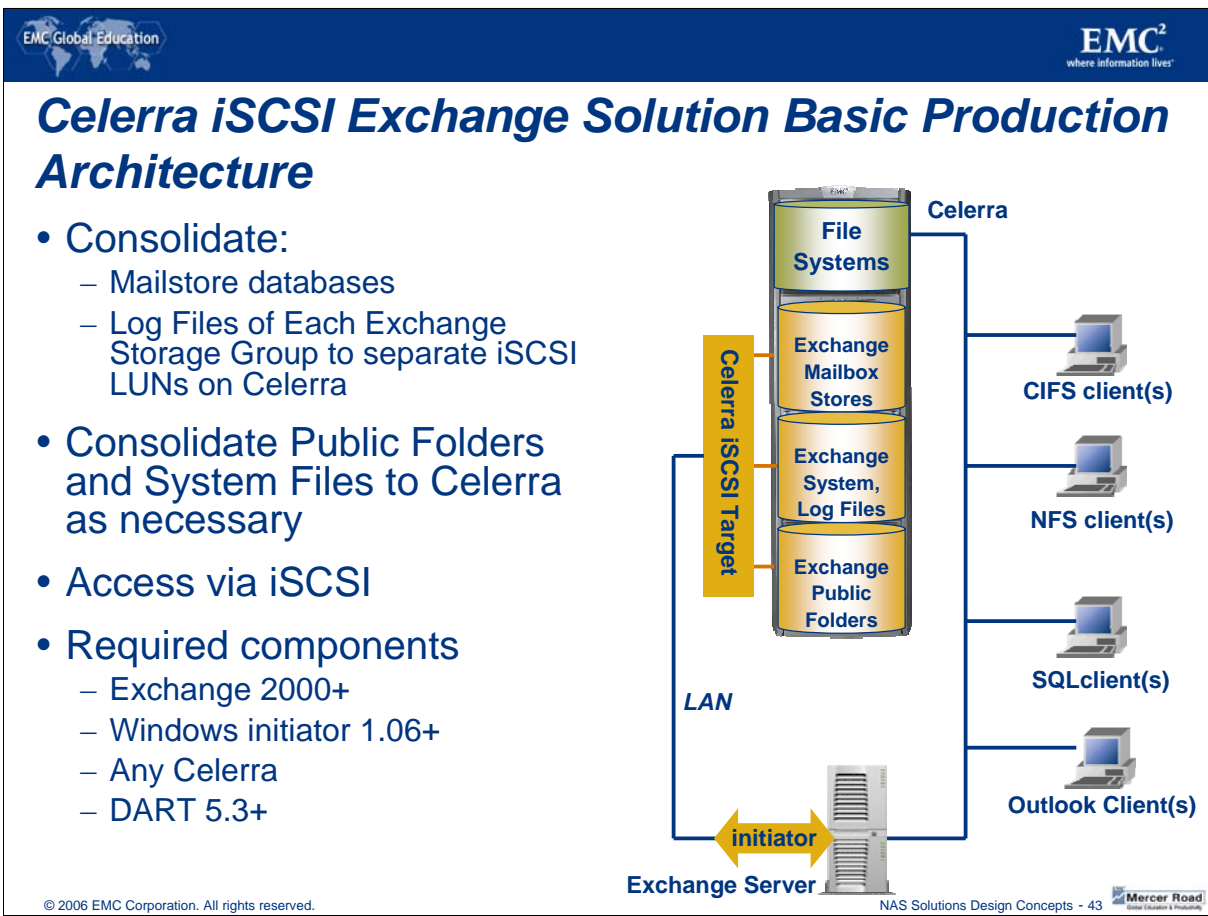
Application Specific Solutions: Exchange with iSCSI

- This is a very specific solution for Exchange implementation to EMC NAS iSCSI devices. Currently, the two devices used are the NS500 and NS700 systems
- Microsoft specific requirements for a viable Exchange solution are:
 - 0.4 IOPS sustained / mailbox
 - 0.5 IOPS peak / mailbox
 - 20 ms read / write latency

Configuration to meet Microsoft's Exchange requirements	RAID 1	RAID 5
DAE quantity per group with AVM	2	4
Spindles used per DAE	8	10
Max mailboxes* per spindle meeting Microsoft requirements (before snaps)	250	150
Medium load mailboxes* for every 4 DAEs	8000	6000
RAID spindle overhead	100%	25%

**Microsoft definition of a medium mailbox is 100 MB*

When deploying MS Exchange over iSCSI to EMC NAS devices, there are several criteria that must be considered. Please take a moment to review the details on this slide (for further information, please see the MS Exchange with iSCSI training materials).



The basic components of the Exchange Celerra–iSCSI solution consist of Windows-iSCSI compliant hardware and software. This includes any Celerra (integrated, gateway, or NSX model), with DART 5.3 or later installed. The permitted hosts include Exchange 2000 on Windows 2000 and Exchange 2003 on Windows 2003 platforms (although we’ll be concentrating on the latte because it is more iSCSI-friendly). Microsoft generally recommends that its latest initiator replace any older version, and eLab has been responsible for testing these (as they are released) for compatibility with the latest version of Celerra-iSCSI targets. Please check the EMC Support Matrix (ESM) for supported versions of initiators.

EMC Celerra Database Solutions

- Database applications traditionally implemented with SAN:
 - Have deterministic performance of a SAN
 - NAS considered appropriate only for low transaction rates with sequential I/O, flat-file like access and/ or non-critical performance
- Things have changed now...
 - NAS is well-suited for different types of I/O patterns: random and sequential, for OLTP and DSS
 - Size of the database is not a determinant for storing DB on NAS
 - Databases on NAS is an established market NAS and expected to grow at a faster rate than fiber channel
 - Oracle has endorsed NAS for databases, including Celerra

In the Industry, database implementations have generally stayed away from using Network infrastructures as their connectivity medium. This mindset has changed due to several technology advances, and Oracle, as Oracle has solidly endorsed NAS as a repository for their databases.

EMC Celerra Database Solutions

Changes that made NAS feasible for databases

- **Network enhancements**
 - Adaptation of Gigabit Ethernet
 - Network Data Separation
- **Increases in Host processing capabilities**
 - Increases in MHz and Processing Power
 - Decreases in cost of host RAM
- **Industry knowledge of managing database objects**
 - Technology knowledge sharing: web training, Metalink, etc.
 - DBAs are getting better at managing all database objects

The technology advances that have made databases viable on NAS devices are listed here. Please take a moment to review them.

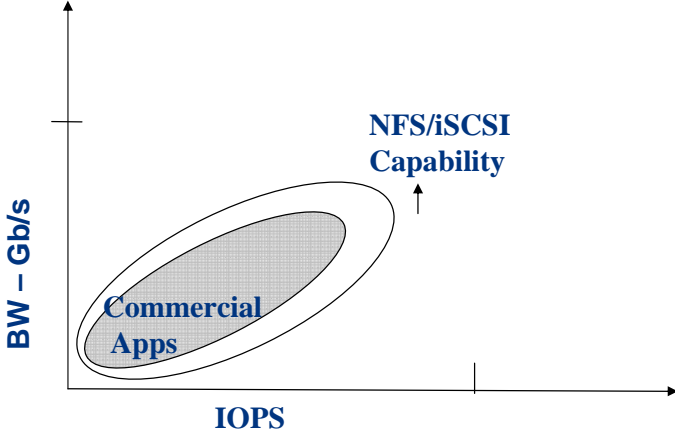

EMC Global Education

EMC²
where information lives™

Is NAS is a Good Fit?

Key Metrics that determine if NAS is a fit:

- Transaction Rates or IOPS
- Bandwidth or Throughput



© 2006 EMC Corporation. All rights reserved.

NAS Solutions Design Concepts - 46

Mercer Road
Data Location & Protection

OLTP (Online Transaction Processing) workloads have lower throughput and higher IOPS (IO's per second) need. They have random reads/writes.

DSS (Decision Support Systems) or Datawarehouse workloads, have higher throughput and lower IOPS needs, and sequential IO's. NAS (NFS+iSCSI) IOPS and throughput can meet most commercial application needs.

Just as Fibre Channel (FC) pulls significantly ahead in the performance race with 4Gb/sec speed, IP/Ethernet is about to leapfrog FC with 10Gb/sec Ethernet over copper wire. Analysts see two uses for the greater storage networking speeds: disk-to-disk backup and archiving, and storage port consolidation.

EMC Celerra Database Solutions

Databases on EMC NAS offer the following advantages:

- **Lower TCO:** Lower acquisition, administration and maintenance costs than equivalent DAS or SAN
- **Greater Manageability:** Easier Implementation, provisioning and volume management
- **High Availability:** NAS clusters provide very high levels of data availability
- **Increased Flexibility:** Easy to re-deploy data with other servers. Simplicity in making databases, or copies of database, available via remounts to other servers.
- **Improved Protection:** Integrated backup and availability
- **Superior Price/Performance:** High performance for commercial apps
- **Benefits of EMC ILM:** Tiered Storage solutions

Justification for NAS databases are listed here. Please take a moment to review this slide.

Course Summary

Key points covered in this course:

- Important technical data to gather about the use of NAS
- How to collect technical data for the NAS environment
- How to interpret and comprehend the gathered data
- Parameters to set and tools used to control and manage the NAS environment
- Best practices for configuring and deploying NAS

These are the key points covered in this training. Please take a moment to review them.