Storage Networking
Cisco MDS 9148 Beta Testing

Prepared by:
Nuno Filipe Alves Ferreira
Technology Consultant / Architect
MTI
Email: nuno.ferreira@mti.com

Document V1.1

Document Control

Revision History			
Version Number	Revision Date	Summary of Changes	Changes Marked
1.0	11-Feb-10	Document creation.	----------
1.1	02-Mar-10	General Release	----------

Approvals			
Name	Signature	Title	Date of Issue
T Conway		MTI VP	11-Feb-10
N Ferreira		MTI Technology Architect	11-Feb-10
M Quinn		MTI Product Management	02-Mar-10

Distribution		
Name	Title	Date of Issue
Tony Conway	MTI VP	11-Feb-10
Cisco Bogota Team	Several	11-Feb-10
Gary Gray	MTI's Cisco AM	11-Feb-10

Table of Contents

1 Contact Details

Nuno Filipe Alves Ferreira
Technology Consultant / Architect
MTI

Email
nuno.ferreira@mti.com

Office Address:
Riverview House
Catteshall Lane
Godalming
Surrey
GU7 1XE

Telephone
01483 520 336

Mobile
07523 652 469

2 Overview

Cisco invited MTI to beta test their new 48* 2/4/8-Gbps port switch – MDS 9148 Multilayer Fabric Switch, and MTI happily accepted the challenge by proposing 3 testing scenarios to cover the features of the switch.

2.1 Product Overview

The Cisco® MDS 9148 Multilayer Fabric Switch (Figure 1) provides an affordable, highly capable and scalable storage networking solution for small, midsize, and large enterprise customers. The switch provides line-rate 8-Gbps ports based on a purpose-built "switch-on-a-chip" application-specific integrated circuit (ASIC) with high performance, high density, and enterprise-class availability. The switch offers outstanding value by providing flexibility, high-availability, security, and ease of use at an affordable price in a compact one-rack-unit (1RU) form factor. With the flexibility to expand from 16 to 48 ports in 8-port increments, the Cisco MDS 9148 offers the densities required to scale from entry-level departmental switch to top-of-the-rack switch to edge connectivity in enterprise SANs. The Cisco MDS 9148 offers nonblocking architecture, with all 48 2/4/8-Gbps ports operating at line rate concurrently.

2.1.1 Figure 1 - Cisco MDS 9148 Multilayer Fabric Switch

3 Analysis

MTI proposed the following scenarios for testing. The scenarios were based in the functionalities that MTI considers important to our deployments (and customer needs) and also based in the new functionalities the switch offers.

3.1 Equipment used:

a. EMC² Celerra NS-120
b. Cisco MDS 9148 Multilayer Fabric Switch
c. Cisco MDS 9222i Multilayer Modular Switch
d. Emulex LPE12000-M8 HBAs

The 3 scenarios proposed with the brief configuration overview were the following:

3.2 Normal 8G connectivity to our new 8G capable Storage

e. We used a Server with dual 8G HBAs and configured the switch for dual fabrics by using 2 VSANS – 10 and 20
f. The server was loaded with EMC PowerPath so that we could make use of the 2 paths to the Storage
g. Iometer was used to generate traffic
h. Switch ports configured for dedicated mode
i. IVR was not used as is not available yet

Network design / Topology below in Figure 2

3.3 NPV Mode

j. We changed the connection of the host to the Storage to a Cisco 9222i switch to act as the NPV CORE
k. Enabled NPiV in the NPV CORE Switch
l. Configured the 9148 in NPV mode and rebooted (NPV EDGE)
m. Configured the NPV EDGE (9148) ports that connect to the NPV CORE (9222i) as Trunking NP Ports
n. Configured the NPV CORE (9222i) ports that connect to the NPV EDGE (9148) as Trunking F-Ports
o. We then channelized both ports in a Trunking Port-Channel where we allowed both our VSANS

Network design / Topology below in Figure 3

3.4 NPV + NPiV Mode

p. We took as base the configurations described in Scenario 2
q. We enabled NPiV in the NPV EDGE Switch (9148)
r. We extended the zoning configuration in the NPV CORE Switch to accommodate the needs of NPiV
s. Most of the remaining configuration was done in the ESX server and in the Storage Array

Network design / Topology below in Figure 4

4 Topologies

4.1 Figure 2 - Normal 8G connectivity Topology (back to analysis)

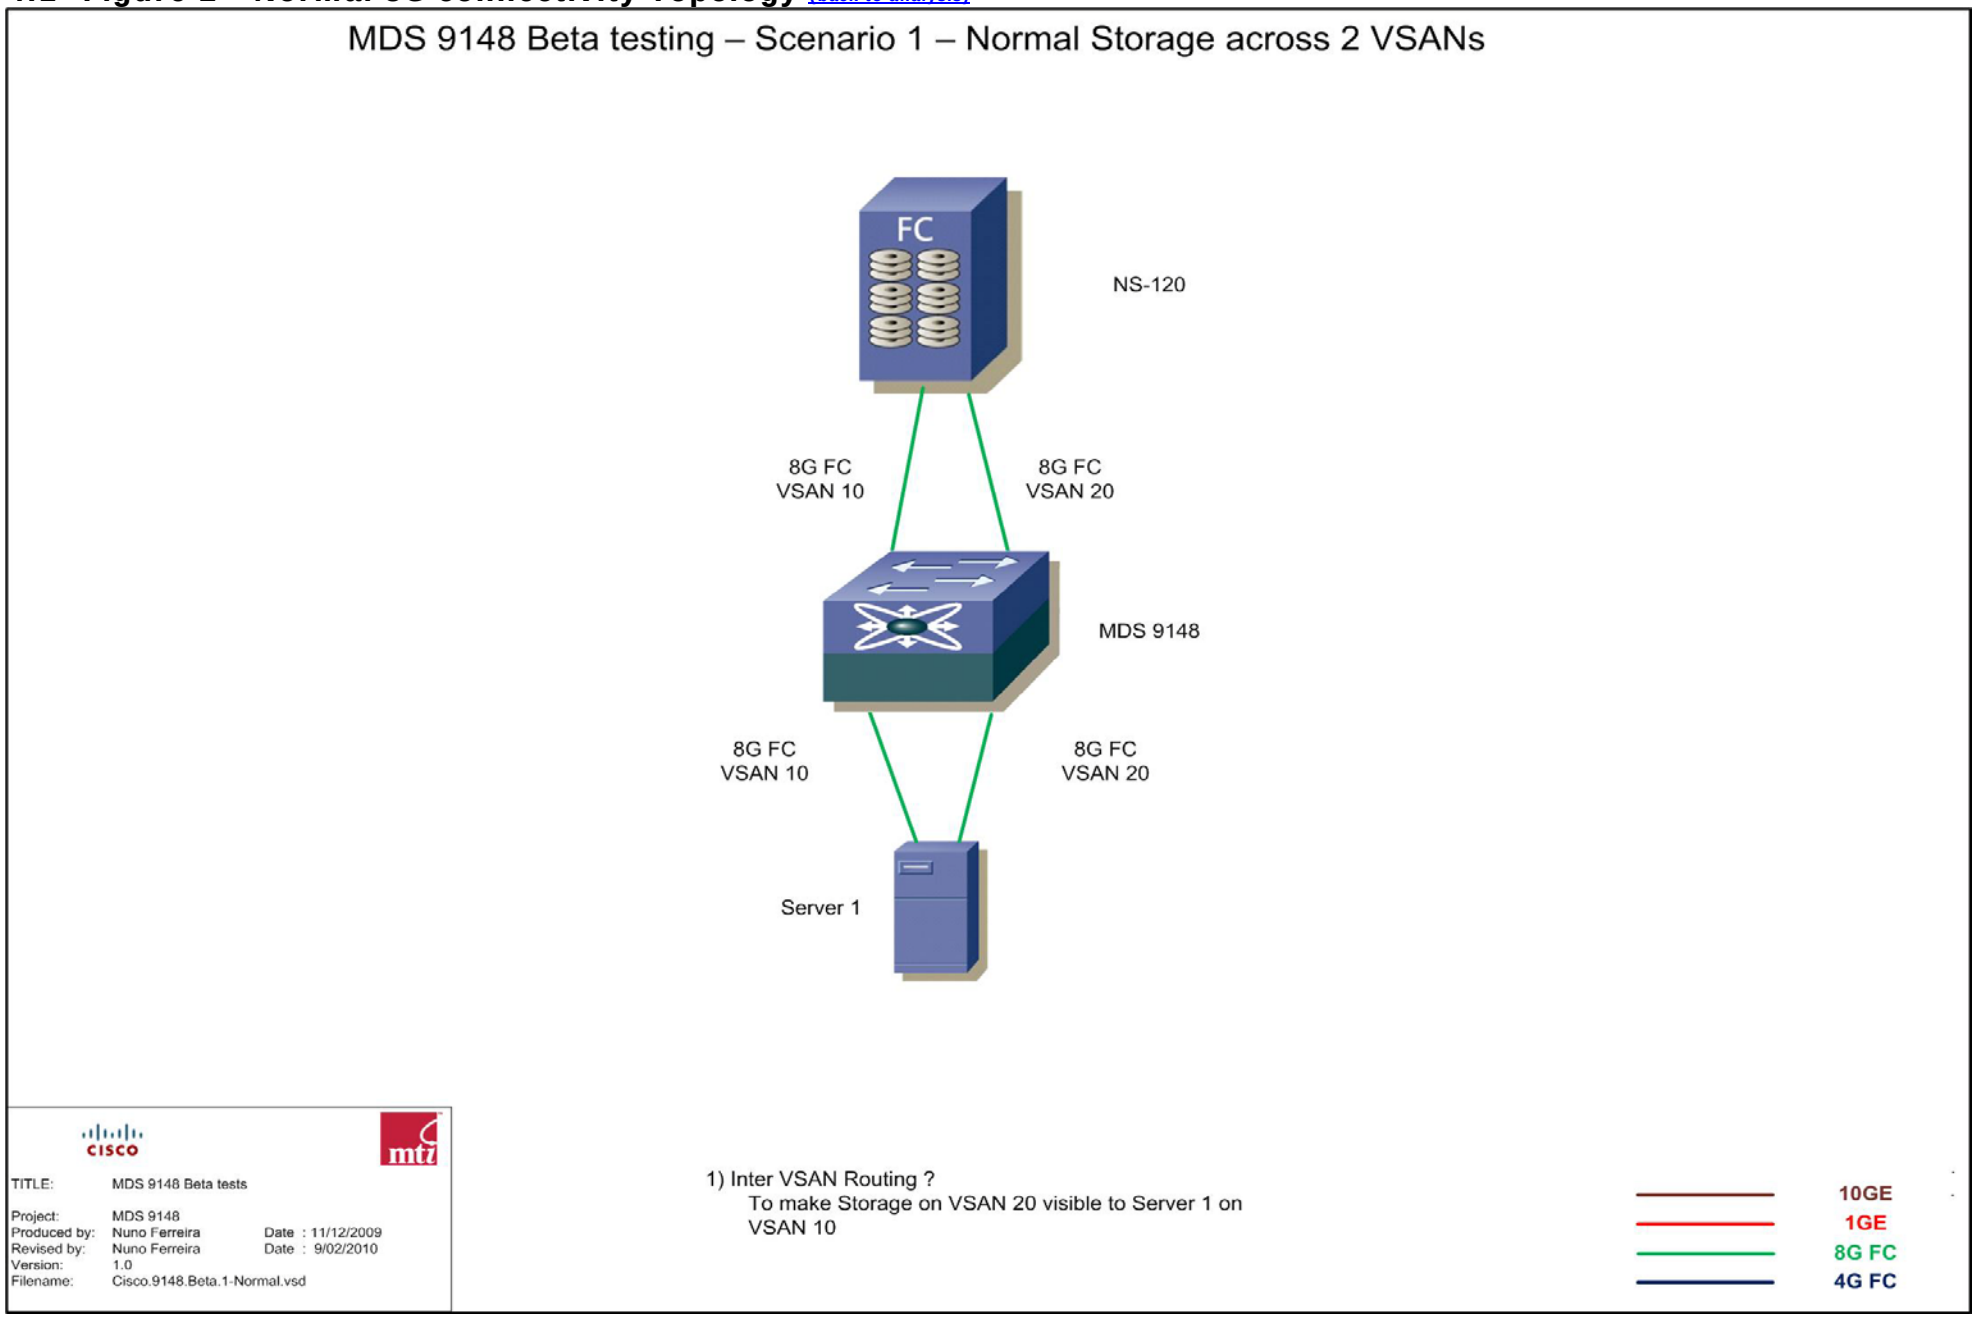

4.2 Figure 3 - NPV Mode Topology (back to analysis)

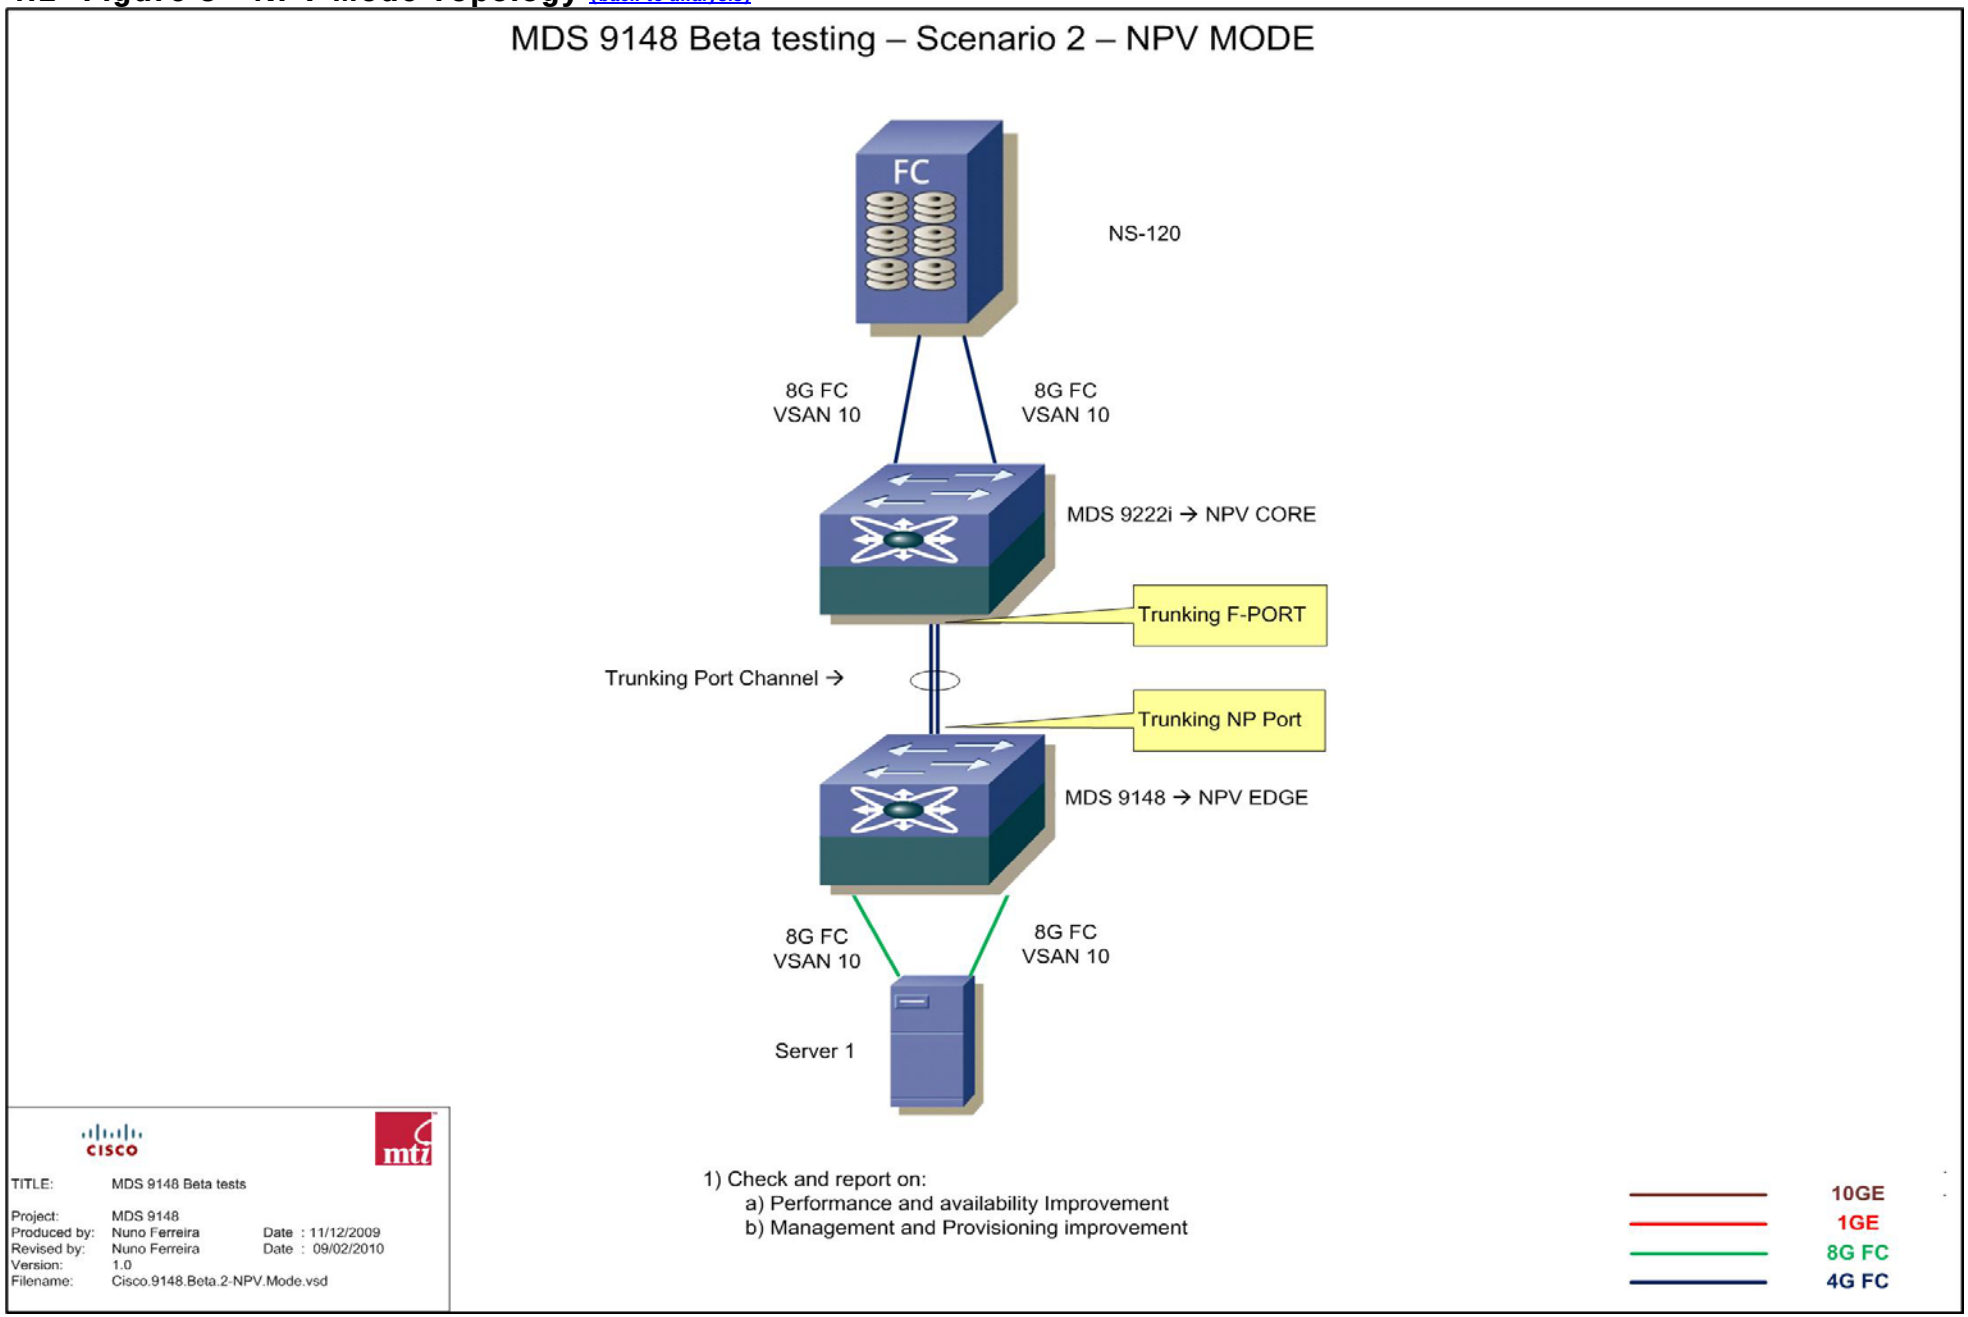

4.3 Figure 4 - NPV + NPiV Mode Topology (back to analysis)

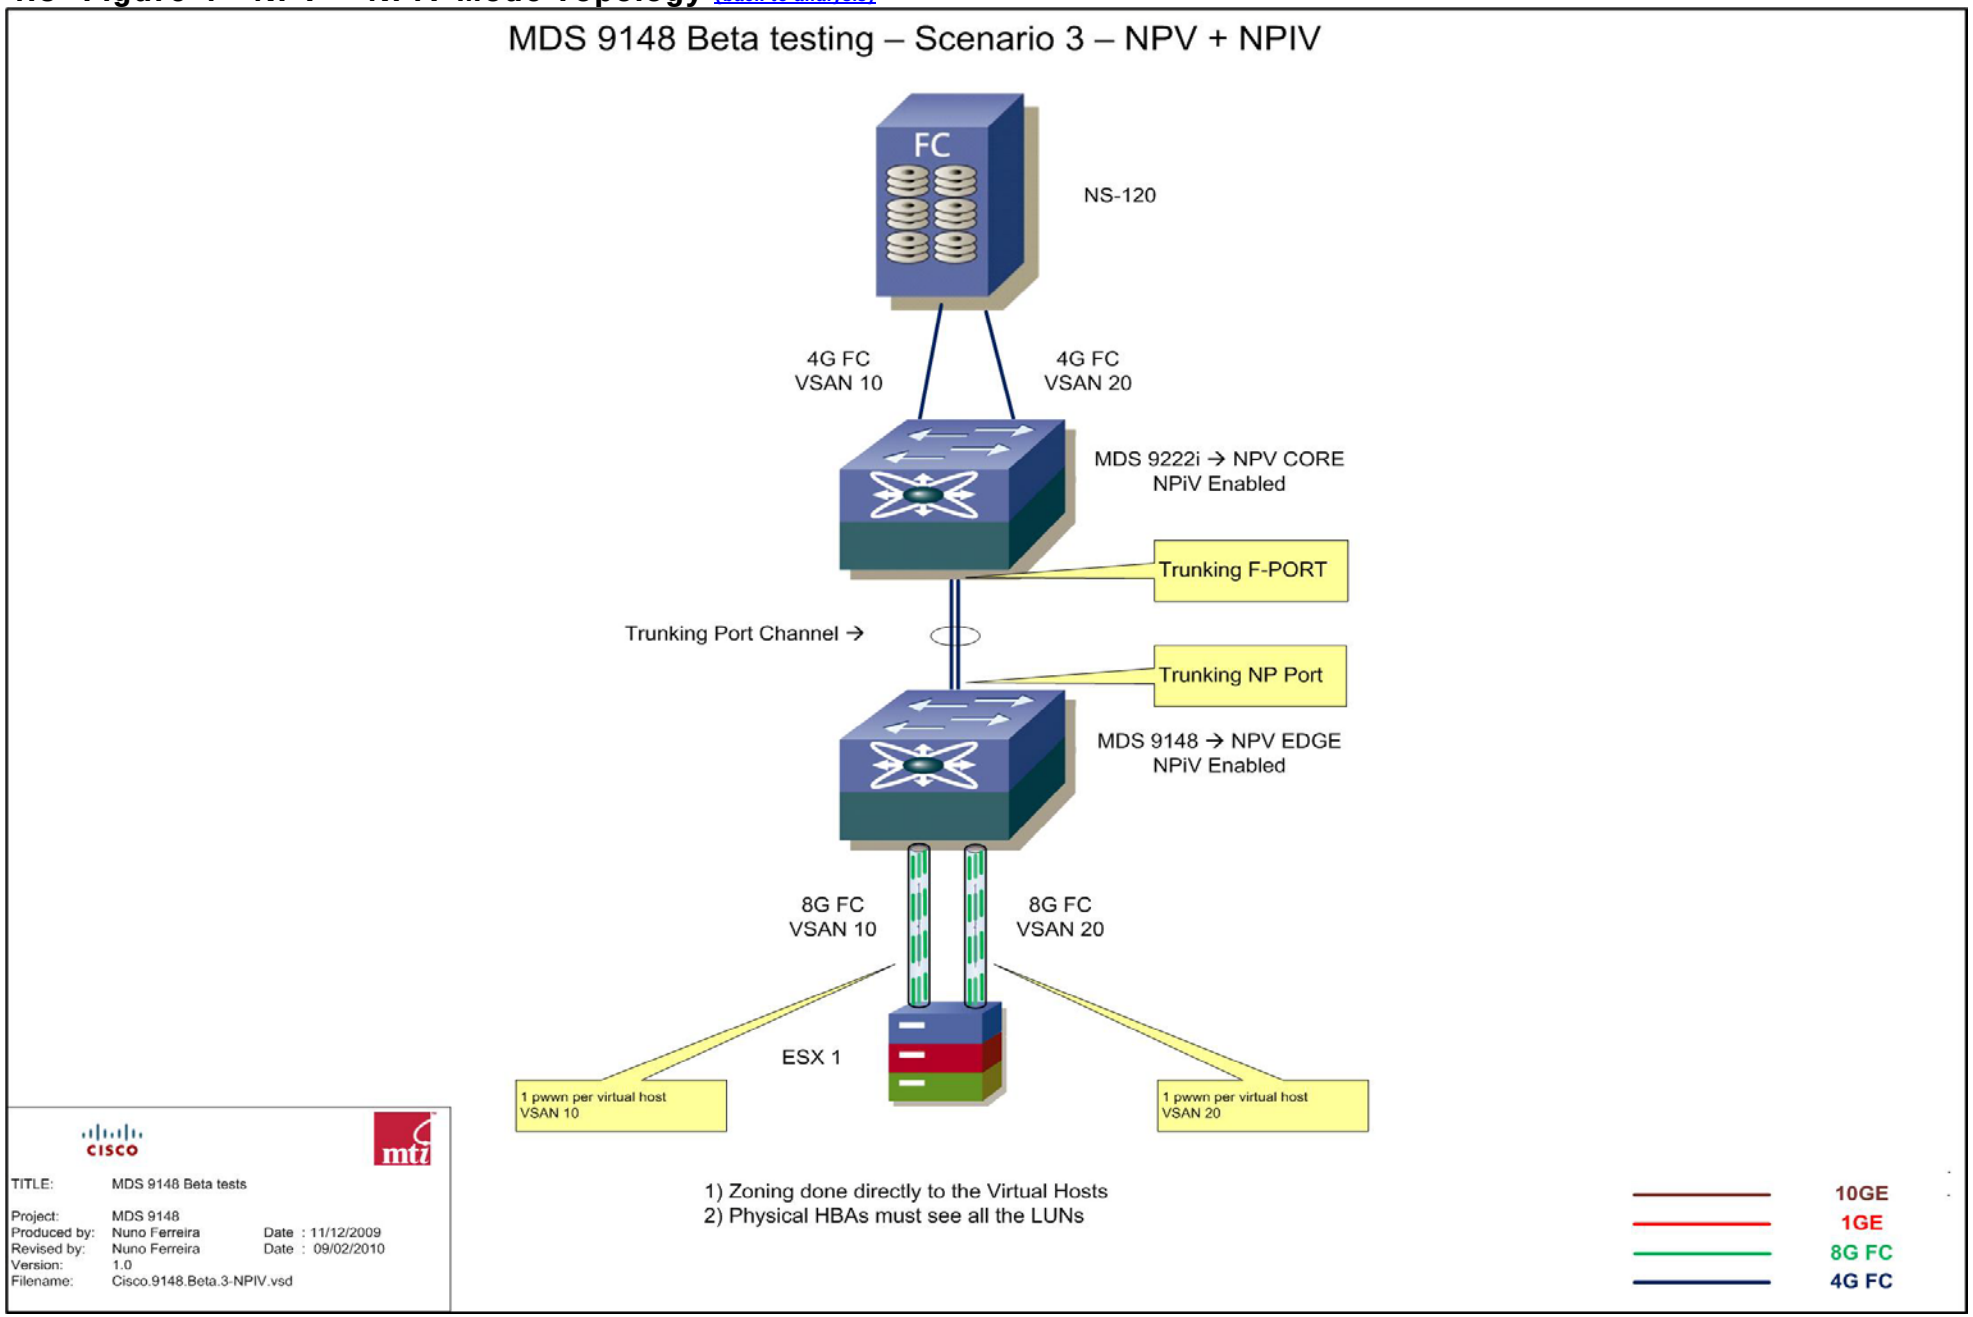

5 Results

The testing of the switch went pretty well and the results were very good.

Below the results are presented with the captures of tests done and a brief description of what the capture sees.

5.1 Preparation

We started up by doing the normal OS and Management software upgrades to the latest versions. Versions used as follows:

Fabric Manager → Version 5.0(1)
Cisco MDS 9148 NX-OS → Version 5.0(1)
Cisco MDS 9222i NS-OS → Version 4.2(3)

5.1.1 Figure 5 – Update requirements

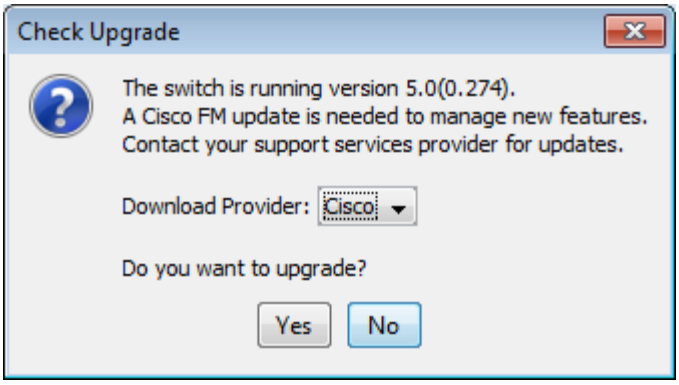

5.2 Normal 8G connectivity Results

The first topology to be deployed and configured was just the normal connectivity using the 8G on HBAs, Switch Ports and Storage.

With this test we would like to see the improvement in terms of speed that the higher bandwidth would give.

We knew that the Cisco Switch was non-blocking (Full Line Rate) so we used these tests also to find out where the bottlenecks were (they were found to be in the Storage Arrays).

The following images shown in Figures 6, 7 and 8 show in more depth how the tests and configurations were done.

5.2.1 Figure 6 – Setup layout and physical interface map

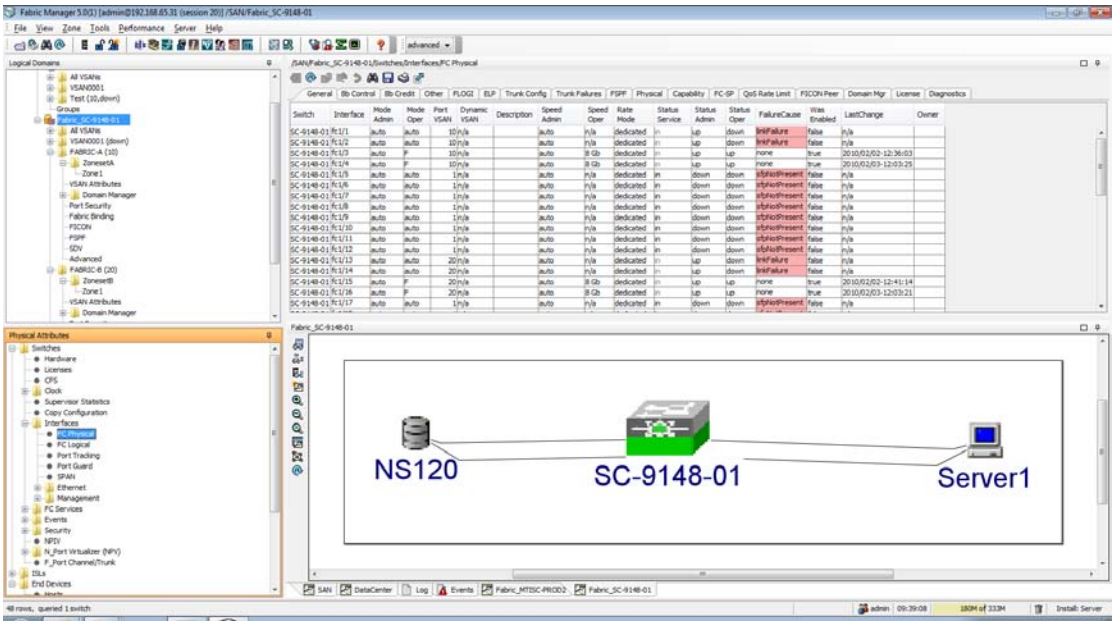

5.2.2 Figure 7 – Connected Hosts

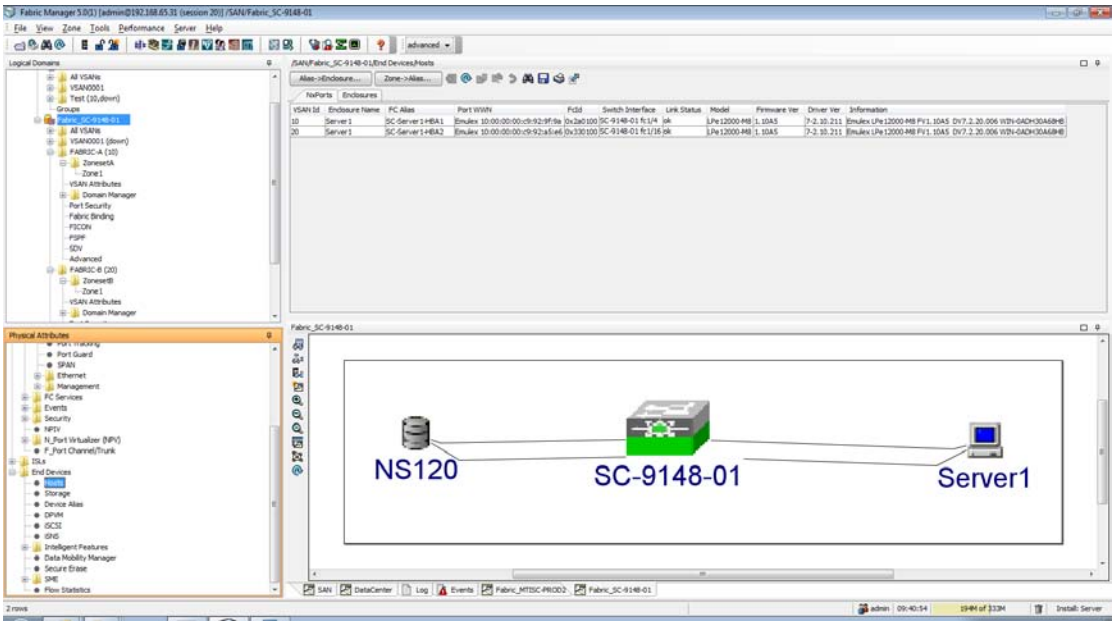

5.2.3 Figure 8 – Connected Storage

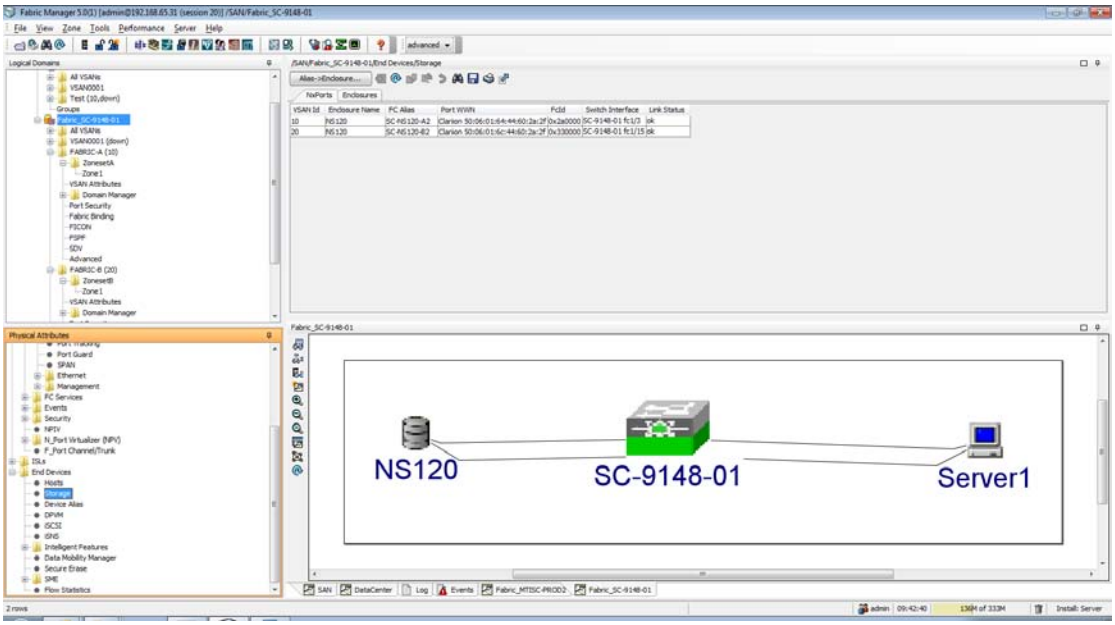

After the configuration was complete we could see our disks in our host as per the below image.

5.2.4 Figure 9 – Disks

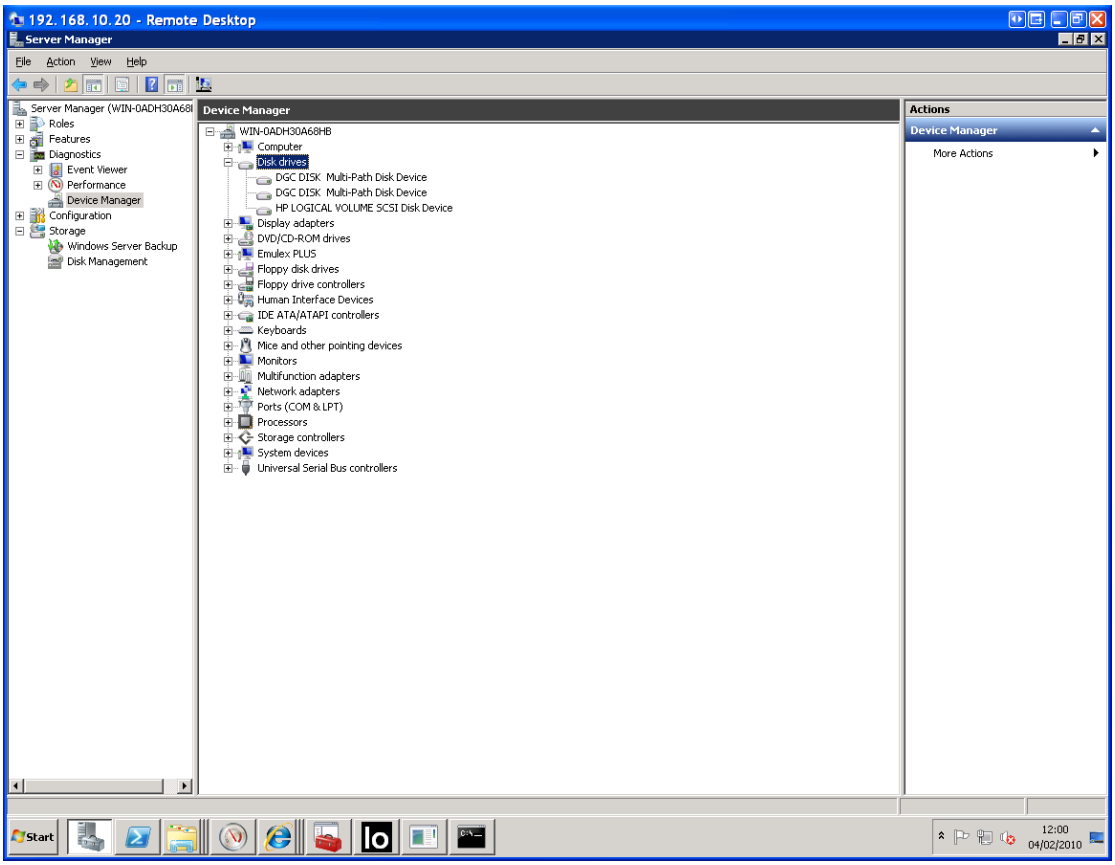

Then we started to generate some traffic to get a feel of what the performance would be like and also as stated above to find out what and where would be the eventual bottlenecks.

5.2.5 Figure 10 – Traffic Generation and Stats

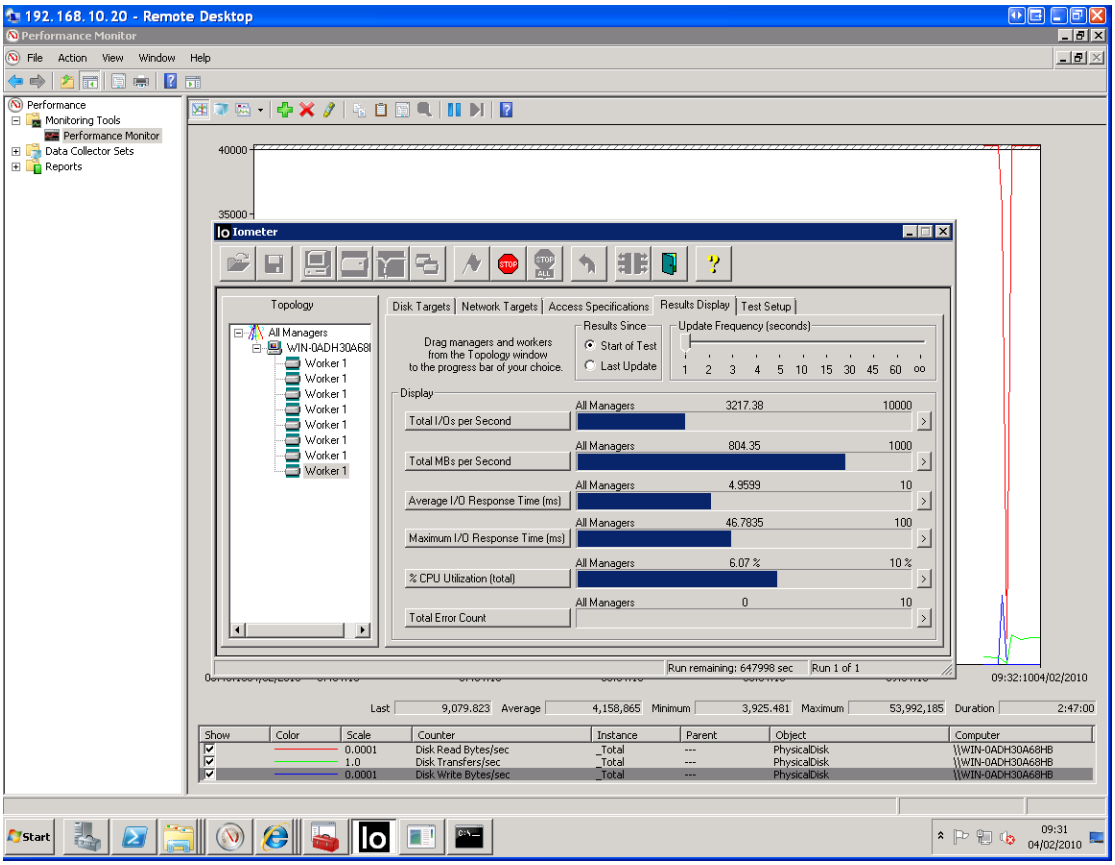

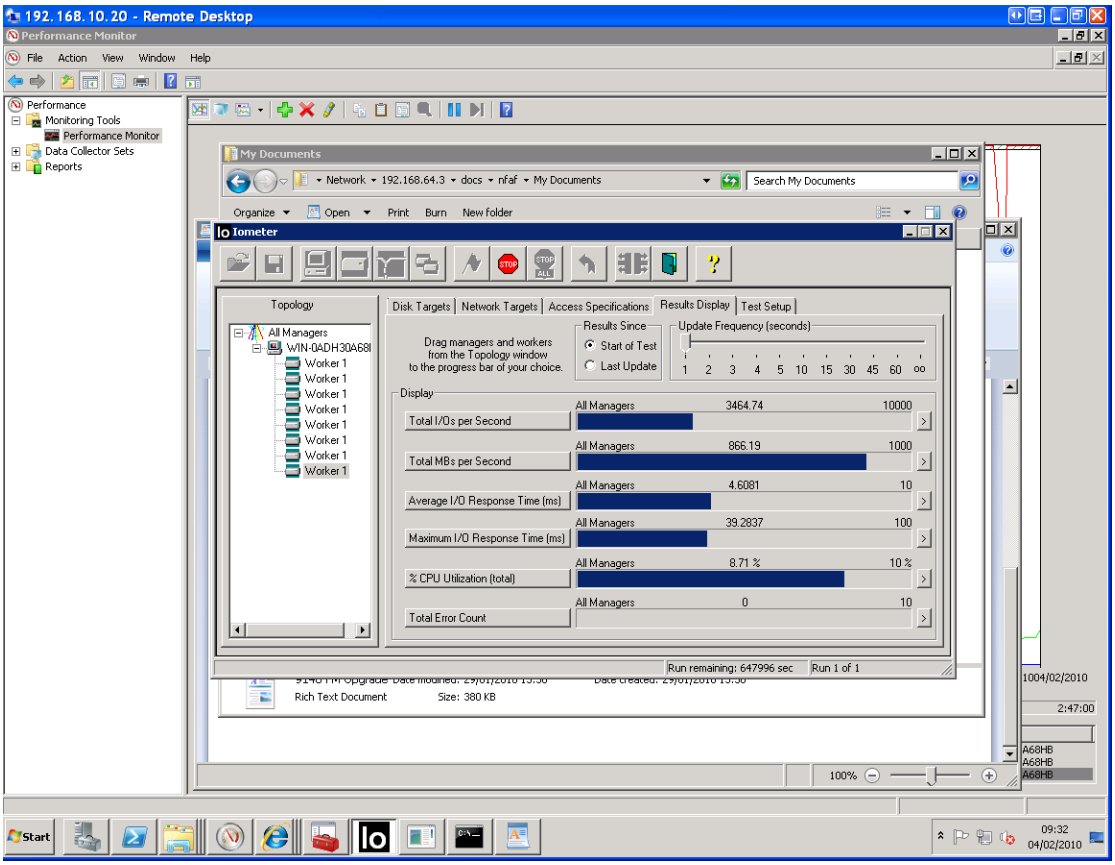

To finish we used the Fabric Manager Web Client to configure a set of collections so we could have a graphical view of what was flowing through the switch as displayed in the following image:

5.2.6 Figure 11 – Traffic and Performance Overview

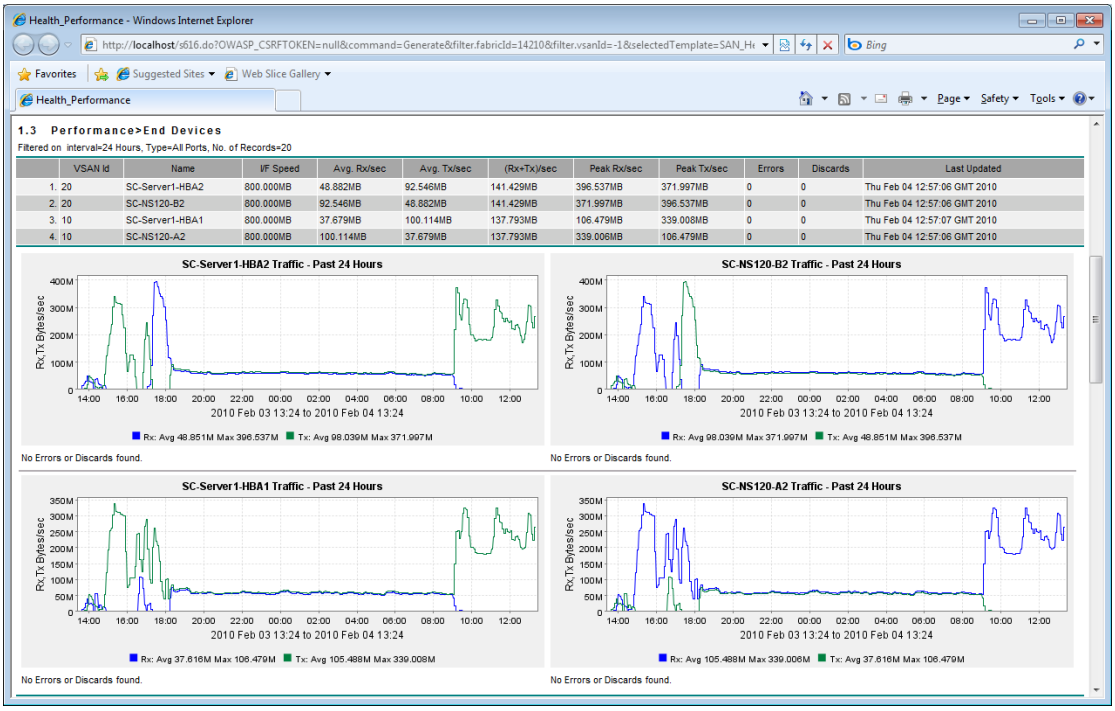

Overall the results were quite good (as can be seen) and we could definitely see the benefits of the increased speed. Hopefully the Storage Vendors will soon be providing faster backplanes.

5.3 NPV Mode Results

We then added our 9222i to act as the NPV CORE Switch and enabled NPiV on it.

We enabled the NPV mode on the 9148 and restarted so the switch could initialize in this mode.

We then configured the ports of our NPV EDGE Switch (9148) connecting to the NPV CORE (9222i) in Trunking NP ports and the same ports on the NPV CORE (9222i) in Trunking F-Ports as shown in the following image:

5.3.1 Figure 12 – Setup layout and physical interface map

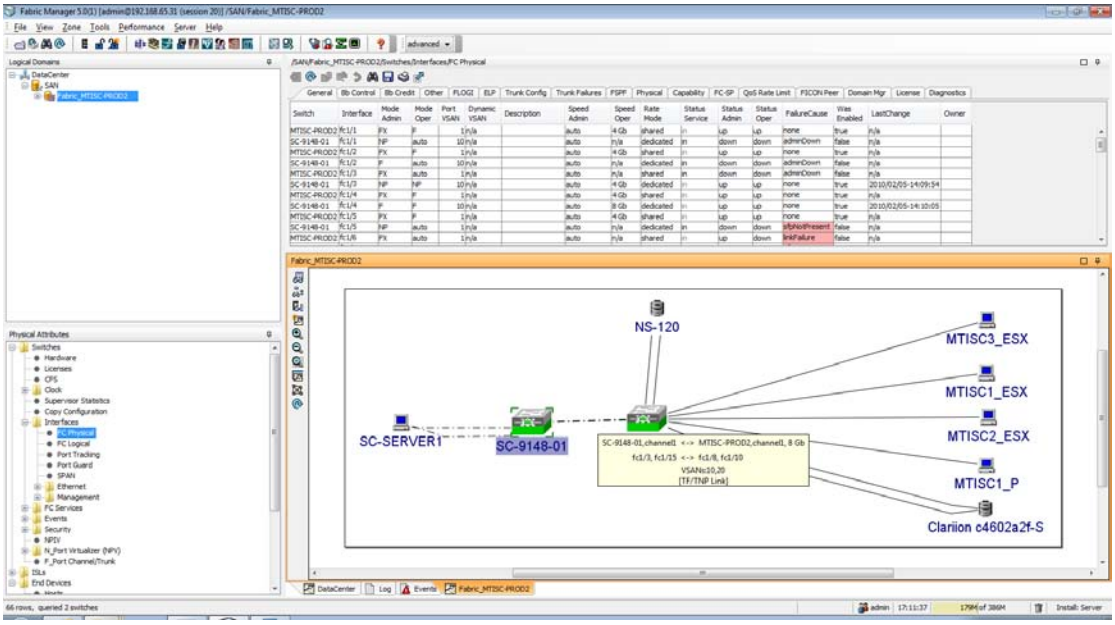

Then we checked connectivity and everything was there. We could see as expected traffic being forwarded to the NPV CORE for Fabric Logins and our 9148 not wasting any Domain ID.

We did some traffic and bandwidth monitoring as well just to see the difference in throughput (because now we had 4G links in the middle).

5.3.2 Figure 13 – Traffic Generation and Stats

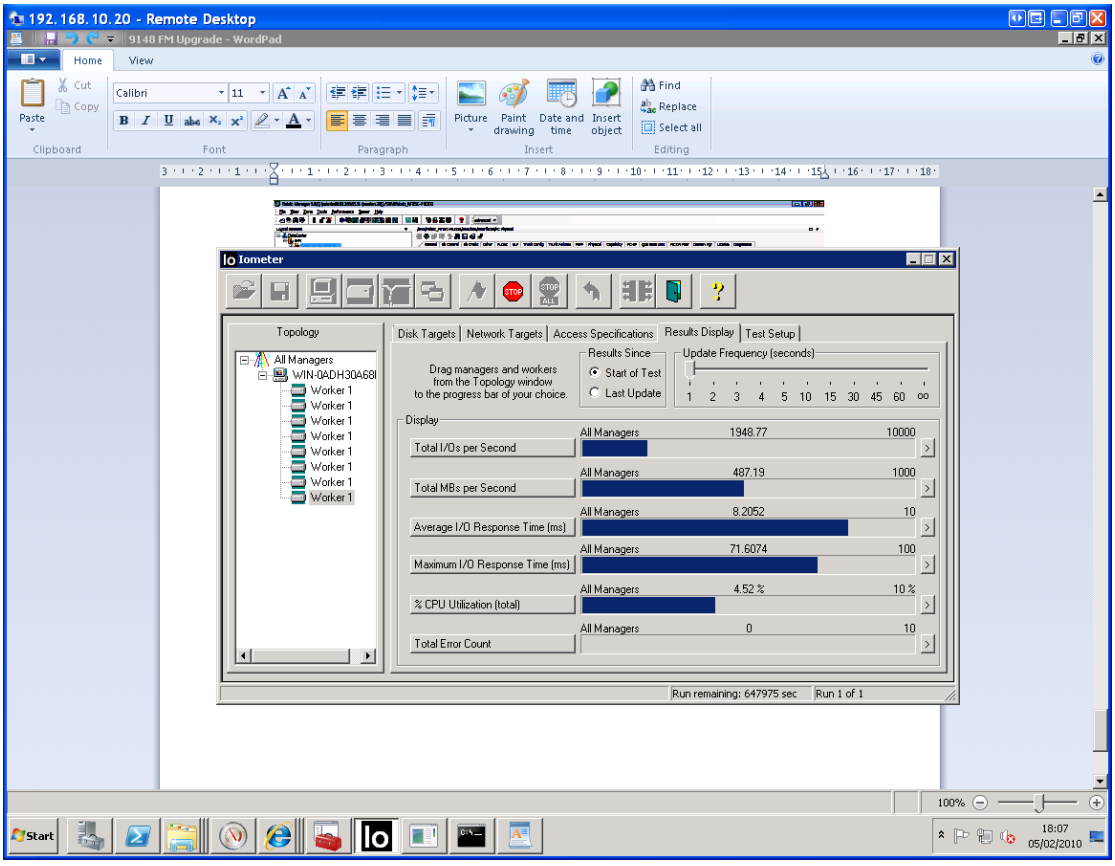

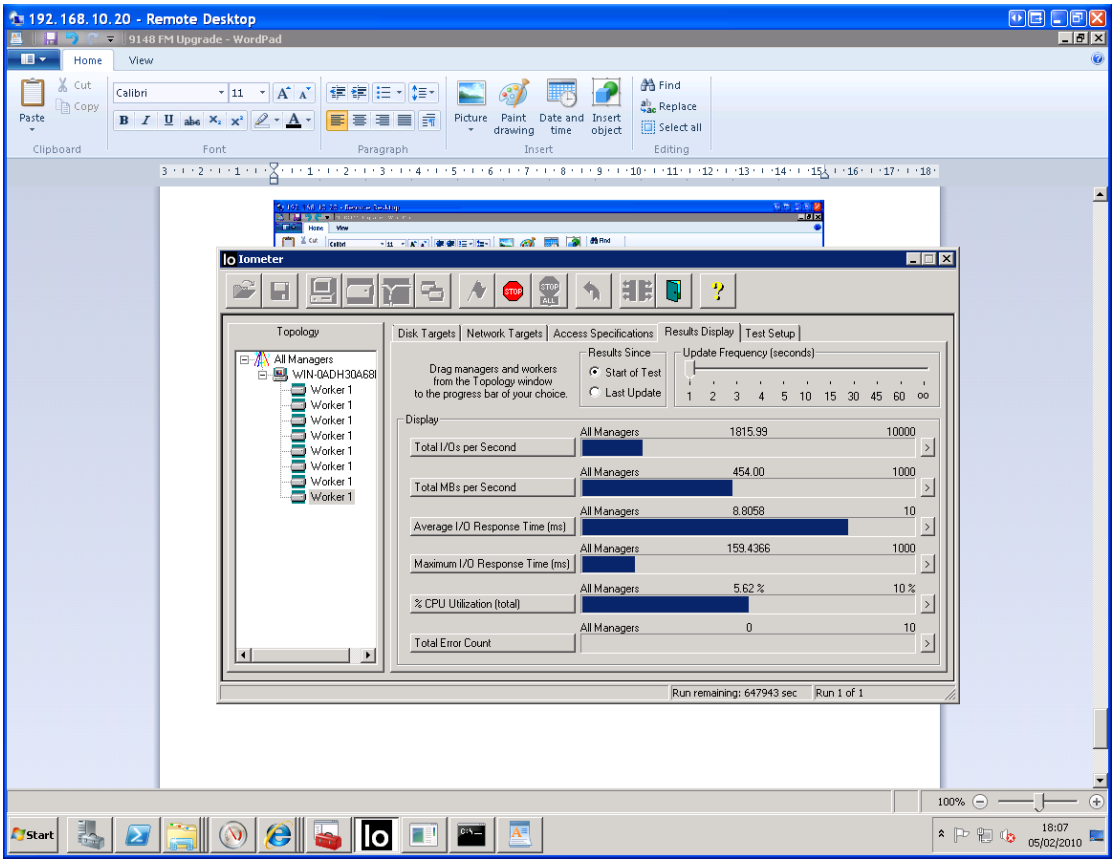

To finish we used the Fabric Manager Web Client to configure a set of collections so we could have a graphical view of what was flowing through the switch as displayed in the following four images that make up figure 14:

5.3.3 Figure 14 – Traffic and Performance Overview

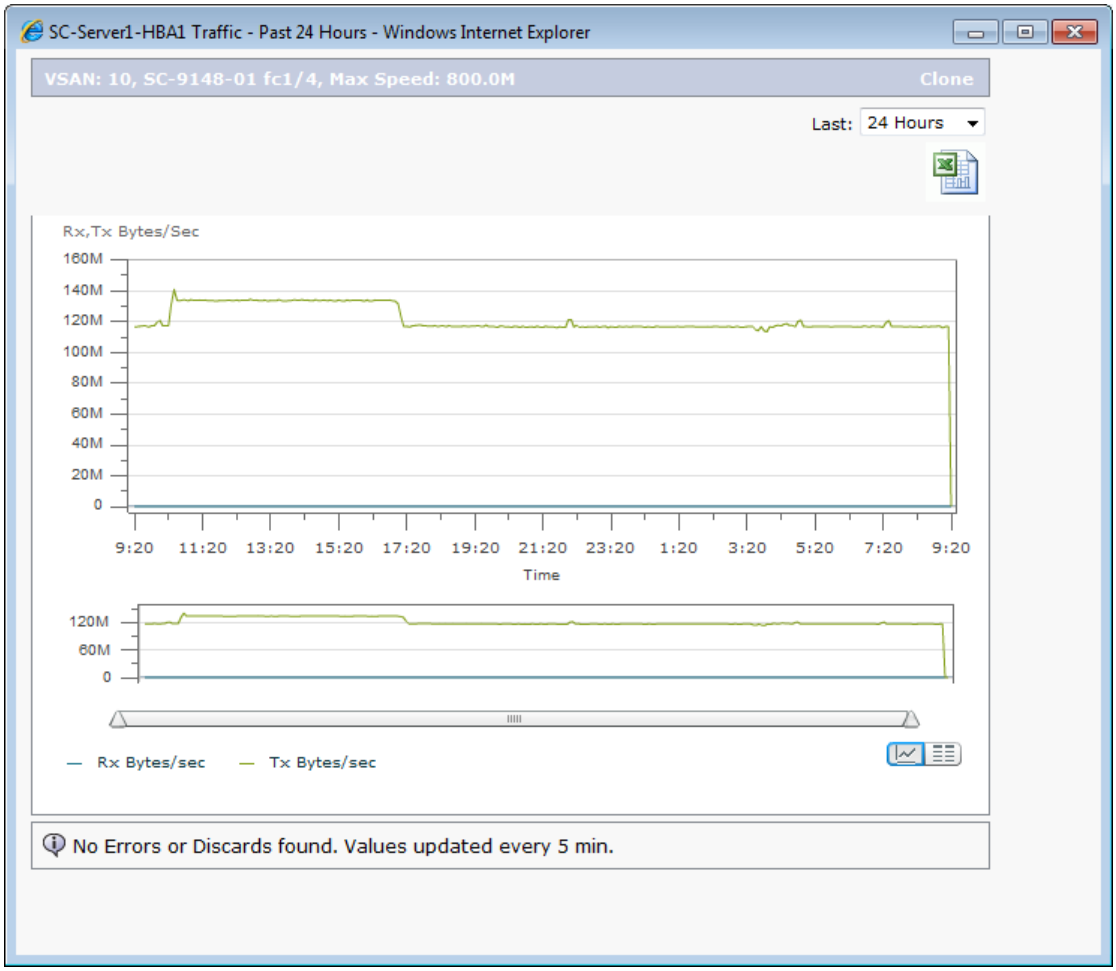

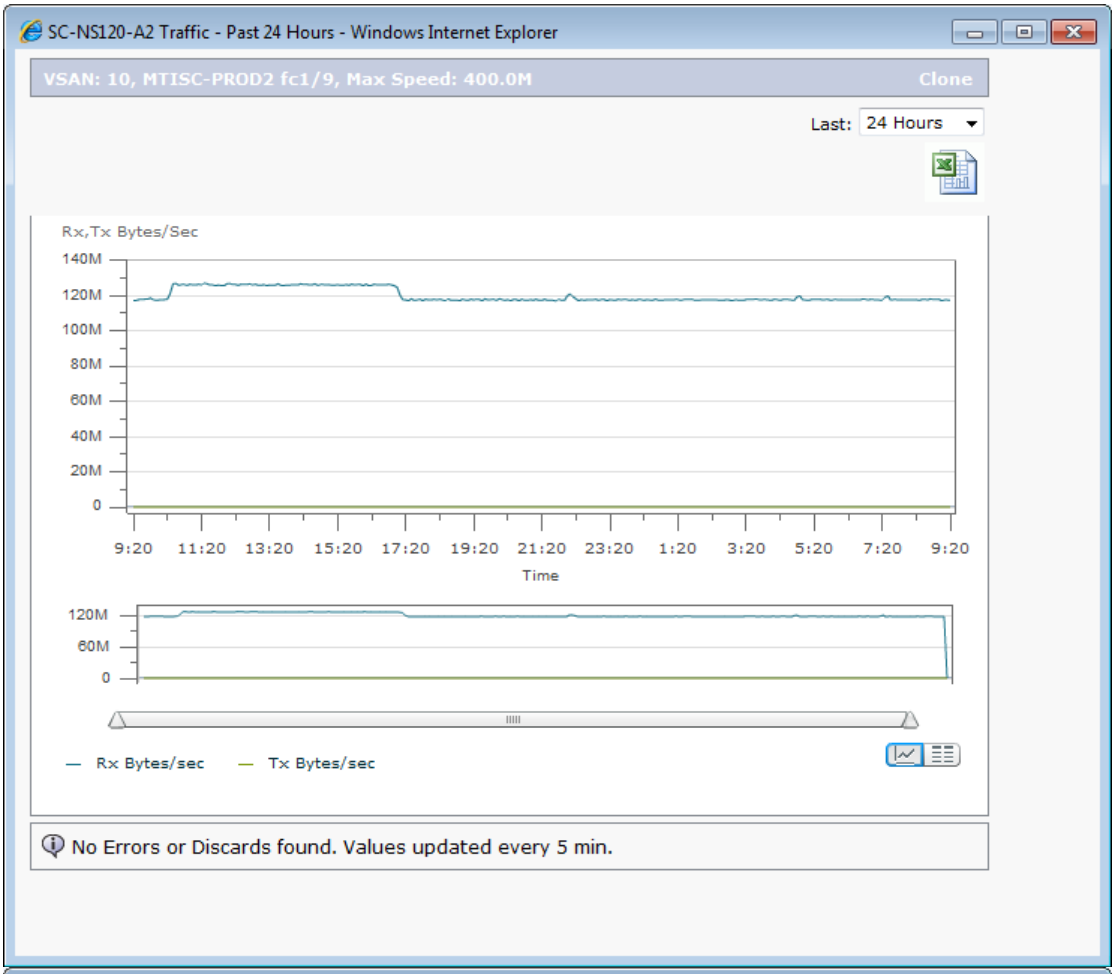

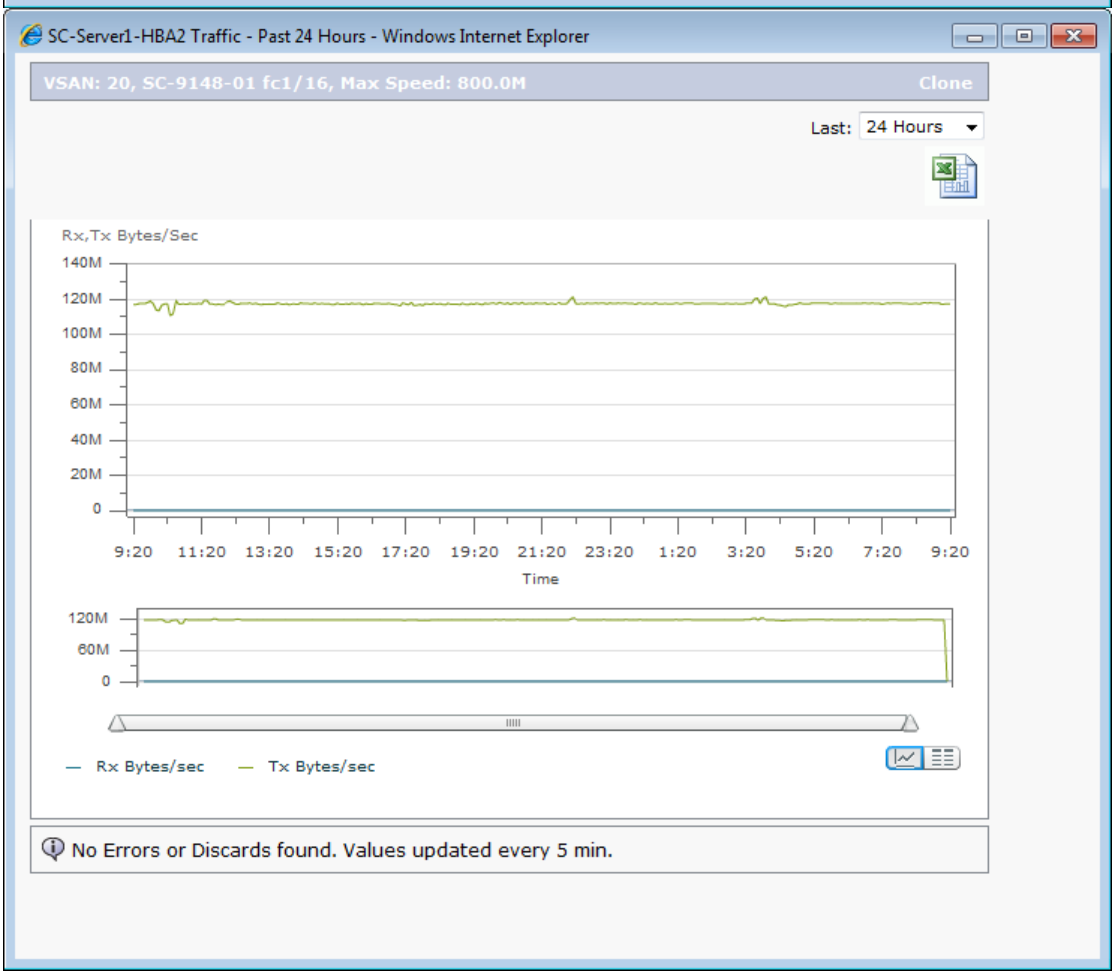

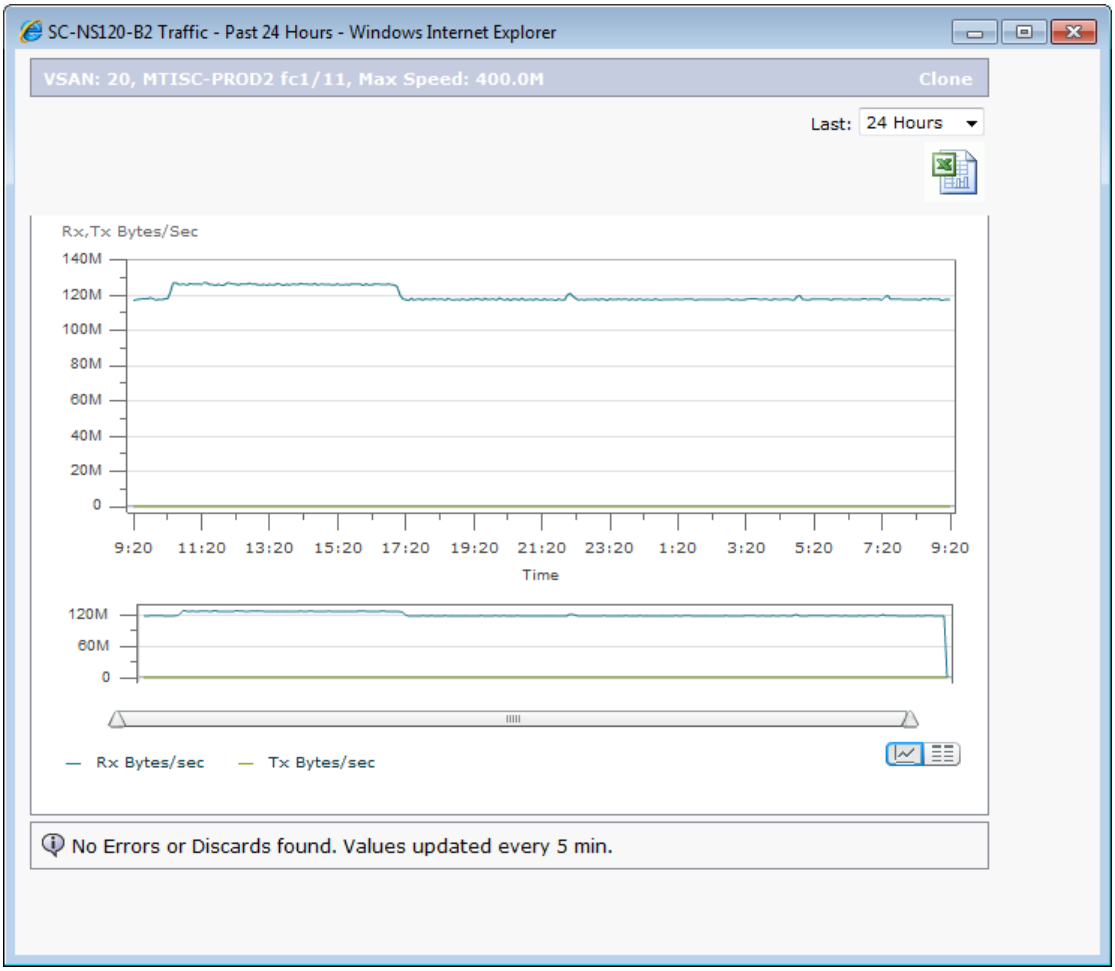

5.4 NPV + NPiV Mode Results

Virtualization is KEY for MTI and therefore this last setup was one of our most important ones.

We wanted to test having the Virtual Machines with their own pWWN so we could ease hardware requirements and create LUNS destined to Virtual Machines.

This also allows SAN Administrators to monitor and route storage access on a per Virtual Machine basis.

So to achieve this we enabled NPiV on our NPV Edge switch. We then installed VMWare vSPHERE (or ESX 4.0) on our Server and created 2 Virtual Machines.

When creating the Virtual Machines we decided to add one of our LUNS as a second disk to our Virtual Machine 1 but our second LUN would be the unique disk of our Virtual Machine 2 – Thus enabling booting from SAN and having a SAN only disk.

Setup layout is shown in figure 15:

5.4.1 Figure 15 – Setup layout and physical interface map

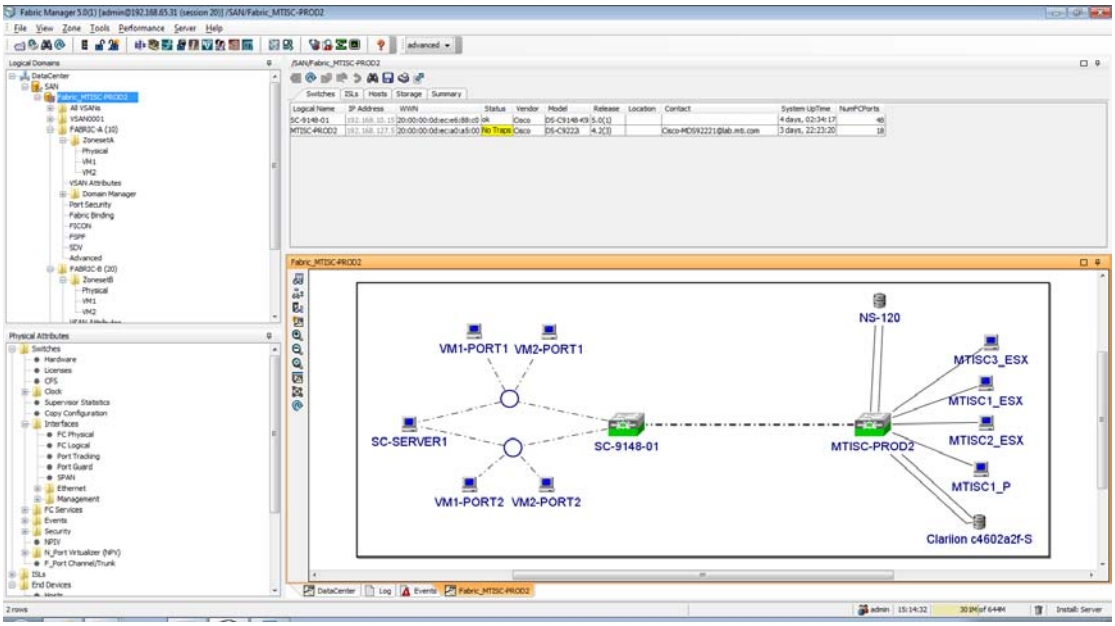

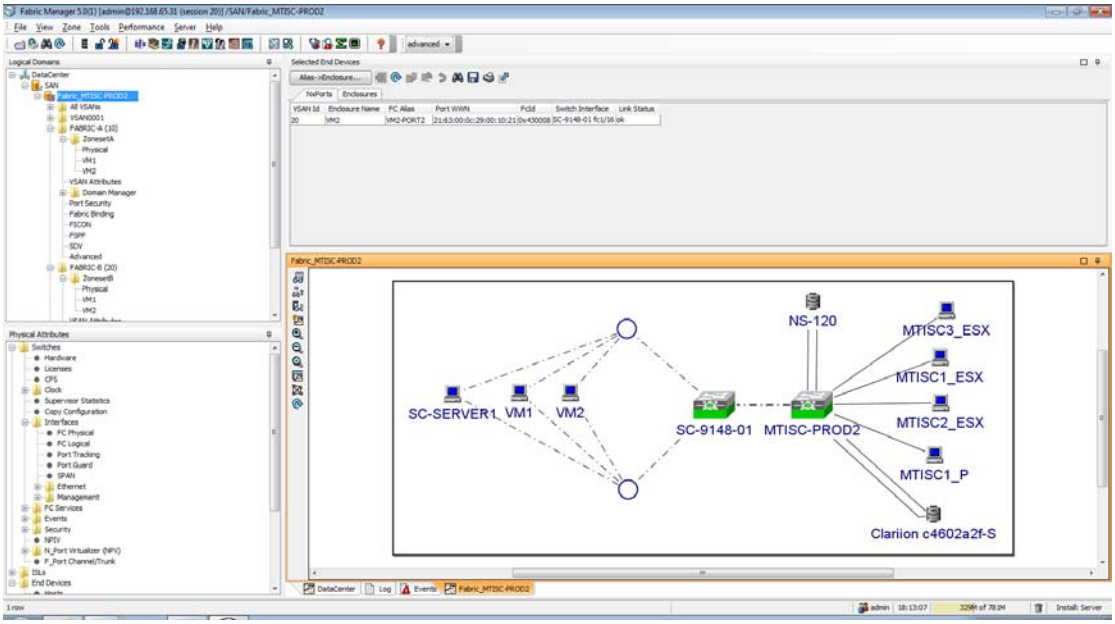

The VM2 was the machine configured with only one disk from our Storage Array. Booting from SAN can be observed in figure 16

5.4.2 Figure 16 – NPiV booting from SAN

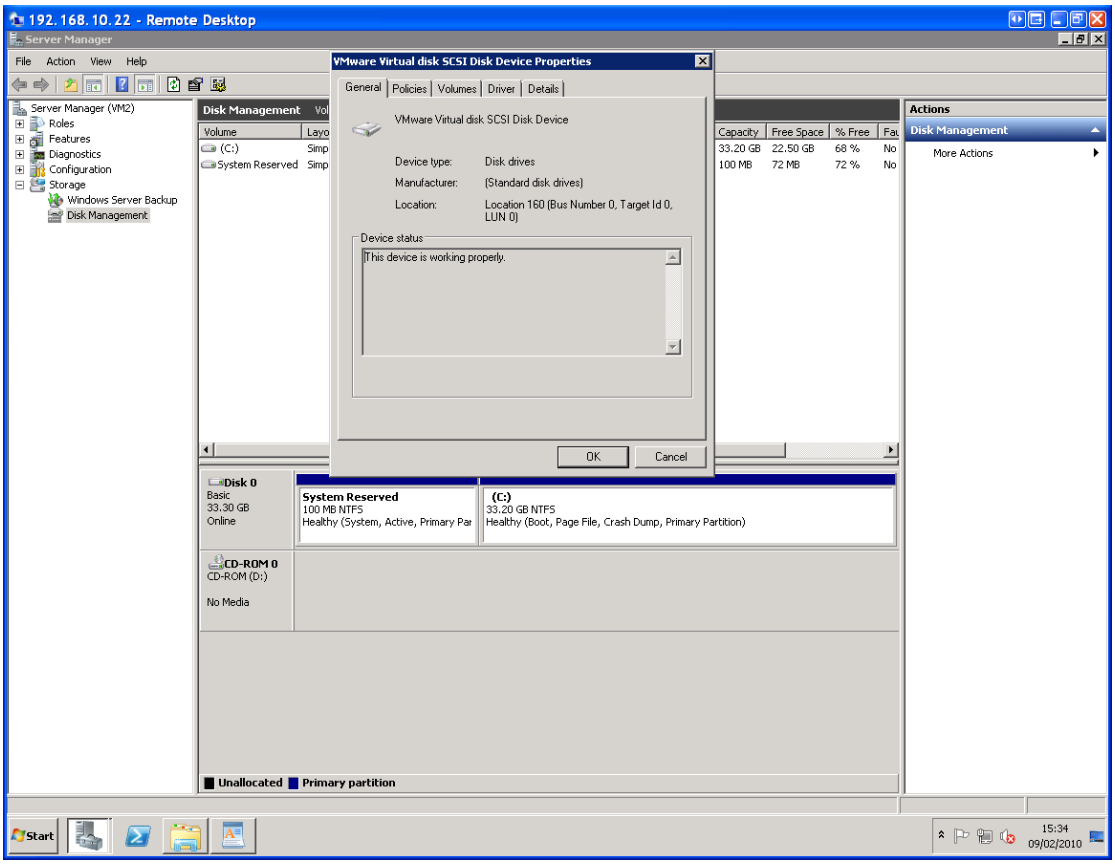

In the following two images, we can see the Zonesets including the interfaces needed for the correct zoning.

5.4.3 Figure 17 – NPiV Zonesets

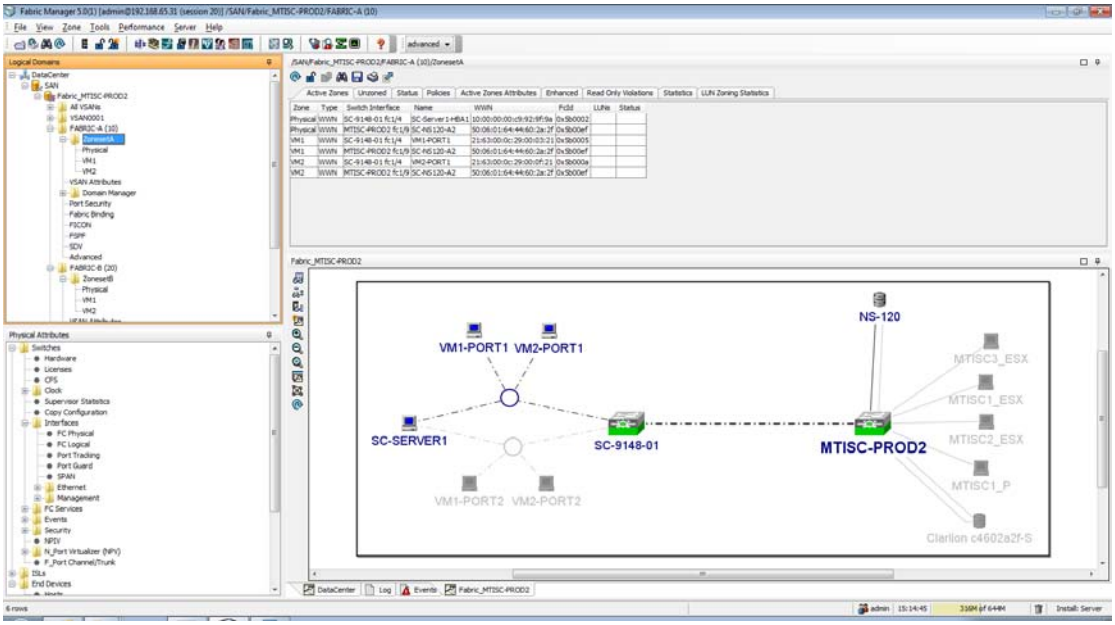

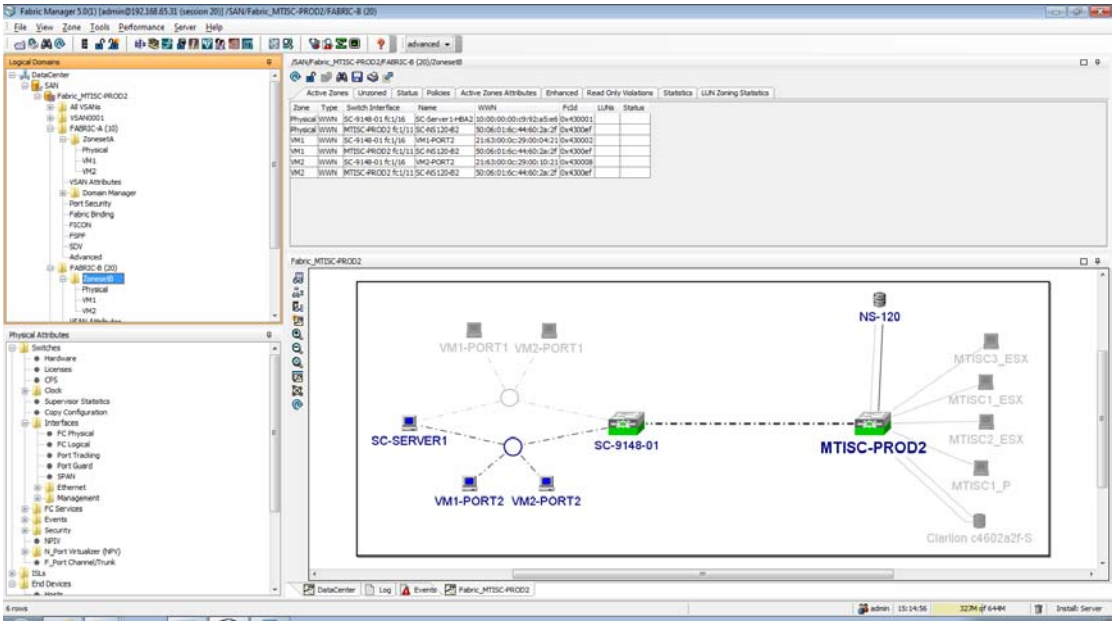

To finish, just a show of the Virtual Machines doing the Fabric Logins through the NPV EDGE Switch to the NPV CORE Switch in figures 18 and 19.

5.4.4 Figure 18 – Fabric Logins – NPV EDGE Switch

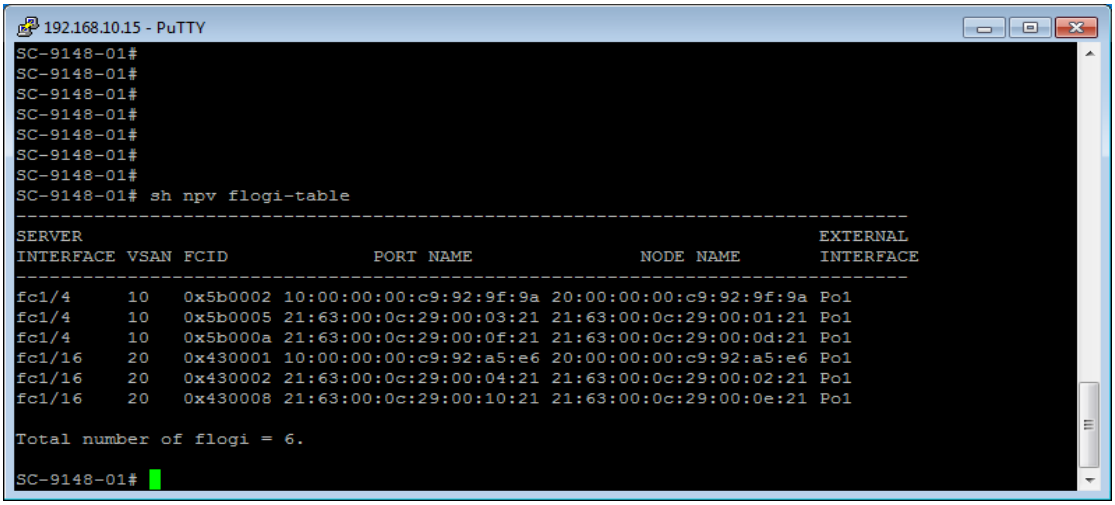

5.4.5 Figure 19 – Fabric Logins – NPV CORE Switch

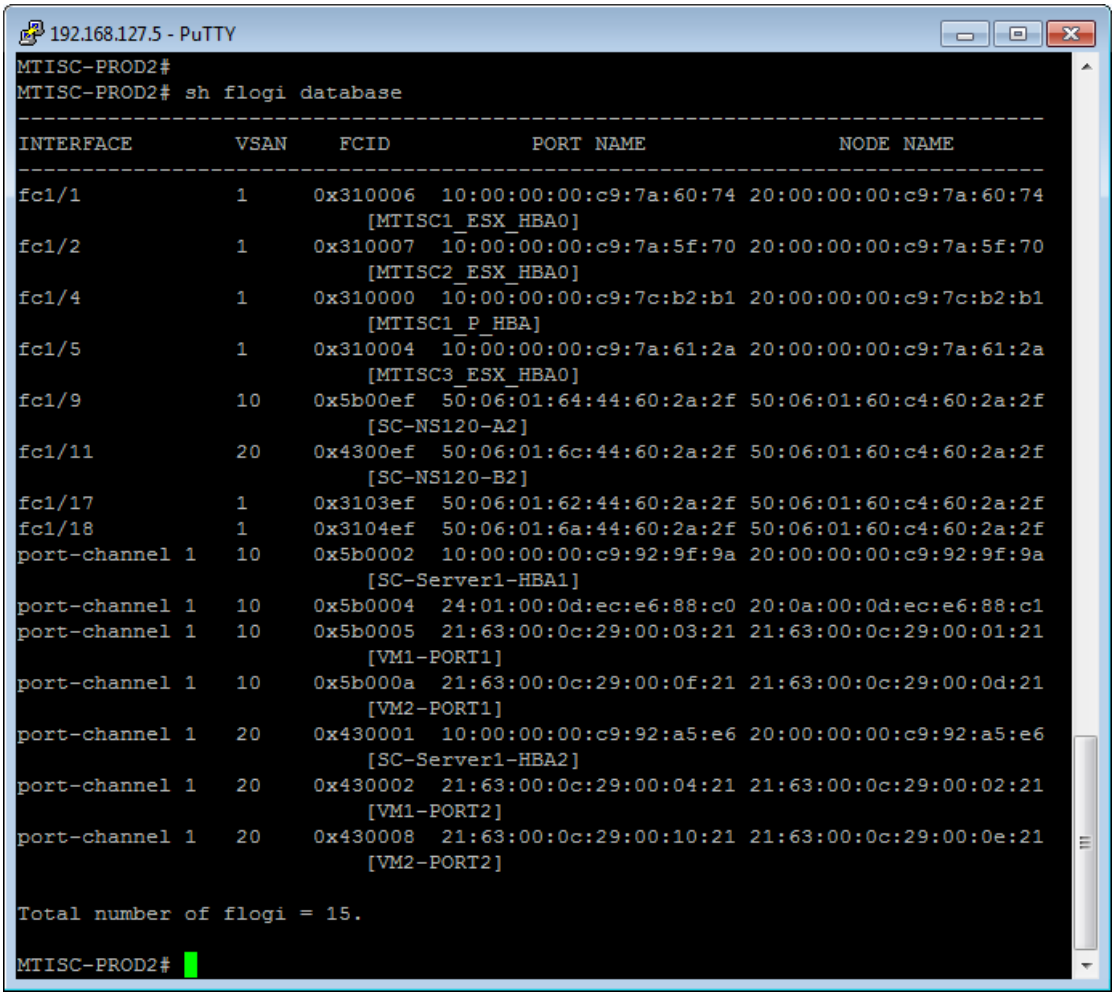

6 Next steps

Put the switch on the market ☺