Compellent

Dell Compellent FS8600

Networking Best Practices Guide

Dell Compellent Technical Solutions Group
January 2015

THIS BEST PRACTICES 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.

Table of Contents

1 Preface

1.1 Audience

This document is intended for systems, networking, storage or backup administrators who are responsible for the day-to-day management of a Dell Compellent FS8600 environment.

Correct management of an FS8600 system 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 and procedures 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 basic 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) or support@compellent.com , 24 hours a day, 7 days a week, 365 days a year for current customers

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)

The roles of these networks are described in the following sections.

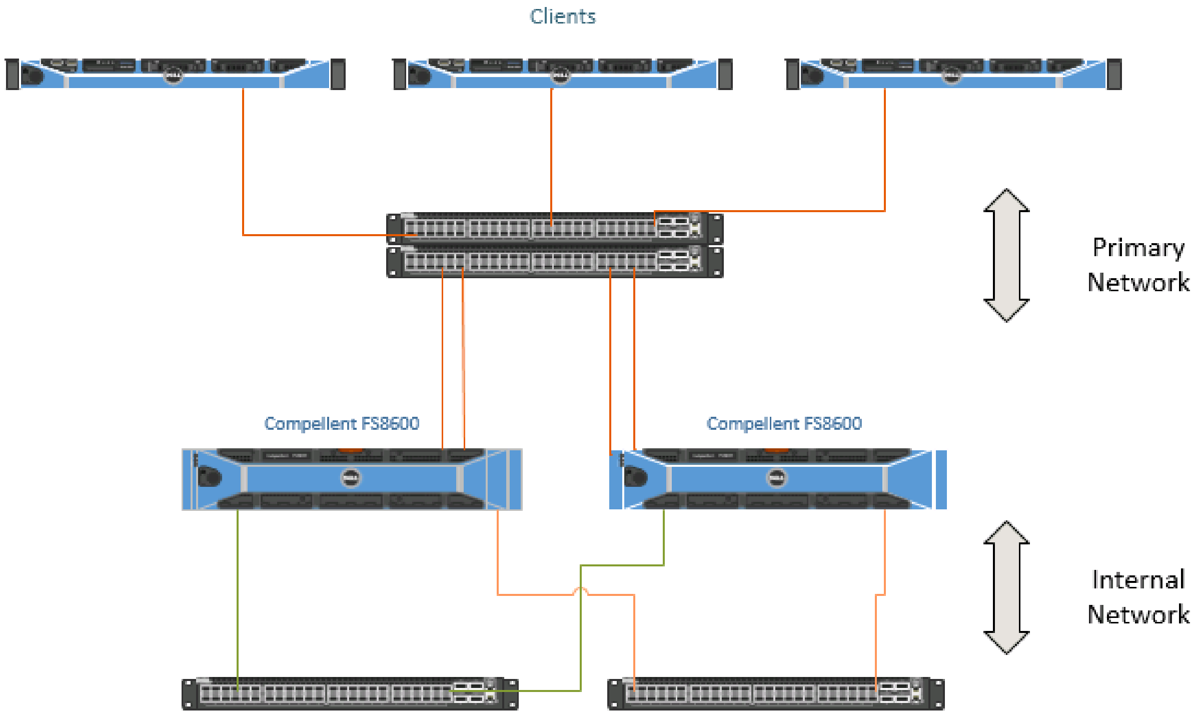

Figure 1: FS8600 networks diagram

2.2 FS8600 Internal Network

The interconnect network's role is to provide communication between the cluster controllers to enable internal data transfer and the heartbeat mechanism, while maintaining low latency and maximum throughput. In version 3 and earlier, there was both IPv4 and IPv6 traffic on the Internal Network. As of FluidFS Version 4, there is no longer a need for IPv4 internal addresses as all traffic utilizes IPv6.

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

In a single-appliance configuration, the interconnect network is comprised of two back-to-back links. To achieve full high-availability and facilitate future 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 it is used 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 the model purchased) multiple ports.

The "primary network" provides access to data from NFS and CIFS clients via virtual IP (VIP).

2.4 FS8600 IPv6 Support

As of FluidFS v4, all networks fully support IPv6. So IPv6 subnets can be created and can access IPv6-based DNS, NTP, or any other external services. IPv6 subnets are created in exactly the same way as IPv4 subnets.

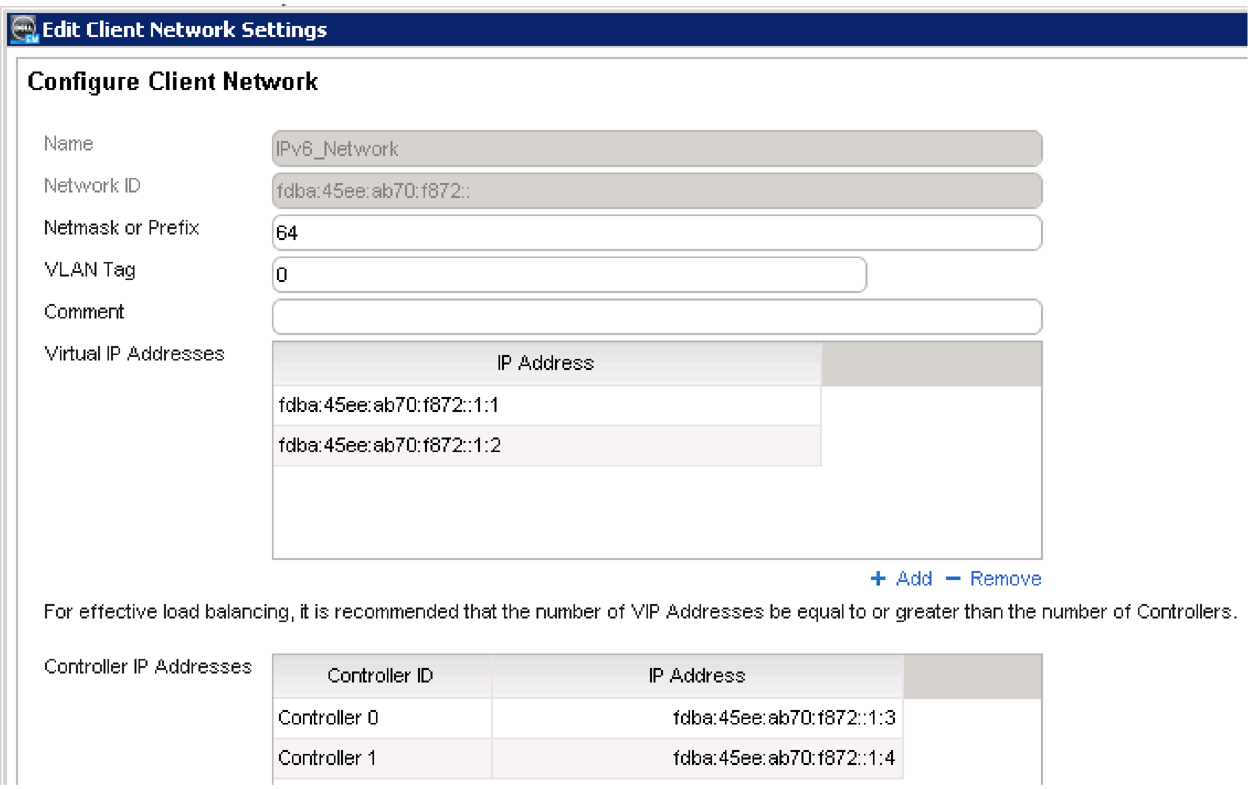

Edit Client Network Settings

Configure Client Network

Name	IPv6_Network
Network ID	fdba:45ee:ab70:f872::
Netmask or Prefix	64
VLAN Tag	0
Comment	

Virtual IP Addresses

IP Address
fdba:45ee:ab70:f872::1:1
fdba:45ee:ab70:f872::1:2

For effective load balancing, it is recommended that the number of VIP Addresses be equal to or greater than the number of Controllers.

Controller IP Addresses

Controller ID	IP Address
Controller 0	fdba:45ee:ab70:f872::1:3
Controller 1	fdba:45ee:ab70:f872::1:4

Figure 2

3 Network Switch Configuration

3.1 Internal Switch Configuration

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

IPv6

In all versions of FluidFS, IPv6 must be enabled for internal traffic.

VLAN

The interconnect network must reside on a dedicated VLAN that is isolated from all other traffic, and must be solely dedicated to 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", "edge" or "edge-port" 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

```
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)#no shutdown
s60(conf-if-range-gi-0/0-47)#exit
s60(conf)#exit
s60#wr
s60#reload
```

PowerConnect series (the same commands are applied to Dell PowerConnect series)

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

VLAN Tagging is supported but not required. In FluidFS v4, there is no longer a requirement to allow untagged traffic for internal communication over the Client Network Interfaces. This means if all FluidFS subnets are tagged, you do not need to allow untagged traffic. This allows for a somewhat simpler network configuration.

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

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 (the default is 1500) on the cluster from Enterprise Manager, Right Click on the cluster and edit Properties.

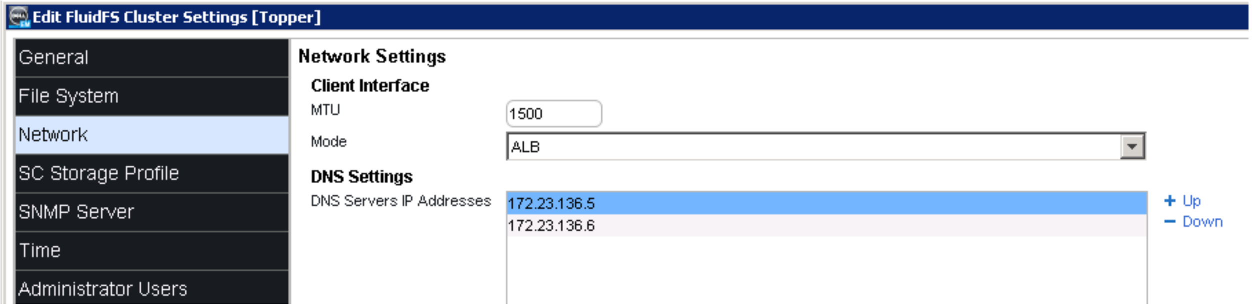

Edit FluidFS Cluster Settings [Topper]

Network Settings

Client Interface

MTU 1500

Mode ALB

DNS Settings

DNS Servers IP Addresses

172.23.136.5
172.23.136.6

Figure 3

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 the appropriate Storage Center System Manager documentation.

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 in a routed environment. See Routed Network Configuration for additional information.

Setting up LACP mode

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

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

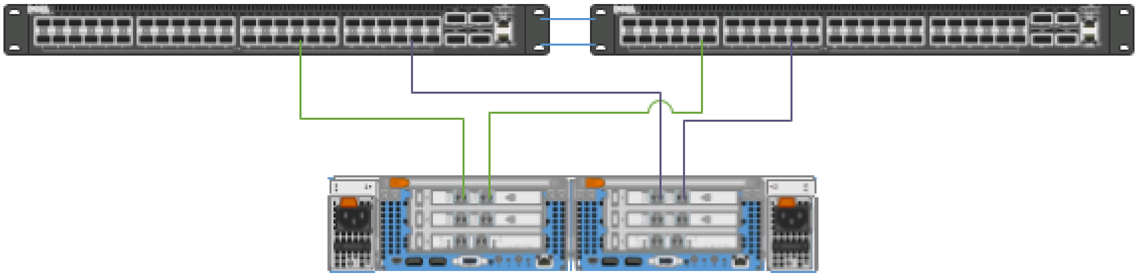

Figure 4

To change the default bonding mode from ALB to LACP within Enterprise Manager, Right Click on the FluidFS Cluster and choose edit settings. Then click on the Network Tab:

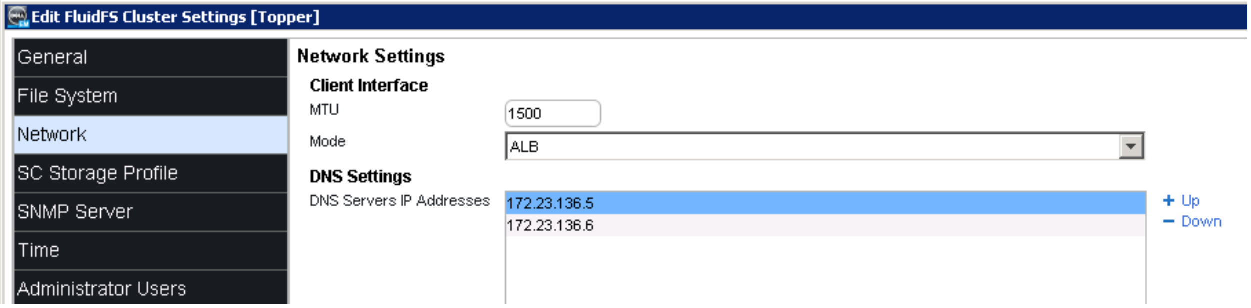

Edit FluidFS Cluster Settings [Topper]

Network Settings

Client Interface

MTU 1500

Mode ALB

DNS Settings

DNS Servers IP Addresses 172.23.136.5
172.23.136.6

Figure 5

Load Balancing

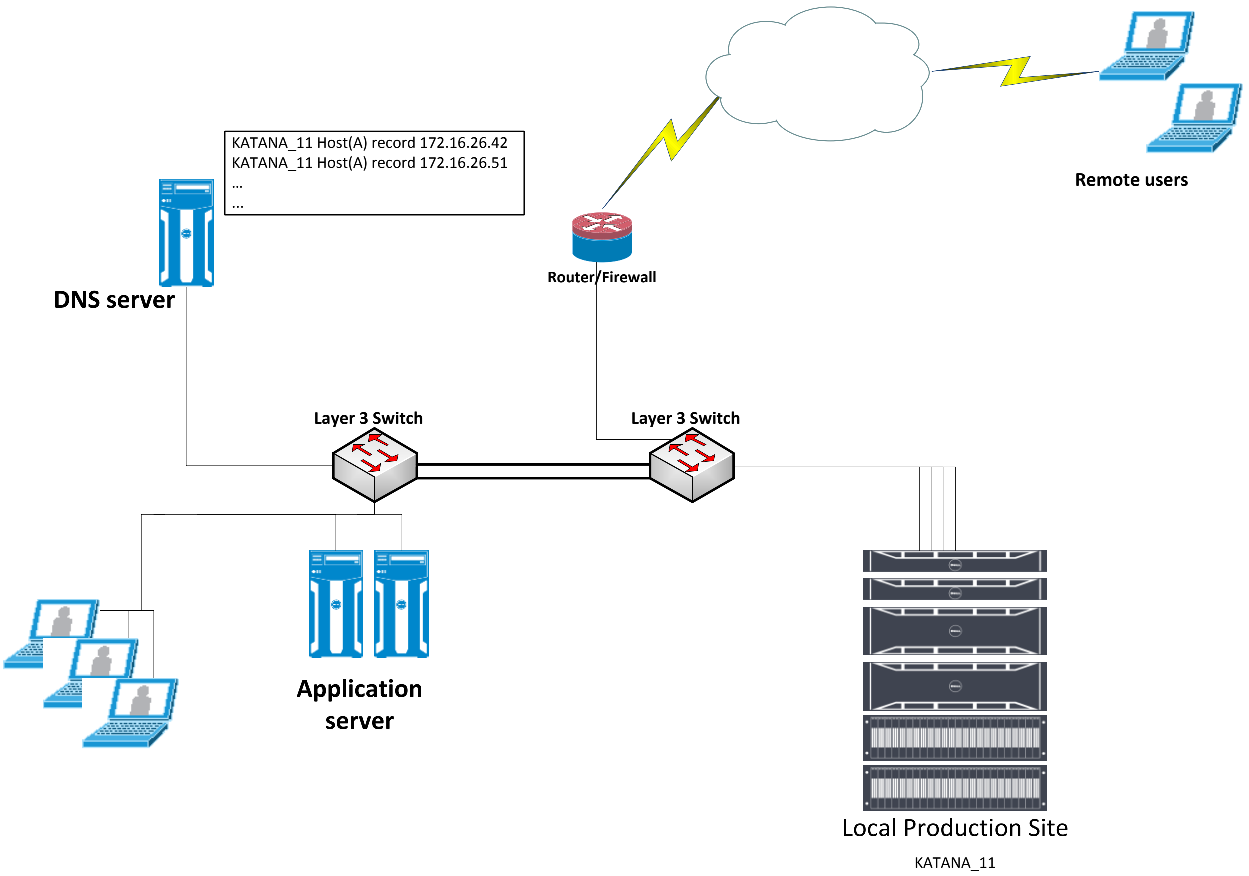

Figure 6

3.3 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 virtual IP(s) and a proxy ARP. This workload balancing method supports both inbound and outbound workloads. In a Flat network, a single VIP is sufficient to successfully load balance all the clients. However, if there are any clients accessing the FluidFS via a routed network, it is recommended to follow the routed network best practice.

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 an 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, in

Enterprise Manager Choose System > Connections > Client Connections.

3.4 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.

If client traffic is routed, it is recommended to configure a VIP per client NIC in an ALB configuration (default) or per controllers in a LACP configuration. Once the VIPs have been identified, create the corresponding DNS entries for FluidFS. For example, a single 10Gb appliance (which has 2 controllers with 2 client network NICs in each controller) will have 4 VIPs if using the default ALB configuration, or 2 VIPs if using LACP. Add a DNS Record for each VIP and associate it to the same Network Name. The virtual IP addresses are then distributed and balanced among the nodes and interfaces.

Round Robin DNS is utilized when there are multiple A records created for the same host name.

```
fluidfs   IN  A  10.0.0.1

fluidfs   IN  A  10.0.0.2
```

The first client DNS query of FluidFS will receive the order above. The second client query will receive:

```
fluidfs   IN  A  10.0.0.2

fluidfs   IN  A  10.0.0.1
```

For each additional DNS query, the IP address that the client uses will rotate indefinitely, allowing for load balancing. This can be seen by running ‘nslookup fluidfs’.

3.5 Routed Network and Static Routes

Routed networks provide an opportunity to enhance performance through static routes. Static routes allow 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 Enterprise Manager, Choose System > Network and either right click on ‘Network or click the link in the top right.

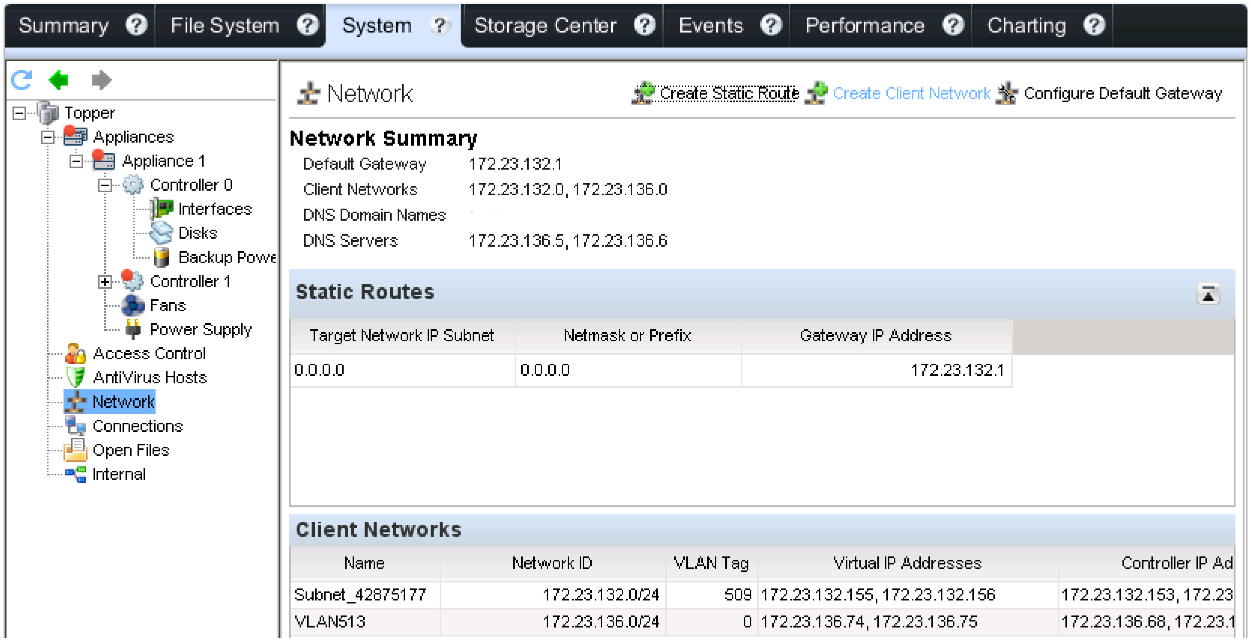

Network Summary

Default Gateway	172.23.132.1
Client Networks	172.23.132.0, 172.23.136.0
DNS Domain Names	
DNS Servers	172.23.136.5, 172.23.136.6

Static Routes

Target Network IP Subnet	Netmask or Prefix	Gateway IP Address
0.0.0.0	0.0.0.0	172.23.132.1

Client Networks

Name	Network ID	VLAN Tag	Virtual IP Addresses	Controller IP Ad
Subnet_42875177	172.23.132.0/24	509	172.23.132.155, 172.23.132.156	172.23.132.153, 172.23
VLAN513	172.23.136.0/24	0	172.23.136.74, 172.23.136.75	172.23.136.68, 172.23.1

Figure 7

4 Additional Subnets

4.1 Additional Subnets

The FS8600 can support or be part of a large number of separate or distinct logical networks, also referred to as subnets. This can be beneficial for 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. Where possible, networks should be isolated in a one-to-one fashion with VLANs.

To add 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)

Adding Subnets can be done in Enterprise Manager under System > Network, as seen in the screenshot above

5 Network Services

5.1 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 learn how 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
If both the configured gateway and the DNS hosts are unreachable, the NTP hosts are used for cluster quorum.. 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.

5.2 DNS

It is recommended that a DNS server be located in the same physical site and network as the FluidFS cluster, or as close to the cluster as possible. This will limit DNS request latency as well as quorum response latencies.

6 Network Security

6.1 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 allow traffic.

The list of Network Ports utilized by FluidFS are located in the FluidFS Support Matrix. This can be found via the FluidFS Documentation link below.

7 Additional Resources

Below are some links to additional resources:

Dell Compellent Documentation

- FluidFS Documentation
- Compellent Knowedge Center