Dell Compellent FS8600

Network-Attached Storage (NAS)

Networking Best Practices Guide

Dell Compellent Technical Solutions Group
July 2013

Table of Contents

Document Revisions

Table 1: Document revisions

Date	Revision	Author	Comments
7/2/13	1.0	Mordi Shushan	Initial Release

1 Preface

1.1 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 an FS8600 requires administrators or teams of administrators capable of managing and configuring enterprise-class 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.

1.2 Purpose

The purpose of this document is to cover specific implementation concepts or specifics related to the Dell Compellent FS8600 networking environment. It is not intended to be a primer or Dell Compellent FS8600 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.

1.2.1 Disclaimer

The information contained within this best practices document is intended to provide general recommendations only. Actual configurations in customer environments may vary due to individual circumstances, budget constraints, service level agreements, applicable industry-specific regulations, or other factors. Configurations should be tested before implementing them in a production environment.

1.3 Customer Support

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

2 Introduction

2.1 FS8600 networks

The FS8600 utilizes three separate logical networks and two physical network segments to achieve maximum performance and security.

- Internal Network
  - Interconnect
  - Management
- Primary Network (Client Network)

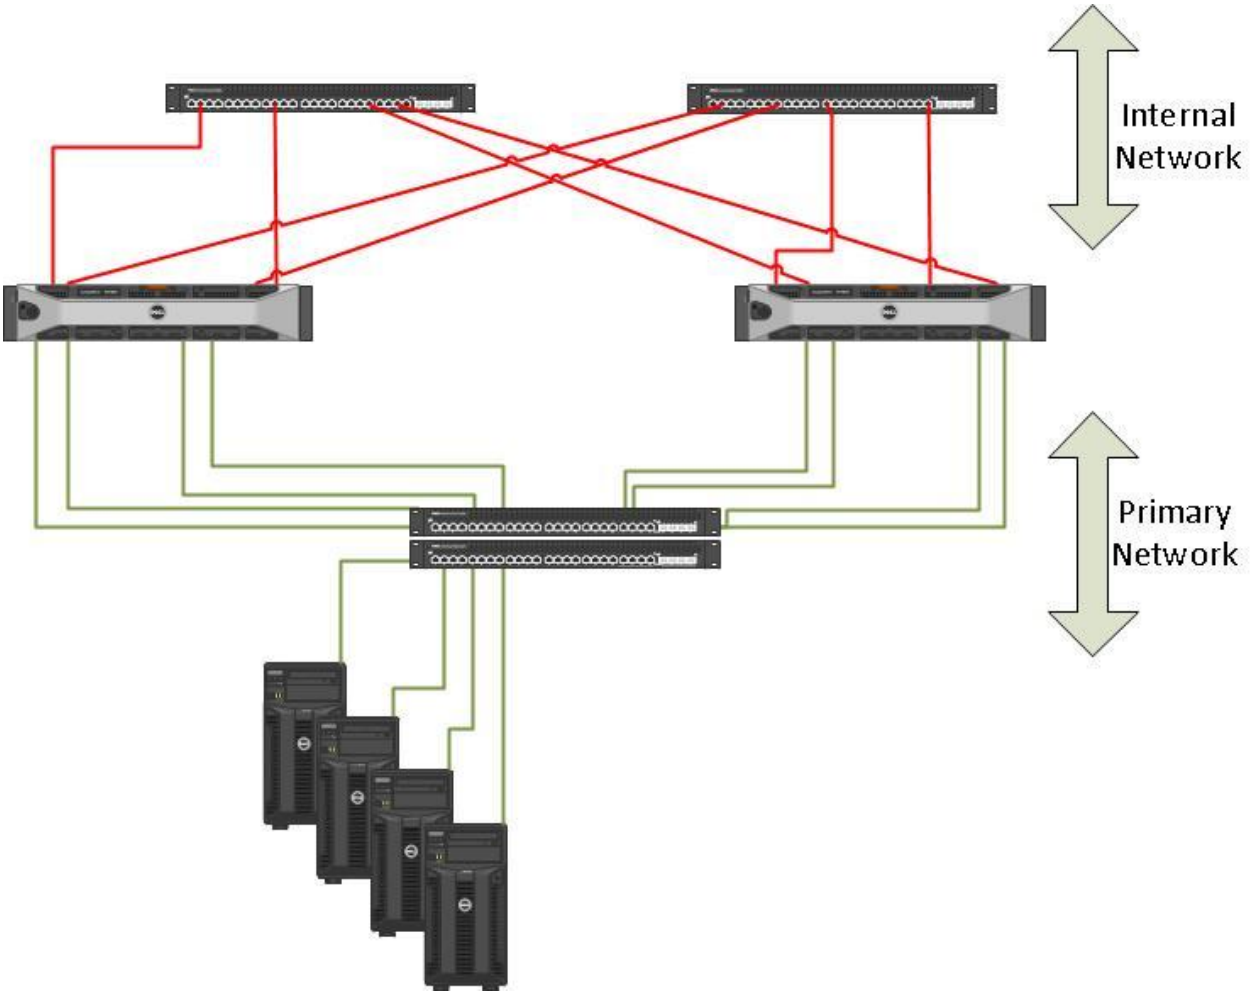

Figure 1: FS8600 networks diagram

2.2 FS8600 Internal Network

The interconnect network's role is to provide communication between the cluster controllers to enable internal data transfer and heartbeat mechanism while maintaining low latency and maximum throughput.

The interconnect segment is comprised of redundant gigabit Ethernet (1GbE) switches or 10GbE (depending on a model purchased) with dual connection to each controller creating a fully redundant mesh.

In one appliance configuration the interconnect network is comprised of two back to back links. To achieve full high-availability and future easy scalability it's recommended to use redundant external switches.

The management network's role is to provide internal logging, syslog mirroring and in some cases option for remote access.

2.3 FS8600 Primary Network (Client network)

The "primary network", also known as the client network, connects the FS8600 system to the customer network via gigabit Ethernet or 10GbE (depending on a model purchased) multiple ports.

The "primary network" provides access to data from NFS and CIFS clients and also access to the NAS administration virtual IP (VIP).

3 Network Switches Configuration

3.1 Internal Switch Configuration

The FS8600 requires the internal network switch configuration to support the following configuration:

VLAN

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. This VLAN must be untagged on all switch ports connected to the 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".

The following is an example of how to configure Dell Force10 switches and Dell PowerConnect series (the same commands are applied to Dell PowerConnect series)

```
s60>enable
s60#configure
s60(conf)#interface range gigabitethernet 0/0 – 47
s60(conf-if-range-gi-0/0-47)#mtu 9216
s60(conf-if-range-gi-0/0-47)#switchport
s60(conf-if-range-gi-0/0-47)#flowcontrol rx on
s60(conf-if-range-gi-0/0-47)#spanning-tree rstp edge-port
s60(conf-if-range-gi-0/0-47)#no shutdown
s60(conf-if-range-gi-0/0-47)#exit
s60(conf)#exit
s60#wr
s60#reload
```

3.2 Primary Switch Configuration

The FS8600 requires the primary network switch configuration to support the following configuration:

VLAN

The FS8600 primary network 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.

VLAN Tagging

If only tagged VLANs will be used for management and data, the FluidFS controllers need to be able to communicate with each other on the Client Network using untagged traffic. In other words, even if there is not a subnet configured to use untagged traffic, the switch must still allow untagged traffic/packets to pass. Blocking untagged traffic can cause issues with FluidFS during maintenance activities such as service pack upgrades, and controller replacement.

The switch must be configured properly in its present state, otherwise the untagged and/or tagged VLAN would not be accessible.

On Dell PowerConnect switches, this correlates to setting all of the frontend network ports to general mode, and then allowing the specific VLAN tags that are to be allowed to pass through.

On Force10 switches, this correlates to using hybrid port mode. General mode (Dell PowerConnect) and hybrid mode (Dell Force10) will permit untagged traffic to pass through in addition to the tagged VLANs you specifically allowed. DO not use trunk mode, as trunk mode will only pass tagged traffic, and will disallow untagged traffic.

MTU

The FS8600 primary network supports Ethernet jumbo frames.Environments can expect a degree of performance improvement, particularly where throughput is concerned.

To change the MTU value (Defaults is set to 1500) on the cluster from the FS8600 CLI, run the following command:

```
CLI> networking client-network-interface edit -MTU 9000
```

Flow Control

It is recommended, but not required, that all switch ports connected to the FS8600 "primary network" interfaces have flow control "enabled".

For more information regarding flow control and Storage Center, consult any Copilot 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 performing 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

ALB and LACP Mode

An FS8600 appliance equipped with multiple network interface cards are bonded to sustain seamless link failure.

FS8600 supports two bonding modes:

- Adaptive Load Balancing (ALB)
- Link Aggregation Control Protocol (LACP)

The default bonding mode for the FS8600 is ALB which requires no switch configuration and exposes all bonded MAC addresses.

For LACP, some switch configuration is required and only a single MAC address is exposed.

The main benefit of using LACP in FluidFS is to reduce the number of required VIPs while maintaining efficient load balancing.

Setting up LACP mode

When setting LACP you will need to perform the following configuration changes:

- In FluidFS, set one primary network VIP per FS8600 controller.
- In FluildFS, set up the bonding mode to "LACP".
- In the "primary network" switch, all switches must be stacks and define one LACP trunk per FS8600 controller.

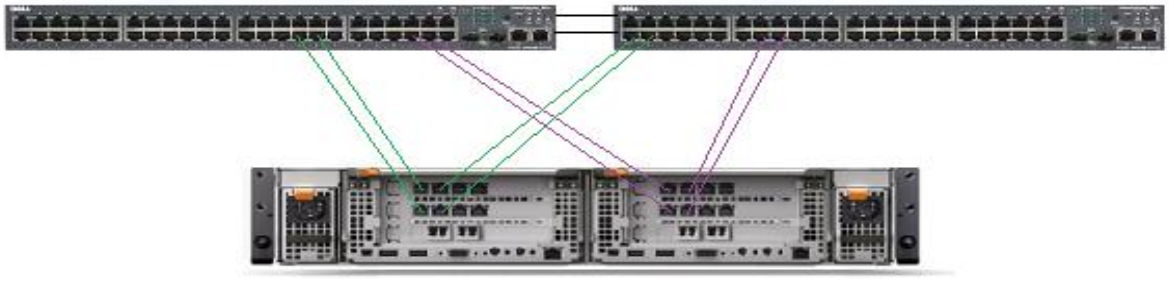

Figure 2: FS8600 LACP configuration.

To change the default bonding mode from ALB to LACP from the FluidFS CLi:

```
CLI> networking client-network-interface edit -Mode LACP
```

4 Load Balancing

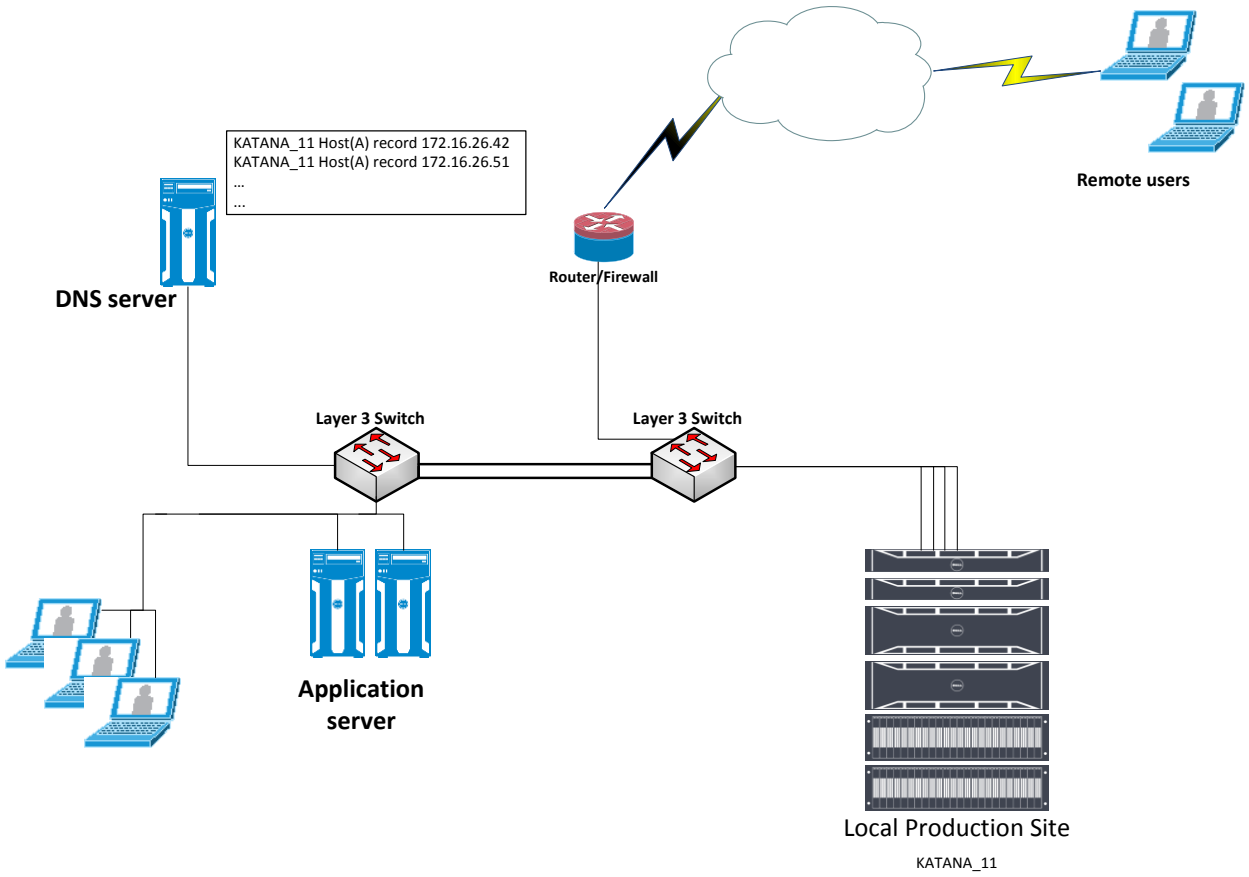

Figure 3: FS8600 in a routed network diagram

4.1 Flat Network Configuration

In case of a flat network, meaning there are no Layer 3 switches or routers between the FS8600 and the clients on the primary network, the FS8600 will use an internal ARP-based load balancing mechanism.

To achieve load balancing across multiple interfaces, the system uses a virtual IP and proxy ARP. This workload balancing method supports both inbound and outbound workloads.

The proxy-ARP protocol enables a host to provide ARP responses on behalf of other hosts. In order to support this, the destination hosts need to be configured with the proxy IP address, but should not answer broadcast ARP requests for that address. The system uses arp-filter and arp-tables to "hide" the proxy address so unicast ARP and other traffic are still served.

To verify that clients connecting to the FS8600 are balanced across all system controllers from the FS8600 CLI, run the following command:

```
CLI> networking client-load-balancing list
```

```
CLI> networking client-load-balancing list
.----------------.---------------.-----------------.-------------------.-----------------.----------------.-----------------.----------.---------------------------.
| Client IP      | Access IP     | Current Node Id | Current Interface | Assigned Node Id | Pinned Node Id | Pinned Interface | Protocol | Is Required Manual Failback |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
| 172.23.136.1   | 172.23.136.75 | 1               | eth1              | 1                | -1             |                  | Other    | No                        |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
| 172.23.136.30  | 172.23.136.74 | 1               | eth1              | 1                | -1             |                  | NFS4     | No                        |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
| 172.23.136.32  | 172.23.136.74 | 0               | eth0              | 0                | -1             |                  | NFS4     | No                        |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
| 172.23.136.32  | 172.23.136.75 | 1               | eth0              | 1                | -1             |                  | NFS4     | No                        |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
| 172.23.136.90  | 172.23.136.75 | 0               | eth1              | 0                | -1             |                  | NFS4     | No                        |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
| 172.23.136.100 | 172.23.136.74 | 1               | eth0              | 1                | -1             |                  | Other    | No                        |
|----------------|---------------|-----------------|-------------------|------------------|----------------|------------------|----------|---------------------------|
```

To verify a specific client connection to the FS8600, from the FS8600 CLI, run the following command (you will need to provide the client IP and the virtual IP used):

```
CLI>networking client-load-balancing view 172.23.136.1 172.23.136.75
```

```
CLI> networking client-load-balancing view 172.23.136.1 172.23.136.75
Client IP                  = 172.23.136.1
Access IP                  = 172.23.136.75
Current Node Id            = 1
Current Interface          = eth1
Assigned Node Id           = 1
Pinned Node Id             = -1
Pinned Interface           =
Protocol                   = Other
Is Required Manual Failback = No
```

4.2 Routed Network Configuration

Since ARP load balancing is only able to balance local network clients, the FS8600 utilizes DNS load balancing to balance clients connected across Layer 3 switches and routers.

In essence the FS8600 storage system will serve several virtual IP addresses, based on the number of interfaces attached to the client network. The virtual IP addresses are distributed and balanced among the nodes and interfaces.

The administrator defines these addresses as a round-robin group in the DNS server in the site. Clients use DNS lookup and thereby are balanced across all of the available nodes and NICs.

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.

4.3 Routed Network and Static Routes

Routed networks provide an opportunity to enhance performance through a feature called static routes. This feature allows you to configure the exact paths in which the FS8600 storage system communicates with various clients on a routed network.

To add a static route via the FS8600 CLI run the following command:

```
CLI>networking static-routes  add <DestinationNetworkID> <DestinationNetMask> <GatewayIP>
```

5 Additional Subnets

5.1 Additional 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 and replication 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.

To add additional subnets you will need the following information:

- Subnet mask
- VLAN information, if applicable
- Private IPs – one IP per FS8600 controller
- Public IP – will be used as Virtual IP (VIP)

CLI>networking subnets add 255.255.255.0 -VLANTag 0 -PrivateIPs 192.168.200.1,192.168.200.2 -PublicIPs 192.168.200.3

```
CLI> networking subnets list
 .--------------.--------------.----------.------------------------------.------------------------------.
 | Network Id   | Netmask      | VLAN Tag | PrivateIPs                   | PublicIPs                    |
 |--------------|--------------|----------|------------------------------|------------------------------|
 | 172.23.132.0 | 255.255.252.0 | 509     | 172.23.132.140,172.23.132.141 | 172.23.132.142,172.23.132.143 |
 |--------------|--------------|----------|------------------------------|------------------------------|
 | 172.23.136.0 | 255.255.255.0 | 0       | 172.23.136.70,172.23.136.71  | 172.23.136.74,172.23.136.75  |
 |--------------|--------------|----------|------------------------------|------------------------------|
 | 192.168.200.0 | 255.255.255.0 | 0      | 192.168.200.1,192.168.200.2  | 192.168.200.3                |
 '--------------'--------------'----------'------------------------------'------------------------------'
```

6 Network Services

6.1 NTP

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 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.

To configure the time server for the Active Directory environment please refer to the following link:
http://technet.microsoft.com/en-us/library/cc784800%28v=WS.10%29.aspx

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.

7 Network Security

7.1 Firewalled Environments

Many enterprise organizations have various layers of security within their network, often at their border, across different branches and work groups. Firewalls are used to implement access restrictions and Individual rules

When deploying Dell F8600 in a firewalled environment we need to make sure that specific ports are allowing traffic.

The list of ports are divided to two groups:

- Enterprise Manager and Data Collector ports
- FluidFS ports

Data Collector Inbound Ports

Port	Protocol	Name	Purpose
514	UDP	syslog	Receiving logs forwarded from Storage Centers
3033	TCP	Web Server port	Receiving: • Communication from all clients, including the Enterprise Manager client and Storage Replication Adapter (SRA) 5.5.3 • Alerts from FluidFS clusters
8080	TCP	Legacy web services port	Receiving: • Communication from Server Agents • Alerts forwarded from Storage Centers • Communication from Storage Replication Adapter (SRA) 3.3.1
7342	TCP	Legacy client listener port	• Communicating with the remote Data Collector • Providing automatic upgrade functionality for previous versions of the Enterprise Manager client
5988	TCP	SMI-S over HTTP	Receiving unencrypted SMI-S communication
5989	TCP	SMI-S over HTTPS	Receiving encrypted SMI-S communication

Data Collector Outbound Ports

Port	Protocol	Name	Purpose
25	TCP	SMTP	Sending email notifications
443	TCP	Storage center	Communicating with: • Managed Storage Centers • Managed zNAS servers
514	UDP	syslog	Forwarding Storage Center logs to syslog servers
1433	TCP	MS SQL server	Connecting to an external Microsoft SQL Server database
3306	TCP	MySQL	Connecting to an external MySQL database
8080	TCP	Vmware SDK	Communicating with VMware servers
27355	TCP	Server agent socket listening port	Communicating with Server Agents
35451	TCP	FluidFS	Communicating with managed FluidFS clusters

Enterprise Manager Client Ports

Port	Protocol	name	purpose
33	TCP	Web Service	Data Collector communication

FluidFS Basic Ports

Port	Protocol	Name
445	UDP	CIFS/SMB
427	TCP and UDP	SLP
2049 - 2049+(domain number - 1)	TCP and UDP	NFS
5001 - 5001+(domain number - 1)	TCP and UDP	mount
5051 - 5051+(domain number - 1)	TCP and UDP	quota
4050 - 4050+(domain number - 1)	TCP and UDP	NLM (lock manager)
4000 - 4000+(domain number - 1)	TCP and UDP	statd
111	TCP and UDP	portmap
44421	TCP	FTP
22	TCP	SSH
80	TCP	HTTP Web management
443	TCP	HTTPS Web management
53	UDP	DNS

FluidFS Additional Ports (Depends on NAS environment and services):

Port	Protocol	Name
138	UDP	NetBIOS
139	TCP	NetBIOS
88	TCP	Kerberos
88	UDP	Kerberos
464	TCP	Kerberos v5
464	UDP	Kerberos v5
543	TCP	Kerberos login
544	TCP	Kerberos remote shell
749	TCP	Kerberos administration
749	UDP	Kerberos administration
135	TCP	AD - RPC
711	UDP	NIS
714	TCP	NIS
389	TCP	LDAP
389	UDP	LDAP
3268	TCP	LDAP global catalog
3269	TCP	LDAP global catalog over TLS/SSL
636	TCP	LDAP over TLS/SSL
123	UDP	NTP
161	UDP	SNMP Agent
162	TCP	SNMP trap
10000	TCP	NDMP
10560-10568	TCP	Replication
1344	TCP	Antivirus - ICAP
8004	TCP	ScanEngine server WebUI (AV host)

8 Additional Resources

Below are some links to additional resources:

Dell Compellent Documentation

- Dell Compellent Replay Manager 7 Administrator's Guide
- Dell Compellent Enterprise Manager 6.3 Users Guide
- Dell Compellent Storage Center 6.3 Users Guide

http://kc.compellent.com