


DELL SCALABLE FILE SYSTEM

A Dell Technology White Paper
Version 1.1



THIS TECHNOLOGY WHITE PAPER 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.

© 2011 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.

Apr. 2011

Contents

- 1. A storage management breakthrough2
- 2. Empowering the Future with DSFS4
- 3. Dell Scalable File System Features6
- 4. Dell PowerVault NX3500 - Unified Storage Solution with DSFS7
- 5. PowerVault NX3500 and Dell Scalable File System Technical Specifications8

1. A Storage Management Breakthrough

The relentless growth of unstructured data is accelerating the need for network file storage systems. As a result, organizations are confronted with several challenges, including data silos, migration complexities, backup and DR complexity, file system scalability limitations, and the administrative cost of deploying and managing multiple storage systems and file servers.

Dell Scalable File System (DSFS) - a key component in Dell's Fluid Data strategy - is a high-performance clustered file system that represents an exceptional level of innovation in distributed systems technology. DSFS eliminates the scalability limitations associated with traditional file systems, supports both scale-out performance and scale-up capacity expansion, and provides a single global namespace for easy administration. DSFS presents a storage pool as a single file system with a single virtual-IP address to the storage client(s).


Its flexible architecture enables DSFS to deliver tremendous price-performance advantages for a wide variety of applications, from home-shares in a small enterprise, to high performance computing, to the world's largest scale-out data repositories. DSFS offers an optimal combination of performance and scalability, making it an excellent choice for a wide range of use cases and deployment environments, including file-intensive user shares, highly available NAS and unified storage for smaller scale environments, and virtual server environments with extensive NFS data.

DSFS is being incorporated in a number of Dell products that serve the needs of large and small organizations. Dell's first DSFS-based product, the PowerVault NX3500 Unified Storage Platform, provides an easy-to-manage solution for both file and block-based applications in small to midsize deployments. The PowerVault NX3500 delivers advanced clustered file system technologies without the complexity or expense of traditional advanced clustered platforms.

High Availability Scale-out Cluster Architecture

A DSFS cluster typically consists of DSFS nodes (deployed in pairs) and the underlying storage arrays. DSFS nodes consist of x86 commodity hardware, offering customers Dell's economies-of-scale as the leading vendor of server hardware to Enterprise customers. DSFS stores file data on conventional storage arrays, from JBOD to full SAN configurations, on multiple platforms, enabling customers to balance price/performance and reliability to meet their requirements. In a DSFS cluster, any single node can fail without affecting data availability or causing data loss, even if write operations were in-flight.

Cross-cluster reliability is achieved through a variety of mechanisms including a high speed cluster interconnect, write cache mirroring, failsafe journaling, and data integrity checks to insure data store consistency. A DSFS Fault Management System (FMS) uses a variety of mechanisms for insuring cluster consistency, and, in the event any node exhibits anomalous behavior, FMS will prevent split-brain scenarios that might disrupt system operations or integrity. Other services monitor the health of the server platform, including temperature and power condition. In some implementations, including the PowerVault NX3500, cache mirroring may be supplemented by internal or external battery backup that supports cache dumping to disk in the event of node loss.



DSFS System Architecture

Scale-out Parallel Performance

Because all nodes in a DSFS clustered system support active I/O, customers benefit from high intrinsic performance without exotic protocols or the need to distribute application load across multiple filers. Network traffic is load balanced across the cluster. Load balancing is implemented natively within DSFS, and customers may optionally deploy network-based load balancing techniques.

Storage platforms using DSFS are also load- and capacity-balanced in the back end. For example, write traffic is load balanced across LUNs, and capacity is monitored to insure even balancing across LUNs.

Scaling Performance and Capacity On Demand

In a scale-out implementation, customers may add nodes to transparently increase performance on-demand. In a scale-on-demand model, customers don't need to provision excess capacity in anticipation of future growth, which makes DSFS-based storage platforms ideal for customers that expect rapid growth over time or phased consolidation of applications. DSFS-based products support the transparent addition of capacity. Additional arrays can be added to a cluster, and those LUNs will be seamlessly mapped into the DSDS virtual volume without downtime.

2. Empowering the Future with DSFS

High Utilization of Storage Hardware

The design of Dell Scalable File System separates users' data and access from the underlying hardware configuration so that servers, CPUs, cache memory and disk drives are optimally utilized. As data gets written to the virtual server, it is distributed across internal file servers, and eventually to all disks connected to the storage cluster.

Share Files Seamlessly Among Heterogeneous Clients

With different client platforms accessing the same shared file system, DSFS provides fully interoperable multi-protocol file sharing for UNIX, Linux, and Windows clients using standard CIFS and NFS file access protocols and authentication methods (NIS, AD, LDAP). Multi-protocol support ensures that a single file share can support both CIFS and NFS access. DSFS includes support for mapping users and permissions among users, and for authentication domains.

Perform On-Demand Virtual Storage Provisioning

DSFS volumes enable administrators to provision storage as needed, so that capacity can be allocated independently of physical storage configuration. The large pool of storage (i.e., the aggregate of all disks seen by the nodes) is split into smaller, virtual containers, each providing administrators with a full set of policies (security style, quotas, snapshots, and alerting).

View User-accessible Point-In-Time Snapshots

Snapshots provide the first level of data protection by letting both end-users and administrators recover data instantly from an unlimited number of recovery points. DSFS incorporates redirect-on-write snapshots instead of copy-on-write solutions typical of other file systems. This approach requires only one I/O operation and delivers higher write performance. Each NAS volume has its own snapshot policy, wherein point-in-time "frozen" images are available to the end user as read-only views. DSFS creates a point-in-time backup of the data and applies volume-based policies that can be set for each volume.

Protect Your Data with Flexible Backup/Replication

DSFS allows fast and reliable replication of any number of volumes to a peer (local or remote) site, enabling continuous access to data in the case of a disaster or site failure to assure business continuity. User-accessible snapshots, snapshot-based asynchronous replication, NDMP backup and anti-virus support help protect all data stored in the system. DSFS also features seamless Active Directory integration and support for native Windows access control lists (ACLs) in addition to NFS User Identifiers (UIDs) and Group Identifiers (GIDs).

Restore Large Volumes of Data Quickly

DSFS provides the ability to restore very large data sets that need to be recovered as a whole to a particular point in time. It helps administrators to recover large data sets (terabyte scale) easily, eliminating long file copies and the need for free space for the recovery process. It gives the ability for client users to restore previous versions of files directly without contacting IT.

Simple and Easy Management

Managing terabytes of NAS storage is simple with the administrative functions supported by DSFS. From installation and initial configuration to ongoing monitoring and storage operations, all functionality is provided via easy-to-use screens and wizards. The management interface can easily be used to set up policies, quotas, snapshots, and replication and engage NDMP backup.

A DSFS cluster is managed as a single NAS system regardless of how many nodes are in the cluster. After DSFS nodes are discovered and added to the cluster, there are no node-specific operations for the administrator. Volumes are virtual entities that span underlying LUNs and provide a context for setting policies related to access, quotas, snapshots, replication, etc. When new LUNs are added, the NAS volume grows accordingly.

A DSFS cluster is accessed through a single VIP or network name, which means that as storage scales, customers don't need to worry about managing multiple mounts, balancing data across them, and redesigning applications to accommodate a fragmented namespace.



3. Dell Scalable File System Features

- **Client load balancing** - Load balancing sends client requests automatically to the node with the least-current workload. It also helps with complex network topologies and storage administration, and makes migration to failover nodes transparent.
- **NFS and CIFS protocol support** - Multi-protocol support enables clients to access a single file share using either protocol for file system sharing services. DSFS supports CIFS and NFS.
- **Distributed file system** - The distributed file system presents all the storage in the deployment as a single file system with a single IP address to the client(s). Multiple nodes can access the file system concurrently
- **Cache management** - DSFS uses patented, highly optimized caching technology across multiple storage controllers. Writing or modifying files first occurs in the cache. Next, data is mirrored to the peer controller's cache to ensure that all transactions are duplicated and secured while eliminating hot spots and bottlenecks.

- **Continuous data protection (CDP)** - Continuous data protection is applied to metadata to help with versioning and prevent corruption in the storage system. All metadata is check summed and written, and subsequent changes are also retained. CDP also helps optimize file system performance and efficiency.
- **Clustered, high availability architecture** - An active-active software architecture supports redundant controller pairs in a clustered configuration. This supports real-time maintenance and helps ensure continuous hardware health monitoring and no disruption to data access.

4. Dell PowerVault NX3500 - Unified Storage Solution with DSFS

The PowerVault NX3500 is the first in a series of products based upon Dell Scalable File System that delivers enterprise class file services to Microsoft Windows and Linux clients. It works with PowerVault MD32x0i and MD36x0i storage arrays to help provide affordable unified storage with iSCSI, CIFS and NFS access to block and file data.

Customers can use the PowerVault NX3500 to consolidate user data as well as other file and block applications into a single, easy-to-manage unified storage system with best-of-breed data management and scaling capabilities. Its scale-up architecture delivers a flexible, load-balanced pool of high performance storage, making it easy to grow capacity up while avoiding the scalability constraints and challenges of managing separate block and file systems. With its dual active-active file controllers and backup power supply, the PowerVault NX3500 gives you data protection and excellent performance with no single point of failure. Find out more about the NX3500 at Dell.com/NX3500.



5. PowerVault NX3500 and Dell Scalable File System Technical Specifications

Specification	Max Value (2-node system)
Max system size	192 TB
Max file size	4 TB
Max files	~32 billion
Number of directories	~34 billion
Max NAS volumes	512
Max snapshots per volume	512
Max snapshots per NX3500 system	10,000
Cache size per NX3500 system with 2 quad-core CPUs	24 GB/12 GB per controller
Max LUNs	16
File name length	255 bytes
Max NFS mounts	1024
Max CIFS shares	1024
Max CIFS client concurrent connections	200
Max local users per NX3500 system	300
Max Quota rules per NX3500 system (user quotas)	65,536
Max quota rules per volume	256
Max block level replication policies	256
Max directory depth	1,024