

DELL POWERVAULT NX3500

A Dell Technical Guide
Version 2.0



THIS TRANSITION GUIDE IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, *PowerConnect*, and *PowerVault* are trademarks of Dell Inc. *Microsoft*, *Windows*, *Windows Server* and *Hyper-V* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Contents

1. Product Overview	2
2. PowerVault NX3500 Clustered Solution	3
3. PowerVault NX3500 Solution Architecture	4
4. PowerVault NX3500 NAS Volumes	4
5. PowerVault NX3500 Physical Configuration	5
7. Controllers.....	6
8. Cache.....	6
9. Backup Power Supply	6
10. PowerVault NX3500 Software Tools	6
11. NAS Replication.....	8
12. PowerVault NX3500 hardware and technical specifications	9
13. PowerVault NX3500 and Dell Fluid File System Technical Specifications	10
14. Networking	11
15. PowerVault NX3500 Configuration Overview.....	12
16. PowerVault NX3500 System Glossary.....	13

1. Product Overview

The PowerVault NX3500 adds scale-up NAS capability to PowerVault MD32x0i and MD36x0i storage to create a scale-up Unified Storage solution. It is an easy-to-use solution for smaller-scale deployments with increasing iSCSI, NFS and CIFS-based data storage needs. With the PowerVault MD backend, you can scale storage of CIFS, NFS and iSCSI-based data up to 389TB of capacity. As your storage needs increase, you can add or modify block and file capacity dynamically as needed without disrupting your existing applications and storage systems.

The PowerVault NX3500 incorporates Dell Fluid File System, a high performance, high availability file system based on proven technology. The PowerVault NX3500 also offers affordable data protection features, including user-restorable snapshots, asynchronous replication, and NDMP backup capabilities.

The PowerVault NX3500 Difference

Today, a majority of organizations store and manage block and file data on separate, non-integrated storage systems. Both SAN (Storage Area Network, primarily Fibre Channel, and iSCSI) and NAS (Network Attached Storage, primarily CIFS and NFS) can be used to store unstructured data, but traditionally IT environments have selected NAS for unstructured data. NAS storage provides a consistent file system that allows the same storage space to be accessed by multiple heterogeneous clients.

Unified Storage combines SAN and NAS protocols in the same hardware platform and often has a single management interface. As unstructured (file) data continues to grow, several challenges with this approach have emerged. Spikes in NAS storage demand can cause performance bottlenecks. Unified Storage typically is priced like SAN, which costs much more in \$/GB than traditional NAS. Also, traditional Unified Storage typically has rigid architectural boundaries that limit file size, number of files, and size of the file system. Forklift upgrades to the next generation of Unified Storage hardware require tedious data migration and often weeks of application and storage network downtime.

A file system is a software layer running on the controller that manages how the data is stored and accessed. Because the file system manages data access, and functions like an operating system for the controller, its architecture and design directly impact storage system scalability and performance.

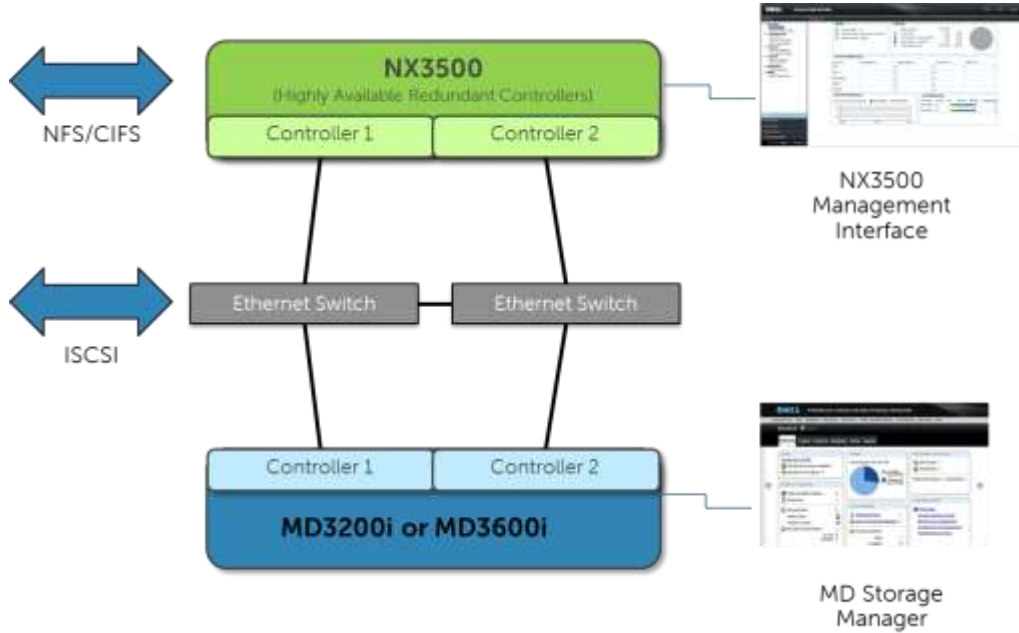
Dell Fluid File System (FluidFS), incorporated in PowerVault NX3500, is uniquely scalable. Based on a scale-out architecture with full redundancy and no single point of failure, it is designed to present a storage pool as a single file system with a single IP address to the client(s) in a flat network topology, or with multiple IP addresses to clients in a routed network. Dell Fluid File System is essentially a distributed file system, with volume management, cluster management for transparent failover, and non-disruptive scalability. Its interconnect technology delivers virtually unlimited scalability for both performance and capacity, independent of each other. FluidFS uses cache efficiently to provide fast, highly reliable reads and writes. Overall, FluidFS is designed to optimize file access

performance and hardware utilization while eliminating the capacity constraints of Windows-based file systems.

Because it can store block, CIFS, and NFS protocol data on the same hardware without the constraints of traditional Unified Storage, the PowerVault NX3500/MD system is a flexible and practical alternative to Windows file servers. And, it features a comprehensive and intuitive file management system that can help you reduce the time it takes to create and manage multiple volumes and shares.

2. PowerVault NX3500 Clustered Solution

The PowerVault NX3500 is deployed as a clustered NAS solution that aggregates multiple storage controllers and presents a NAS storage pool to UNIX, Linux and Windows clients as one virtual file server. This is an enterprise-class storage solution, optimized for a wide range of applications and deployment environments.



PowerVault Unified Storage Solution - High Level Logical Diagram

3. PowerVault NX3500 Solution Architecture

The PowerVault NX3500 system is based on a NAS gateway architecture with the hardware residing on three physically separate components:

- PowerVault NX3500 controller0
- PowerVault NX3500 controller1
- Backup power supply

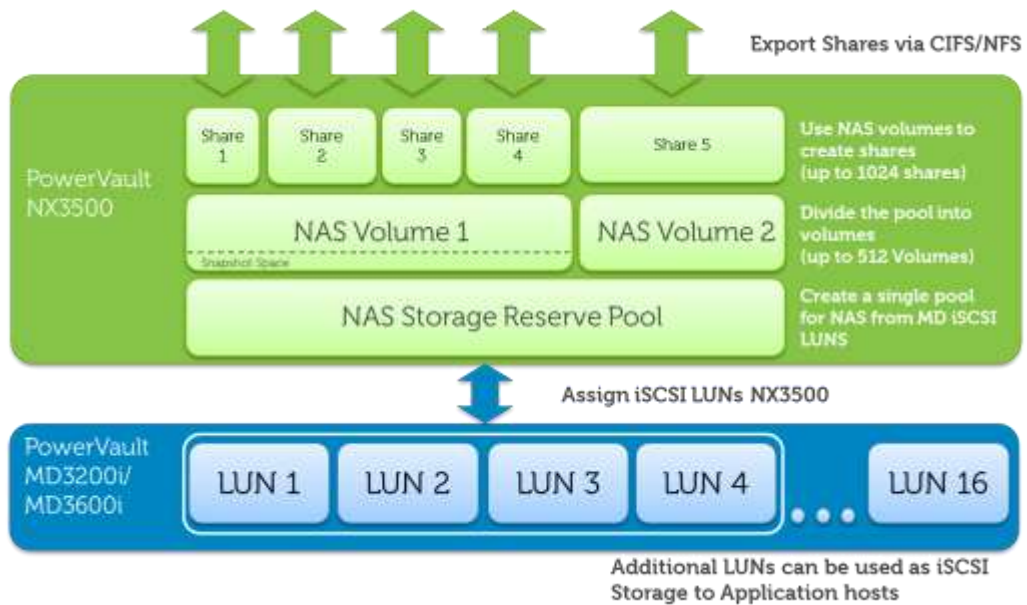
A unified storage system based on the PowerVault NX3500 and MD3xx0i hardware would include:

- Two PowerVault NX3500 controllers
- Backup power supply
- iSCSI-based storage, the NX3500 system currently supports 1 MD32x0i or MD36x0i array and multiple MD120x0 expansion shelves
- iSCSI connections
- Two internal networks
- Client network

4. PowerVault NX3500 NAS Volumes

A NAS volume is a subset of the file system and comprises a virtual unit of storage capacity that is less than or equal to the total available space. The minimum size of a NAS volume is 20 MB and the maximum size is the total amount of available space on the backend storage system. You can increase the capacity of a NAS volume without impacting data access or bringing down your SAN.

You may choose to create one large NAS volume or, alternatively, you can use multiple NAS volumes, creating, growing, and deleting them as required. The PowerVault NX3500 supports up to 512 NAS volumes per solution. The size of a file system can be increased at any time but may not be decreased in the initial release of the product. However, NAS volume size may be increased or decreased anytime. To make NAS volumes available to users, each volume must be shared (CIFS) or exported (NFS) separately. Users need to specifically map or mount each share.



NAS System Diagram

NAS Shares

NAS file shares are the final export of the NAS storage to the user. One NAS volume can have multiple file shares and shares can be either CIFS, NFS or mixed-protocol (CIFS and NFS). File share access can be fully managed to enable default or controlled access based on user name or by defining trusted users and client machines that can access NFS export directories.

5. PowerVault NX3500 Physical Configuration

The NX3500 is available in a 3U rack mount configuration:

1U PowerVault NX3500 controller0

1U PowerVault NX3500 controller1

1U Backup power unit (with two backup power supplies)



PowerVault NX3500

7. Controllers

In the clustered solution, each PowerVault NX3500 controller has two connections to the client network, two connections to the storage area network and two controller back-to-back connections for the cluster's internal network. Each controller is protected by a backup power supply (BPS), which protects data during power failure.

PowerVault NX3500 software, installed on each of the controllers, consists of an operating system, volume management, Dell Fluid File System, and Dell's advanced clustering technology. A virtual IP address provides access to the PowerVault NX3500 system.

The PowerVault NX3500 is equipped with a built-in traffic load balancing mechanism. Read/write operations are handled through mirrored Non-Volatile RAM (NVRAM). Cache data is mirrored between the paired NAS controllers. This assures a quick response to clients' requests while maintaining data integrity. Data from the cache to permanent storage is transferred asynchronously according to a variety of optimized data-placement schemes.

8. Cache

Each controller is equipped with 12 GB of RAM, most of which is used for caching to boost performance. Writing or modifying files occurs first in the cache. Data is then mirrored to the peer controller's cache. This feature ensures that all transactions are duplicated and secured.

9. Backup Power Supply

Each controller receives its power from a dedicated BPS and from the power grid. The BPS has sufficient battery power to allow the controller to execute orderly shutdown, thus enabling the cache to be used as NVRAM. Each controller regularly monitors the BPS battery status and requires the BPS to maintain a minimum level of charge for normal operation.

10. PowerVault NX3500 Software Tools

The PowerVault NX3500 solution includes configuration and management tools, as well as snapshot and replication software.

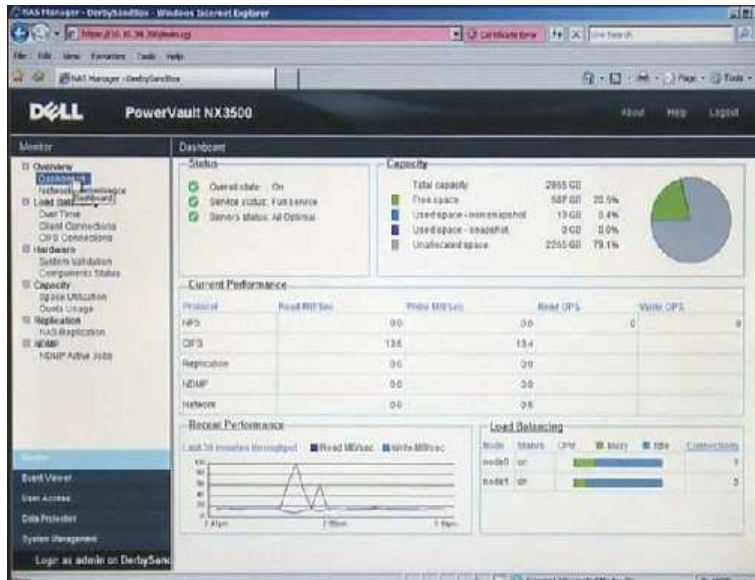
Configuration and management of the PowerVault NX3500 solution uses the following three tools:

PowerVault Modular Disk Storage Manager (MDSM) handles block storage configuration and management, including the provisioning of MD3xx0i storage for NAS.

PowerVault NAS Configuration Utility (NASCU) guides you through network configuration and pairing the PowerVault NX3500 NAS controllers together. It will also start the process of pairing the system to the PowerVault MD3xx0i storage appliance. It is recommended that administrators determine their network configuration and IP allocation for their NAS controllers before executing NASCU.

Technical Guide for PowerVault NX3500

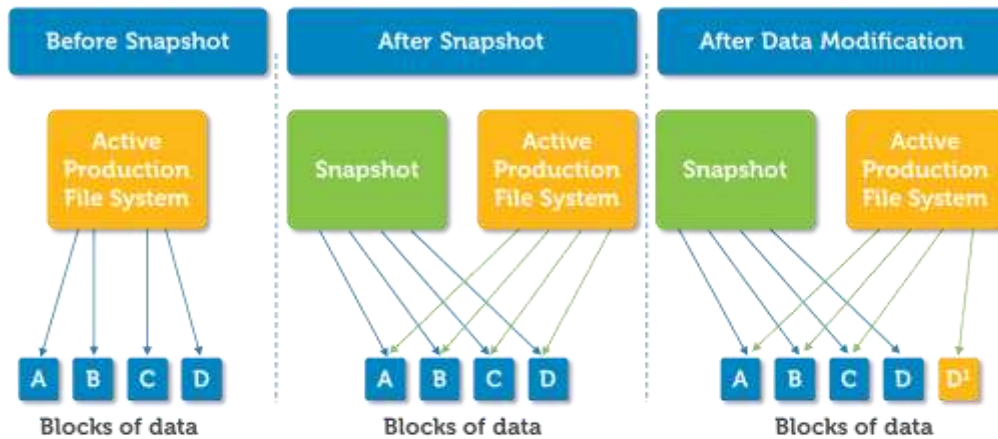
PowerVault NX3500 NAS Manager lets you configure NAS network IP addresses and create and manage volumes and shares. You also use the NAS Manager to set up snapshots and replication, engage NDMP backup, and monitor NAS file system performance.



PowerVault NX3500 NAS Manager

Dell Fluid File System Snapshots

You can use snapshots to restore files or volumes back to a previous state. Snapshots are also used with data protection features such as backup and replication to provide consistent copies of a production file system without stopping access to the production file system. Snapshot policies are applied at the NAS volume level. The first snapshot of a NAS volume is created, and then all snapshots created after the baseline snapshot will be a delta from the previous snapshot. FluidFS uses a Redirect-on-Write Snapshot design: Each snapshot creates a thin copy by saving only the differences rather than a full copy of the file system.



NX3500 uses a Redirect-on-Write Snapshot Design: Provides higher performance snapshots than older copy-on write snapshot designs.

FluidFS Snapshots

11. NAS Replication

Block-based NAS replication is a feature that uses snapshot technology in the PowerVault NX3500 file system. Replication is volume-level and can be used to replicate volumes on the same or another PowerVault NX3500 system. After the initial replication, changes to the stored data are recorded to the target volume.

When replicating a volume to another PowerVault NX3500 system, you must set up the other system as a replication partner. Once a partner relationship is established, replication is bi-directional: one system can hold target volumes for the other system as well as source volumes that are replicated to the other system. You can set up replication policies to run on various schedules as well as on demand.

12. PowerVault NX3500 hardware and technical specifications

Features	Dell™ PowerVault™ NX3500
Protocol support	Server: CIFS (SMB 1), NFS v3, NDMP V4, SNMP Client: NTP, iSCSI, DNS, Active Directory, LDAP, NIS (Network Information Service)
Storage arrays supported	Dell PowerVault MD3200i series (MD3200i, MD3220i) Dell PowerVault MD3600i series (MD3600i, MD3620i)
Expansion capability	Scale up to 192 drives with MD iSCSI arrays and expansion enclosures. Refer to PowerVault MD storage array specifications for expansion support.
Storage controllers	Dual controllers operate in an active-active environment mirroring each other's cache. Each controller contains 12GB memory protected by the PowerVault NX3500 backup power supply.
Management	PowerVault NAS Manager GUI, CLI interface
Front-end NAS connectivity	Two 1Gb IP ports per controller for front-end CIFS/NFS connectivity
Backend storage array connectivity	Two 1Gb IP ports per controller for back-end iSCSI connectivity
Snapshot capability	Redirect-on-write snapshots
Replication capability	Asynchronous 1-to-1, 1-to-Many replication for disaster recovery
Power	AC power supply (per controller power supply): Wattage: 400 W Voltage: 100-240 VAC, 50/60 Hz Heat dissipation: 1666 BTU/hr maximum Maximum inrush current under typical line conditions and over the entire system ambient operating range may reach 25 A per power supply for 10 ms or less. Backup Power Supply: Wattage: 500 W (output) Voltage: 120V or 230V models
Physical (each controller)	Height: 4.29 cm (1.69 in), Width: 43.4 cm (17.09 in) Depth: 61.26 cm (24.12 in) Weight: 15 kg (33.02 lbs)
Physical (backup power supply, 120 V and 230 V Models)	Height: 4.2 cm (1.65 in.) Width: 43.4 cm (17.09 in.) Depth: 72.3 cm (28.46) Weight: 30.0 kg (66.1 lb)
Environmental	Temperature Operating 10° to 35° C (50° to 95° F) with a maximum temperature gradation of 10° C per hour NOTE: For altitudes above 2950 feet, the maximum operating temperature is derated 1° F/550 ft. Storage -40° to 65° C (-40° to 149° F) with a maximum temperature gradation of 20° C per hour Relative humidity Operating 8% to 85% (noncondensing) with a maximum humidity gradation of 10% per hour Storage 5% to 95% (noncondensing) Maximum vibration Operating 0.25 G at 3-200 Hz for 15 min Storage 0.5 G at 3-200 Hz for 15 min

13. PowerVault NX3500 and Dell Fluid File System Technical Specifications

Attribute	Max Value (2-node)
Max system size	389 TB
Max file size	4 TB
Max files	~32 Billion files
Number of directories	~34 Billion directories
Max NAS volumes	512
Max snapshots per volume	512
Max snapshots per NX3500 system	10,000
Memory size per cluster	24 GB / 12 GB per controller
Max LUNs	16
File name length	255 bytes
Max NFS mounts	1024
Max CIFS shares (per controller)	1024
Max CIFS client concurrent connections	200
Max local users per NX3500	300
Max Quota rules per NX3500 (user quotas)	65,536
Max quota rules per volume	256
Max block level replication policies	256
Max directory depth	1,024

14. Networking

iSCSI Connections

iSCSI connections provide the NX3500 controllers access to the storage subsystem through Ethernet switches.

Internal Network

The PowerVault NX3500's internal network uses two independent 1 GbE ports. The internal network is the infrastructure for PowerVault NX3500 clustering, including the heartbeat monitor, transferring data, mirroring information between the controllers' caches and distributing data evenly across all LUNs in the system.

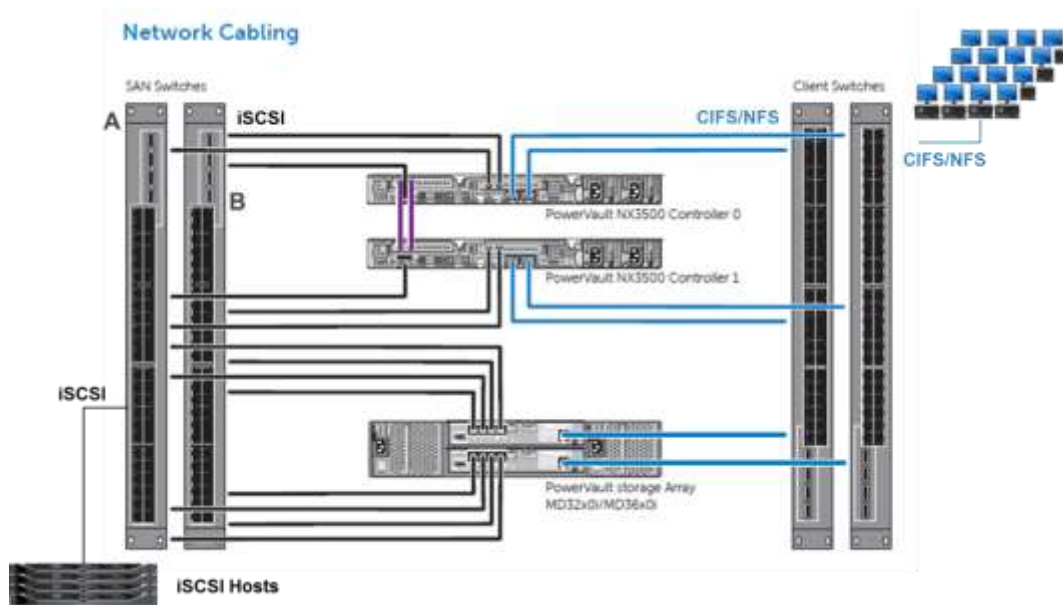
In a clustered solution, PowerVault NX3500 controllers are connected peer-to-peer and employ dual links for redundancy and load balancing. To achieve complete data distribution and to maintain high availability, each of the controllers in the PowerVault NX3500 cluster system must have access to the other controller in the system. The internal network achieves this goal.

The PowerVault NX3500 internal management network (also known as internal network B) connects both controllers. All internal administrative related functions are performed on this network. In case of a suspected "split brain" situation, in which the controllers in a cluster lose communication with each other but continue to function independently, the management network is used to automatically reset the suspected controllers. This prevents a split brain situation while ensuring data integrity.

Client Network

The connection between the PowerVault NX3500 solution and the NAS client network is based on virtual IP addresses (VIP) that serve every controller in the cluster. This enables clients to access the PowerVault NX3500 as a single unit through the virtual IP network and enables the system to load balance the network traffic between controllers, as well as continue service if one controller fails. The client network comprises two GbE NICs from each controller, which connect to client network switches.

The PowerVault NX3500 supports multiple switch topologies. Working with the Dell Services and support teams, you can select the ideal topology for your environment. (Note: Best practices suggest that you separate SAN and LAN traffic.)



Network Cabling

15. PowerVault NX3500 Configuration Overview

The PowerVault NAS Configuration Utility (NASCU) walks you through the steps necessary to set up the network configuration and pair your PowerVault NX3500 controllers. The configuration process involves the following steps:

Storage Array Identification - NASCU performs iSCSI discovery of your storage array. The NASCU wizard asks you to enter the last IP from the range of iSCSI host port IPs you have configured on your MD storage array.

NAS Controller Discovery - You enter the MAC address of the embedded NIC 1 for each PowerVault NX3500 controller. Embedded NIC 1 MAC address can be found on the system identification panel underneath the front bezel of your PowerVault NX3500s or in the front panel LCD menu.

NAS Cluster Identification - You specify the name that will be used to identify your PowerVault NX3500 cluster within the PowerVault NAS Manager web interface.

Primary Network Configuration - You specify the following: NAS Management virtual IP, Client Access virtual IP, and controller IPs (private maintenance IPs for each individual PowerVault NX3500 controller in your cluster; these should not be used by clients directly).

Internal Network Configuration - You specify the IPs to be used for internal communication between your paired PowerVault NX3500 controllers. These IPs should be from a private IP space not reachable from any other system in your network.

SAN Network Configuration - You define unique SAN IPs based on the subnets you configured on your storage array iSCSI host ports. These SAN IP addresses will provide communication between your MD storage array and your PowerVault NX3500 cluster. The IPs for each subnet group must have access to the iSCSI port IPs.

Configuration Summary - The NASCU wizard presents a configuration summary. You may modify or confirm it. When you confirm the configuration, NASCU configures your PowerVault NX3500 cluster and provides the Host Port Identifiers (IQNs) for each controller.

16. PowerVault NX3500 System Glossary

Understanding the terminology related to a Dell Fluid File System based PowerVault NAS system will help you successfully deploy, manage, and maintain your unified storage environment.

Backup Power Supply (BPS) - Provides back up battery power in the event of a power loss.

Client access VIP - Virtual IP addresses that clients use to access CIFS shares and NFS exports hosted by a PowerVault NAS system. The PowerVault NAS system supports multiple client access VIPs.

Controller (NAS controller or nodes) - Server appliance installed with the PowerVault NX3500 Dell Fluid File System software.

Controller pair - Two NAS controllers that are configured as pair in a PowerVault NAS clustered system. Cache data is mirrored between the paired NAS controllers.

Dell Management Application (DMA) - Data Management Application is also known as a Backup Application Server.

Dell PowerVault Modular Disk Storage Manager (MDSM) - The management software that ships with the PowerVault MD32x0i/MD36x0i array.

Dell Fluid File System (FluidFS) - High-performance, scalable file system software installed on NAS controllers.

Host Port Identifier - Unique ID used to identify hosts in a network.

Internal network A (peer connection) - The PowerVault NX3500's internal network consists of two independent Gigabit Ethernet ports. The internal network is the infrastructure for PowerVault NX3500 clustering, including the heartbeat monitor, data transfer, and mirroring information between the controllers.

Internal network B (internal management/IPMI) - The PowerVault NX3500 internal management network (also known as internal network b) connects both controllers. All administrative related functions and controller reboots are performed on this network.

Technical Guide for PowerVault NX3500

LAN/client network (primary network) - The network through which clients access NAS shares/exports. The PowerVault NAS system is connected to customer IT environment and its NAS clients using this network. NAS storage pool Virtual disks created on the PowerVault MD32x0i/MD36x0i storage arrays dedicated to the PowerVault NX3500 system.

NAS volume (NAS container or virtual volume) - A virtualized volume that consumes storage space in the NAS storage pool. Administrators can create CIFS shares and NFS exports on a NAS volume and share them with authorized users. A PowerVault NAS system supports multiple NAS volumes.

NAS replication - Replication between two PowerVault NAS systems or between two NAS volumes.

NAS replication partners - PowerVault NAS systems participating in a replication activity.

PowerVault NX3500 Architecture - The PowerVault NX3500 clustered NAS solution consists of a pair of controllers and the PowerVault Modular Disk (MD) iSCSI storage array. In addition, both controllers are protected by a backup power supply (BPS), which helps protect data during power failure.

Network Data Management Protocol (NDMP) - Network Data Management Protocol is used for backup and restore.

Peer controller - The peer NAS controller with which a specific NAS controller is paired in a PowerVault NAS system.

Power module (battery unit) - One of the battery units in a Backup Power Supply.

PowerVault MD3xx0i - Refers to the PowerVault MD3200i, MD3220i, MD3600i, MD3620i iSCSI storage solutions.

PowerVault NAS Configuration Utility - The setup wizard used to initially discover and configure a PowerVault NAS system. This utility is only used for initial setup.

PowerVault NAS Manager - The web-based user interface, which is part of the PowerVault NX3500 software, used to manage the PowerVault NAS system.

PowerVault NAS system - A fully configured, highly-available and scalable NAS appliance, providing NAS (CIFS and/or NFS) services, which is comprised of a pair of NAS controllers, a BPS, a PowerVault storage subsystem and the NAS Manager.

Standby controller - A server appliance that is not installed with the PowerVault NX3500 FluidFS software. For example, a new or replacement controller from the Dell factory is considered as a standby controller.

SAN network (iSCSI network) - The network that carries the block level (iSCSI) traffic and to which the storage subsystem is connected. NOTE: It is recommended that this network be isolated from the LAN/client network.

Technical Guide for PowerVault NX3500

PowerVault NAS system - A highly-available and scalable NAS system, comprised of pairs of NAS controllers, a BPS, a PowerVault storage subsystem and the FluidFS software, providing NAS (CIFS and/or NFS) services.

Standby controller - “Raw” server appliance, not installed with the FluidFS.

SAN Network / iSCSI Network - The network that carries the block level (iSCSI) traffic and to which the storage subsystem is connected. It is recommended that this network be separated from the client network.