ENSC 427: Communication Networks

Spring 2014

Final Report

Analysis of Applications Through IP VPN

www.sfu.ca/~leetonyl/Ensc427Group12.html

Group 12

Lee, Tony	301111050	leetonyl@sfu.ca
Nguyen, Anthony	301110184	anthonyn@sfu.ca
Truong, Henson	301114646	hensont@sfu.ca

Table of Contents

1.0 Abstract

Network security is becoming a bigger issue as technology gets more advanced. Many network security threats today are spread over the Internet. Some examples of these threats include viruses, spyware, zero-day attacks, denial of service attacks, and identity theft. Common network security components may include anti-virus software to prevent and get rid of viruses, a firewall to block unauthorized access to the network, intrusion prevention system to identify fast spreading threats, and virtual private networks (VPN) for secure remote access.

VPN technology connect computer clients from a private network, across a public network, such as the Internet, to another private network, while still guaranteeing the security, user experience, and integrity of a private network. The use of a VPN connection can cause a non-desirable effect on the performance of applications. We plan to analyze the performance of various applications typically used by corporations using OPNET 16.0 and compare the advantages and disadvantages of using a VPN connection.

2.0 Introduction

The performance of a connection using a VPN connection is not as optimal as a connection without one. A connection with a VPN connection is affected by the public network speed the client is connected to (interconnection between VPN and Internet service provider), processor power to encrypt and decrypt algorithms used with the particular VPN type, and the distance between the client and the VPN server. Businesses and companies usually deploy VPN connections to secure access to their servers and to prevent sensitive information being leaked out.

3.0 VPN Overview

A VPN establishes a temporary connection that connects two individual private networks together using the Internet as a pathway to connect the two networks together.

There are two basic kinds of VPN connections: remote access and site-to-site. A remote access VPN is used for connecting computers that are at a fixed or mobile location to a central resource residing in a separate private network. An example of a remote access VPN is a worker in the field connecting to the central headquarters' server. A site-to-site VPN connects computers that would be at a permanent location to a central resource. An example of a site-to-site VPN would be workstations at a branch office connecting to the central headquarters' server or to a different branch office in a different region that has their own private network.

A VPN is created entirely by software - no special network hardware is required to create one. To establish a VPN connection, 3 main components are required: the VPN client, a firewall, and VPN server. The client creates a tunnel spanning across the Internet and firewall, to the VPN server as shown in figure 1 below.

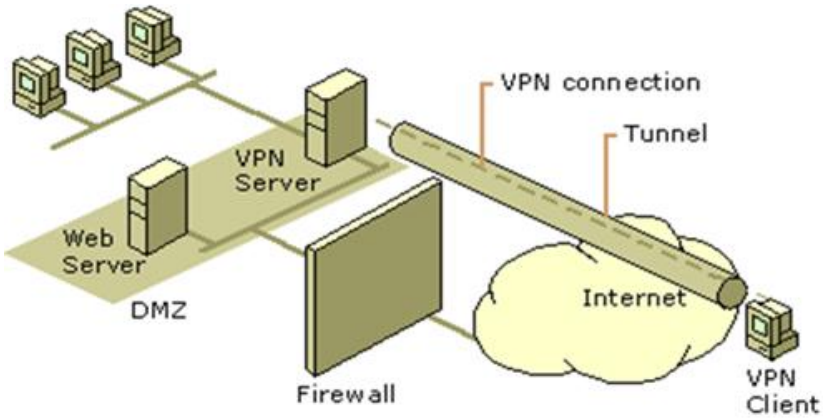

Figure 1: How a VPN connection is established

A VPN connection relies on tunneling to create a secure channel of communication between client and server. Tunneling is the encapsulation of a packet in another packet before transporting it across the internet. There are two types of tunneling: voluntary and compulsory. Voluntary tunneling is when the user sets up a VPN connection to connect to the VPN server; this is used for connecting clients to a VPN server (remote access). Compulsory tunneling on the other hand is when the network server creates the VPN connection and creates a secure tunnel between two or more VPN servers or routers. Compulsory tunneling is used for connecting the local networks of several subsidiaries together (site to site).

Tunneling also provides security as the packet travels through the Internet. A packet gets encrypted when being sent out and decrypted once it reaches its destination. Only the sender and receiver knows the encryption key to encrypt and decrypt the packet. If someone gets a hold of that packet without the encryption key to decrypt that packet, that packet would be useless to them.

The firewall is what gives the VPN network protection. It filters out unwanted and unauthorized access to the network. The firewall may be set to disallow some specific type of traffic going through (IP, HTTP, FTP for example), but having a VPN connection bypasses the firewall's restrictions because the packet going through the firewall is encapsulated in another packet so the firewall does not know what type of traffic the packet going through is, thus allowing it through.

There are many different variations of VPN connections. Typically they differ only by the protocol which the connection uses to tunnel traffic. The protocol chosen for the connection deals with how secure the connection is. The protocol deals how the information inside the packet gets encrypted before getting transported over the internet. The tougher the encryption, the more time resources it takes to encrypt and decrypt each packet. There needs to be a balance between the workload of the machines and the security needed. Some common types of VPN connections are: L2TP, PPTP, and IPsec.

A VPN incorporating IPsec is one of the types which provides greater security than the other VPN types because of its use of two encryption modes: tunnel and transport. Tunnel mode

encrypts the header and the payload of each packet. Transport mode encrypts the payload. PPTP supports multi-protocol VPNs with 40-bit and 128-bit encryption using Microsoft Point-to-Point Encryption. L2TP provides only the tunnel mode security of IPsec but not the transport mode encryption. Balancing security and performance is key to a well-designed VPN connection.

A well designed VPN can provide great benefits for a company including:

- Security
- Extension of network connectivity virtually across the globe
- Simplified network topology
- Reliability
- Improved productivity
- Reduced operational costs compared to wide area networks (WAN)

4.0 OPNET Simulation

The objective of our project is to look at how the response time of applications are affected by using a VPN connection as compared to a regular connection. For this we measured the response time for the applications shown in the table below.

Table 1: Applications and statistics to be measured

Application	**Statistic Measured**
Database Query	Response Time (sec)
Email	Download Response Time (sec)
File Transfer Protocol	Download Response Time (sec)
HTTP	Page Response Time (sec)
Remote Login	Response Time (sec)

4.1 OPNET Objects

The objects used in our simulation include:

- Application Config
- Profile Config
- IP VPN Config (To configure the VPN tunnel)
- ip32_cloud (The Internet)
- ppp_server (Server)
- ppp_workstation (Workstation)
- ethernet4_slip8_gtwy (Router)
- ethernet2_slip8_firewall (Firewall)
- PPP DS1 links

4.2 Topology and Implementation

For our simulation we set up 6 different scenarios:

- Vancouver server to Toronto clients through a regular connections
- Vancouver server to Toronto clients through a VPN connection
- Vancouver server to London clients through a regular connection
- Vancouver server to London clients through a VPN connection
- Vancouver server to Toronto clients through a regular connections with DDOS attackers
- Vancouver server to Toronto clients through a VPN connection with DDOS attackers

Scenario 1: Regular Connection to Toronto Clients

This scenario will act as our baseline for our future scenarios.

The server is placed near Vancouver and is connected through a series of routers and firewalls through the internet to 7 workstation clients in Toronto.

The Applications Config is configured to use the default profile. We add a new application by adding a new row and configured the new row for the Remote Login under light load.

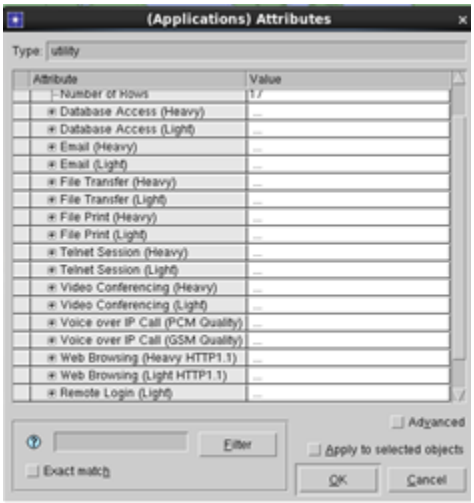

Figure 2: Applications Config

We configured the Profile Config module to use the sample profile which includes the Engineer and Sales Person profile. The Engineer profile is modified to include the Remote Login application.

Figure 3: Engineer Profile

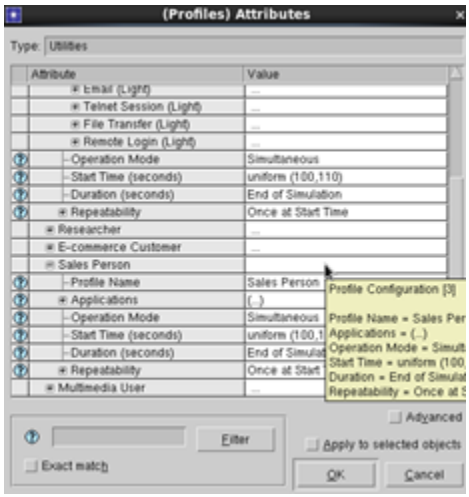

Figure 4: Sales Person Profile

The clients are all set to use the Engineer and Sales Person profiles, which include the applications: Database, Email, FTP, HTTP, and Remote Login.

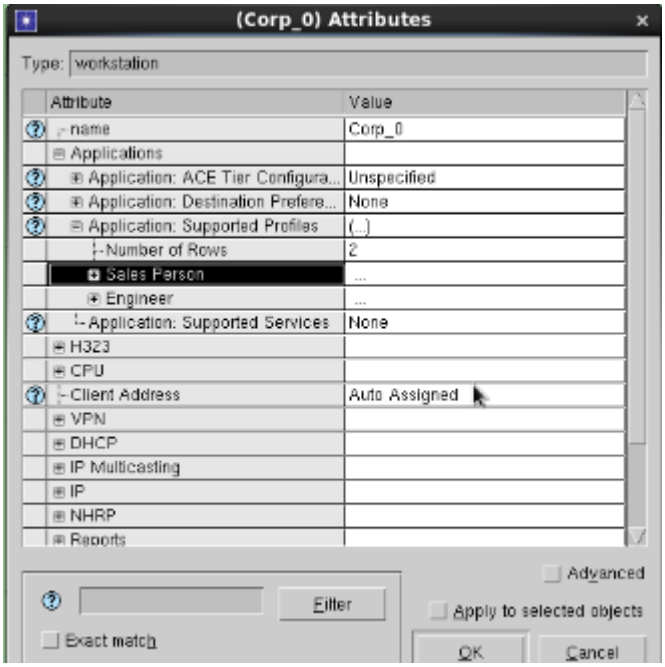

Figure 5: Client Profile Setup

A PPP Server is placed in Vancouver and is connected to a router. From there it is connected to a firewall router, then to the Internet cloud, and finally to a destination router connecting all the clients together. A visually reference is shown in Figure 8

Scenario 2: VPN Connection to Toronto Clients

To protect the server so that only VPN authorized clients are allowed to connect to it we need to disable the proxy server on the firewall for each of the applications. This will prevent any traffic from those applications from going through. We edited the proxy server setting on the server in Vancouver and turned off the proxy server as shown in Figure 6 below.

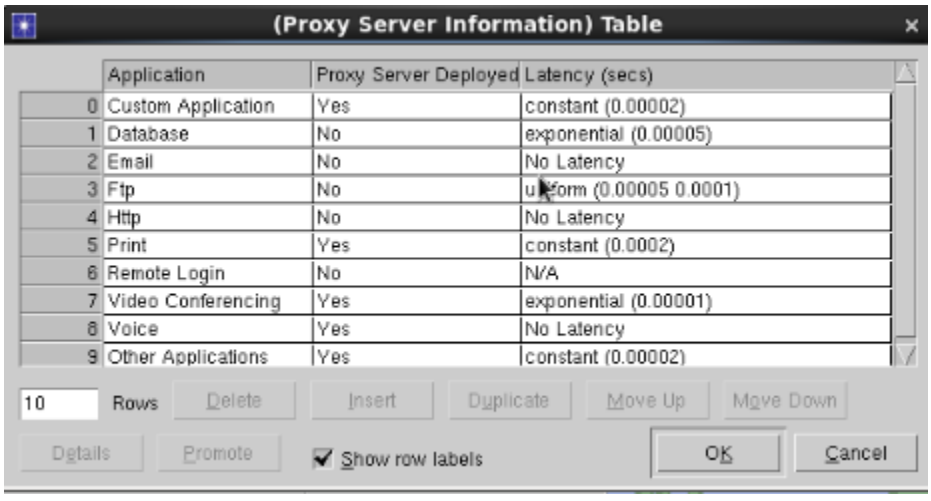

Figure 6: Proxy Server Setup

Now that the applications are being blocked by the firewall we need to establish the VPN tunnel path. This is done in the IP VPN object; the attributes need to be configured to create a tunnel between the two end routers connecting the server and the clients. Additionally each individual client must be added into the client list as shown in Figure 7 below in order for them to send and receive traffic.

Figure 7: IP VPN Config

The final topology for this scenario is indicated in Figure 8 below.

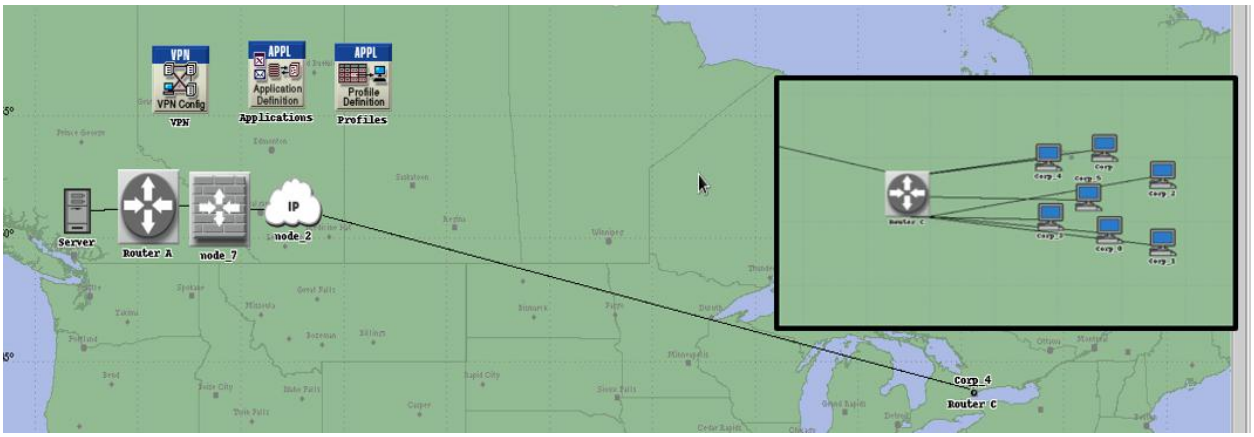

Figure 8: Toronto Topology

Scenario 3: Regular Connections to London Clients

This scenario has the exact settings as scenario 1, where the clients were all situated in Toronto. We duplicated scenario 1 and moved all the clients to be in the London, England area.

Scenario 4: VPN Connection to London Clients

Like scenario 3, this scenario has the same settings as the VPN to Toronto clients' scenario. The VPN to Toronto clients' scenario was duplicated and the clients were moved to London, England area. The full topology for the London client cases is shown in Figure 9 below.

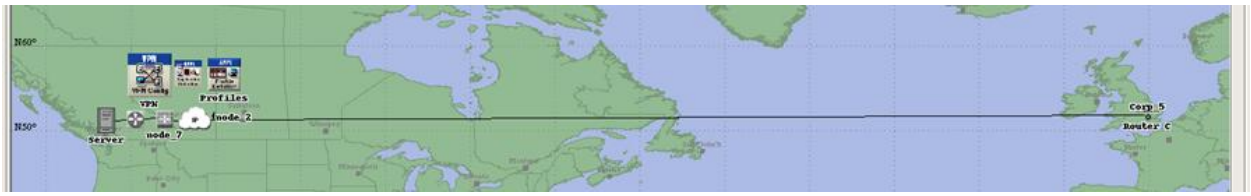

Figure 9: London Topology

Scenario 5: DDOS Through a Regular Connection

The baseline for this scenario is made by duplicating the Vancouver to Toronto clients' scenario. From the duplicated scenario, we added several groups of workstations spread out across North America. For our DDOS scenarios we wanted to have many workstations try to have unauthorized access the Vancouver server, which we will be called "attacking clients".

We edited the attributes of the Profile object to add an additional profile, Attacker, to this scenario, this new profile has the same settings as the sample Sales Person and Engineer profile, except the applications are all configured to be heavy traffic, that is, the size of the packets sent from these clients are much larger than the VPN clients.

Then we connected groups of workstations using the Attacker profile to the Internet. In this scenario, these attacker clients will be able to reach the Vancouver server and request access from the applications. We expect that this setup will cause the server to receive an abundance of application requests which will affect the server's resources.

Scenario 6: DDOS through a VPN Connection

Scenario 5 is duplicated and we implement the same settings for the VPN connection as scenario 2. The firewall will block all traffic that it identifies for the applications that we are testing for. We expect that the traffic from the attacking clients will be stopped at the firewall, not having a chance to reach the server.

The topology for scenario 5 (minus the VPN) and scenario 6 is shown in Figure 10 below.

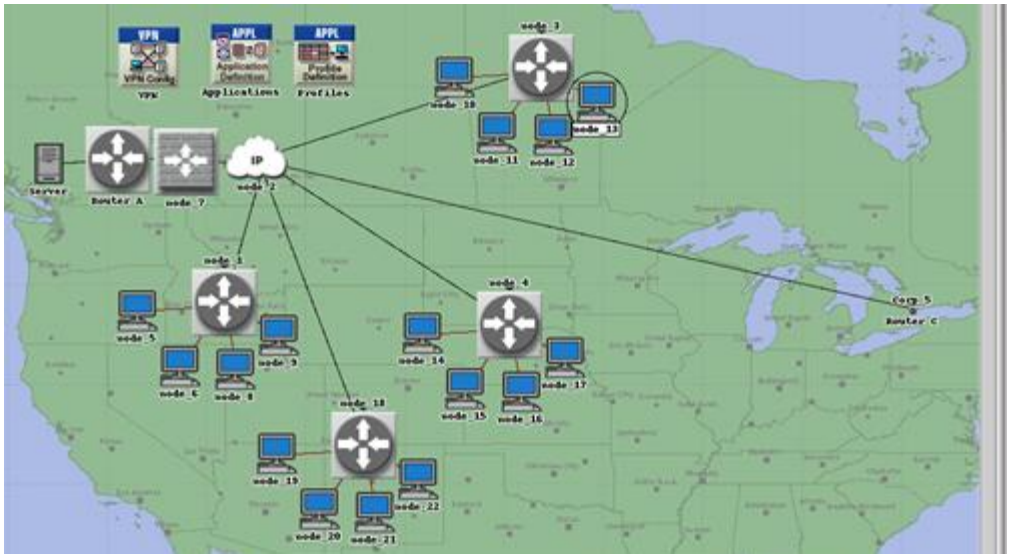

Figure 10: DDOS Topology

4.2 Simulation Results

For the five applications that were analyzed, we collected the statistics that related to some variance of response time for communication between the client and server. The definition for those statistics are listed below.

Response Time: The time elapsed between the client sending a request packet and receiving the response packet.

Download Response Time: The time elapsed between sending a request for email and receiving emails from the email server.

Page Response Time: Time required to retrieve the page with all the objects.

Server Performance Load: The rate at which a request for any application arrives at the server.

This section will go over the numerical findings of our simulation results, the discussion of what these numbers might mean will be talked about in section 5. Our simulation time was set to 2 hours.

For the figures in sections 4.2.1 to 4.2.5, the color of the graphs correspond to the scenarios as follows:

- light blue to scenario 3 (VPN connection to London clients)
- red to scenario 4 (Regular connection to London clients
- green to scenario 2 (VPN connection to Toronto clients)
- blue to scenario 1 (Regular connection to Toronto clients)

4.2.1 Database Query

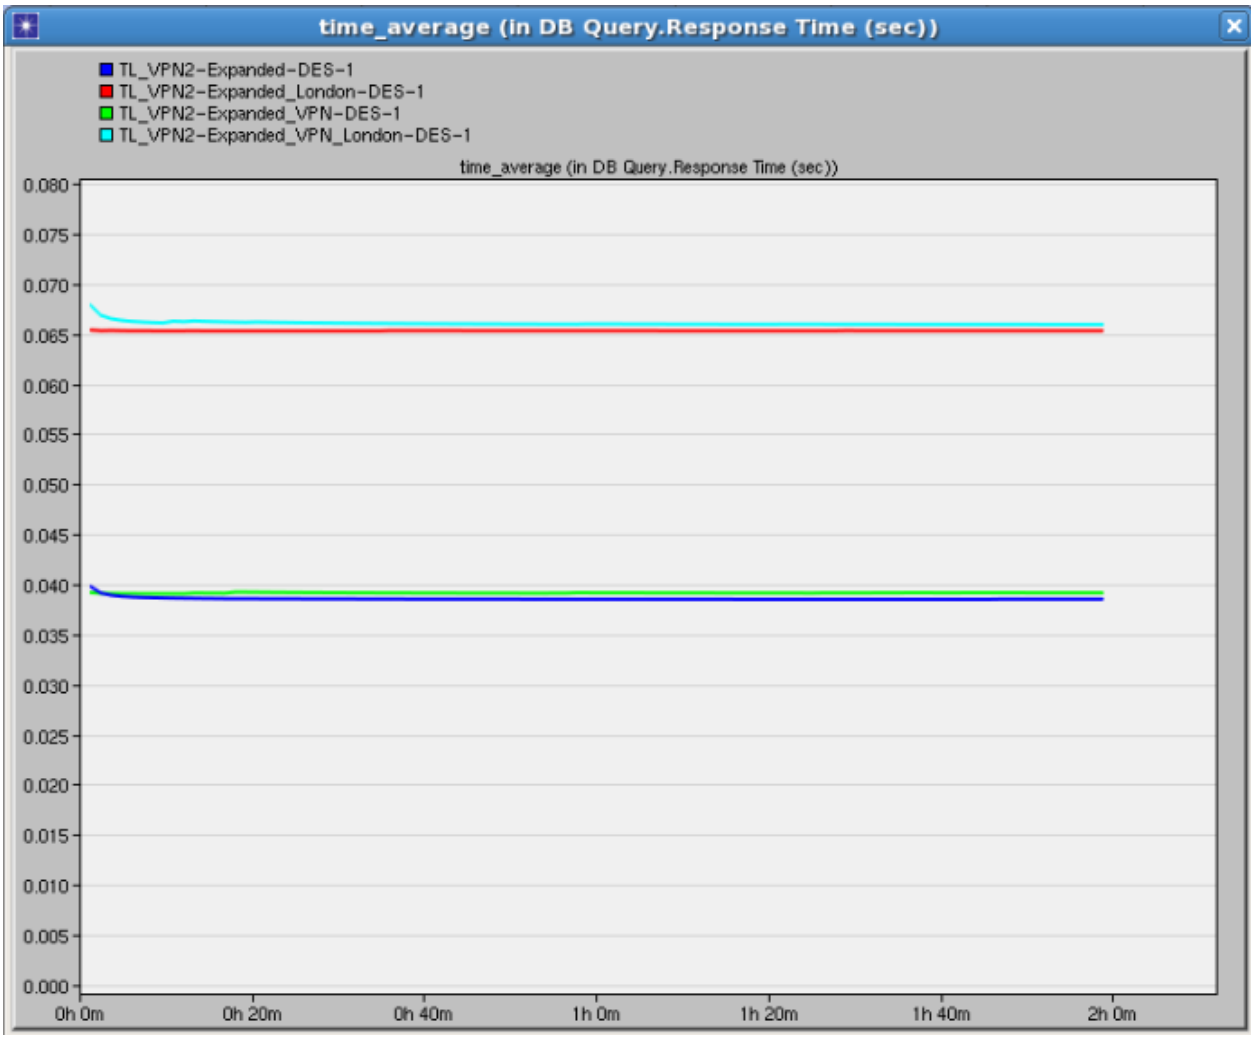

Figure 11: Average Response Time for Database Query

In Figure 11 above, the graph is showing the response time, in seconds, for a database query. The top two trace represent the response time of the clients placed in London, while the bottom two trace represent the response time of the clients situated in Toronto. With regards to the clients connected to the server in Vancouver via a regular connection, the average response time for the clients in London is approximately 65 milliseconds, while the response time for the clients in Toronto is approximately 39 milliseconds. When looking at the VPN connected counterparts of these clients, the response times are approximately 0.7 milliseconds higher respectively. The difference between the response time between London and Toronto is approximately 25 milliseconds.

4.2.2 Email

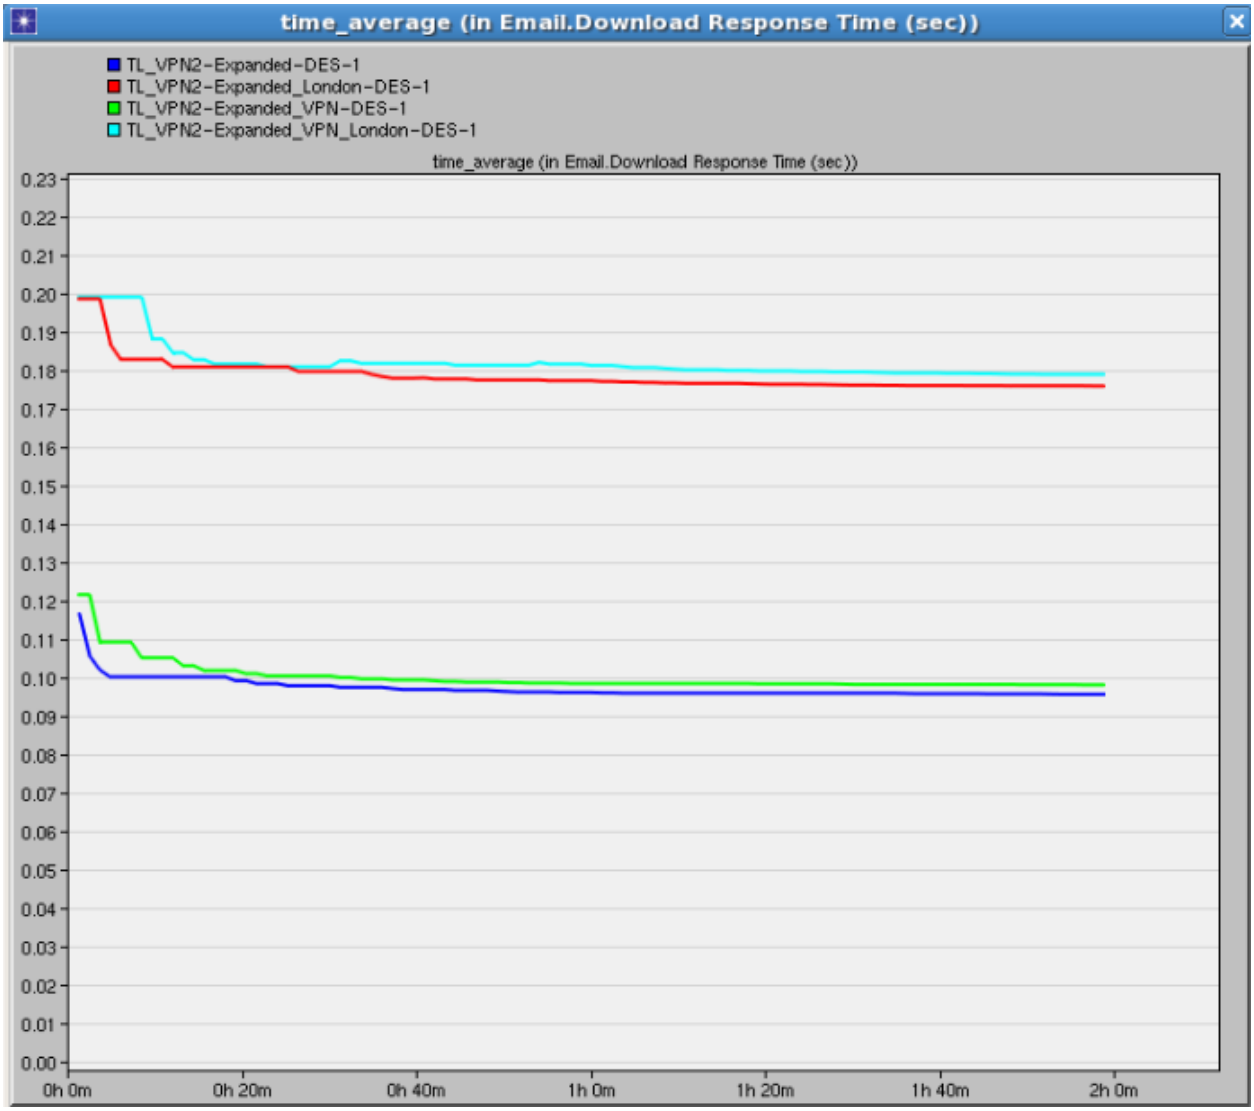

Figure 12: Average Download Response Time for Email

Figure 12 shows the download response time for the email application. The graph is structured the same as for the database application. The top two trace show the clients in London, and the bottom two trace show the clients in Toronto. The response time for the regular connections are approximately 175 milliseconds for London, and 95 milliseconds for Toronto, with the VPN connections adding on about 3 milliseconds more. The increase in response time between London and Toronto is 80 milliseconds for this application.

4.2.3 File Transfer Protocol

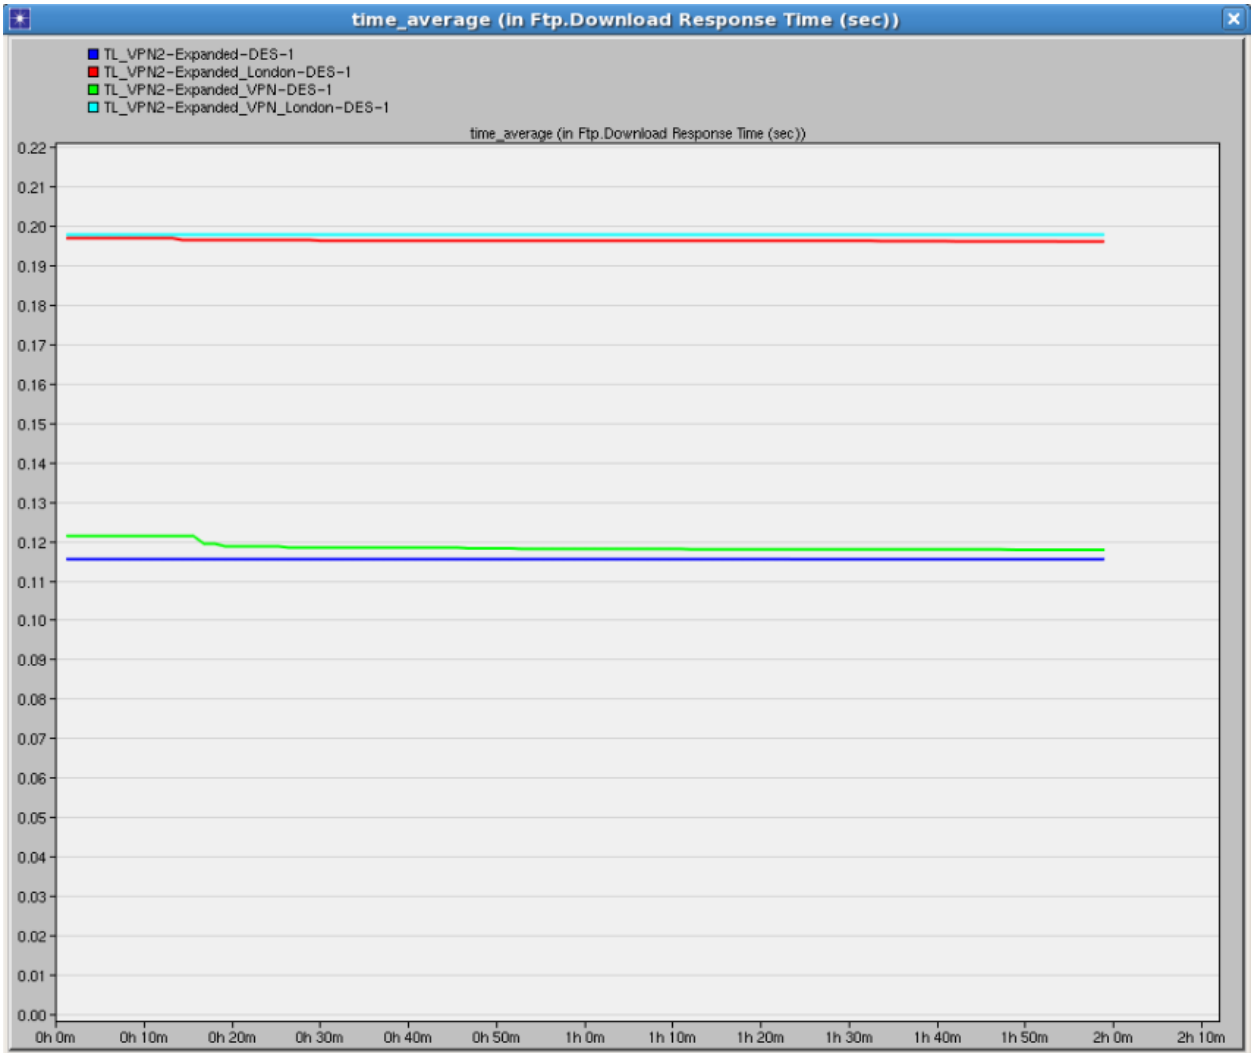

Figure 13: Average Download Response Time for FTP

Like the email application, we collected the download response time for the FTP application, which is shown in the above figure, Figure 13. For this application the download response time for the clients in London is averaged to be about 195 milliseconds, and for Toronto, is about 115 milliseconds. Additionally the increased in response time for the VPN specific case adds an extra 2 milliseconds to the regular connection. Finally the download response time difference between the two locations for this application is about 80 milliseconds.

4.2.4 HTTP

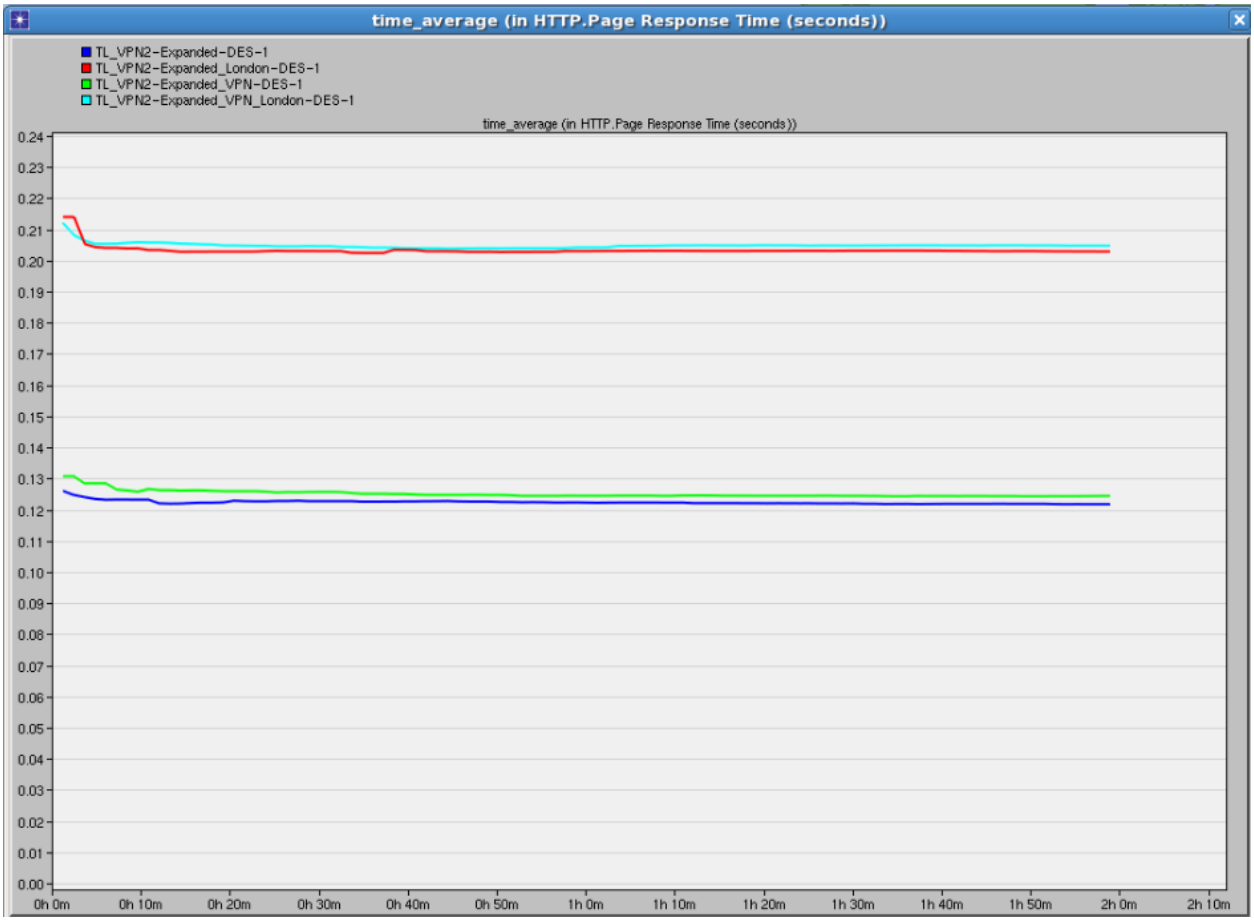

Figure 14: Average Page Response Time for HTTP

For our response time analysis on the HTTP application, we opted to collect the page response time statistics. The page response time for regularly connected clients in London is seen in the above figure to be around 205 milliseconds, and for the corresponding clients in Toronto, the page response time is about 125 milliseconds. Looking at both location's VPN connected clients, the increase in response time for them is about 2 milliseconds higher than their clients through a regular connection. The last numerical note here is that the difference between the page response times for London and Toronto for this application is approximately 80 milliseconds.

4.2.5 Remote Login

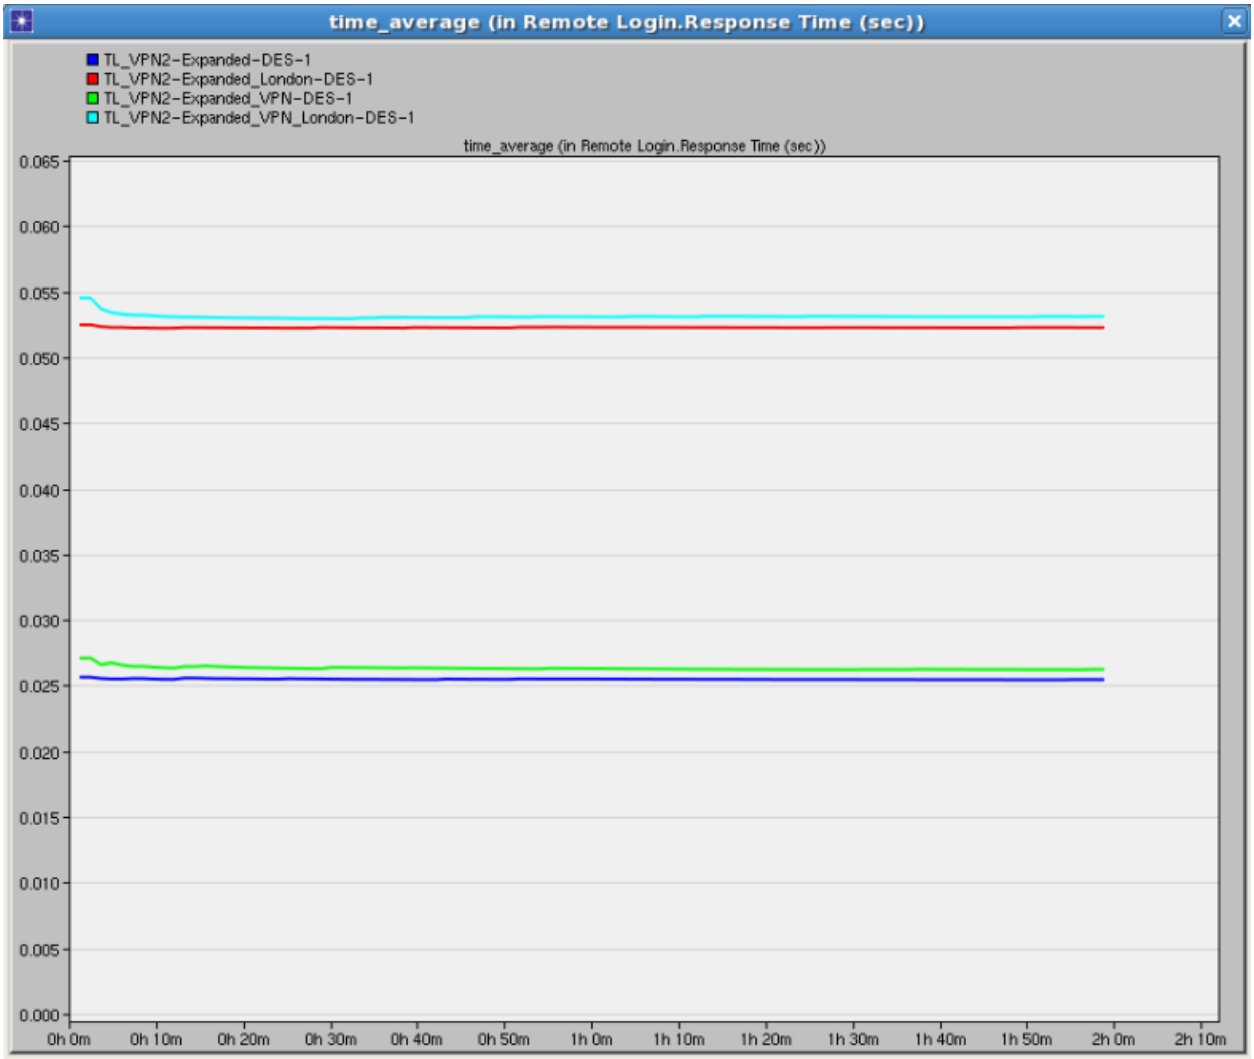

Figure 15: Average Response Time for Remote Login

The final application that was analyzed was remote login. Remote login refers to applications where a user is accessing a server through a virtual interface, that is, the user is not working on the hardware of the accessed server directly. One example of a remote login application would be Microsoft's Remote Desktop Connection. From the results seen in Figure 15, the regular connection to London had a response time of about 53 milliseconds and a response time of about 26 milliseconds for Toronto. The respective VPN connections for these two locations had an increased response time of about 1 millisecond. This difference in response time between the two locations for this application is approximately 30 milliseconds.

4.2.6 Server Performance Load

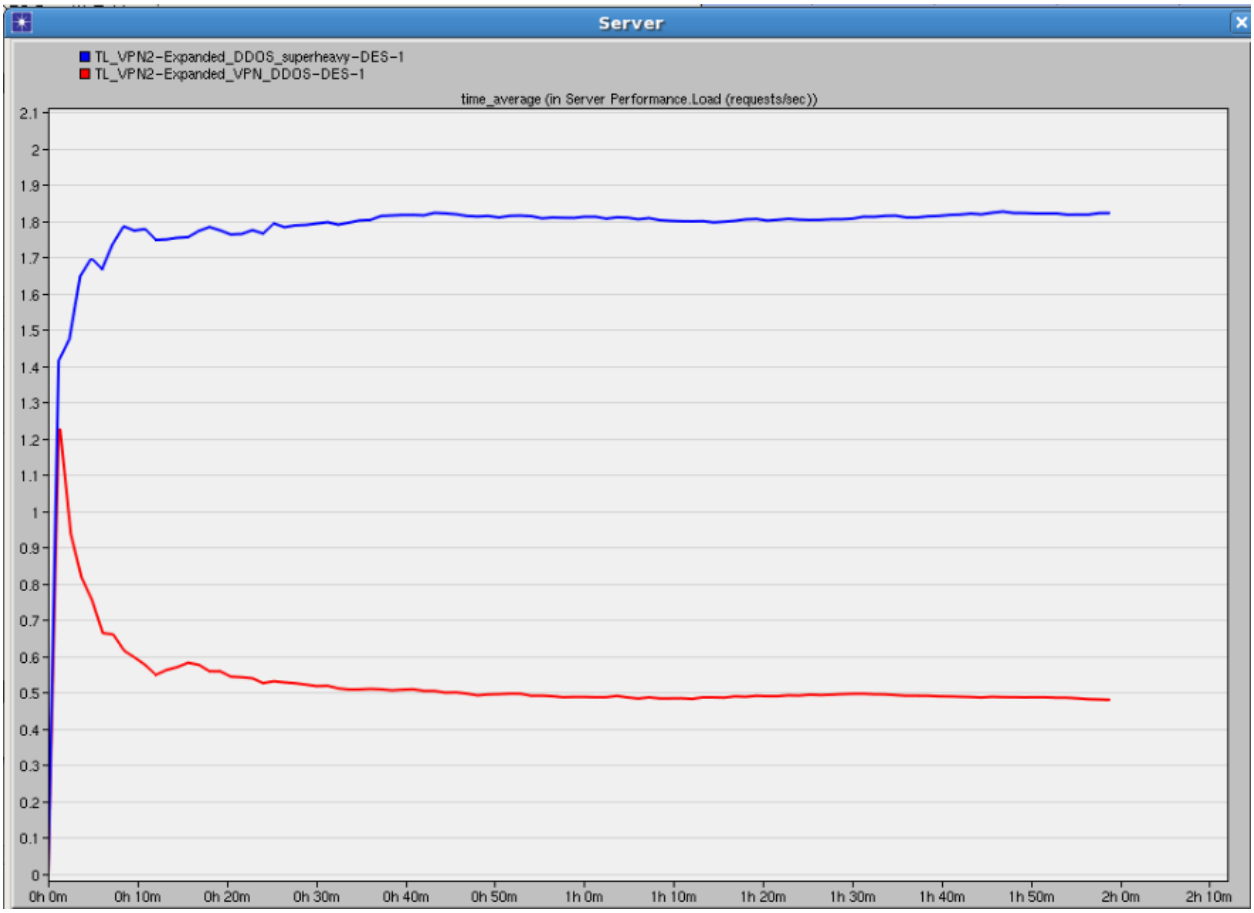

Figure 16: Average Server Performance Load

For our DDOS test cases (scenarios 5 and 6) we collected the server performance load statistics, which is measured in requests per second. In the above figure, the blue trace represents scenario 5, where the firewall is allowing traffic for all our applications to the server, including the attackers' application. The combination of requests averages to about 1.8 requests per second for the server. The red trace (scenario 6) where the firewall blocks recognized traffic for our applications shows only the requests to the server by the established VPN clients in Toronto. As expected, the load for the VPN only scenario is about 0.5 requests per second.

5.0 Discussion

The results from sections 4.2.1 to 4.2.5 is tabulated in the table below.

Application	**Time difference between London and Toronto (ms)**	**Toronto: Time difference between VPN and regular connection (ms)**	**London: Time difference between VPN connect and regular connection (ms)**
Database	25	0.7	0.7
Email	80	3	3
File Transfer Protocol	80	2	2
HTTP	80	2	2
Remote Login	30	1	1

The point of this project was to see how much of a negative impact a VPN connection would have on standard applications used by corporations. From our results, we can see that the increase in response time is in the range of a few milliseconds, which is insignificant to the response time increase caused by the client's distances.

Another note is that the increase in the response time for VPN connections is pretty much the same the clients in Toronto and London, even though there was a substantial increase in the response time through a regular connection.

We noticed that for the applications, E-mail, FTP, and HTTP, the response time difference between the two locations was much higher than for the database and remote login applications, this may be possibly due to the applications sending bigger data packets in general. Because those applications tend to deal with transferring files from the server to the client.

As suggested by the results, the distance of the clients from the server plays a more prominent role in affecting the response times of the applications compared to the VPN connection. This is to be expected, as the traffic going to farther clients would need to pass by more network devices to reach its destination.

6.0 Difficulties and Future Work

Originally we had wanted to go one step deeper into the protocols surrounding VPN, specifically we wanted to compare the protocols Point to Point Tunneling (PPTP) vs. Layer 2 Tunneling Protocol (L2TP). However after 2 weeks of researching we couldn't find a model or information that would help us create a model of PPTP or L2TP. For our alternative option, we switched to analyzing applications through a VPN connection over a WiMAX connection. The reason we

wanted to use WiMAX was because it was a fast wireless connection, suitable for employees doing field work. These employees wouldn't have a guaranteed internet connection when working remotely. The idea was to see if there was any negative effect of a VPN connection over WiMAX. For our WiMAX implementation attempt, we weren't able to establish a VPN connection over WiMAX, possibly because the WiMAX subscriber stations did not have the capabilities to perform point-to-point encryption necessary for creating the tunnel for a VPN connection. Based on this observation, we did not check whether VPN would work for LTE as well for it would have delayed our project even more.

For future work, we would like to once again approach the idea of using a WiMAX network as a medium to establish a VPN connection with remote clients, or use similar long distance wireless solutions such as LTE. We would also like to create a more realistic topology for analyzing VPN connections. This includes having a mix of wired and wireless connections, multiple server and clients workstations, and more attackers trying to access the server at different time within the simulation.

7.0 Conclusion

We were able to successfully implement a VPN network and analyze the response time for the applications: Database, E-mail, FTP, HTTP, and Remote Login. Through our analysis, we can conclude that a VPN connection results in a minor delay in the applications that we looked at. From this we conclude that the minor delay is not significant enough to discourage the use of a VPN connection. The more influencing factor in the response time for the applications is the distance of the client to the server. As the increase in response time for clients using a VPN only connection to the server is on the scale of a few milliseconds, it would be beneficial for companies to implement VPN only connections for server available for external access.

8.0 References

[1] (12 March 2014) F. Parkar and K. Wong, "Analysis of IP VPN Performance." [Online]. Available:http://www2.ensc.sfu.ca/~ljilja/ENSC427/Spring12/Projects/team12/ENSC427_Group12_FinalReport_Spring2012.pdf.

[2] (22 March 2014) "Firewalls and VPN Networks," [Online]. Available:http://www.eng.tau.ac.il/~netlab/resources/booklet/lab11.pdf
.
[3] (22 March 2014) "How much does VPN slow my Internet down?," Cactus VPN, [Online]. Available: www.cactusvpn.com/vpn/vpn-slow-internet-connection.

[4] (4 April 2014) P. Ferguson, Cisco Systems and G. Huston, Telstra, "What Is a VPN?," Cisco, [Online]. Available:http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_1-1/what_is_a_vpn.html.

[5] (12 February 2014) S. Hussein and A. Hadi, "The Impact of Using Security Protocols in Dedicated Private Network and Virtual Private Network," *International Journal of Scientific and Technology Research*, [Online]. 11(2), pp. 170-175. Available: http://www.ijstr.org/final-print/nov2013/The-Impact-Of-Using-Security-Protocols-In-Dedicated-Private-Network-And-Virtual-Private-Network.pdf.

[6] (9 April 2014) B. Marcel “Secure Data Access for Home Users”, Tom’s Hardware, [Online]. Available: http://www.tomshardware.com/reviews/secure-remote-access,1803-2.html.

[7] (2 April 2014) “What is VPN?,” [Online].
Available: technet.microsoft.com/en-us/library/cc731954(v=ws.10).aspx

[8] (2 April 2014) X. Bau and F. Zhang, “The Application of VPN Technology in the University’s Library,” *2011 IEEE 3rd Conference on Communication Software and Networks (ICCSN)*, pp. 563–566, May 2011.

[9] (3 April 2014) “How Virtual Private Networks Work,” (13 October 2008), Cisco, [Online]. Available: http://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/14106-how-vpn-works.html.