Department of Computing Science

University of Alberta

MINT 709 Master of Science in Internetworking

Virtualization Performance Analysis

Dated: 13th July, 2012

Submitted To :

Professor Dr. Mike MacGregor

Submitted By :

Sheraz Ahmad

UNIVERSITY OF ALBERTA
MASTER OF INTERNETWORKING PROGRAM

AUTHORITY TO DISTRIBUTE REPORT

Title:

Virtualization Performance Analysis

The Master of Internetworking Program, University of Alberta may make available or authorize others to make photo/microfilm or soft copies of this report without restrictions after July 13th, 2012.

Signature of Author: Sheraz Ahmad

Date: July 13th, 2012

Table of Contents

LIST OF SYMBOLS AND ABBREVIATIONS

MPLS Multi-Protocol label switching

SAN Storage Area Network

VNC Virtual Network Computing

VPN Virtual Private Network

FTP File Transfer Protocol

iSCSI Internet Small Computer Systems Interface

OS Operating System

CentOS Community Enterprise Operating System

NFS Network File System

SMB Server Message Block

LDAP Lightweight Directory Access Protocol

BGP Border Gateway Protocol

IGP Interior Gateway Protocol

LIST OF FIGURES

LIST OF TABLES

Acknowledgement

Rome was not built in a day. To compile a thesis on the "Virtualization Performance Analysis" is not a child play. I feel myself lucky to work with Prof. Dr. Mike MacGregor –Professor and chair, department of computing science, University of Alberta, for his valuable suggestions, unparalleled nice way of criticism and encouragement.

I also extend my deepest sense of gratitude and appreciation to my nice teacher, ShahNawaz for his superb way to natural guidance, solid and ideal ideas, kind help and sincere involvement throughout my project. I recall in my mind the old sayings of wise person that a good teacher is an everlasting flower of all the generations. I express my profound thanks to my parents specially my father who is the solely responsible for enabling me to reach this milestone and achieve my Goals from the very beginning of my academic career.

I offer my sincere thanks to my younger brother Bilal. For his lovely habits which he extended to freshen me up whenever I felt tired. I'm also thankful to my class fellows and colleagues for their dreamlike company, legendary friendship and their professional conduct.

Finally, I would especially like to acknowledge the pearl like moments spent at MINT lab and student union hall with all my friends.

EXECUTIVE SUMMARY

This report is to present “**Virtual Server Performance Analysis over an MPLS Network”**. It compares the server's CPU and network performance under normal load and high load. It also aims to explain the basics along with the performance of SANs from a customer perspective. It also goes through overall configuration including security. A part of this report sheds light on benchmarking. Some of the other parts briefly discuss the Openfiler Operating System.

1.0 INTRODUCTION

This report contains the results achieved after configuring virtual servers over an MPLS network through utilization of tools like iozone and the built-in performance reporting mechanism in ESXI-4.1. It also demonstrates several physical server specifications, virtual server specifications, SAN configurations, performance analysis and benchmarking.

This report includes an analysis of network performance under normal conditions and when a large amount of data is being transferred. It explains the use of MPLS/VPN and sham link with emphasis on SANs and the basic configuration of the Openfiler OS.

1.1 Server Specifications

The servers used in my tests are as specified in Table 1, Figure 1 and Figure 2.

Table.1. Server Specifications

Feature	**Server-PE-1**	**Server-PE-2**
CPU type	Intel Xeon	Intel Xeon
Memory	4GB	16GB
Processor	3.36GHz	3.36GHz
Hard disk	300GB	200GB
L2 cache	1MB	1MB
System bus	800MHz	800MHz
OS support	64bit	64bit
LAN cards	Gigabit supported	Gigabit supported

Figure 1. Server-PE-1 Summary

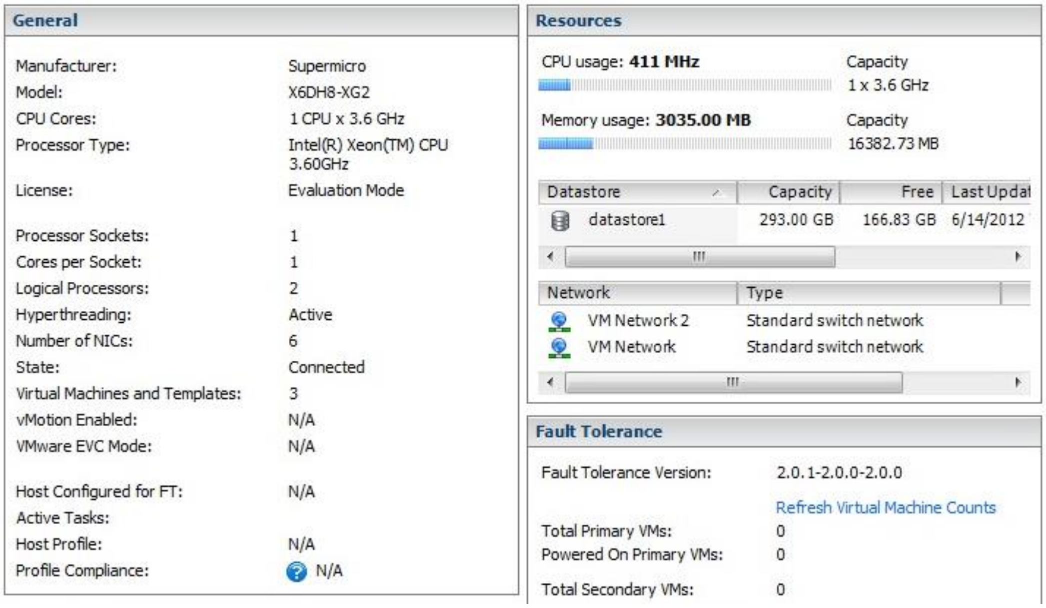

Figure 2. Server-PE-2 Summary

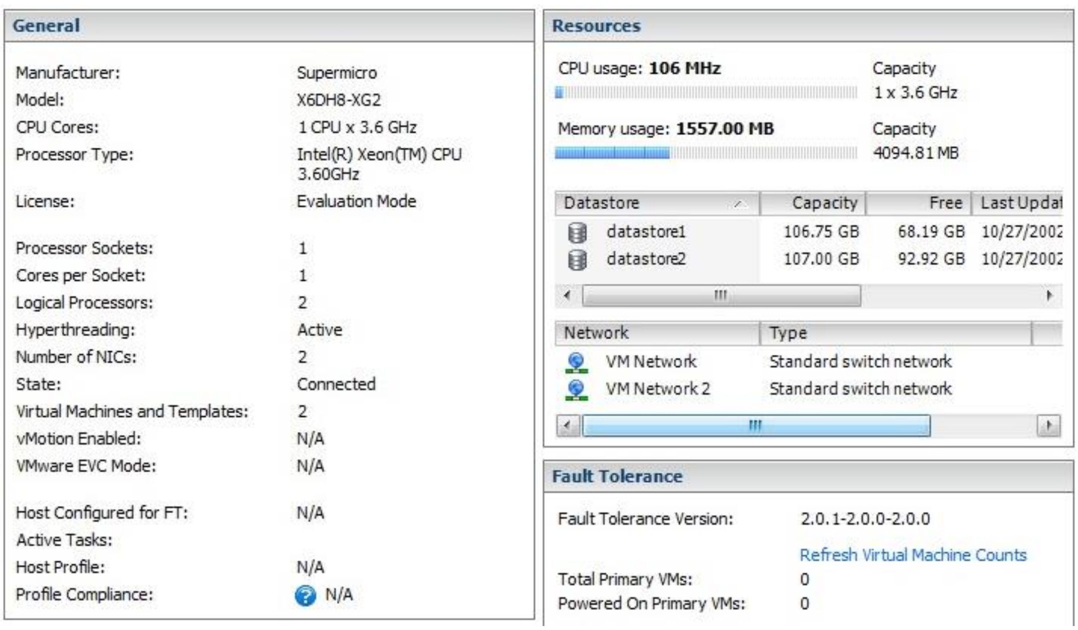

1.2 Network Device Specifications

The network devices used in my tests are as specified in Table 2.

Table. 2. Network Device Specifications

Features	2900-Routers	2800-Router	2600-Router
Brand	Cisco	Cisco	Cisco
NVRAM	255Kb	240Kb	32Kb
Flash	254464Kb	62720Kb	32768Kb
Ports	Gigabit	Gigabit	Fast

1.3 Performance Strategy

To improve the performance of the servers, the following guidelines were used.

- All RAM chips available in the lab were utilized to allocate more RAM per server.
- Dedicated a larger fraction of the physical server to the virtual servers.
- Used SCSI instead of IDE, because SCSI is faster than IDE.
- Used fixed size Virtual Hard Disk.
- For performance analysis, I used iozone and some built-in tools in VMware.

1.4 Network Diagram

Figure 3 shows an MPLS-enabled network. CE-1 is configured with OSPF to the service provider network (PE-1). PE-1 is configured with OSPF. PE-1 is also configured with BGP. An IBGP session is configured between PE-1 and PE-2. For provider network routing, an IGP is configured. For the purpose of this network, OSPF is used as the IGP. The provider network is configured with OSPF.

CE-2 is also configured with OSPF. A sham link is configured between PE-1 and PE-2 for redundancy. A layer 3 VPN is configured between CE-1 and CE-2 clients.

There are two physical servers connected to routers PE-1 and PE-2. Both servers and the CE clients are in the same VPN so that they can communicate with each other. This network design is scalable. We can add more servers for additional clients in different VPNs.

I have configured server PE-1 with virtual FTP, VNC and SAN servers while in PE-2 server, FTP and VNC have been configured. The SAN server is providing services to both CE-1 and CE-2 clients.

Figure 3. Network Design

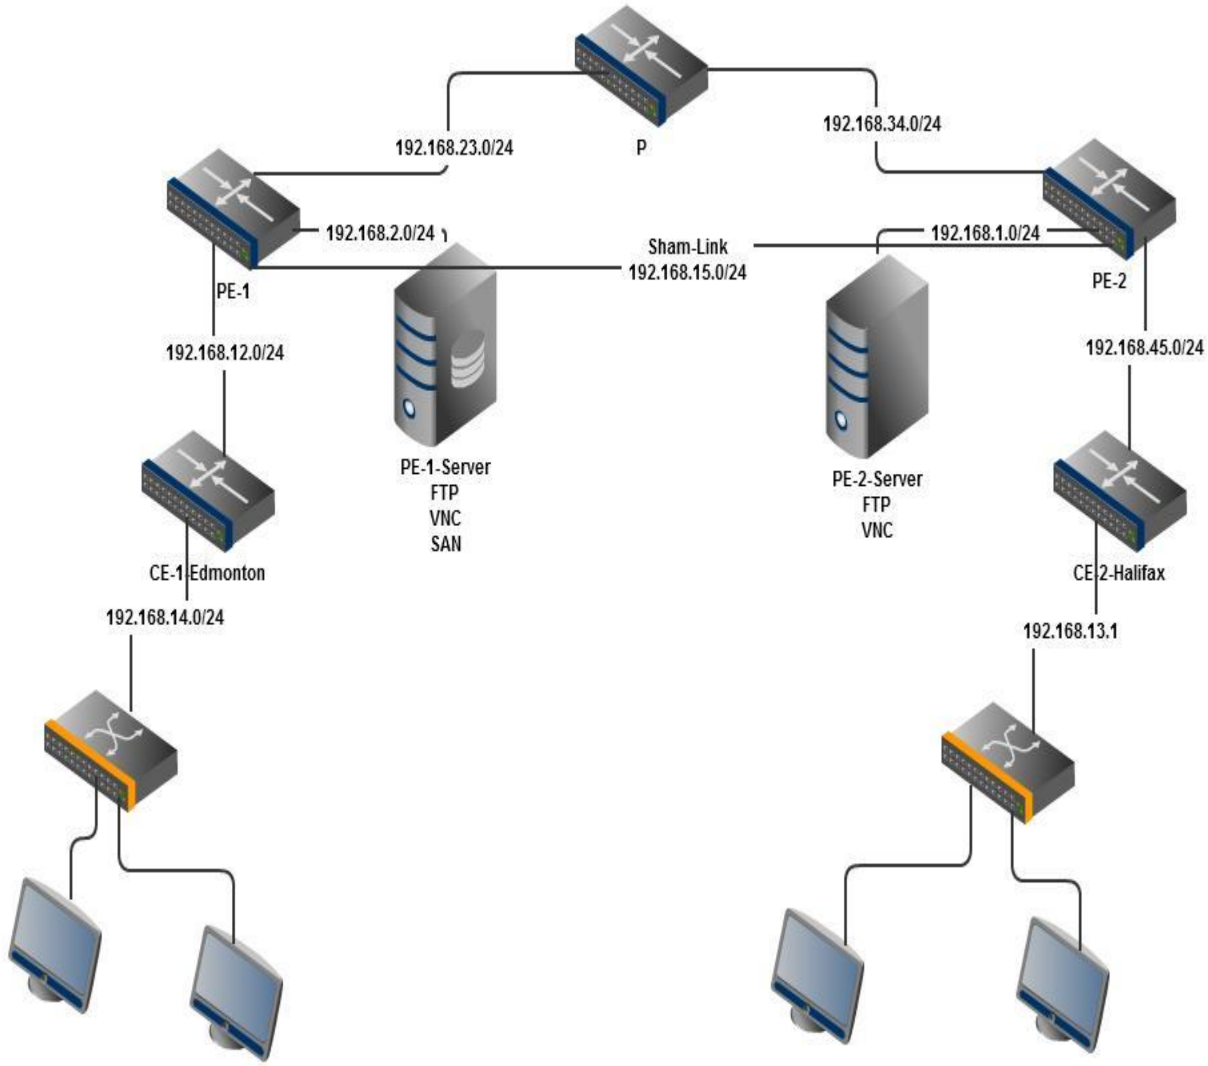

2.1 Virtual Servers Specifications

Below are the technical specifications of FTP, VNC and SAN servers. Figure 4 through Figure 8 show memory and other resource allocations to the virtual servers.

2.1.1 File Transfer Protocol Server

Resources allocated to the FTP server are as shown in Figure 4.

Figure 4. FTP Server

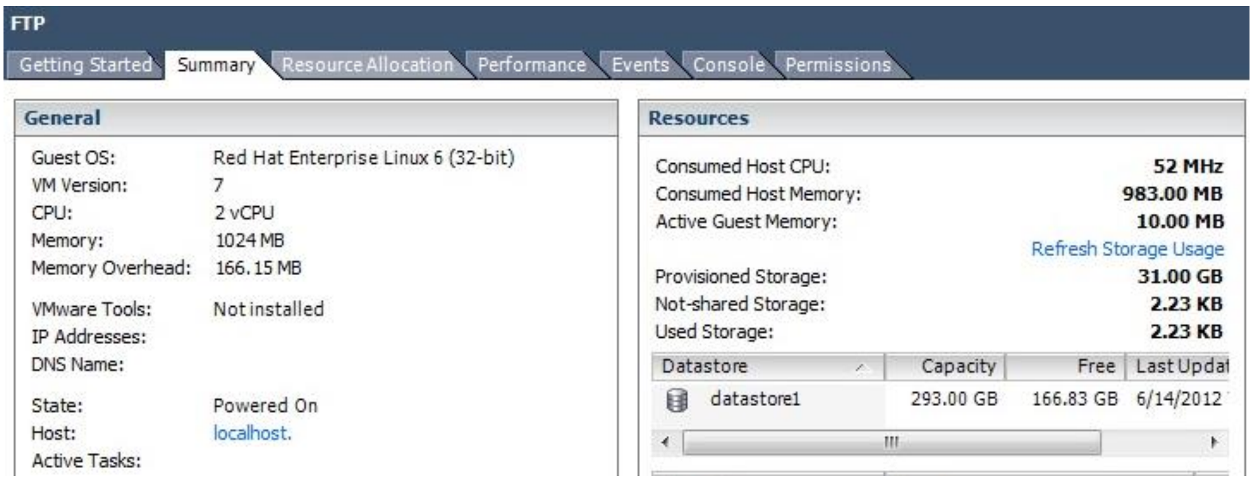

2.1.2 Virtual Network Computing Server

Resources allocated to the VNC server are as shown in Figure 5.

Figure 5. VNC Server

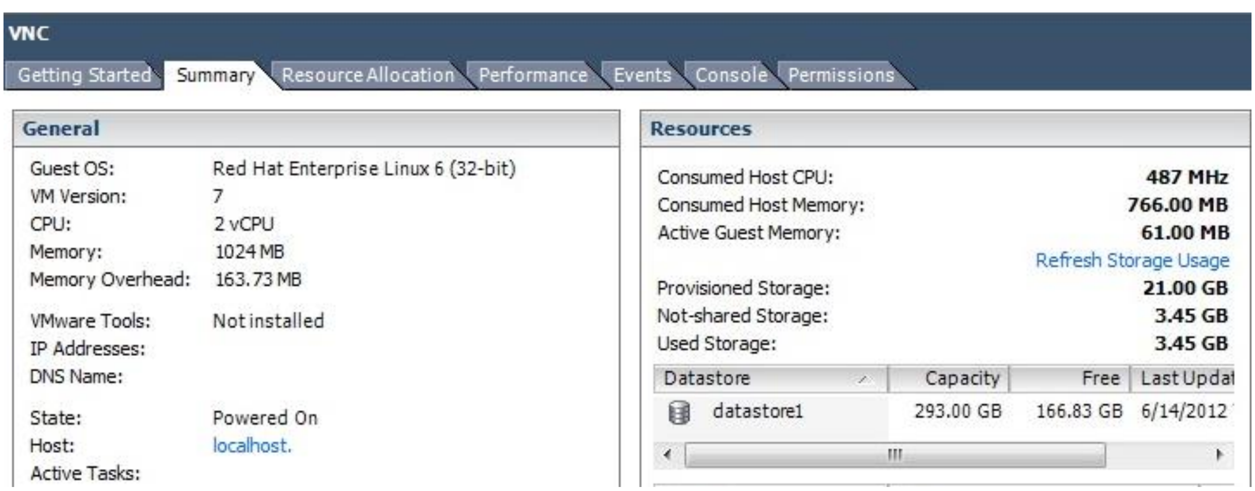

2.1.3 Storage Area Network Server

Resources allocated to the SAN server are as shown in Figure 6.

Figure 6. SAN Server

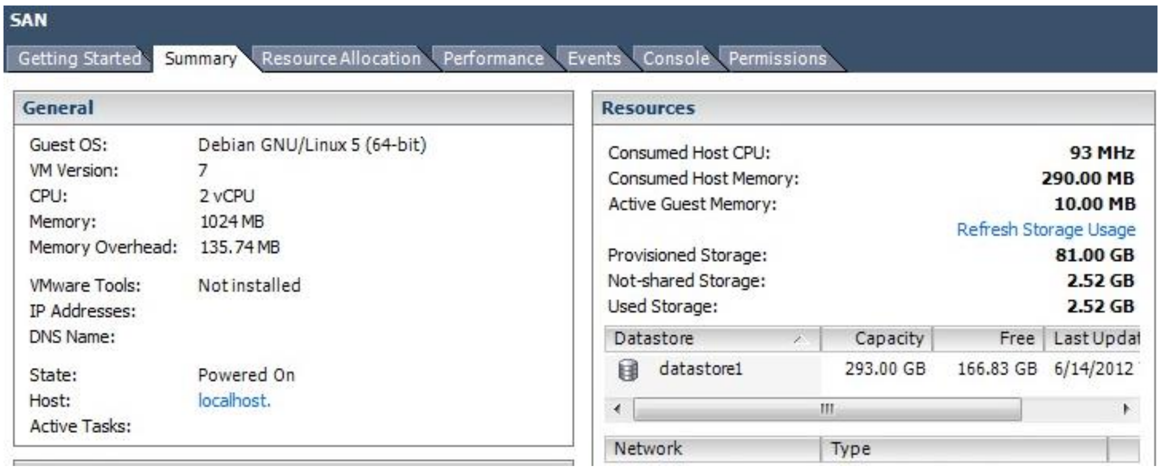

PE-2-Server-FTP

Resources allocated to the FTP server are as shown in Figure 7.

Figure 7. FTP Server

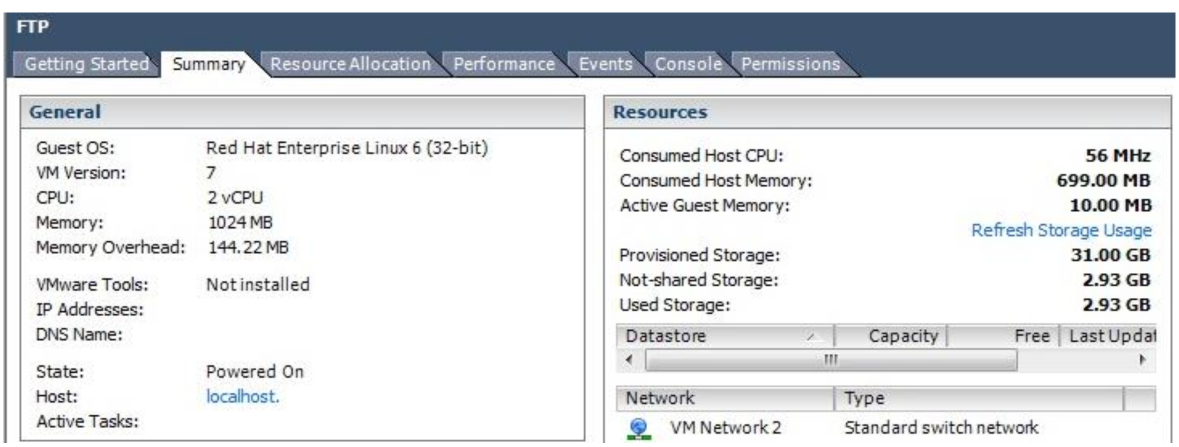

PE-2-Server-VNC

Resources allocated to the VNC server are as shown in Figure 8.

Figure 8. VNC Server

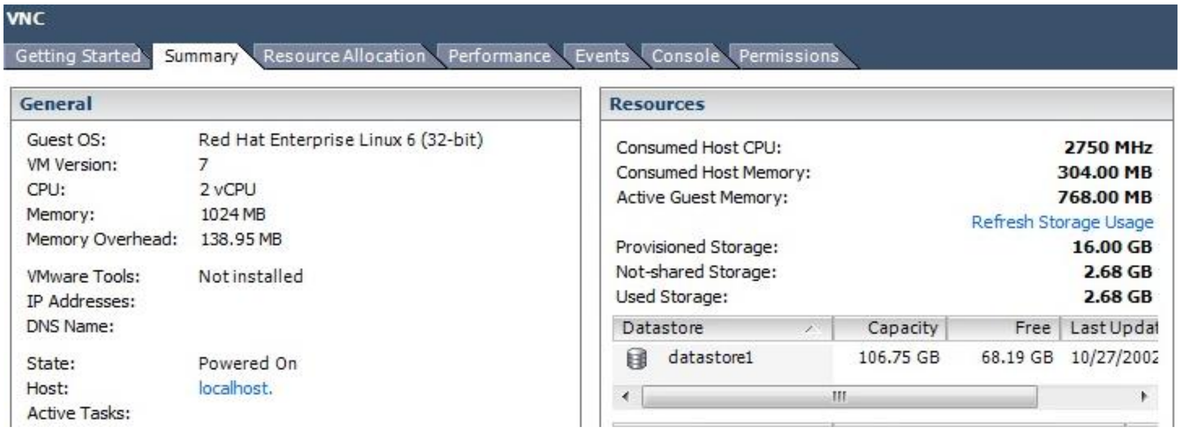

2.2 Network Configuration Strategy

Different technologies have been implemented to enhance network performance. These technologies are discussed below.

2.2.1 Multi-Protocol Label Switching Technology (MPLS)

MPLS is a packet forwarding technology that uses labels for forwarding decisions. With MPLS, the layer-3 header analysis is done just once to enter in the MPLS domain. MPLS could be used for the following reasons:

Virtual Private Networking (VPN)

VPN is a network technology that creates a secure network connection over a public network such as the Internet or a private network owned by a service provider. Large corporations, educational institutions and government agencies use VPN technology to enable remote users to securely connect to a private network. A VPN can connect multiple sites over a large distance just like a WAN. VPNs are often used to extend intranets worldwide to disseminate information and news to a wide user base. Educational institutions use VPNs to connect campuses that can be distributed across the country or around the world.

Traffic Engineering

Traffic engineering is a method of optimizing the performance of a telecommunications network by dynamically analyzing, predicting and regulating the behaviour of data transmitted over that network. Traffic engineering is also known as teletraffic engineering and traffic management. The techniques of traffic engineering can be applied to networks of all kinds, including the PSTN, LANs, WANs, cellular telephone networks, proprietary business and the Internet.

QoS enablement

One of the primary benefits of MPLS-based services is the ability to support QoS. It plays a vital role for companies that provide voice and video services.

Cost Savings

Depending on the specific mix of applications and network configuration, MPLS-based services can reduce costs by 10% to 25% over comparable data services (frame relay and ATM). As companies add voice and video traffic, cost savings can rise to as much as 40% network wide.

Improved Performance

Because of the “any-to-any” nature of MPLS services, network designers can reduce the number of “hops” between network points, which translates directly to increased response time and improved application performance.

Disaster Recovery

MPLS-based services improve disaster recovery in a variety of ways. First and foremost, data centers and other key sites can be connected in multiply redundant ways to the cloud (and thus to other sites on the network). Secondly, remote sites can quickly and easily reconnect to backup locations if needed, unlike with ATM and frame networks, in which either switched or backup permanent-virtual-circuits are required.

In this project i have focused on the use of ***VPNs***.

2.2.2 Layer-3 VPN

A VPN is a private data network that makes use of the public network infrastructure while maintaining privacy and reservation through the use of a tunnelling protocol based on an MPLS network. In my project I designed the VPN in such a way that all servers and customer edge routers are in the same VPN so that the servers can be reachable from every client behind the customer routers.

<u>Ping result from CE-1-Client station to :</u>

(a) PE-1-Server

Figure 9 illustrates the ping result from CE-1 client to PE-1 server.

Figure 9. Ping PE-1 Server

```
C:\Users\Administrator>ping 192.168.2.3

Pinging 192.168.2.3 with 32 bytes of data:
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60

Ping statistics for 192.168.2.3
    Packets: Sent = 4, Received [illegible]st = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 1ms, Average = 1ms
```

(b) PE-2-Server

Figure 10 illustrates the ping result from CE-1 client to PE-2 server.

Figure 10. Ping PE-2 Server

```
C:\Users\Administrator>ping 192.168.1.5

Pinging 192.168.1.5 with 32 bytes of data:
Reply from 192.168.1.5: bytes=32 time=2ms TTL=62
Reply from 192.168.1.5: bytes=32 time<1ms TTL=62
Reply from 192.168.1.5: bytes=32 time<1ms TTL=62
Reply from 192.168.1.5: bytes=32 time<1ms TTL=62

Ping statistics for 192.168.1.5:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 2ms, Average = 0ms
```

<u>Ping Result from CE-2 Client Station to :</u>

(a) PE-1-Server

Figure 11 illustrates the ping result from CE-2 client to PE-1 server.

Figure 11. Ping from CE-2 to PE-1 Server

```
C:\Users\Administrator>ping 192.168.2.3

Pinging 192.168.2.3 with 32 bytes of data:
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60
Reply from 192.168.2.3: bytes=32 time=1ms TTL=60

Ping statistics for 192.168.2.3
    Packets: Sent = 4, Received [illegible]t = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 1ms, Average = 1ms
```

(b) PE-2-Server

Figure 12 illustrates the ping result from CE-2 client to PE-1 server.

Figure 12. Ping from CE-2 to PE-1 Server

```
C:\Users\Administrator>ping 192.168.1.5

Pinging 192.168.1.5 with 32 bytes of data:
Reply from 192.168.1.5: bytes=32 time=2ms TTL=62
Reply from 192.168.1.5: bytes=32 time<1ms TTL=62
Reply from 192.168.1.5: bytes=32 time<1ms TTL=62
Reply from 192.168.1.5: bytes=32 time<1ms TTL=62

Ping statistics for 192.168.1.5:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 2ms, Average = 0ms
```

2.2.3 Sham Link

A sham link in OSPF is used in case MPLS network goes down. When we configure a backdoor sham link, OSPF gives priority to that link, which we actually don’t need when we have an MPLS network. This link should be configured in such a way that it becomes active only if the MPLS network goes down.

Figure 13. Sham link configuration at PE-1 Router

```
!
!
router ospf 2 vrf customer
 domain-id 0.0.0.1
 area 0 sham-link 22.22.22.22 44.44.44.44
 redistribute static
```

Figure 14. Sham link configuration at PE-2 Router

```
router ospf 2 vrf customer
 domain-id 0.0.0.5
 log-adjacency-changes
 area 0 sham-link 44.44.44.44 22.22.22.22
 redistribute static
```

3.0 Introduction to Storage Area Networks

A SAN's primary purpose is to transfer data between computer systems and storage elements. A SAN consists of a communication infrastructure, which provides physical connection, and a management layer that organizes the connection, storage elements and computer systems so that data transfer is secure and robust. The term SAN is usually identified with block I/O service rather than file access service.

A SAN can also be a storage system consisting of storage elements, storage devices, computer systems and/or applications plus all control software. In this project I designed a SAN with an MPLS network rather than a traditional LAN. The SAN server is connected to the PE router and is accessible throughout the MPLS Network.

3.1 Openfiler

To configure the SAN, I used Openfiler, which is an operating system. Openfiler can use a lot of characteristics to share files. When you have set up the application you have the possibility to develop disk volumes and share information from the SAN.

For example, this configuration can serve the objective of creating an iSCSI SAN. Numerous users use this software with dedicated machines in VMware workstation and ESX server. It is supplied in 32-bit, 64-bit, ISO, VMware ESX server and XEN appliances. It is a totally free application that can be put in every single laptop or computer ranging from the previous ones to a high efficiency tower. It supports LDAP and has volume-management, enabling simple storage management.

Openfiler is a storage management approach, and is a standalone Linux distribution. It supports all the main network directories. It is also easy to use. Once installed, we can access it via a Web interface and configure it graphically according to our needs.

3.2 Openfiler Installation and Configuration

The step by step installation procedure is:

1. Download ISO for Openfiler and connect to virtual server. Upon restarting the 1st screen that appears is shown in Figure 15.

Figure 15. Setup step-1

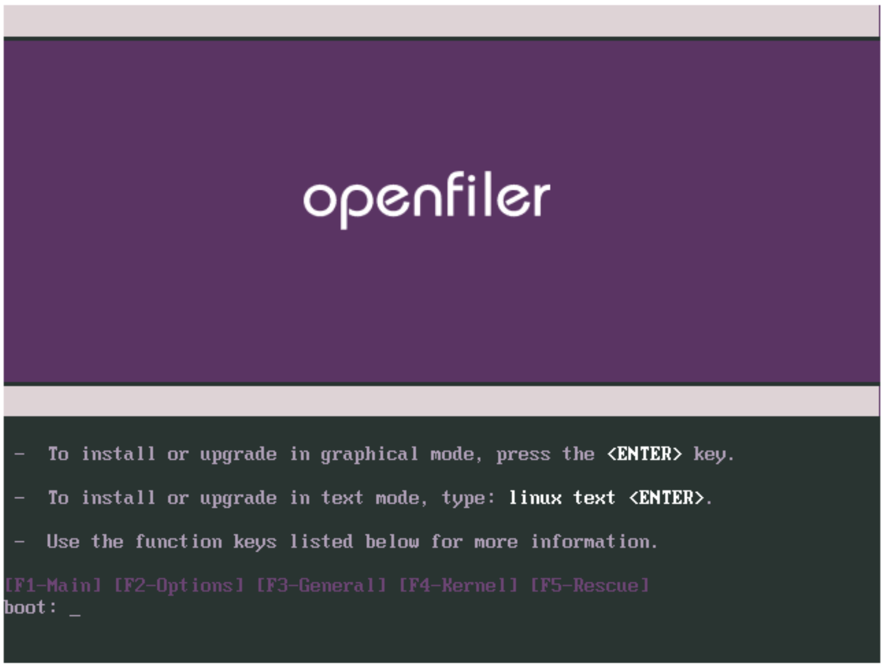

2. Press ENTER.

Figure 16. Setup step-2

3. Click next and select a keyboard.

Figure 17. Setup step-3

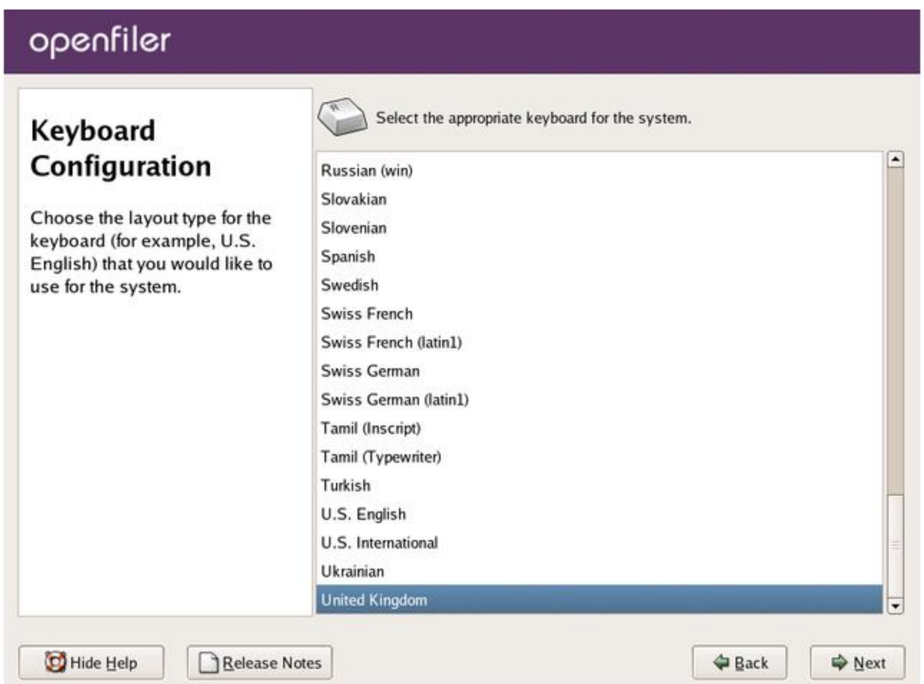

4. Click next and configure partitions.

Figure 18. Setup step-4

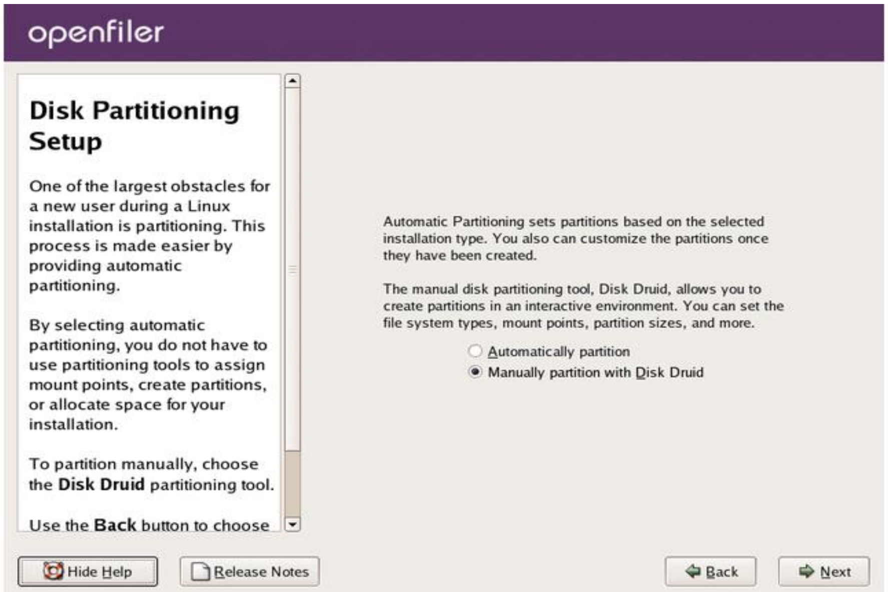

5. Click next and configure networking.

Figure 19. Setup step-5

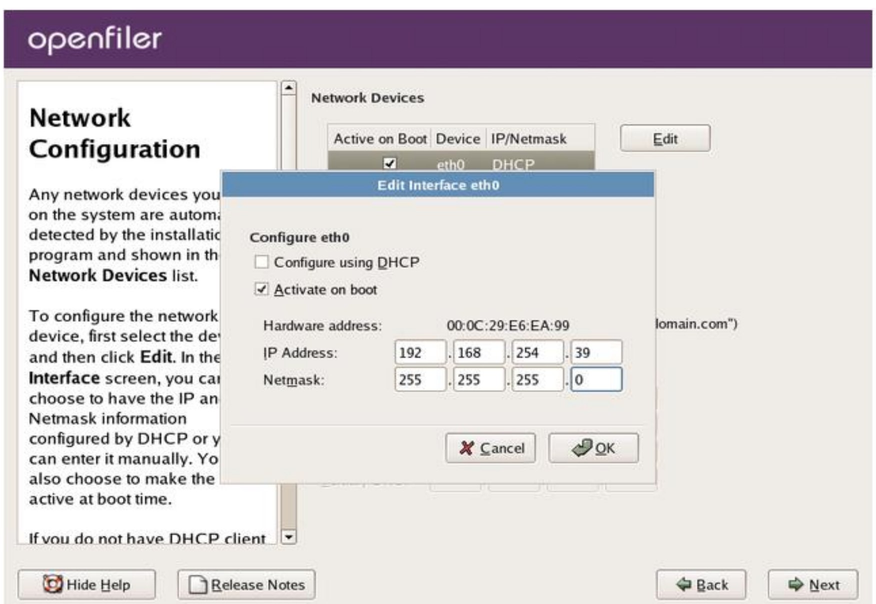

6. Click next and set Root password.

Figure 20. Setup step-6

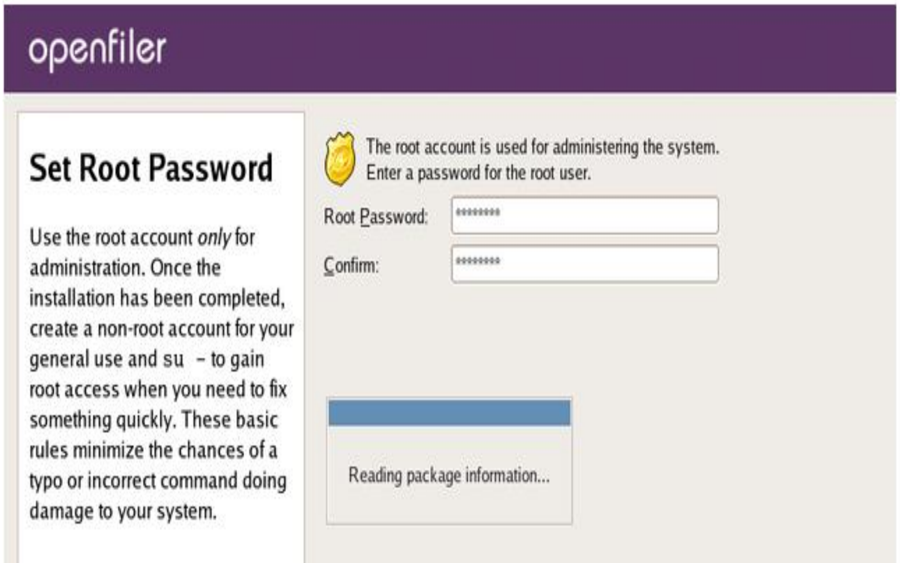

7. Make sure to disconnect .ISO before clicking reboot.

Figure 21. Setup step-7

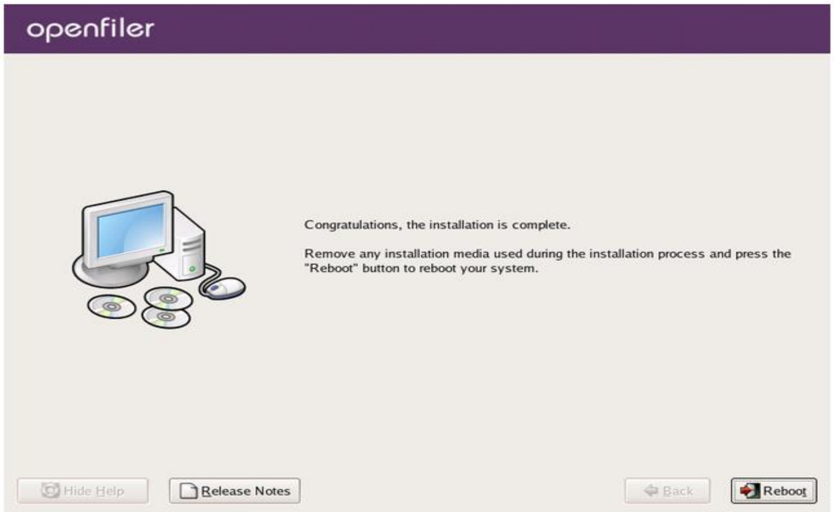

8. Now we can access Openfiler via web interface by entering Openfiler IP address and port 446.

Figure 22. Setup step-8

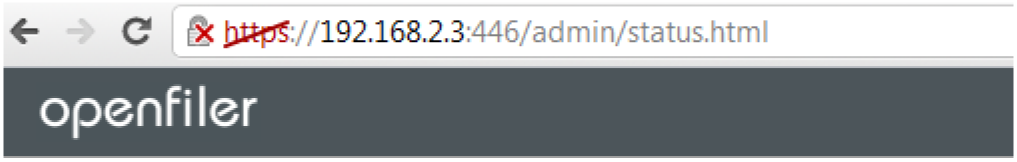

Configuration:

1. Initial Information/Status page.

Figure 23. System information

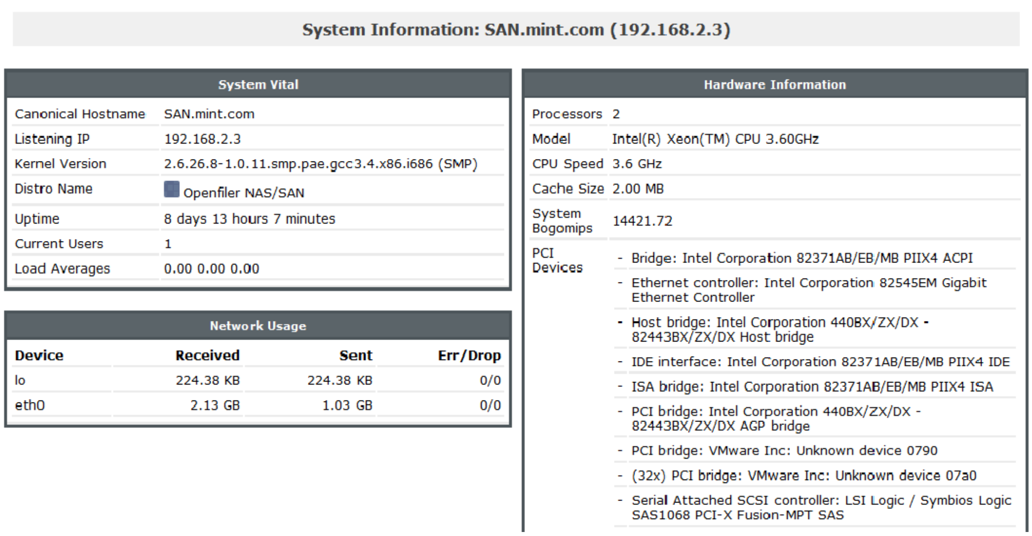

System Information: SAN.mint.com (192.168.2.3)

System Vital	
Canonical Hostname	SAN.mint.com
Listening IP	192.168.2.3
Kernel Version	2.6.26.8-1.0.11.smp.pae.gcc3.4.x86.i686 (SMP)
Distro Name	Openfiler NAS/SAN
Uptime	8 days 13 hours 7 minutes
Current Users	1
Load Averages	0.00 0.00 0.00

Network Usage

Device	Received	Sent	Err/Drop
lo	224.38 KB	224.38 KB	0/0
eth0	2.13 GB	1.03 GB	0/0

Hardware Information	
Processors	2
Model	Intel(R) Xeon(TM) CPU 3.60GHz
CPU Speed	3.6 GHz
Cache Size	2.00 MB
System Bogomips	14421.72
PCI Devices	- Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI - Ethernet controller: Intel Corporation 82545EM Gigabit Ethernet Controller - Host bridge: Intel Corporation 440BX/ZX/DX - 82443BX/ZX/DX Host bridge - IDE interface: Intel Corporation 82371AB/EB/MB PIIX4 IDE - ISA bridge: Intel Corporation 82371AB/EB/MB PIIX4 ISA - PCI bridge: Intel Corporation 440BX/ZX/DX - 82443BX/ZX/DX AGP bridge - PCI bridge: VMware Inc: Unknown device 0790 - (32x) PCI bridge: VMware Inc: Unknown device 07a0 - Serial Attached SCSI controller: LSI Logic / Symbios Logic SAS1068 PCI-X Fusion-MPT SAS

80 GB virtual hard disk was used.

Figure 24. Disk Information

Edit Disk	Type	Description	Size	Label type	Partitions
/dev/sda	SCSI	VMware Virtual disk	80.00 GB	msdos	9 (view)

2. Two volumes were created. The 1st one is shown below.

Figure 25. Disk Vol-1

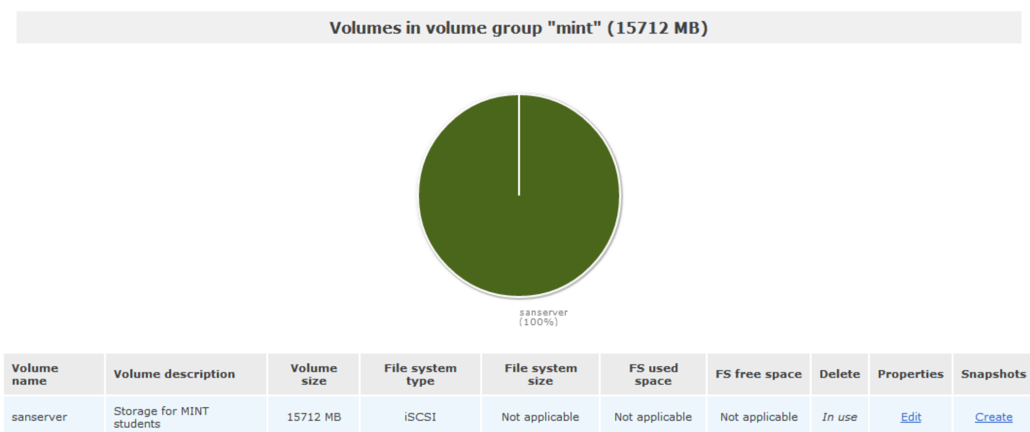

Volume name	Volume description	Volume size	File system type	File system size	FS used space	FS free space	Delete	Properties	Snapshots
sanserver	Storage for MINT students	15712 MB	iSCSI	Not applicable	Not applicable	Not applicable	*In use*	Edit	Create

The 2nd volume is.

Figure 26. Disk Vol-2

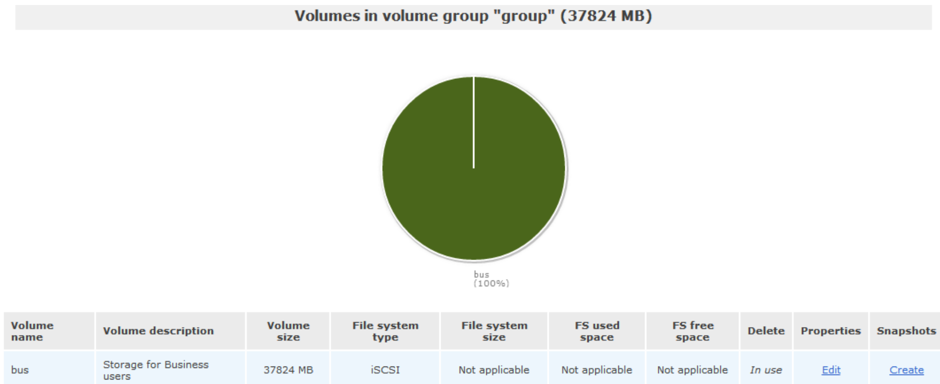

Volumes in volume group "group" (37824 MB)

Volume name	Volume description	Volume size	File system type	File system size	FS used space	FS free space	Delete	Properties	Snapshots
bus	Storage for Business users	37824 MB	iSCSI	Not applicable	Not applicable	Not applicable	In use	Edit	Create

3. RAID is configured.

Figure 27. Raid Vol

Software RAID Management

Array	Level	Array Size	Device Size	State	Synchronization	Manage	Add	Used In	Delete
/dev/md0	RAID-5	15.35 GB	7.68 GB	Clean	Synchronized	View members	All RAID partitions are used	mint VG	In use

3. Before setting up an iSCSI target, make sure that service is enabled under "SERVICES" as shown in figure below.

Figure 28. SAN Services

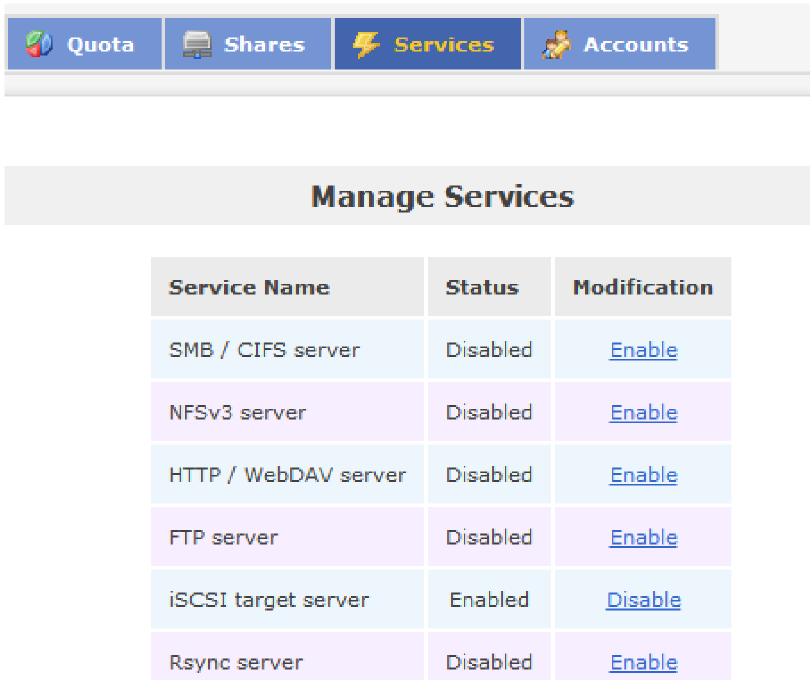

4. Now click on iSCSI target setup and configure it.

Figure 29. iSCSI Target Configuration

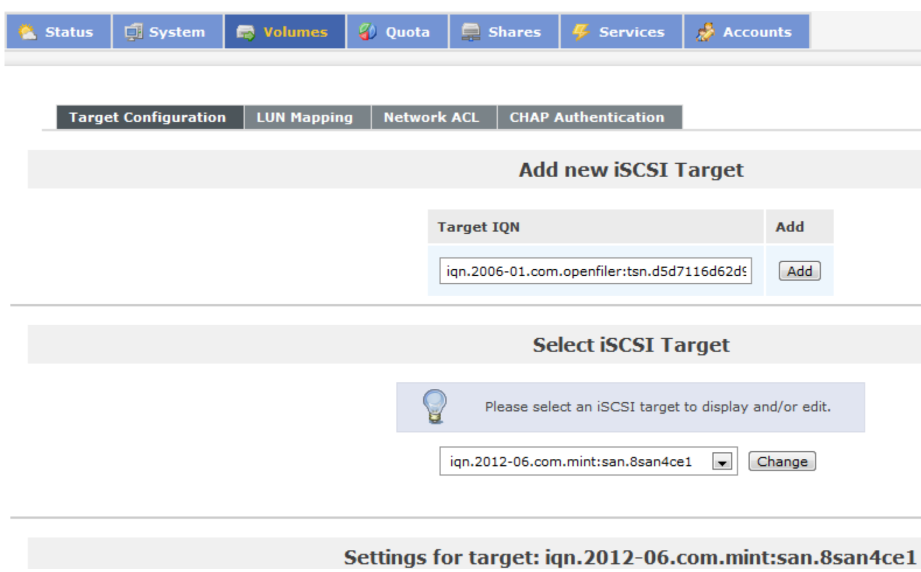

5. Click on Lun Mapping and map it.

Figure 30. Lun Mapping

Map New LUN to Target: "iqn.2012-06.com.mint:san.8san4ce1"

Name	LUN Path	R/W Mode	SCSI Serial No.	SCSI Id.	Transfer Mode	Map LUN
Storage for MINT students	/dev/mint/sanserver	write-thru	3dH4Ej-D5oK-9NJq	3dH4Ej-D5oK-9NJq	blockio	Map
Storage for Business users	/dev/group/bus	write-thru	6U6ngD-zvZI-LIFd	6U6ngD-zvZI-LIFd	blockio	Map

For security purposes CHAP authentication is added.

Figure 31. CHAP Setup

Now it's time to setup client as an iSCSI initiator. In my case I used Windows-7 and Ubuntu as an iSCSI initiator.

In Windows-7 click **start** to get the administrative tools and then click on iSCSI initiator.

Figure 32. ISCSI Initiator

When the screen in Figure 33 appears, write down the SAN server IP and click connect.

Figure 33. IP Configuration

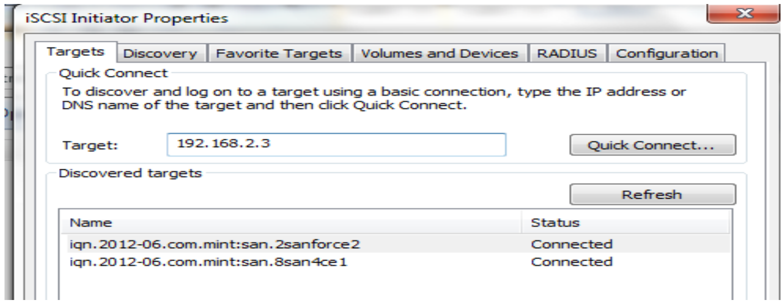

Now right click on the My Computer icon. Then click on manage and select "Disk Management". The screen shown in Figure 34 will appear.

Figure 34. Disk Management

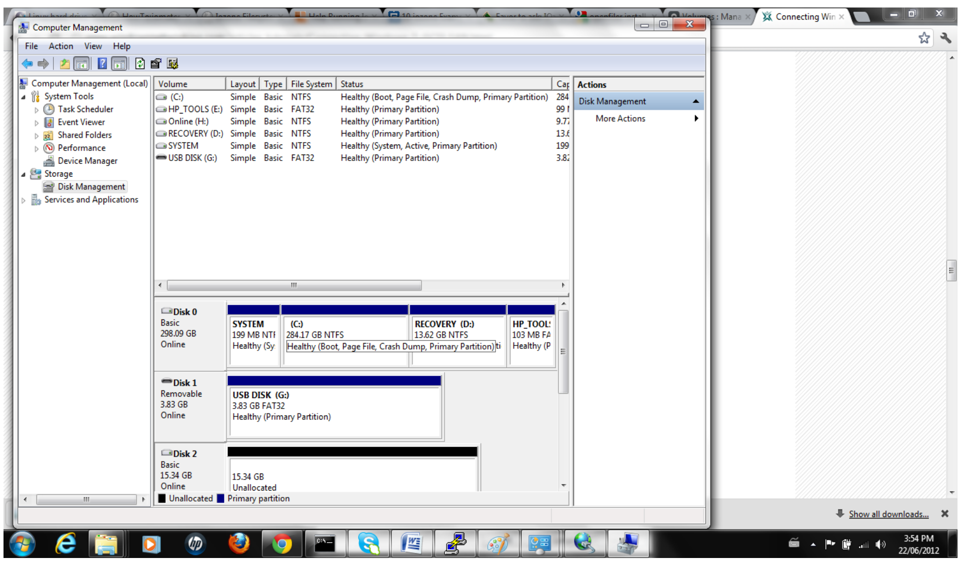

On this screen we can create as many drives as we want. They will all look like they are locally attached. In my case I created drive H named "online" for testing purposes.

Figure 35. My Computer

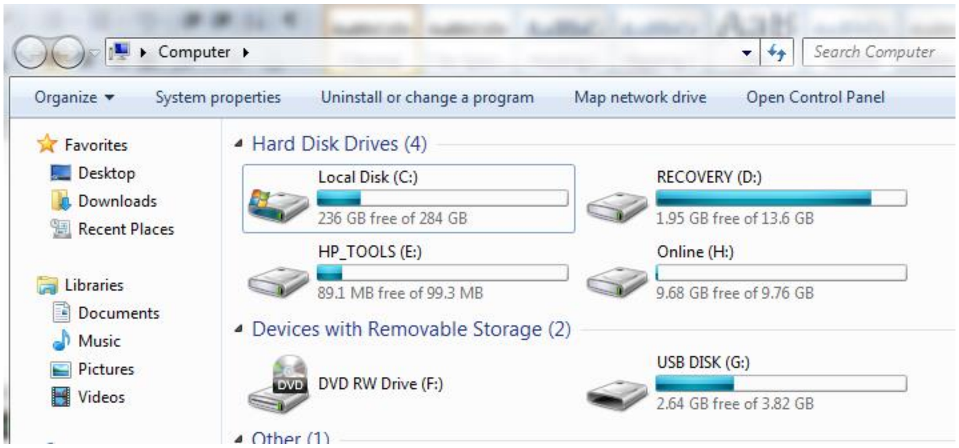

3.3 Setting up Ubuntu as an iSCSI initiator

Open a terminal window and issue the following command:

iscsiadm –m discovery sendTargets –p 192.168.2.3

Once discovered, we can create partitions same way as we did in Linux.

4.1 Performance Analysis

In general performance means the accomplishment of a given task measured against present known standards of accuracy, completeness, cost and speed.

4.1.1 Client’s Perspective

From the client’s perspective, performance means how quickly and accurately the requested services are provided by the server.

4.1.2 Server’s Perspective

From the server’s perspective, performance means how quickly and accurately the requested services are being performed by the server. Server performance is said to be satisfactory if it performs many requests in a short amount of time.

4.2 Servers-PE-1 Performance Analysis

Server performance has been analyzed under normal and high load of data transfer and is discussed below.

4.2.1 Under Normal Conditions

(i) File Server: File is transferred at speed of 10.3 Mbps as shown in Figure 36.

Figure 36. File Transfer Status

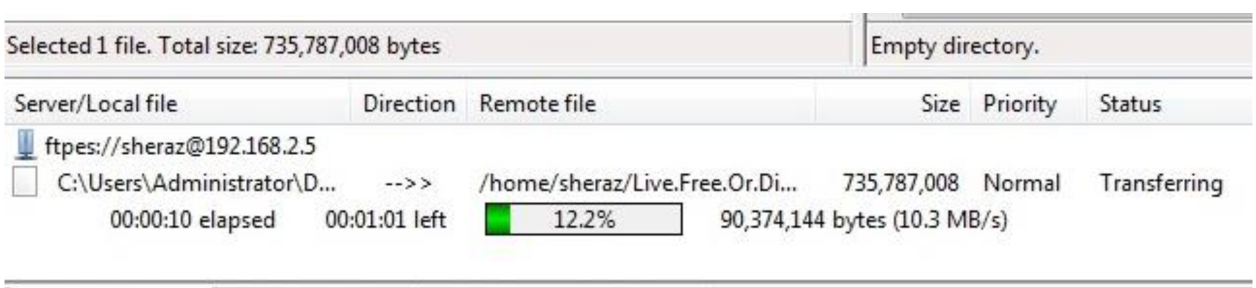

Figure 37 shows that CPU usage is almost 45% while 7 GB was transferred.

Figure 37. CPU Performance

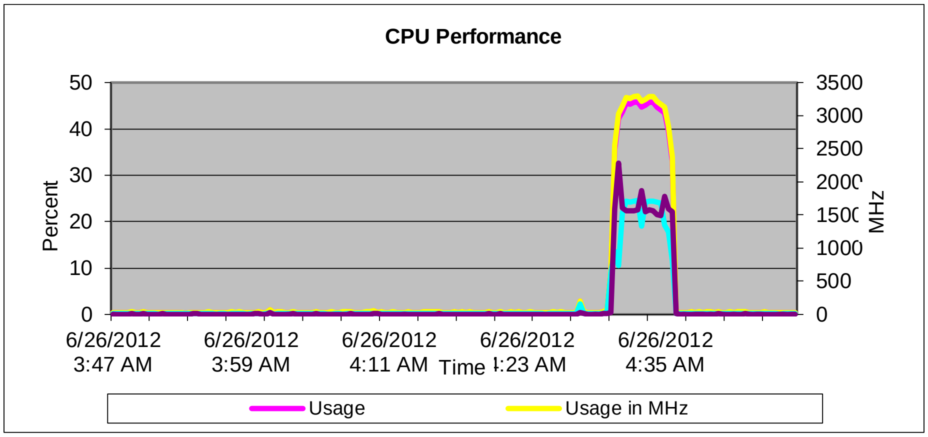

Figure 38 demonstrates that data was received at the rate of 2 Mbps, with a maximum of 8.5 Mbps.

Figure 38. Network Performance

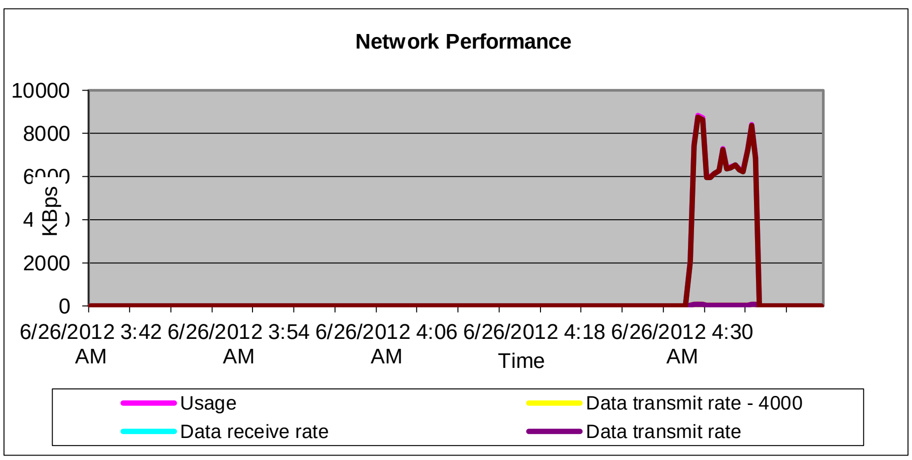

Table 3 elaborates that the data transfer started at 4:31AM. Initially data was received at the rate of 2 Mbps which reached a maximum of 8.5 Mbps at the time slot between 4:31:20 to 4:31:40. This means the data rate reached a maximum in 10 seconds. This is documented in Table 3.

Table 3. Data Transmission

Time	Usage	Data transmit rate - 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
6/26/2012 4:30:40 AM	0	0	0	0	0
6/26/2012 4:31:00 AM	2086	26	2059	26	2059
6/26/2012 4:31:20 AM	7512	70	7442	70	7442
6/26/2012 4:35:00 AM	6356	31	6324	31	6324
6/26/2012 4:36:40	0	0	0	0	0

(ii) Virtual Network Computing (VNC) Server

In this case two VNC connections were established to the VNC server from the Windows-7 client. Figure 39 illustrates the CPU usage, which is almost 50%. Even at this load, the key strokes and mouse responses was still very quick. Response time was totally acceptable for users.

Figure 39. CPU Performance

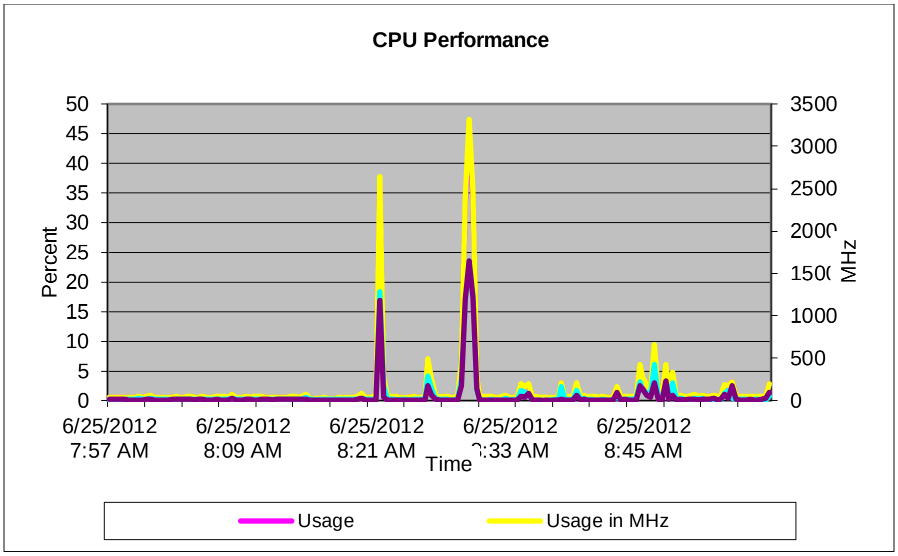

Figure 40 illustrates no conspicuous load on network when establishing VNC connections. Initially the rate at which the server received data was only 2 kbps. This reached a maximum of 6 kbps.

Figure 40. Network Performance during VNC

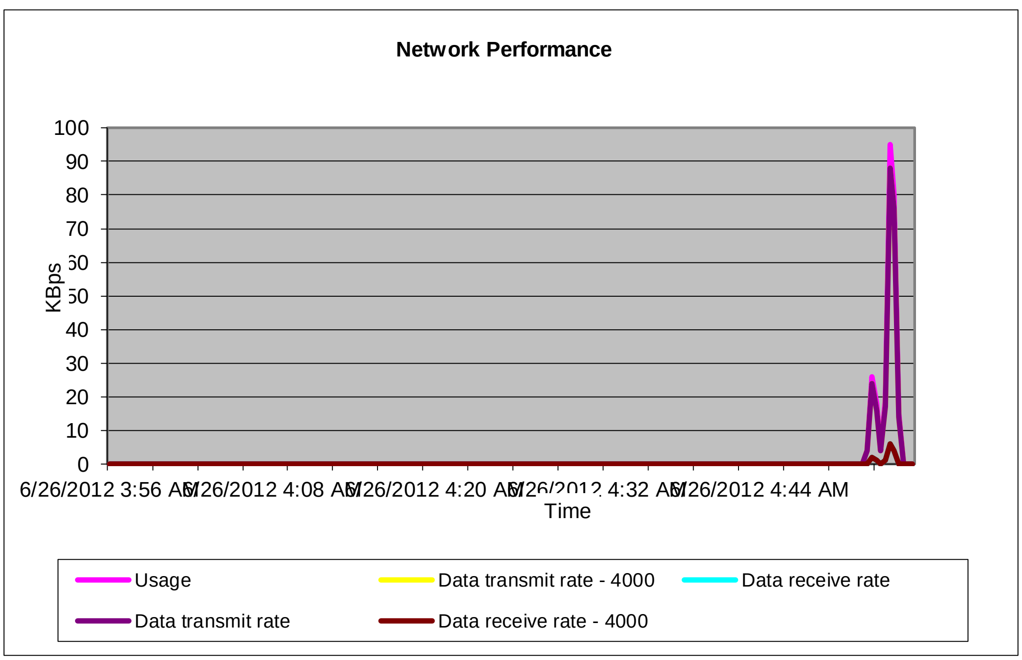

Table 4 shows that a few kilobytes were transferring, and network usage is 95 kbps.

Table 4. Data Transmission

Time	Usage	Data transmit rate - 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
6/26/2012 4:52:20 AM	26	24	2	24	2
6/26/2012 4:52:40 AM	18	16	1	16	1
6/26/2012 4:53:20 AM	19	17	1	17	1
6/26/2012 4:53:40 AM	95	88	6	88	6

(iii) Storage Server

Figure 41 shows a file transfer from client to SAN server.

Figure 41. Data Transfer Rate

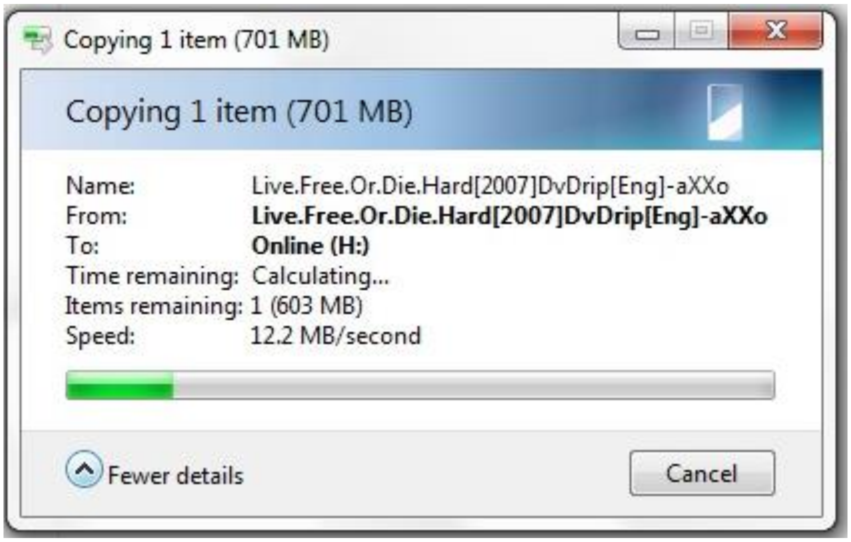

Figure 42 demonstrates a clear 10% CPU usage when copying a 7 GB file from client to SAN.

Figure 42. CPU Performance

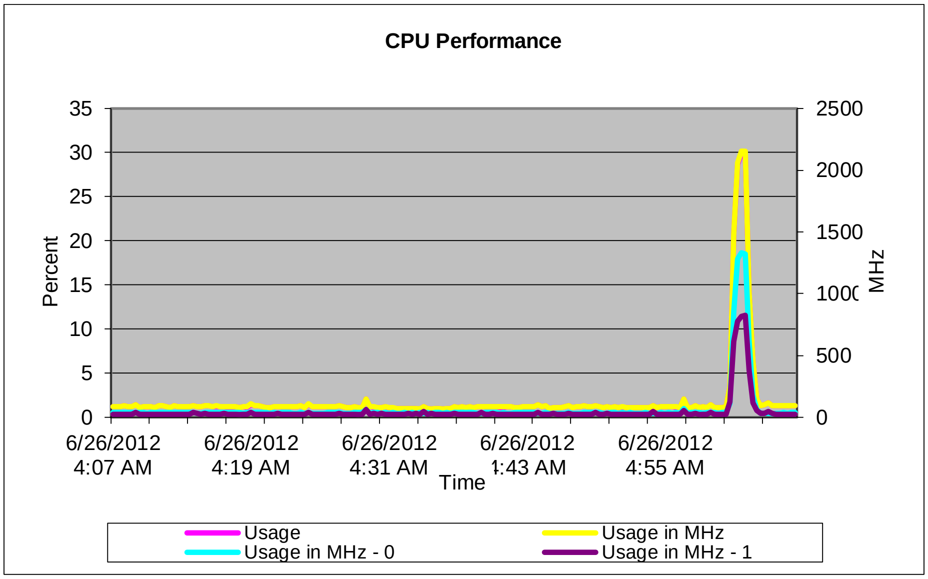

Figure 43 proves the maximum rate at which the SAN server received data was almost 9 MB/sec.

Figure 43. Network Performance

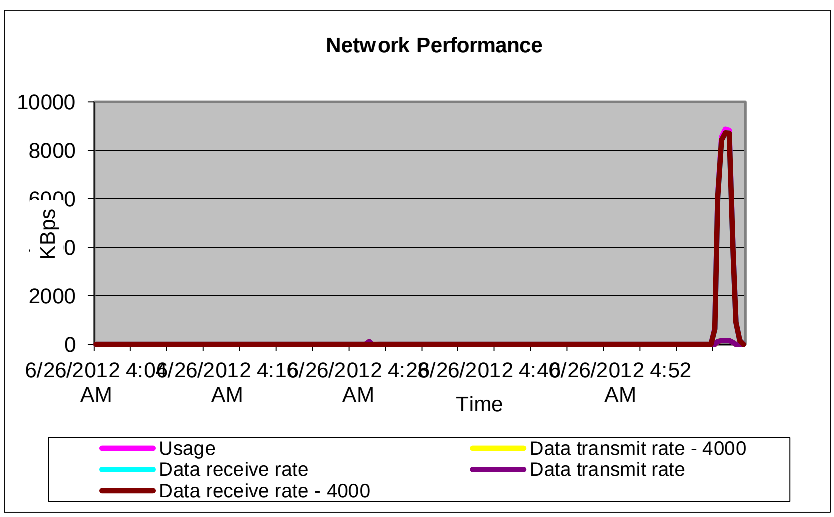

Table 5 depicts the maximum network usage was 8.8 MB/sec.

Table 5. Data Transfer

Time	Usage	Data transmit rate - 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
6/26/2012 5:00:20 AM	0	0	0	0	0
6/26/2012 5:00:40 AM	653	12	640	12	640
6/26/2012 5:01:00 AM	6140	102	6037	102	6037
6/26/2012 5:01:20 AM	8567	139	8428	139	8428
6/26/2012 5:01:40 AM	8878	155	8722	155	8722

4.2.2 Under High Load

Virtual server performance analysis is discussed below.

(i) File Server

Figure 44 shows CPU performance when more than 20 GB data were transferred to the FTP server. Under high load, CPU usage was 46.76% which was 0.76% less than the CPU usage under normal load.

Figure 44. CPU Performance

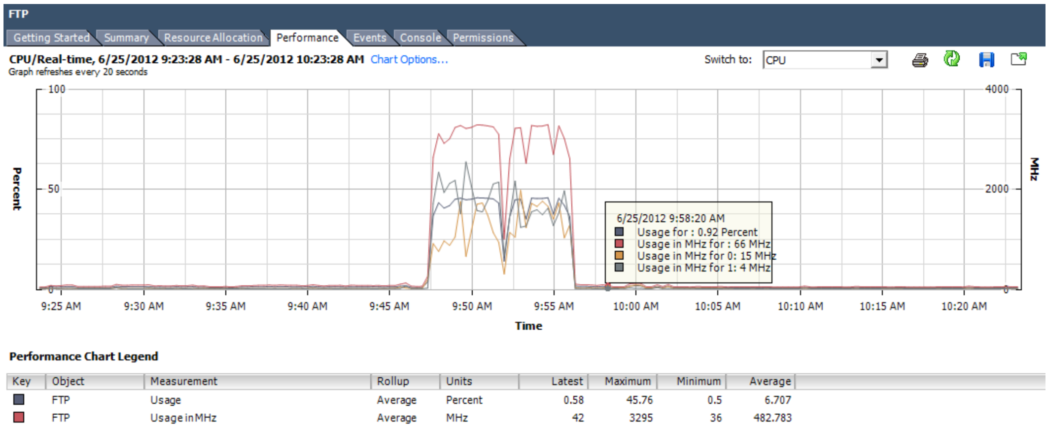

Under high load, when 20 GB data was being transferred to FTP server, the maximum rate at which the server received data was 1.5 Mbps and at that time network usage was 1.2 Mbps.

Figure 45. Network Performance

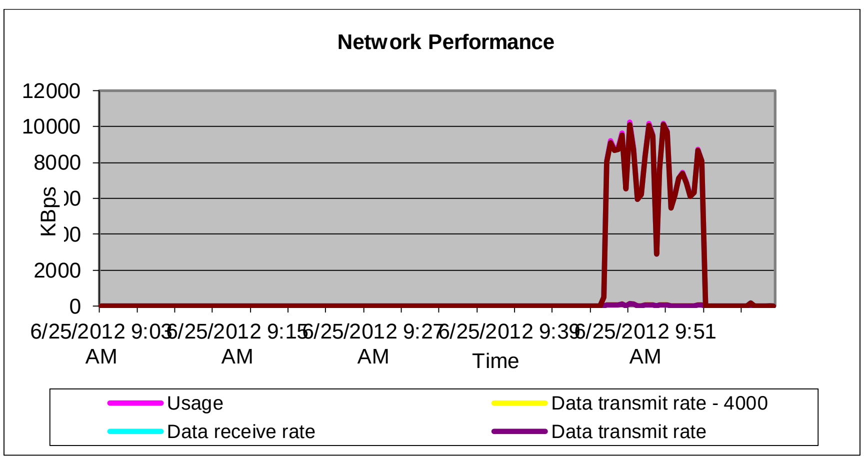

Figure 45 gives a graphical representation of network performance, whereas Table 6 gives the tabular version of network performance.

Table 6. Data Transmission Rate

Time	Usage	Data transmit rate – 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
6/25/2012 9:47:00 AM	0	0	0	0	0
6/25/2012 9:47:20 AM	512	8	504	8	504
6/25/2012 9:53:20 AM	5509	25	5484	25	5484
6/25/2012 9:53:40 AM	6205	32	6173	32	6173
6/25/2012 9:54:00 AM	7135	33	7101	33	7101
6/25/2012 9:54:20 AM	7422	35	7387	35	7387
6/25/2012 9:54:40 AM	6895	32	6862	32	6862
6/25/2012	6158	28	6129	28	6129

(ii) Virtual Network Computing

There was no evidence that VNC performance slowed down while the server was being accessed by many users at the same time. I tested 10 VNC connections on different ports and the performance was almost the same as when accessed by only 2 users.

(iii) Storage Server

Figure 46 illustrates the CPU usage, which rises to 28% when transferring 20 GB of data to the SAN.

Figure 46. CPU Performance

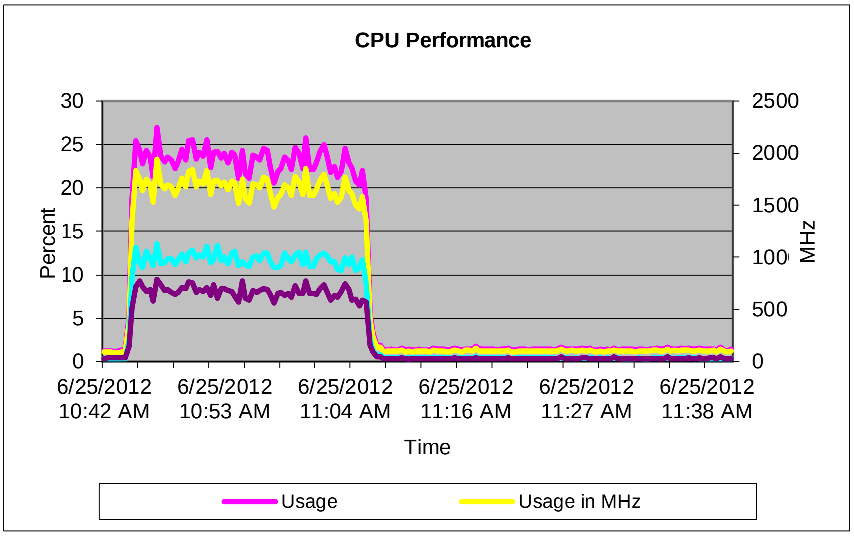

Figure 47 shows that under high data load, network usage reached a maximum of 8.1 Mbps.

Figure 47. Network Performance

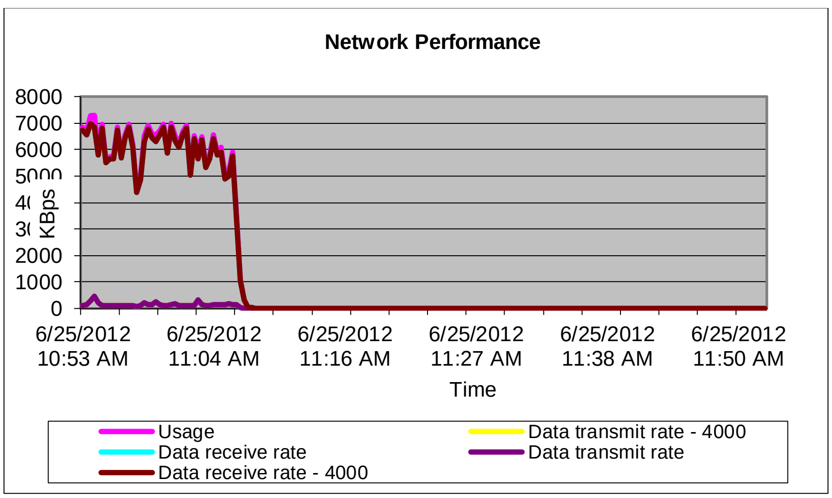

Table 7 shows that the maximum rate at which server received data was almost 7.5 Mbps.

Table 7 Data Transmission Rate

Usage	Data transmit rate - 4000	Dat a receive rate	Data transmit rate	Data receive rate – 4000
0	0	0	0	0
1161	19	1141	19	1141
5429	94	5334	94	5334
8140	300	7839	300	7839
7788	283	7504	283	7504
7294	162	7132	162	7132
7459	222	7236	222	7236
7001	163	6837	163	6837
6360	155	6205	155	6205
7253	170	7082	170	7082

4.3 Server-PE-2 Performance Analysis

The aerver connected to router PE-2 has been analyzed under the conditions that are explained below.

4.3.1 Under Normal Conditions

Different servers have been analyzed under normal conditions.

(i) File Server

Figure 48 shows a 15% CPU usage under normal conditions.

Figure 48. CPU Performance

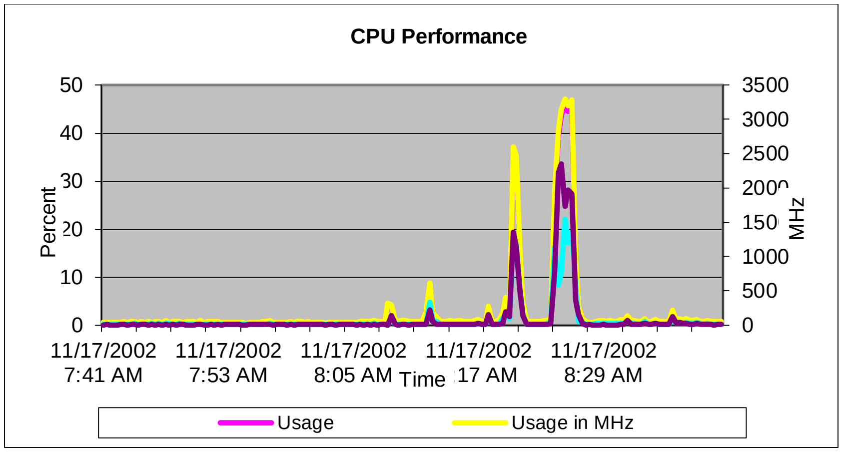

Figure 49 shows the maximum rate at which the server received data was 6.6 Mbps; network usage was almost 7 Mbps.

Figure 49. Network Performance

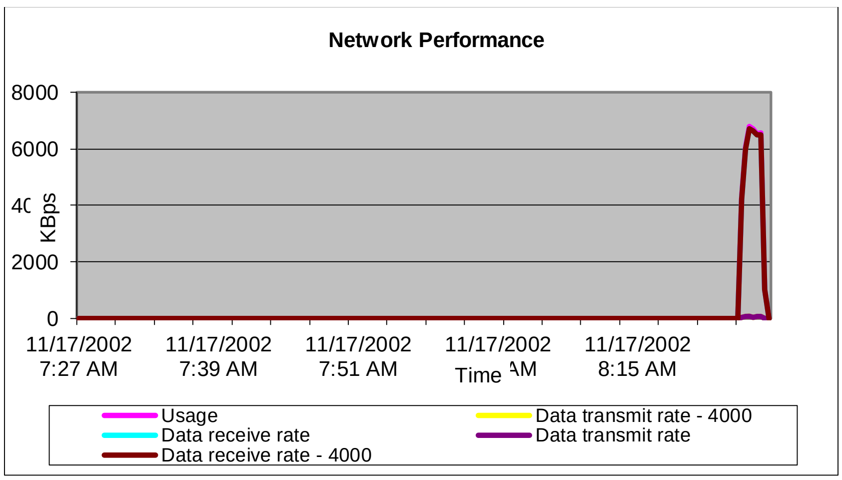

Table 8 shows the tabular form of the graph in Figure 49.

Table 8. Data Transmission Rate

Usage	Data transmit rate - 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
0	0	0	0	0
4276	55	4221	55	4221
6065	75	5990	75	5990
6790	79	6711	79	6711
6693	55	6638	55	6638
6543	60	6482	60	6482
6569	57	6511	57	6511
1034	5	1029	5	1029
0	0	0	0	0

(ii) VNC Server

Figure 50 shows a remarkable rise in CPU usage, which is 45% during connection establishment.

Figure 50. CPU Performance

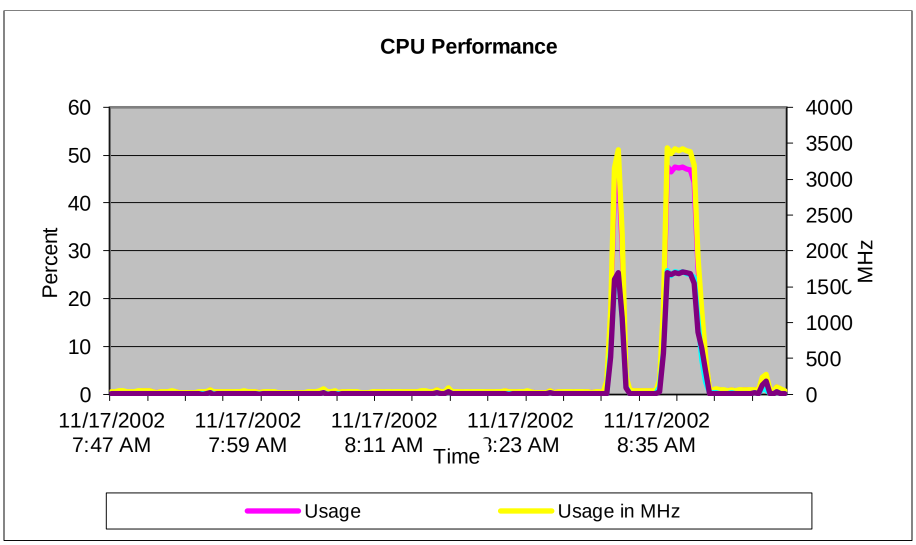

Figure 51 shows performance graph during VNC connection establishment. The maximum rate at which the server received data was only 3 kbps, and network usage at that time was 30 kbps.

Figure 51. Network Performance

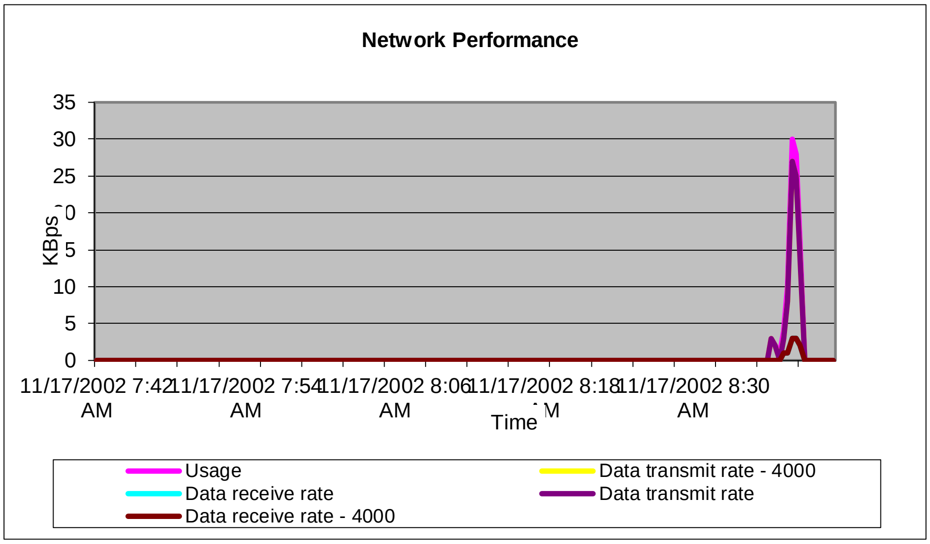

Table 9 is the tabular representation of VNC connection establishment.

Table 9. Data Transmission Rate

Usage	Data transmit rate - 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
0	0	0	0	0
3	3	0	3	0
2	2	0	2	0
1	0	0	0	0
4	3	1	3	1
10	8	1	8	1
30	27	3	27	3
28	25	3	25	3
16	14	2	14	2
0	0	0	0	0

4.3.2 Under High Load

In this section, analysis of virtual servers has been carried out under high load of data transfer.

(i) File Server

Figure 52 illustrates the maximum data rate at which server received data was 6.5 Mbps and network usage at that time was 6 Mbps.

Figure 52. Network Performance

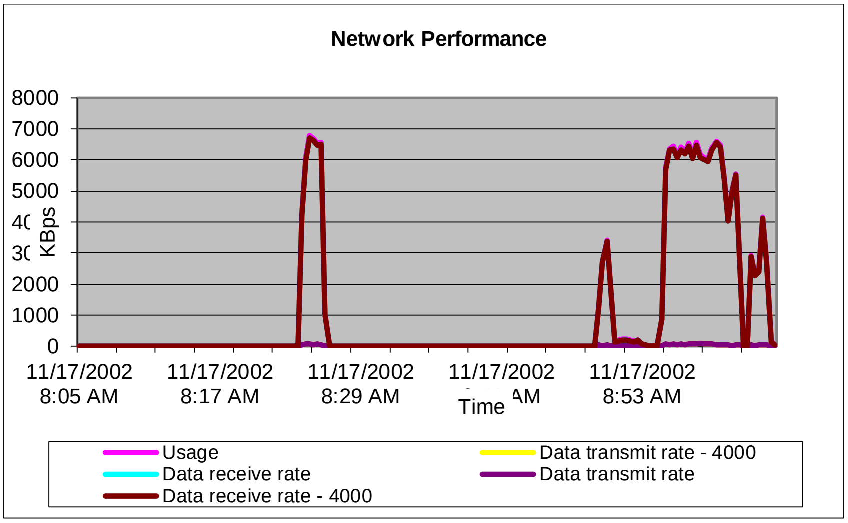

Table 10 represents the data of Figure 52 in tabular form.

Table 10. Data Transmission Rate

Usage	Data transmit rate - 4000	Data receive rate	Data transmit rate	Data receive rate - 4000
0	0	0	0	0
18	8	9	8	9
3	1	1	1	1
877	10	867	10	867
5760	62	5697	62	5697
6354	41	6313	41	6313
6443	80	6363	80	6363
6124	50	6074	50	6074
6418	83	6334	83	6334
6229	45	6184	45	6184

(ii) VNC Server

It has been found that the performance of the VNC server is not noticeably affected by establishing many connections on different ports. All key strokes and mouse clicks were responding with an acceptable latency.

5.0 Benchmarking

Benchmarking is the process of evaluating the performance of servers running different tests.

5.1 Benchmarking for SAN

In the SAN sever, two hard drives were utilized and then two logical volumes were created. In one volume, a 15 GB RAID-5 was configured. I named it media-2. The other volume, which I named media-1, was 30 GB, normally partitioned. It is used as Windows partitions for clients. The software used for benchmarking is called iozone, which is a command-line oriented file system benchmark tool. The benchmark generates and measures a variety of file operations. The command used to generate the output displayed in the screen shots below was:

iozone -Rb –l5 –u5 – r 4k –s 100m.

This command generates an Excel file with 5 records per process in which every write and read is 1 GB with 4kb record size.

Read – Indicates the performance of reading a file that already exists in the file system.

Write – Indicates the performance of writing a new file to the file system.

Re-read – After reading a file, this indicates the performance of reading a file again.

Re-write – Indicates the performance of writing to an existing file.

5.2 Benchmarking from CE-1 Client

5.2.1 To Media-1

Figure 53. Benchmarking CE-1 to Media-1

```
Excel chart generation enabled
Record Size 4 KB
File size set to 102400 KB
Command line used: iozone -Rb -l5 -u5 -r 4k -s 100m
Output is in Kbytes/sec
Time Resolution = 0.000000 seconds.
Processor cache size set to 1024 Kbytes.
Processor cache line size set to 32 bytes.
File stride size set to 17 * record size.
Min process = 1
Max process = 5
Throughput test with 1 process
Each process writes a 102400 Kbyte file in 4 Kbyte records

Children see throughput for  1 initial writers   =  244898.38 KB/sec
Parent sees throughput for  1 initial writers    =    5890.67 KB/sec
Min throughput per process                       =  244898.38 KB/sec
Max throughput per process                       =  244898.38 KB/sec
Avg throughput per process                       =  244898.38 KB/sec
Min xfer                                         =  102400.00 KB

Children see throughput for  1 rewriters         =  582269.31 KB/sec
Parent sees throughput for  1 rewriters          =    5673.66 KB/sec
Min throughput per process                       =  582269.31 KB/sec
Max throughput per process                       =  582269.31 KB/sec
Avg throughput per process                       =  582269.31 KB/sec
Min xfer                                         =  102400.00 KB

Children see throughput for  1 readers           =  715243.12 KB/sec
Parent sees throughput for  1 readers            =  702961.34 KB/sec
Min throughput per process                       =  715243.12 KB/sec
Max throughput per process                       =  715243.12 KB/sec
Avg throughput per process                       =  715243.12 KB/sec
Min xfer                                         =  102400.00 KB

Children see throughput for 1 re-readers         =  687131.69 KB/sec
Parent sees throughput for 1 re-readers          =  664822.63 KB/sec
Min throughput per process                       =  687131.69 KB/sec
Max throughput per process                       =  687131.69 KB/sec
Avg throughput per process                       =  687131.69 KB/sec
Min xfer                                         =  102400.00 KB
```

5.2.2: To Media-2

Figure 54. Benchmarking CE-1 to Media-2

```
Excel chart generation enabled
Record Size 4 KB
File size set to 102400 KB
Command line used: iozone -Rb output.wks -l5 -u5 -r 4k -s 100m
Output is in Kbytes/sec
Time Resolution = 0.000000 seconds.
Processor cache size set to 1024 Kbytes.
Processor cache line size set to 32 bytes.
File stride size set to 17 * record size.
Min process = 5
Max process = 5
Throughput test with 5 processes
Each process writes a 102400 Kbyte file in 4 Kbyte records

Children see throughput for  5 initial writers   =  130079.97 KB/sec
Parent sees throughput for  5 initial writers    =    1596.69 KB/sec
Min throughput per process                       =   17646.49 KB/sec
Max throughput per process                       =   43698.76 KB/sec
Avg throughput per process                       =   26015.99 KB/sec
Min xfer                                         =   40832.00 KB

Children see throughput for  5 rewriters         =   25275.76 KB/sec
Parent sees throughput for  5 rewriters          =    2457.36 KB/sec
Min throughput per process                       =    3994.77 KB/sec
Max throughput per process                       =    6433.94 KB/sec
Avg throughput per process                       =    5055.15 KB/sec
Min xfer                                         =   63588.00 KB

Children see throughput for  5 readers           = 1062813.04 KB/sec
Parent sees throughput for  5 readers            =  990008.35 KB/sec
Min throughput per process                       =  130472.29 KB/sec
Max throughput per process                       =  297043.97 KB/sec
Avg throughput per process                       =  212562.61 KB/sec
Min xfer                                         =   45116.00 KB

Children see throughput for 5 re-readers         = 1165854.58 KB/sec
Parent sees throughput for 5 re-readers          = 1070423.40 KB/sec
Min throughput per process                       =  173599.47 KB/sec
Max throughput per process                       =  339025.81 KB/sec
Avg throughput per process                       =  233170.92 KB/sec
Min xfer                                         =   52704.00 KB
```

5.3 Benchmarking from CE-2 Clients

5.3.1 To Media-1

Figure 55. Benchmarking CE-2 to Media-1

```
Excel chart generation enabled
Record Size 4 KB
File size set to 102400 KB
Command line used: iozone -Rb output.wks -l5 -u5 -r 4k -s 100m
Output is in Kbytes/sec
Time Resolution = -0.000000 seconds.
Processor cache size set to 1024 Kbytes.
Processor cache line size set to 32 bytes.
File stride size set to 17 * record size.
Min process = 5
Max process = 5
Throughput test with 5 processes
Each process writes a 102400 Kbyte file in 4 Kbyte records

Children see throughput for  5 initial writers  =   51934.52 KB/sec
Parent sees throughput for  5 initial writers   =    9519.14 KB/sec
Min throughput per process                      =    7635.56 KB/sec
Max throughput per process                      =   12667.68 KB/sec
Avg throughput per process                      =   10386.90 KB/sec
Min xfer                                        =   61844.00 KB

Children see throughput for  5 rewriters        =  489356.37 KB/sec
Parent sees throughput for  5 rewriters         =   11048.26 KB/sec
Min throughput per process                      =   64760.96 KB/sec
Max throughput per process                      =  132335.20 KB/sec
Avg throughput per process                      =   97871.27 KB/sec
Min xfer                                        =   50176.00 KB

Children see throughput for  5 readers          = 1157613.44 KB/sec
Parent sees throughput for  5 readers           = 1067430.12 KB/sec
Min throughput per process                      =  113831.44 KB/sec
Max throughput per process                      =  310056.72 KB/sec
Avg throughput per process                      =  231522.69 KB/sec
Min xfer                                        =   37424.00 KB

Children see throughput for 5 re-readers        = 1190798.91 KB/sec
Parent sees throughput for 5 re-readers         = 1097157.69 KB/sec
Min throughput per process                      =  143536.03 KB/sec
Max throughput per process                      =  304086.28 KB/sec
Avg throughput per process                      =  238159.78 KB/sec
Min xfer                                        =   48072.00 KB
```

5.3.2 To Media-2

Figure 56. Benchmarking CE-2 to Media-2

```
Processor cache size set to 1024 Kbytes.
Processor cache line size set to 32 bytes.
File stride size set to 17 * record size.
Min process = 1
Max process = 5
Throughput test with 1 process
Each process writes a 102400 Kbyte file in 4 Kbyte records

Children see throughput for  1 initial writers  =  237103.95 KB/sec
Parent sees throughput for  1 initial writers   =    4656.53 KB/sec
Min throughput per process                      =  237103.95 KB/sec
Max throughput per process                      =  237103.95 KB/sec
Avg throughput per process                      =  237103.95 KB/sec
Min xfer                                        =  102400.00 KB

Children see throughput for  1 rewriters        =  627396.62 KB/sec
Parent sees throughput for  1 rewriters         =    4724.27 KB/sec
Min throughput per process                      =  627396.62 KB/sec
Max throughput per process                      =  627396.62 KB/sec
Avg throughput per process                      =  627396.62 KB/sec
Min xfer                                        =  102400.00 KB

Children see throughput for  1 readers          =  698116.56 KB/sec
Parent sees throughput for  1 readers           =  686561.91 KB/sec
Min throughput per process                      =  698116.56 KB/sec
Max throughput per process                      =  698116.56 KB/sec
Avg throughput per process                      =  698116.56 KB/sec
Min xfer                                        =  102400.00 KB

Children see throughput for 1 re-readers        =  694367.81 KB/sec
Parent sees throughput for 1 re-readers         =  683140.75 KB/sec
Min throughput per process                      =  694367.81 KB/sec
Max throughput per process                      =  694367.81 KB/sec
Avg throughput per process                      =  694367.81 KB/sec
Min xfer                                        =  102400.00 KB
```

6.0 CONCLUSION & RECOMMENDATIONS

In this project virtual servers using an MPLS network under normal conditions and under high data transfer loads have been tested to assess their performance. For testing purposes, a concept of connecting virtual servers to provider edge routers and keeping them in one layer-3 VPN has been introduced. An MPLS sham link is used for redundancy.

All clients and servers were in the same VPN to prevent unreachability. I used two physical Intel servers, installed ESXI-4.1, and configured virtual servers. An FTP server and a VNC server were configured on both physical servers while Openfiler for storage was configured on the PE-1 server. In addition to the above, an SSL certificate has been configured for secure data transfer. The VNC server was configured in such a way that all users can VNC only via a secure channel.

CentOS 6.2 is considered to be the preeminent server OS, because it provides a free enterprise class computing platform to anyone. It is free, stable and there are many online free resources available. For reliability and security purposes, the SAN was configured in the open source operating system called Openfiler. This is a modified version of Linux which provides an iSCSI target for iSCSI initiators like VMware ESX and Windows.

As with a physical computer, the overall objective in improving virtual server performance is to eliminate any performance bottlenecks. In general, the more processors, RAM and hard disk space that exists, the better the virtual server's performance will be. It is recommended to overestimate the amount of RAM required to run the host OS, virtual servers and the guest OS. Performance drops dramatically when a virtual machine runs out of available RAM and starts paging. In addition, poor host OS performance turns directly into poor virtual machine performance. It is also recommended to use iSCSI and SANs to enhance virtual hard disk performance.

Recommendations for improving disk performance:

- Use a hard disk solution that allows fast access, such as a SCSI hard disk, a redundant array of RAID, or a SAN.
- Put each virtual hard disk on a dedicated volume, SCSI hard disk, RAID or SAN. It is easiest to put virtual hard disks together with their associated virtual machine configuration files on a RAID or SAN because this keeps everything in one place.
- Put virtual hard disks on a different physical disk than the host OS.
- Defragmentation should be done on a regular basis, if a dynamically expanding virtual hard disk is configured. A fixed-size virtual hard disk uses a reserved block of storage space, which means that data is less likely to be fragmented as it is stored.
- In order to free up more physical disk space it is necessary to "compact" the hard drive.

Recommendations for improving network performance:

- Distribute the networking load. If several instances of a virtual server are running, the networking load should be distributed between them in the same manner as the physical servers.
- Physical network adapters should be added. Dedicating at least one physical network adapter to each virtual machine is highly recommended for best performance.
- Configure a separate VLAN to host the virtual server.
- The network should be monitored regularly so that corrective steps may be taken to improve network performance.

REFERENCES

[1] J. Reagan, CCIP MPLS Study Guide.

[2] **http://www.redbooks.ibm.com/redbooks/pdfs/sg245470.pdf**, accessed in June 1st, 2012.

[3] http://www.networkworld.com/research/2007/040207-mpls-migration-reasons.html, accessed in June 2nd, 2012.

[4] http://www.windowsnetworking.com/articles_tutorials/pros-cons-virtual-server.html, accessed in June 7th, 2012.