# Dell Compellent FS8600

**Dell Compellent Best Practices**

August 2012

# Contents

## Tables



## Figures

## Document Revision

**Table 1. Revision History**

| Date | Revision | Description |
|---|---|---|
| June 7, 2012 | A | Initial draft. |
| June 26, 2012 | B | Updated Replication Section |
| August 06, 2012 | C | Updated NDMP Section. Replacing 3-way with “Remote NDMP” terminology. Sync’d Switch Topology verbiage on pages 7 and 9. |

# Preface

## Audience

The audience for this document is intended to be systems, networking, storage or backup administrators who are responsible for the day-to-day management responsibilities of a Dell Compellent FS8600 environment. Proper management of a FS8600 requires administrators or teams of administrators capable of managing and configuring enterprise-grade Fibre Channel SAN and Ethernet networks, any enterprise-grade backup software intended to be used, the Dell Compellent Storage Center itself as well as general purpose NAS administration.

## Purpose

The purpose of this document is to cover specific implementation concepts or specifics related to the management of a Dell Compellent FS8600 environment. It is not intended to be a primer or introductory resource for any of the subject matters involved, and it assumes at least introductory knowledge of many of the subjects covered in this document. This document should be used in conjunction with other Dell Compellent resources, such as the Dell Compellent Storage Center Connectivity Guide, FS8600 Admin Guide and Hardware Manual, Enterprise Manager 6 User Guide, or any other available documentation resources.

## Customer Support

Dell Compellent provides live support 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

# Introduction

As unstructured file data continues to grow at challenging rates, IT file server infrastructure supporting business and mission critical systems can easily become strained and fragmented. This can manifest in a variety of painful forms such as rising management costs, lost storage efficiencies, complicated or fickle backup architectures and challenges to overall system availability.

The Dell Compellent FS8600 is uniquely suited to overcome the challenges presented by explosive file data growth and file sprawl. With the combination of the Dell Fluid File System's (FluidFS) scale-out single namespace and Compellent Storage Center's Data Progression and Dynamic Capacity, IT teams can reduce management overhead as well as leverage the platform's unique storage efficiencies and cost of ownership attributes while providing class-leading availability and data protection.

It should be noted that the FS8600 is a member of the Dell FluidFS NAS product family. It is in many ways conceptually similar to the FluidFS products available on the Dell PowerVault product line as well as the Dell EqualLogic product line. While many of the concepts, particularly front-end NAS concepts are similar or even identical, other subjects, such as management of interaction with backend storage, may be radically different. Relevant resources for each FluidFS platform should be leveraged for the appropriate product.

# Connectivity

The Dell Compellent FS8600 leverages industry-standard connectivity protocols and concepts, namely Ethernet and Fibre Channel connectivity, and the proper implementation and configuration of the systems involved is a paramount concern. Particularly in large-scale or otherwise complex deployments, considerable care and planning of the connectivity systems involved is required.

An FS8600 environment is dependent on several conceptually isolated networks, each of which may contain other subdivisions of logically independent networks. While each logically independent network may meet its own requirements, each network can have a substantial impact on the requirements or overall behavior of the other networks involved. Therefore, each individual network as well as the network configuration as a whole must be carefully understood.

## Ethernet connectivity

FS8600 Ethernet connectivity is the connectivity medium that is used to carry client or interconnect traffic. Client network connectivity is used to expose file resources, backup protocols and management interfaces to users, administrators and other systems. Interconnect Ethernet connectivity is used in FS8600 clusters of two or more FS8600 appliances to carry cluster traffic between NAS controllers.

As stated in the “Purpose” section of this document, it is assumed that the reader has already consulted the FS8600 Admin Guide and Hardware Manual.

For environments with networks living on multiple switch vendors/models (Cisco Catalyst<->Brocade Foundry), and particularly environments adopting a new vendor switch as part of the FS8600 deployment (for example Dell Force10 into an established Cisco network), all network features should be thoroughly tested before entering the deployment into production.

***Monitoring***

It is also strongly suggested that a system for monitoring and trending all ports on all switches in an environment be preemptively set up prior to the production implementation of an FS8600 environment, as this can be a substantial resource in detecting potential bottlenecks or protecting against future performance or availability issues.

Each FS8600 interface can also be monitored through both SNMP traps and SNMP polling after SNMP has been configured through either the “Start Configuration Wizard” or through the appropriate configuring interfaces in the web interface or CLI. The interfaces are exposed to SNMP polling through the standard net-snmp OIDs.

***Client Connectivity***

The FS8600 client network connectivity is responsible for exposing and carrying NAS resource traffic to clients and other systems. In any FS8600 deployment, the client connectivity Ethernet ports are reserved to the left most network interfaces on each Ethernet card when observing the rear of the FS8600 appliance.

***Separate Networks/Subnets***

The FS8600 is capable of supporting or living on large numbers of separate or distinct networks, also referred to as subnets. This can be beneficial for reasons such as consolidating established or legacy file servers, as well as following best practices for providing dedicated resources to specific client groups or environments. NDMP backups should also be isolated to a dedicated network, and it is also suggested that where possible, networks should be isolated in a one to one fashion with VLANs (see below).

The scale-out architecture of FluidFS ensures that each host or client that accesses NAS resources is evenly distributed across all interfaces in an FS8600 cluster. Generally, this means only one IP address is needed per network. The FluidFS platform balances clients across interfaces in the following way:

Client-A tries to access NAS resources with the DNS provided hostname of cl0-fs8600.local, which is a DNS A record to 10.10.30.5. Not having any previous network information for 10.10.30.5, Client-A will issue an ARP request on its available interfaces. Seeing the request, the FS8600 will respond saying it owns 10.10.30.5 through its first NIC, and the client and FS8600 will proceed to communicate over that interface. Client-B then tries to access NAS resources through the cl0-fs8600.local hostname, and following DNS, similarly reaches out to the network with an ARP request for 10.10.30.5. In order to balance the connection from Client-B, the FS8600 will respond to Client-B through its second NIC, and communication between Client-B and the FS8600 will continue through that interface. Going forward, the FS8600 will continue to answer new ARP requests through all of the interfaces available in the cluster.

There are reasons why multiple IP addresses per network may be needed, such as supporting IP addresses or legacy hostnames from consolidated legacy file servers. In some environments, there may

also be performance reasons for maintaining multiple IP addresses per network. Take, for example, an NFS client that accesses NFS exports “cl0-fs8600.local:/vol1/mail cl0-fs8600:/vol1/htdocs/ and cl0-fs8600:/export/ora1”. The first time the client requests NAS resources (for example for /vol1/mail), it will reach out with an ARP, just as any other client. However, the second time it requests NAS resources (for example /vol1/htdocs), the client already knows which MAC address it needs to reach the NAS resources, and it will simply use that established lane of communication, meaning that going forward, the client will only communicate through one NIC on the FS8600. For high performance clients, this may not be ideal. In these cases, mechanisms such as round robin DNS or the use of a hosts file can specify multiple hostnames for multiple IPs, such that “cl0-fs8600-1” resolves to 10.10.30.5 and “cl0-fs8600-2” resolves to 10.10.30.6. Then each time a client attempts to access both “cl0-fs8600-1:/vol1/mail” and “cl0-fs8600-2:/export/ora1”, the client will ARP for each resource, and hence be distributed across NICs for each resource.

Similarly, some network devices, such as VPNs or routers, may cache ARP/MAC addresses in such a way as to only distribute traffic for a destination IP through one FS8600 NIC. For resources that are heavily accessed through such devices, mechanisms such as round robin DNS should be used in order for clients to request multiple destinations through the VPN or router. This will balance traffic across the FS8600 NICs.

When adding subnets, it is suggested to set aside sufficient IP address per subnet to accommodate the addition of more FS8600 appliances as part of the cluster at a later point in time.

It should be noted that any client connectivity, including NDMP, should always target a client VIP, and client traffic should never target the management IP.

***VLANs***

The FS8600 is VLAN-aware, meaning it is capable of understanding and communicating with VLAN tagged Ethernet frames, allowing it to address multiple networks across multiple VLANs at a time. The management network is not VLAN-aware and must be an untagged VLAN network to the FS8600 NICs. Similarly, any installation of Enterprise Manager must be in a VLAN capable of reaching the FS8600 management network. Similarly, the network gateway for the FS8600 must be reachable, preferably through the management/untagged network, because reaching the gateway is used as part of the FluidFS cluster quorum logic.

It is generally suggested to use and isolate networks to a VLAN and network per purpose where possible (for example, a certain data is only accessed by isolated group of clients). This is also true for functions such as replication or NDMP backup traffic.

The management network cannot be VLAN-isolated, and it requires that the switch ports “untag” the VLAN associated with the management network to the FS8600 interfaces.

***MTU***

The FS8600 is Ethernet Jumbo Frames (MTU 9000+) capable, and environments can expect a degree a performance improvement, particularly where throughput is concerned, by enabling Jumbo Frames on the FS8600 platform as well as the client connectivity network topology and clients. It must be noted that the MTU setting for client connectivity on the FS8600 is a global variable for all networks, so environments with mixed MTU topologies need to determine whether they can enable Jumbo Frames on the FS8600 and associated network infrastructure without impacting the non-Jumbo Frames network infrastructure. It is also worth noting that not all vendors agree on the exact definition of Jumbo Frames, and as such, documentation from each vendor should be consulted.

### *Spanning Tree*

The FS8600 requires that spanning tree on all switch ports connected to client connectivity ports must be disabled. For many network vendor implementations, this will be called “portfast” or “edge” mode within the switch configuration semantics.

For networks with sprawled network topologies or otherwise complicated/collapsed/converged topologies, it is strongly recommended that a modern implementation of rapid spanning tree protocol, multiple spanning tree protocol or other similar technology be deployed, as the impact to client connectivity in the event of a failure will be determined by network convergence delays. In many environments, the time required for spanning tree to detect and react to a topology change will be the primary source of observed time for a client to recover from a network or NAS controller failure.

For networks that make use of VLANs, it is suggested that the spanning tree implementation support a mechanism for an instance of spanning tree per VLAN.

### *Load Balancing and ALB vs. LACP*

There are few or no technical or rational reasons why LACP should be chosen over ALB, and ALB should be the assumed connectivity mode for all FS8600 environments.

FluidFS actively monitors client connectivity across the cluster and maintains a policy list for what NAS protocols to migrate in the event that the cluster experiences a failover event that causes lopsided client connectivity across the cluster. It is strongly suggested that the default policies be left intact unless a known client connectivity issue requires a change.

### *Flow Control*

It is recommended but not required that all switch ports connected to the FS8600 client side interfaces have flow control enabled. Outside of the switch ports directly connected to the FS8600, Flow Control can be enabled or disabled as required.

Note this in contrast to the Dell Compellent Storage Center, for which it is strongly suggested to disable flow control on switch ports connected to iSCSI I/O cards. For more information regarding Flow Control and Storage Center, consult any Co-Pilot Support Tech Alerts for the Storage Center configuration in question.

### *Switch Topology*

An FS8600 appliance or cluster can only be as highly available or high-performance as the switch infrastructure supporting it. Architecturally, three guiding principles can be used to construct a best practice switch connectivity topology:

- Avoid any single points of failure
- Ensure sufficient inter-switch throughput
- Make every client connectivity port in an FS8600 cluster available to any potential client

It is also important to remember that the client network is responsible for all traffic of exposed NAS resources, including general client connectivity, management, replication and backups. Accordingly, the throughput and impact of backup traffic between the FS8600 and the backup server should be accounted for.

The FS8600 and the FluidFS architecture are fundamentally Ethernet/Switch topology agnostic as long as the topology meets the stated best practice requirements. This means that technologies such as switch stacking or other vendor proprietary switch trunking technologies are valid topologies, so long as they behave as open standards Ethernet to the FS8600 interfaces.

### *Understanding failure scenarios*

This section assumes that an FS8600 is being supported by a correctly configured highly available switch topology. In a correctly configured, highly available environment, a failure of any individual system, including an FS8600 controller, should be able to be survived while maintaining data availability.

In the event of a switch failure, clients who had a connected session to an FS8600 interface connected through that switch will begin their TCP/IP timeout/ARP cache counter and expire the now invalid ARP cache for that FS8600. At that point, the client will issue another ARP request, which will be answered to by one of the FS8600 interfaces available through another switch. In these scenarios, the performance of the spanning tree implementation in question is likely to be the primary concern for recovery times.

In the event of an FS8600 controller failure, the surviving controller will issue a gratuitous arp to the clients of the failed controller. This gratuitous arp instructs those clients to renegotiate their arp cache, which will be answered by one of the interfaces on another controller in the cluster.

It should be noted that clients who lose connectivity to an FS8600 controller during the middle of a CIFs data operation will error out and will have to re-establish a CIFs session with the surviving controller.

### *Interconnect Connectivity*

The FS8600 is dependent on the ability of each NAS controller in an FS8600 cluster to communicate through a dedicated network, called the interconnect network, to any other controller in the cluster. The interconnect network is used both for inter-node communication to satisfy client I/O across FS8600 appliances, as well as cluster quorum and availability communication. The interconnect network reserves the rightmost ports on a NAS controller.

An example of a client side I/0 that generates interconnect traffic would be a client read request that is best resolved by querying another controller's cache or file system allocation than satisfying the I/O with an expensive operation to disk. This type of behavior can cause the overall FluidFS architecture to be relatively "chatty" over its interconnect network, particularly in high-performance environments.

For single appliance deployments, the interconnect ports on each controller are cross connected directly to each other, bypassing any switch. For deployments with multiple FS8600 appliances in a single NAS cluster, the interconnect ports must be connected to a switch infrastructure. For more information, please consult the FS8600 Hardware Manual.

### *VLANs*

The interconnect network must live on a dedicated VLAN that is isolated from all other traffic and solely purposed for an individual cluster's interconnect traffic. Similarly to the management network, this VLAN must be untagged on all switch ports connected to FS8600 interconnect connectivity ports.

***MTU***
All switch ports connected to FS8600 interconnect connectivity ports must be enabled for Jumbo Frames (MTU equal to or greater than 9000).

***Spanning Tree***
Spanning tree must be disabled for all FS8600 interconnect ports. For many network vendor implementations, this will be called “portfast” or “edge” in the switch configuration semantics.

***Flow Control***
The FS8600 requires that all switch ports connected to the FS8600 interconnect interfaces have flow control enabled.

Note this is in contrast to the Dell Compellent Storage Center, for which it is strongly suggested to disable flow control on switch ports connected to iSCSI I/O cards. For more information regarding Flow Control and Storage Center, consult any Co-Pilot Support Tech Alerts for the Storage Center configuration in question.

***Switch Topology***
An FS8600 appliance or cluster can only be as highly available or high-performance as the switch infrastructure supporting it. Architecturally, three guiding principles can be used to construct a best practice switch connectivity topology:

- Avoid any single points of failure
- Ensure sufficient inter-switch throughput
- Make every client connectivity port in an FS8600 cluster available to any potential client

Because of the chatty nature of the Interconnect network, throughput between switches responsible for the Interconnect network may require sufficient inter-switch throughput as to match the aggregate client connectivity throughput.

Through the use of VLANs, the interconnect network can reside on the same physical switch infrastructure as the client connectivity network. The interconnect network can also reside on dedicated switch infrastructure solely purposed for the interconnect network, which may be ideal for larger or higher performance deployments, so long as all controller interconnect ports are be able to reach all other controller interconnect ports.

## Fibre Channel Connectivity
The FS8600 is unique in the FluidFS family for its usage of Fibre Channel based connectivity to provide access to its backend supporting storage, a Dell Compellent Storage Center or in some cases two Storage Centers. Similarly to the front end Ethernet connectivity, the FS8600 is only as highly available and high performance as the supporting backend SAN switch environment and topology.

Architecturally, the following guiding principles can be used to construct a best practice Fibre Channel connectivity topology:

- Avoid any single points of failure

- Ensure sufficient ISL throughput
- Make all NAS controller Fibre Channel ports able to reach all Storage Center ports per fault domain
- Avoid or minimize impact to other hosts in the SAN
- Adhere to established Dell Compellent Storage Center best practices such as those outlined in the Dell Compellent Storage Center Connectivity Guide.

***Logically isolated fabrics***

Not all environments or deployments warrant the implementation of a Fibre Channel infrastructure that avoids any single point of failure for the SAN. For these environments, it is suggested that a virtual fabric technology (Cisco's VSANs or Brocade's logical switch features) be implemented in order to logically isolate port and zoneset configurations into two separate logical fabrics. While this may introduce conceptual complexity, it provides the benefit of isolating one logical fabric from administrative errors in the other. This allows for the survival of the FS8600 environment even if there is a critical configuration error in one of the logical fabrics.

Each individual port in an FS8600 controller should be associated with a unique, logically isolated fabric, and the same logically isolated fabrics should be used for each FS8600 controller in a cluster.

***Storage Center Fault Domains***

As stated previously, it is strongly suggested that the relevant Dell Compellent Storage Center documentation be consulted when implementing a Fibre Channel architecture that involves the FS8600.

A Storage Center fault domain should be created per logically isolated fabric, physical or virtual. In the event that a logically isolated fabric is being introduced to an established architecture as part of an FS8600 implementation, the Storage Center connectivity into that fabric and its relevant Fibre Channel fault domain configuration should be adjusted accordingly.

It should be noted that the Storage Center logically links Fibre Channel over Ethernet with Fibre Channel fault domains. As the FS8600 is not currently supported on FCoE-based connectivity, the proper configuration of FCoE vs. FC fault domains needs to be considered.

***Legacy vs. Virtual Port Mode***

As stated previously, it is strongly suggested that the relevant Dell Compellent Storage Center documentation be consulted when implementing a Fibre Channel architecture that involves the FS8600.

The Dell Compellent Storage Center has two modes of front-end Fibre Channel connectivity configuration - Legacy Port Mode and Virtual Port Mode, each of which has its own rationale for implementation. Both Legacy Port Mode and Virtual Port Mode are supported by the FS8600. Generally speaking, the primary concern that drives any Storage Center to Legacy Port Mode is the existence of legacy HP-UX and Solaris systems or current AIX systems. Legacy Port Mode can become conceptually complex very quickly, and in-depth coverage of Storage Center Fibre Channel connectivity falls outside the scope of this document.

It should be noted that converting from Legacy Port Mode to Virtual Port Mode is a "one way" move. While a Storage Center in Legacy Port Mode can be converted to Virtual Port Mode, a Storage Center in Virtual port Mode cannot be converted back to Legacy Port Mode.

***Zoning***

Traditionally, Dell Compellent recommends implementing WWN-based zoning policies for Storage Center SANs but also supports port-base zoning out of recognition that some customers may prefer this method. However, with the FS8600, the recommendation is reversed. Dell Compellent recommends zoning the FS8600 controllers using port-based zoning but also supports WWN-based zoning out of recognition that some customers may prefer this method. The rationale for suggesting port-based zoning is that in the rare event that an FS8600 controller or Fibre Channel card needs to be replaced, the use of port zoning avoids the need to change the relevant zone set.

For large SANs, it may be advantageous to group the front-end Storage Center Fibre Channel ports into a Fibre Channel alias on the fabric and create zones by using the FC Alias for the appropriate Storage Center. For Legacy Mode Storage Centers, both the Front-End Primary and Front-End Reserve can be included in a single zone.

A zone should be created per NAS controller port to the respective Storage Center front-end ports. No zone should be created that allows a NAS controller port to see any other NAS controller port. No more than one Storage Center should be included per zone. No zone should be created that allows a NAS controller port to see anything other than a Storage Center front-end port.

# Network Services

All FS8600 environments will be dependent on various network services that, while providing crucial services, also introduce complexity that needs to be accounted for.

## Domain Name System

***Round Robin DNS***

Route-aware devices (routers/gateways, VPNs etc) may cache IP to MAC Address destination relationships, which could cause uneven distribution of clients on a limited number of FS8600 interfaces. Because of this, some environments may benefit from the use of Round Robin DNS systems. For a pool of clients, this will iteratively alter the IP address that is returned to the clients, causing the clients to connect through the route-aware device to the FS8600 in a balanced way. When implementing a Round Robin DNS system, it is important to understand the caching relationships of all subsequent DNS systems between the Round Robin DNS server and the clients. For example, if ns0.local and ns1.local are configured for Round Robin DNS but clients use ns0.office.local as their primary DNS target, ns0.office.local could possibly poll ns0.local to satisfy the request but then cache the request and return the same value to the next client, hence removing the benefit of implementing Round Robin DNS. Some environments may see a benefit in keeping a short TTL value for the relevant DNS records for added overall flexibility.

When allocating IPs and creating A record pools for Round Robin DNS, there should be at a minimum as many IPs and associated A records as there are interfaces in the FS8600 cluster.

***Cluster Quorum***

In the absence of a reachable gateway, the ability to reach the configured DNS hosts is used as a cluster quorum voting mechanism. As such, DNS must be similarly highly available and reachable by the FS8600.

***Active Directory***
For FS8600 environments that will be integrated with Active Directory, all relevant DNS entries must be up to date and valid prior to attempting to bind the FS8600 into that Active Directory domain. The DNS servers configured on the FS8600 must include DNS servers responsible for the DNS of the Active Directory domain in question, and one of the configured DNS suffix entries configured on the FS8600 must be the Active Directory domain in question.

***Network Lock Manager***
The specific configuration for NFS NLM is highly platform dependent. However, as a rule of thumb, most client implementations are highly dependent on a valid alignment of:

- The client's perceived hostname for itself (this is because the hostname can be included as part of the NLM request)
- The matching of the client's perceived hostname to visible DNS forwards
- The matching of DNS forwards to the IP address of the client
- The matching of DNS forwards to proper DNS reverse lookups

In many environments and platforms, troubleshooting of NLM will be difficult or impossible without a valid end-to-end DNS configuration.

## Network Time Protocol
The FS8600 can use the NTP to poll time information from authoritative outside sources. Generally, it is recommended that a minimum of two sources be used. Three or more sources is the suggested ideal configuration.

For FS8600 environments that will be integrated with Active Directory, the NTP sources should be the Active Directory Domain Controllers for the domain in question.

***Cluster Quorum***
In the event that both the configured gateway and DNS hosts are unreachable, the ability to reach the configured NTP hosts is used as a cluster quorum voting mechanism. Because this mechanism is used in the event that DNS is unreachable, at least one of the time servers should be configured by IP as opposed to DNS.

## Client Identification Services

***LDAP/NIS***
As both LDAP and NIS are acknowledged to be unsecure authentication protocols, all NIS and LDAP traffic should be confined to secured or otherwise untrusted networks and VLANs. As NFSv3 is also understood to be unsecure, it may be possible that authentication traffic may be able to reside in the same network or VLAN as NFS traffic.

The FS8600 should be pointed to name services that are logically close to the FS8600 in the overall network topology. While the FS8600 will cache some name information for short periods of time, name lookups can introduce unexpected latency in performance or latency-sensitive environments.

As with networking, the FS8600 can only be as highly available as the systems on which it depends. As such, LDAP and NIS services should be configured in such a way as to be highly available.

Both LDAP and NIS are covered in more depth in the Dell Compellent FS8600: NFS File Server Consolidation Guide.

***Active Directory***
As detailed in the Network Time Protocol section, accurate time resolution is required for an FS8600 to participate as part of a cluster.

Active Directory is covered in more depth in the Dell Compellent FS8600: CIFs File Server Consolidation Guide.

# FS8600 User Access Settings

## NAS volumes

NAS volumes represent one of the primary building blocks in the creation of isolated data sets, and are isolated containers within the overall Fluid File System. NAS volumes should be created per data set requirements (for example, CIFs vs NFS data sets) and also to meet other requirements (for example, data protection RPO & RTO requirements as met by snapshot and replication). It should be noted that unlike some other platforms, NAS volumes are a purely logical level of abstraction within the single namespace of the Fluid File System, and there are no performance benefits of spreading data sets across multiple NAS volumes.

***Thin Provisioning***
While it is possible to oversubscribe the capacity of storage presented to a FS8600 cluster using Storage Center's Dynamic Capacity thin provisioning technology, it is not possible to oversubscribe space allocated from the FluidFS file system. The FluidFS file system can only allocate as much storage to NAS volumes as has been allocated from Storage Center to the FS8600.

It should be noted that, while all Storage Center volumes are thinly provisioned by default, as data is written over time to the FS8600, the storage allocated and consumed on the Storage Center for that data cannot be reclaimed. If the overall storage capacity of the Storage Center is oversubscribed, extreme care should be taken to monitor the available capacity in the Storage Center to ensure the system always has sufficient free space available.

Data sets with extremely high change rates that also leverage FluidFS snapshots can threaten to consume their allocated storage very quickly and hence should never be oversubscribed.

***Access time granularity***
The access time granularity setting governs how FluidFS will manage access time request updates to files. Traditionally, very few applications or use cases are dependent on file access time, which in and of itself generates a significant metadata performance burden on both a client and on the NAS platform. By default, tracking of access time is disabled for NAS volumes, but this can be changed both at NAS volume creation time as well as at any time in the future after the initial NAS volume creation.

***File access security style***
The file access security style setting controls which permission schemes are authoritative for that volume. It is generally suggested that data sets that will be exclusively shared through NFS or CIFs reside on different NAS volumes. For NAS volumes that will be exported via NFS, this setting should be set to "UNIX". For NAS volumes that will be shared via CIFs, it should be set to "NTFS". In the case that

a NAS volume may contain data sets that need to be exposed through both protocols, the setting can be set to “MIXED”.

NAS volumes in UNIX mode will always conform to POSIX style permissions for the data set. A NAS volume in NTFS mode will conform to the NTFS ACL scheme for a given data set. As a rule of thumb, a NAS volume in MIXED mode will enforce the permissions of a file from the perspective of the last protocol to manage the permissions of a file. Take, for example, a file that is stored in UNIX permissions scheme, but the permissions allow a CIFs user to edit and manage the permissions of the file. After the CIFs user has managed the file and the file permission data is committed from the CIFs connection, that file will now have the permissions scheme of an NTFS ACL, as was committed by the client. If configured, user mapping will allow for file permissions interaction between LDAP/NIS and Active Directory user data, allowing for the translation of permissions across protocols.

## Exports and Shares

Exports and Shares represent points of isolated visibility into the overall single namespace Fluid File System.

It is suggested that shares do not expose the “/” or root of any individual NAS volume. The “/” folder contains a folder called “.clusterConfig”, which is a stored backup of the cluster configuration and can be used to access backups of the cluster configuration through the file protocols. However, it also contains sensitive information that should not be exposed to unprivileged users.

### *NFS Exports*

Most classical Unix NFS export concepts can be carried over into an FS8600 environment. It is worth noting that exports can be created below export, and the same directory structure of the file system can be exposed through multiple exports.

### *Limit reported size*

The size available and announced through an NFS share will by default be the overall size of the NAS volume. For NAS volumes with many exports, this may be undesirable. Using the “Limit reported size” function under the advanced section of an NFS export, the size reported to the client can be arbitrarily limited at any time. Note that setting a limited size below what the export is already consuming can have disastrous consequences for clients already accessing that export.

### *Secure Port*

There are certain platforms with interpretations of the NFS RFC that cause them to request NFS mount negotiation over what are deemed insecure or unprivileged ports. In order to allow those platforms to access FS8600 exports, the Secure Port setting may need to be enabled. This should not have any impact on other platforms already accessing that export.

### *CIFS Shares*

CIFS shares names must be unique throughout the FS8600. For environments leveraging replication between FS8600s, shares must be unique for each FS8600 environment so that in the event that the NAS volume and associated shares need to be brought online, there is no conflict with established shares.

### *User-based directory tree*

For shares that are intended to provide user home folder data sets, FluidFS does have the capability to emulate a share per CIFS user authenticated against the FS8600 that is limited to that user/work station. For example, when user-based directory trees has been configured and Active Directory user

"bob" connects to the FS8600, the user will be presented with any available general shares, as well as with a share labeled "bob" that is visible to only that workstation.

**Figure 1. Example of User Based Share View for "bob"**

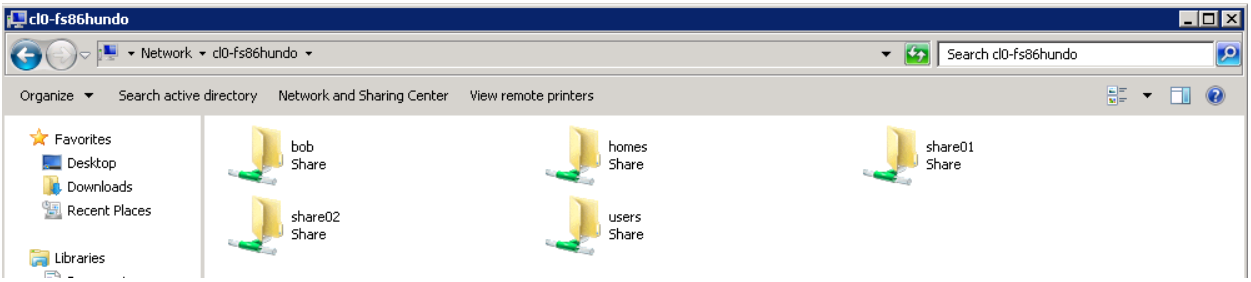


When the Active Directory user "jane" connects to the FS8600, that user will be presented with the "jane" share.

**Figure 2. Example of User Based Share View for "jane"**

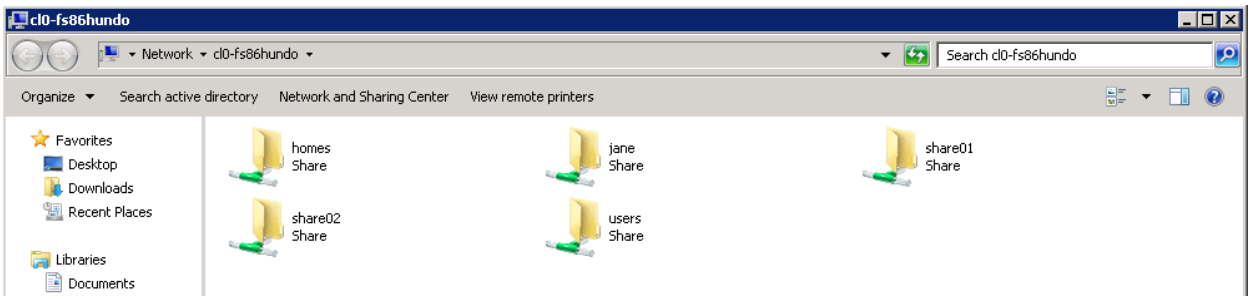


The folder for that user must exist in the file system before the share can be accessed. This is because shares represent points of visibility into the file system, not folders or file systems themselves. Consequently, some administrators may want to first create a dedicated "/users/" share within the NAS volume and modify the permissions of that share so that only administrators can manage the share or the contents of that share. This administrator-only share can also be provisioned as a hidden share (see below). The "homes" share can then be created using the user-based directory tree option, which will automatically live inside of "/users".

**Figure 3. CIFS Share View Example**

Show CIFS Shares for NAS Volume: All NAS Volumes

| Share | NAS Volume | Shared Directory | Accessible | Comment |
|---|---|---|---|---|
| homes | vol1 | /users/ (User) | | |
| share01 | vol1 | /share01 | | |
| share02 | vol1 | /share02 | | |
| users | vol1 | /users | | |

Directory is accessible  Directory is not accessible

As part of a new user creation process, an administrator can then proceed to create the directory in the user's share for that user.

**Figure 4. Example View Within the "users" Share**

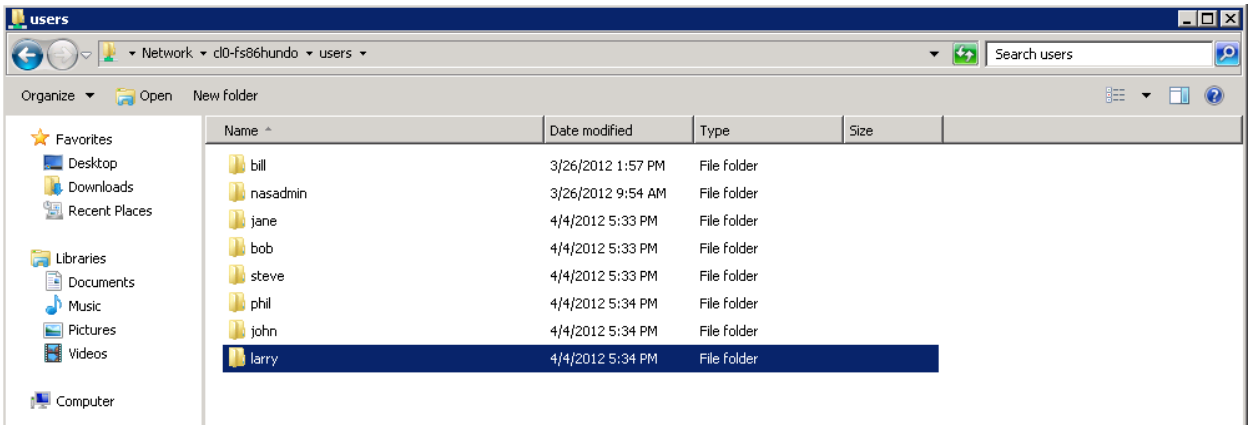


It should be noted that only one share per FS8600 cluster can have user-based directory trees enabled at a time.

***Hidden files and folders***

In the advanced configuration page for any CIFS share, the share can be configured to hide certain files or folders based on matching string names as configured. For NAS volumes that must expose the root of the NAS volume as a share, it is strongly suggested to add "/.clusterConfig/" to obfuscate the cluster configuration directory away from unprivileged users.

***Hidden shares***

When creating a share, a "$" character can be added to the end of the share name, which will hide or obfuscate away the existence of the share from end-user visibility.

**Figure 5. Establishing a Hidden Share**

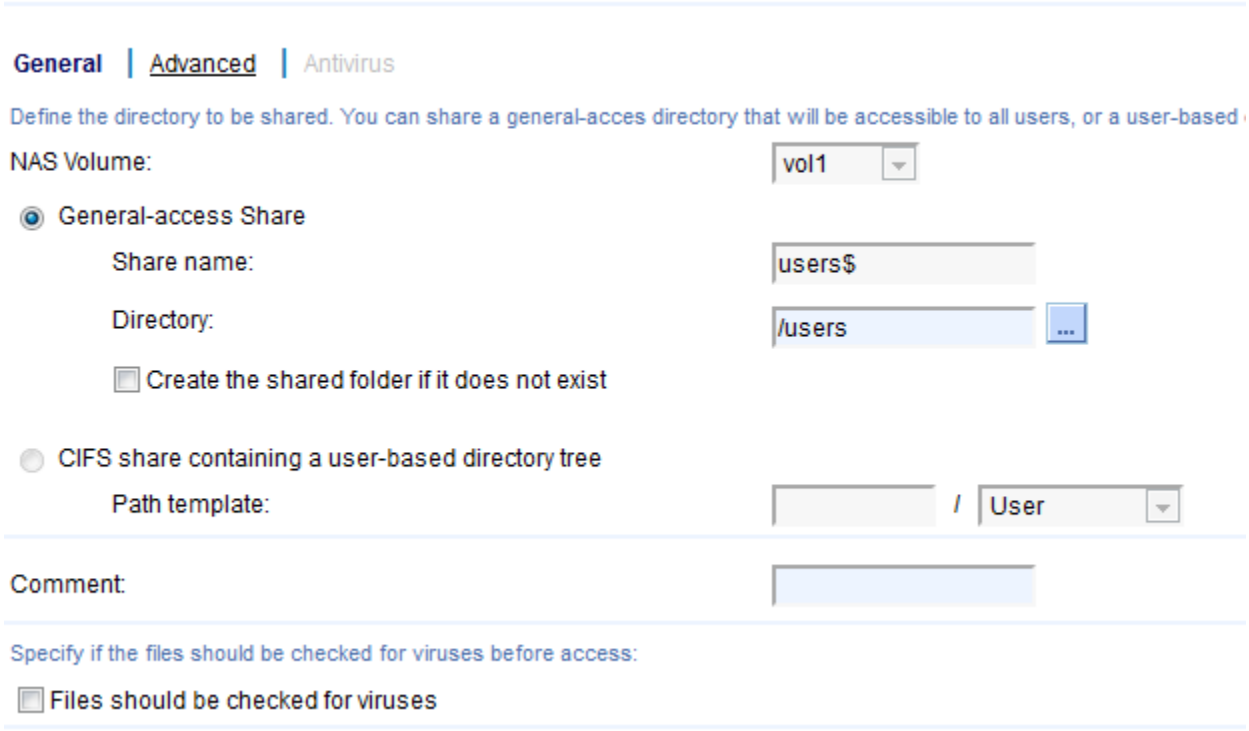


From an end-user perspective, the share will no longer be exposed by default.

**Figure 6. Hidden Share No Longer Visible to Users**

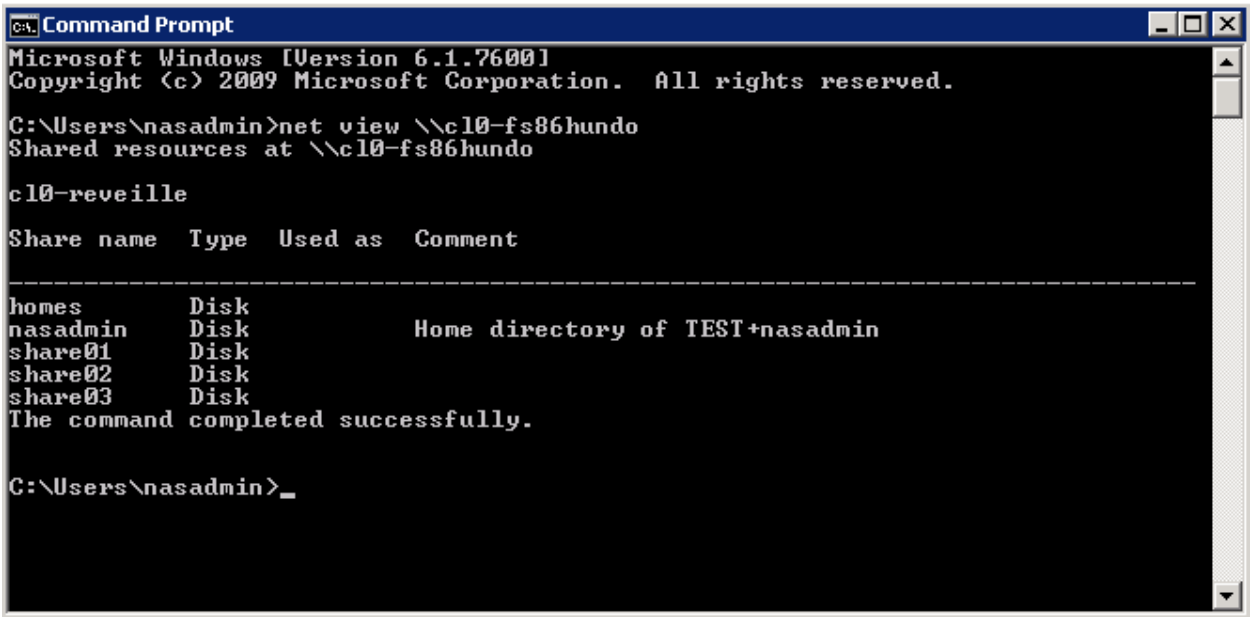


However, the share can still be accessed by entering the specific UNC for the share.

**Figure 7. Hidden Share Accessed Using UNC Path**

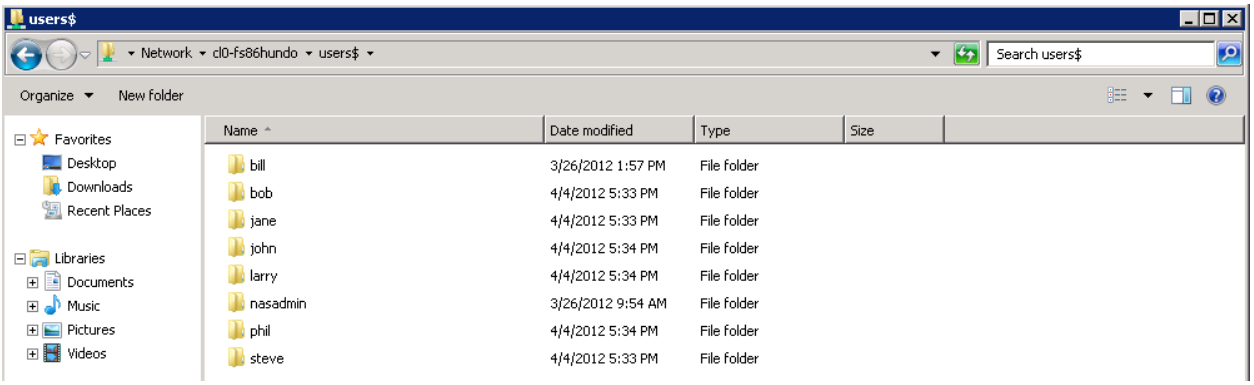


***Share creation and permissions***

When a share is created, it assumes the default permissions are inherited from the file system. Consequently, an administrator will need to manage the permissions of the share before it can be accessed by users. As part of the initial permissions construction, appropriate users/groups should be given access to the share by managing the shares permissions directly. Some clients may also require that the "owner" of the share be altered from the default to a user present within the Active Directory domain before users can traverse into the share.

Users in the Active Directory group "Domain Admins" are granted root-like permissions over every share. When first provisioning a share, a user in this group should be used to easily build share schemas for a share.

Individual users can also be granted "Full Access Rights" through the FluidFS CLI. Because an account that has been granted "Full Access Rights" does have super user privileges over all data sets, it is generally only suggested for accounts used for backups or other maintenance tasks, but it can also be useful for initial share permission schema creation.

Through the CLI, a user can be granted through:

system authentication full-access-account set DOMAIN+username

To print a list of users with Full Access Rights:

system authentication full-access-account view

To remove Full Access Rights from an account:

system authentication full-access-account delete

***Previous Versions***

Windows clients can access the "previous versions" user level file protection feature built into the Windows platform. The items listed in previous versions will return a list of versions as associated with FS8600 file system-level snapshots.

## Quotas

The security style selected for the volume will also determine which user database is used for quota assignments to user objects. NAS volumes configured for NTFS security will look to the configured Active Directory domain as well as to the locally configured accounts for user objects. UNIX will look to

the configured NIS or LDAP services as well as to the locally configured accounts for user objects. Finally, MIXED NAS volumes will look to both Active Directory and the configured NIS or LDAP database as well as to the locally configured accounts for user objects.

Each NAS volume can be configured for a maximum of 512 user group “rules”, where each rule can be a mechanism of applying a quote to a user or user group. Because of this, it is suggested that quotas first be applied using the “Default” quota option per volume, followed by assigning quotas per group, and only assigning quotas to specific users when an individual user has specific requirements that cannot be met otherwise.

Quotas must be assigned to both user objects in the case of mixed protocol data sets.

# FS8600 Data Protection

## Snapshots

The Fluid File System’s redirect-on-write behavior is a central concept of the architecture and provides the foundation for its snapshot functionality. FluidFS snapshots are a powerful and elegant mechanism for data protection. Snapshots introduce no performance overhead and no performance degradation over time. Only blocks of files that are changed are locked in an individual snapshot, and only the data required to maintain the historical delta between snapshots is kept on disk.

Because of the performance and ease of recovery of snapshots, they should be used as the front-line data protection mechanism. Data sets with a high churn rate can be protected with hourly snapshots. For longer-term data protection, such as nightly and weekly backups with long retention periods, slower backup solutions, such as NDMP backup to tape, should be leveraged.

When creating NAS volumes and shares, the data protection requirements of each individual file set should be identified, and the file set should be placed on a NAS volume with the appropriate snapshot policies. For example, mission-critical files with high churn rates can go on a NAS volume that is snapshotted every 30 minutes, whereas archival shares that only need to be snapshotted daily can go on a different NAS volume with a daily snapshot policy.

## Replication

The FluidFS platform offers snapshot-based asynchronous replication between “like-to-like” FS8600 clusters. “Like-to-like” implies that the number of appliances in each FS8600 cluster is identical between the replication source site and the destination. For example, if the source site has two FS8600 appliances in the cluster, the destination must also have two FS8600 appliances in the cluster. The back-end Storage Center array configurations can differ between the source site and the destination as long as there is sufficient capacity presented to the FS8600 for all NAS volumes being replicated.

When a replication is initiated between FS8600 clusters, a temporary snapshot of the current active data is created on the source NAS volume. This temporary snapshot is used to maintain a consistent data set throughout the replication and to encapsulate any active data that is not already managed by a previous snapshot. The replication will leverage the delta changes maintained by any interval snapshots that were taken since the previous replication so that only the data changed between replications is moved.

Snapshots on the source volume will be replicated in such a way as to be accessible on the destination volume, and if the destination volume is promoted to an active volume, the snapshots will appear as if they were local to the destination volume.

It is a strongly suggested best practice to dedicate a NAS volume at the destination FS8600 cluster to the purpose of being a container for each NAS volume on the source FS8600 cluster that is to be replicated. It is also strongly suggested that the destination NAS volume be configured with an identical NAS volume name. This is particularly important for NAS volumes with NFS exports, as the NAS volume name and file structure is part of the NFS export name.

Replication performance is highly dependent on the overall churn rate of the data set, which dictates the amount of data that needs to be replicated. Replication performance is also highly dependent on the performance of the back-end Storage Centers, the available performance capacity of the FS8600 clusters themselves and, most importantly, the available bandwidth between replication sources and destinations. The FS8600 is capable of replicating through one interface per controller.

For NAS volumes with active user quotas, it is strongly suggested that no user quotas be configured at the destination NAS volume. If user quotas must be enabled at the destination volume, there must be no conflicting quota rules at the conflicting destination.

Similarly, there should be no conflicting Windows to Unix User Mappings between the source and destination FS8600 clusters.

It should be noted that neither interaction nor replication between different FluidFS-based products is supported at this time.

## NDMP Backups

The FS8600 supports the industry standard NDMP protocol for backing up NAS data to external backup systems. The FS8600 can perform NDMP backups through “Remote NDMP” over Ethernet NDMP architecture. In a Remote NDMP architecture, the FS8600 appliance authenticates to the backup server solution using the NDMP protocol. Data is ingested from the FS8600 over Ethernet, and written to the backup media. The FS8600 uses the NDMP v4 protocol.

When an NDMP session is created on an FS8600 and a backup is initiated, a hidden snapshot is created of the active data set. This hidden snapshot is used to maintain consistency throughout the backup. When data transfer begins, all controllers in the FS8600 cluster participate in harvesting data from disk. However, the backup client will only ever be connected through one interface per FS8600 cluster. To stream NDMP data across multiple interfaces, individual data sets can be broken out into separate NDMP targets in the backup software, each target with a different IP address.

FluidFS is capable of Direct Access Restore (DAR) NDMP functionality, which can dramatically shorten restore times for large file sets backed up over NDMP. The best practices for each backup vendor’s DAR should be followed closely and thorough testing of individual file recovery from large file sets should be conducted, as a poor DAR implementation can have an overall negative impact on restore timelines.

It should be noted that each backup vendor is likely to suggest its own best practices and impose its own limitations for NDMP backups, and those best practices and limitations should be understood before implementing the overall NDMP architecture. As an example of these limitations, many backup vendors restrict NDMP restore functionality to an NDMP platform of the same type from which the

backup data was taken. Consequently, if data that was backed up from an FS8600 must be recovered without the availability of a target FS8600 to restore to, recovery of that data could be very complex.

NDMP restores from file data backed up from other FluidFS product lines are not supported.

# Storage Center best practices for the FS8600

From a purely technical perspective, a Dell Compellent Storage Center treats an FS8600 as a Fibre Channel-attached server, which means that many Storage Center features, such as Dynamic Capacity and Data Progression, are automatically applied to FS8600 data volumes. Similarly, many Storage Center best practices, such as monitoring data progression and management of alerts from Enterprise Manager, also apply to FS8600 environments.

## Data Progression

By default, expected Data Progression behavior is closely associated with Data Instant Replay technology, in that Data Progression will by default only move data to Tier 3 or RAID 5/6 that is marked as “read-only” in association with a given Storage Center Replay.

For FS8600 environments with minimal change rates, it is suggested to simply allow Storage Center to Replay FS8600 volumes with a scheduled Replay profile that maintains minimal historical Replays, allowing the default storage profile to automatically tier the data within those volumes according to Storage Center default policies.

Some environments with high change rates or otherwise more complicated requirements may need to create custom storage profiles for FS8600 volumes. It is strongly suggested that, unless specific requirements demand the need for altered storage profiles, the storage profile settings should be left to the defaults. Poor implementations of storage profiles can lead to significant performance and capacity issues.

## Data Instant Replay, Remote Instant Replay and Live Volume

Each FS8600 NAS controller has a write cache that is mirrored with its peer controller in the appliance. Because of this write cache, it is possible that at any given time, there may be a write operation that is in FS8600 cache but has not yet been placed to disk by the Storage Center. This means that Data Instant Replay and Remote Instant Replay are not valid mechanisms for data protection of FS8600 data, and similarly there are no current benefits from the Storage Center Live Volume feature for FS8600 volumes.

## Solid state disks

For Storage Centers with Solid State Disks (SSDs), it is a strongly suggested best practice to not allow FS8600 volumes to allocate blocks from SSD tiers, presumably Tier 1. This can be done through the use of storage profiles that only allow allocated storage from Tier 2/3. For systems with SSD and consequential NAS performance requirements, it should be noted that Tier 2 and Tier 3 should have sufficient performance capacity to meet those requirements without the Tier 1 SSD.

SSDs should be avoided because of the potential for high change rate FluidFS data to “crowd” out other volumes from the SSD tier in terms of overall available capacity in the tier. Additionally, the write cache present within the FS8600 NAS appliance accomplishes the majority of the functionality of the SSD tier.

### Managing FS8600 Storage Center Volumes

When an FS8600 is first installed, as part of the file system format procedure, Enterprise Manager will create a number of volumes according to internally determined best practices based on the size of the volumes, the number of FS8600 appliances, etc. It will also create a dedicated folder for FluidFS systems as well as the cluster volumes for that specific environment. It is strongly suggested to never move the volumes out of that folder. While there is no technical reason why they cannot be moved, as Storage Center volume folders are purely organizational, keeping the FS8600 volumes separated from other volumes minimizes their exposure to administrative accidents.

## VMware

It is Dell Compellent's strongly suggested best practice to use Dell Compellent Storage Center as the primary storage platform for VMware virtual machines and virtual machine-attached data. For Dell Compellent best practices with Storage Center and VMware, please consult the Dell Compellent Best Practices with VMware document and associated technical tips.

Should administrators want to use the FS8600 as a tertiary data repository for data such as commonly used ISOs, the established VMware best practices for ESX/ESXi over NFS apply to an FS8600 environment as well.

Should administrators want to expose the snapshot directory through the VMware client GUI, a simple symlink can be created that exposes the otherwise hidden .snapshots folder.

**Figure 8. Using FS8600 NFS Export as Datastore to VMware ESX/ESXi**

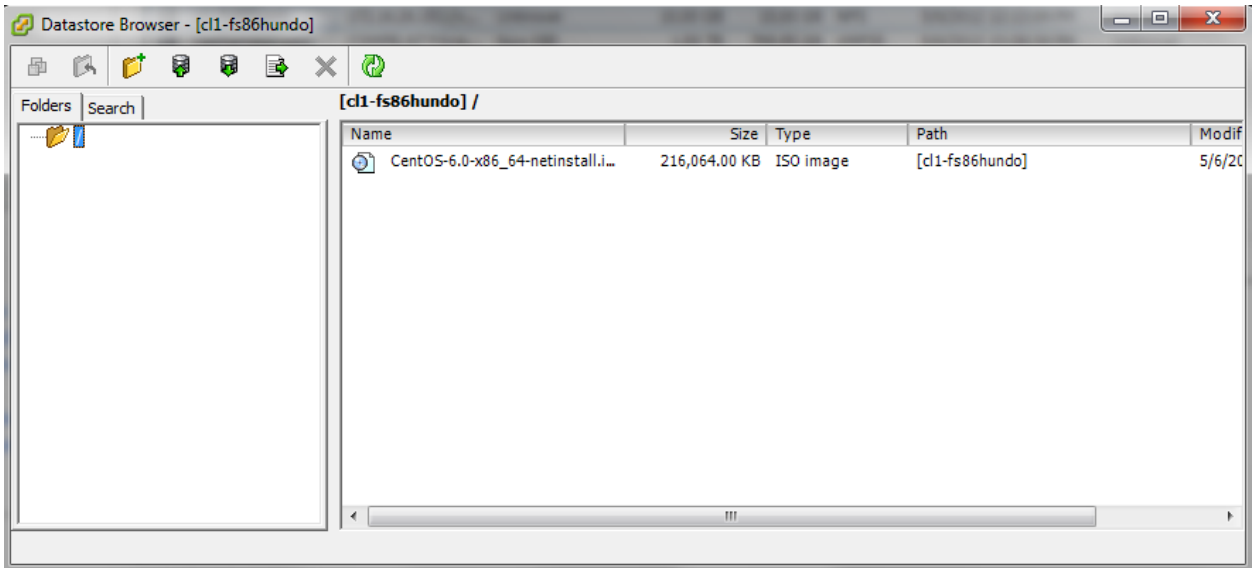

**Figure 9. Creating Symlink to Expose .snapthots**

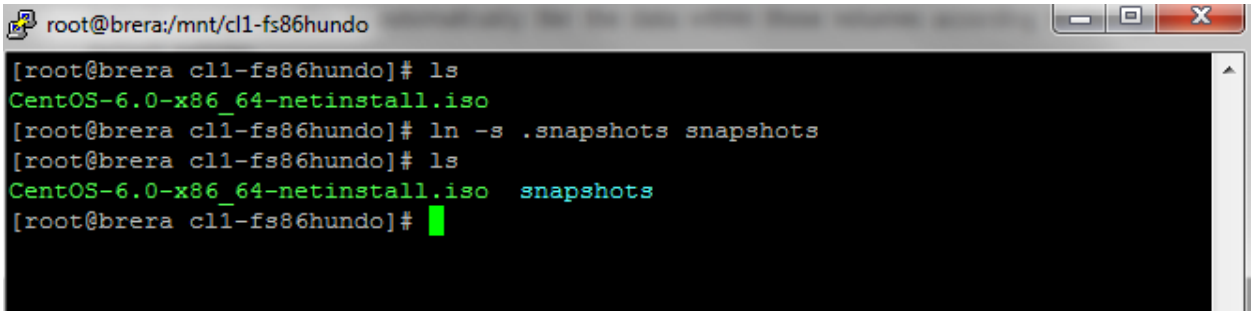


**Figure 10. View of the snapshots Symlink Created**

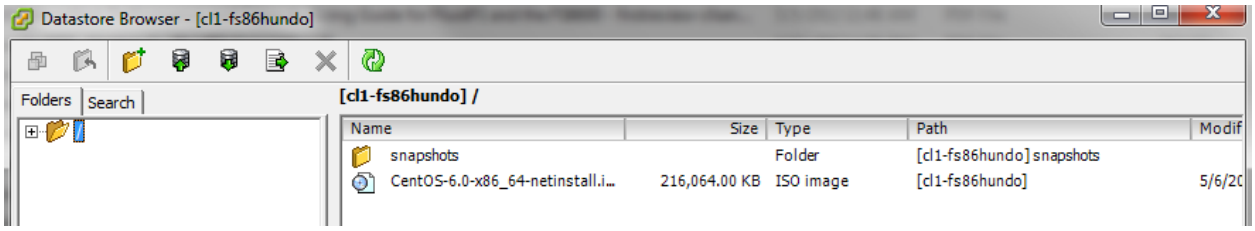


**Figure 11. Viewing Contents of the .snapshots Using the Symlink**