ENSC 427

COMMUNICATION • NETWORKS

Analysis of IP VPN Performance

FINAL REPORT

Parkar, Faiz
# 301044177, fkp@sfu.ca

Wong, Kevin
# 301072117, kcw6@sfu.ca

Submitted on April 15th, 2012 to
Prof. Ljiljana Trajkovic

Simon Fraser University

Project Website: http://www.sfu.ca/~fkp/Index.html

Abstract

A Virtual Private Network (VPN) is a mechanism that allows remote users to gain access to a central network over the Internet. In addition, there can be a firewall that prevents unauthorized use of the VPN tunnel connection to the respective server. As it is an important part of an office environment, the goal of this project is to analyze the performance of some commonly used applications such as Email, File-Transfer-Protocol (FTP) and Remote-Login over a firewall-protected IP VPN connection. OPNET Modeler 16.0 will be used to model and simulate a basic network scenario of an office with minimal number of workstations, as well as a scenario created to model an expansion in the office's network.

Table of Contents

Glossary

VPN	**Virtual Private Network**
FTP	**File Transfer Protocol**
LAN	**Local Area Network**
OPNET	**Optimized Network Engineering Tools**
TCP	**Transmission Control Protocol**
UDP	**User Datagram Protocol**
IP	**Internet Protocol**
IPsec	**Internet Protocol Security**
DES	**Discrete Event Simulator**
MPLS	**Multiprotocol Label Switching**

1. Introduction

1.1 VPN Overview

Virtual Private Network (VPN) is a method of providing secure access between remote or mobile workstations and a central organizational network. It utilizes public telecommunication infrastructures (such as the Internet) as a means of connecting the client to the server. VPNs are usually secured using authentication, firewalls, and encryption technologies. This security is very important as VPN's are typically used for office related work, in which case privacy is a huge concern. VPNs are also capable of supporting all the typical network services, such as data sharing and network resource accessing.

There are two main types of VPNs: remote-access VPN and site-to-site VPN. A remote-access VPN connects the user to a Local Area Network (LAN). An example of this is when an employee connects to the company's private network from a remote location. In order to set up remote-access VPNs, companies will typically outsource to enterprise service providers, who will help them set up network access servers and provide remote users with the necessary software to connect. A site-to-site VPN connects multiple fixed sites together through a public network. This type of VPN requires large-scale encryption and dedicated equipment. Within this VPN category, there are 2 subtypes: intranet-based and extranet-based. Intranet-based site-to-site VPNs connect multiple remote workstations to a single private network. Extranet-based site-to-site VPNs connect the LANs of multiple companies together. In our project, we have chosen to model an intranet-based site-to-site VPN.

1.2 IP VPN Advantages

There are many advantages to using IP VPN, such as:

- Eliminating the need for physical leased lines to connect remote users to a private network, thereby reducing costs
- Enhancing security
- Increasing bandwidth and efficiency
- Accessing resources remotely
- Bypassing Internet filters
- Sharing files for a long period of time
- Maintaining online anonymity

2. OPNET Implementation

2.1 Network Objects

The Object Palette in OPNET was used to select the required link and node models for creating the VPN network. The following objects were selected from the available Internet Toolbox (Figure 1):

- Application Configuration
- Profile Configuration
- IP VPN Configuration
- ip32_cloud
    - Used to model IP traffic and supports up to 32 serial connections
    - Packets are routed using first-come-first-serve
    - Includes support for TCP, UDP amongst others
- ppp_wkstn
    - Models remote client using applications over TCP/IP and UDP/IP
- ppp_server
    - Models the office's internal application(s) server
- ethernet4_slip8_gtwy (**router**)
- ethernet2_slip8_firewall (**firewall**)
- PPP_DS1
    - Link model used to c nodes running IP
    - Data rate of 1.544 Mbps

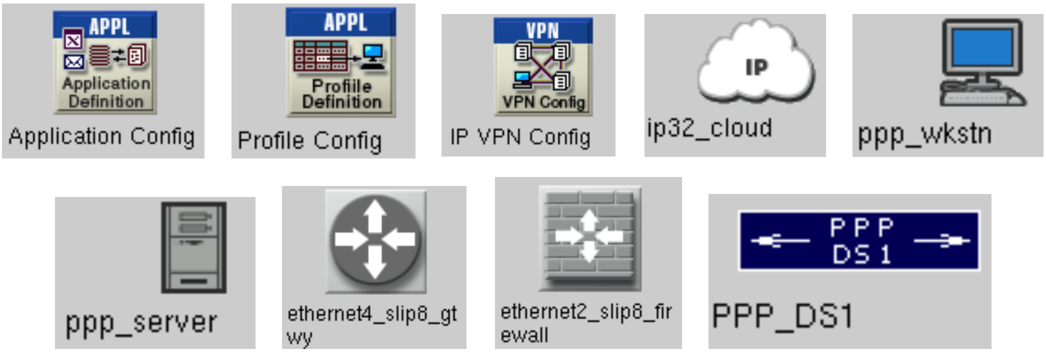

Figure 1: Node & Link Models

2.2 Network Topology

The network model has two scenarios, one with two workstations (Figure 2) and the second with ten workstations (Figure 3). Each of the workstations is set to use the custom VPN user profile, which will be described later in the report. Routers A and B serve as one end of the IP VPN tunnel and are used to connect the workstations to the Internet cloud, which is connected to the Router D through a firewall at Router C. Router D will serve as the other end of the VPN tunnel. Firewall details will be shown later in the report, but one of the important observations of this project is that using the VPN tunnel from A/B to D will bypass any proxy settings made on the firewall due to VPN's built-in security technologies such as IPsec. Finally, the packets are then processed at the application server before returning back to the client workstations through the previously described tunnel.

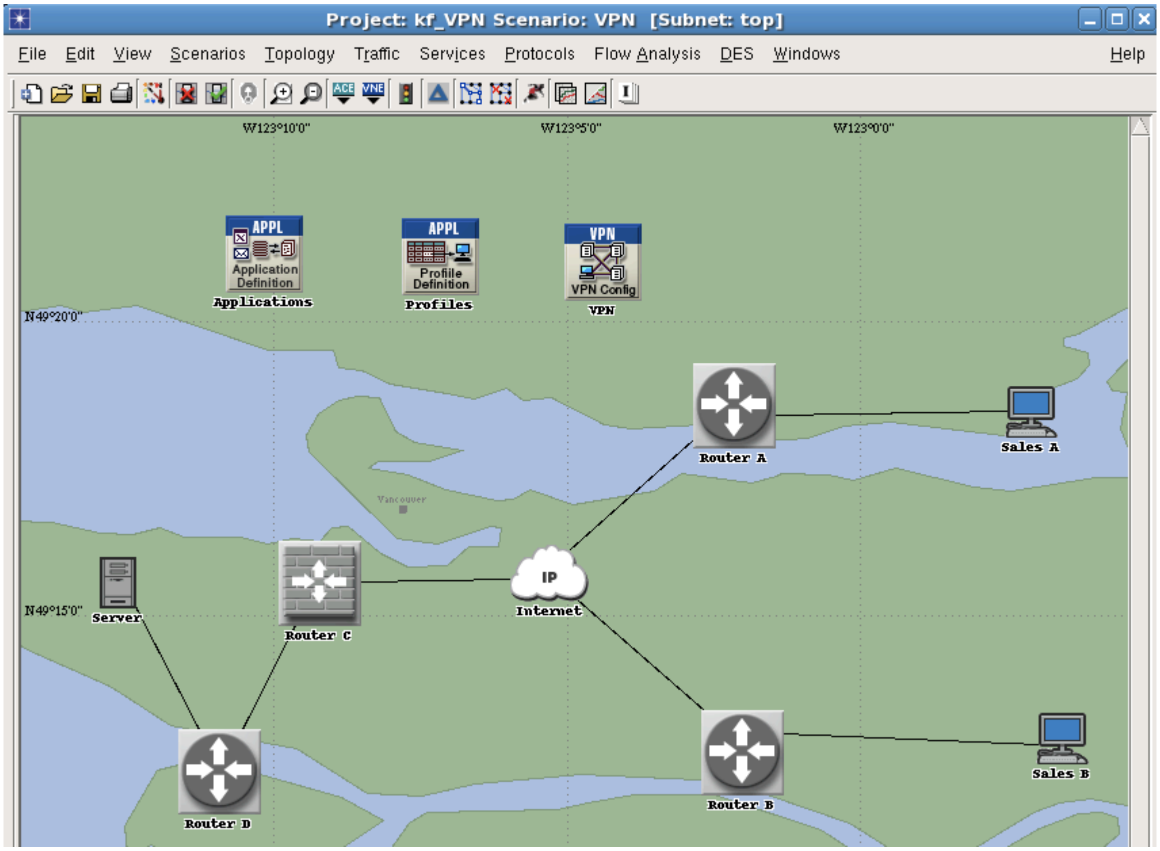

Figure 2: Baseline Network Scenario

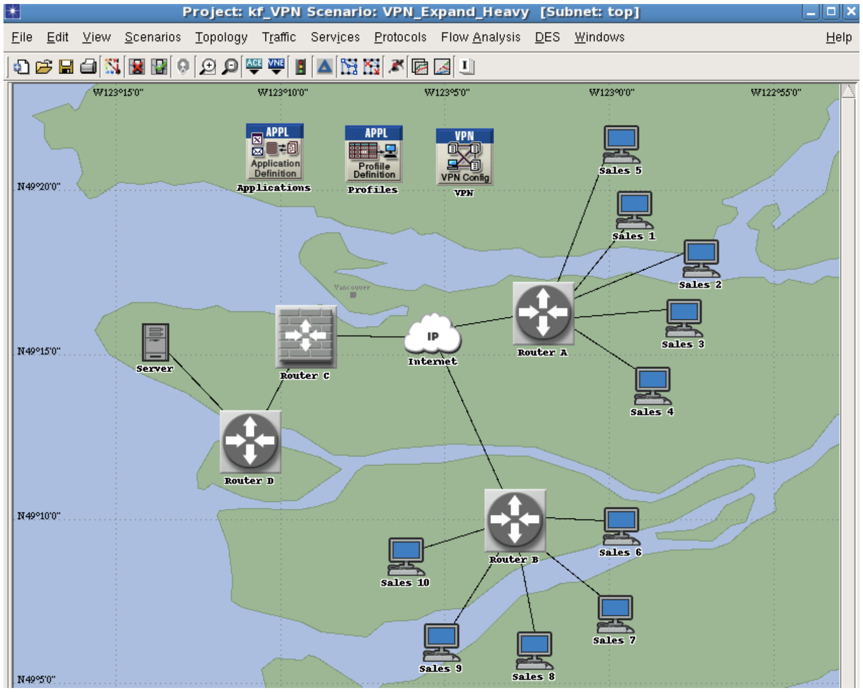

Figure 3: Expanded Network Scenario

2.3 Application Configuration

The application configuration module is used to specify what type of applications will operate in the network. The settings are made using the attributes window (Figure 4). For this project's network, the table has three entries that specify Email, FTP and Remote-Login as the operational applications. Each of these is set to provide heavy usage by setting the description parameter to high load.

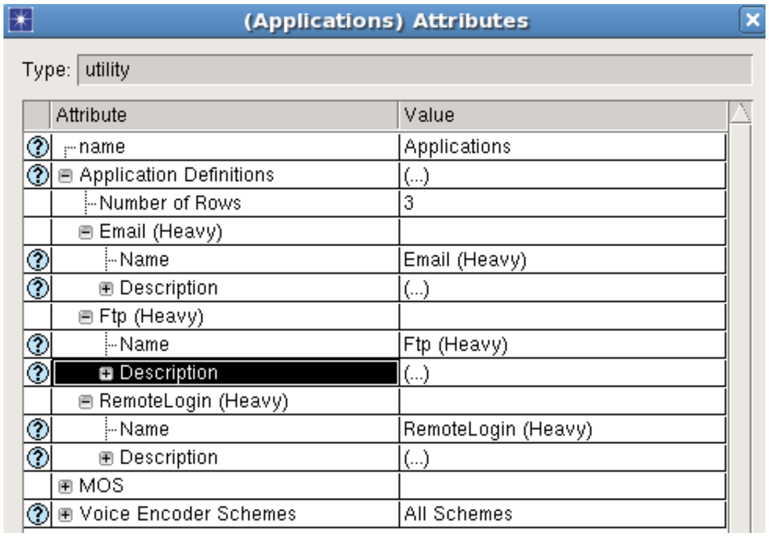

(Applications) Attributes

Type: utility

Attribute	Value
name	Applications
⊟ Application Definitions	(...)
Number of Rows	3
⊟ Email (Heavy)	
Name	Email (Heavy)
⊞ Description	(...)
⊟ Ftp (Heavy)	
Name	Ftp (Heavy)
⊞ Description	(...)
⊟ RemoteLogin (Heavy)	
Name	RemoteLogin (Heavy)
⊞ Description	(...)
⊞ MOS	
⊞ Voice Encoder Schemes	All Schemes

Figure 4: Application Config. Attributes

2.4 Profile Configuration

As mentioned in section 2.1, the profile ‘VPN User’ is assigned to each workstation so the clients operate as required for this network. The profile contains information about which of the applications from the Application configuration will be used. To simulate a realistic scenario of a client using the applications in intervals, each of the selected applications have parameters for start time offset, duration of user, and finally repeatability. Assuming the profile begins using applications at time zero, all applications are given a normal distribution with the means at one-minute intervals (Figure 5). This provides peaks of application traffic at 60, 120 and 180 seconds are start time. Profile configuration allows multiple profiles, which is useful in VPN networks to give each user different levels of application(s) access. This aspect is not considered for this project, as the end goal is to analyze VPN performance based on office expansion and not authorization variance. Hence, the single profile ‘VPN User’ is used for all workstations. Finally, the attributes are set to have this profile run only once in the simulation with each application continuing it’s operation till the end of the profile.

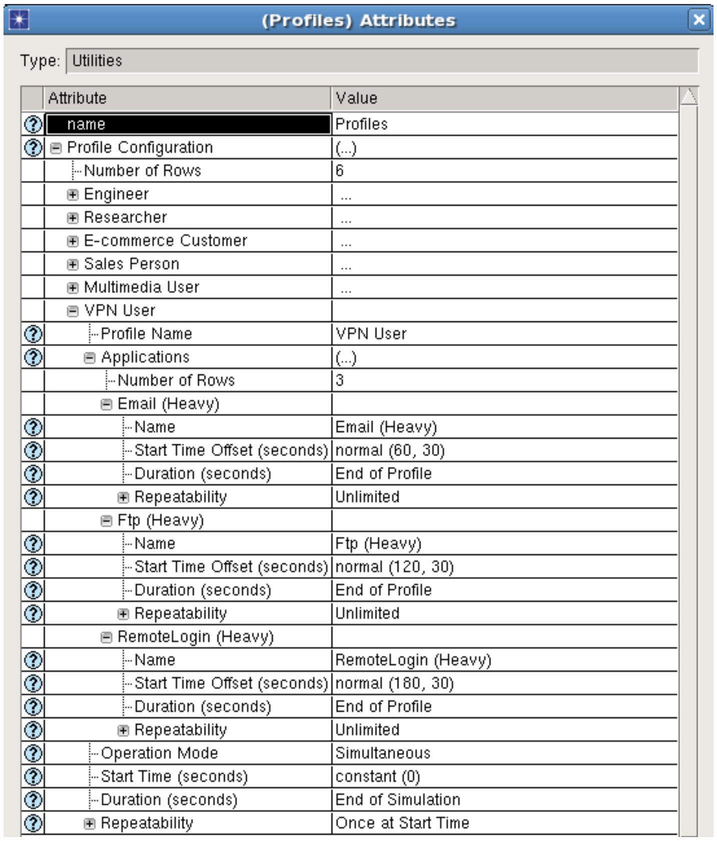

Type: Utilities

	Attribute	Value
?	name	Profiles
?	⊟ Profile Configuration	(...)
	Number of Rows	6
	⊞ Engineer	...
	⊞ Researcher	...
	⊞ E-commerce Customer	...
	⊞ Sales Person	...
	⊞ Multimedia User	...
	⊟ VPN User	
?	Profile Name	VPN User
?	⊟ Applications	(...)
	Number of Rows	3
	⊟ Email (Heavy)	
?	Name	Email (Heavy)
?	Start Time Offset (seconds)	normal (60, 30)
?	Duration (seconds)	End of Profile
?	⊞ Repeatability	Unlimited
	⊟ Ftp (Heavy)	
?	Name	Ftp (Heavy)
?	Start Time Offset (seconds)	normal (120, 30)
?	Duration (seconds)	End of Profile
?	⊞ Repeatability	Unlimited
	⊟ RemoteLogin (Heavy)	
?	Name	RemoteLogin (Heavy)
?	Start Time Offset (seconds)	normal (180, 30)
?	Duration (seconds)	End of Profile
?	⊞ Repeatability	Unlimited
?	Operation Mode	Simultaneous
?	Start Time (seconds)	constant (0)
?	Duration (seconds)	End of Simulation
?	⊞ Repeatability	Once at Start Time

Figure 5: Profile Config. Attributes

2.5 IP VPN Configuration

To ensure that VPN is being used in the path between the workstations and application server, the attributes in this module allow you to set tunnel parameters. The window shown below is for the expanded network but the method applies to both scenarios since the difference is only in remote client numbers. Given that the network has two routers connecting the workstations to the Internet cloud, a tunnel must be declared from Router A to D, as well as Router B to D. Furthermore, the remote client list option is used to specify which workstations are allowed to access the VPN tunnel. This mimics the credential-authentication required for employees to remotely connect using VPN. For this project they are arbitrarily named with the convention “Sales #”.

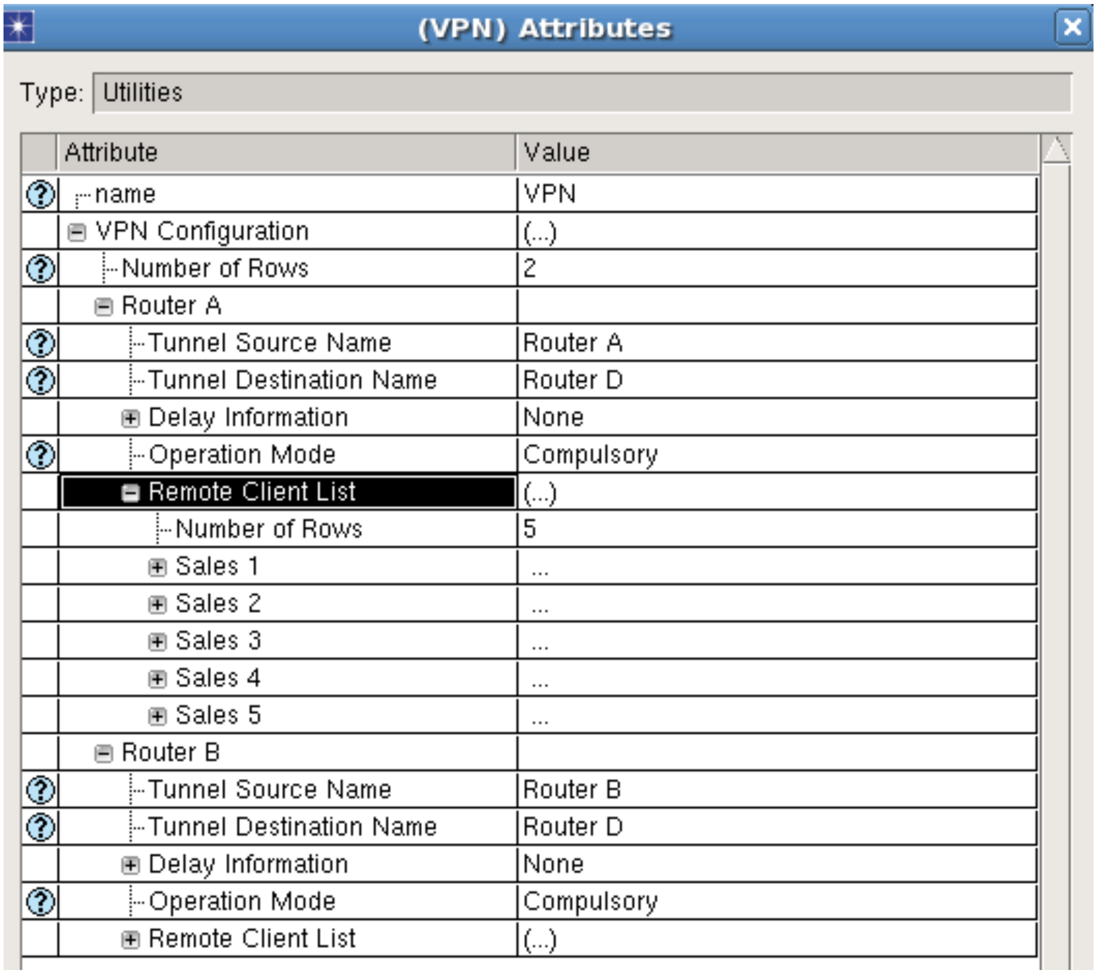

(VPN) Attributes

Type: Utilities

	Attribute	Value
?	name	VPN
	⊟ VPN Configuration	(...)
?	Number of Rows	2
	⊟ Router A	
?	Tunnel Source Name	Router A
?	Tunnel Destination Name	Router D
	⊞ Delay Information	None
?	Operation Mode	Compulsory
	⊟ Remote Client List	(...)
	Number of Rows	5
	⊞ Sales 1	...
	⊞ Sales 2	...
	⊞ Sales 3	...
	⊞ Sales 4	...
	⊞ Sales 5	...
	⊟ Router B	
?	Tunnel Source Name	Router B
?	Tunnel Destination Name	Router D
	⊞ Delay Information	None
?	Operation Mode	Compulsory
	⊞ Remote Client List	(...)

Figure 6: IP VPN Config. Attributes

2.6 Firewall / Server / Workstations

The remaining nodes that require operational settings are the firewall, workstations and server. The firewall's attribute of importance is the proxy server, where each application can have the proxy deployed based on what traffic needs to be blocked or allowed to pass through. Figure 7 shows how Email and Remote Login have no proxy deployed, but FTP does. This would mean that in a non-VPN remote connection, any FTP application traffic would be blocked in the network. This is not what is required in this project's network, hence the VPN deployed from Router A/B to D will bypass this firewall using IPsec and other security protocols included in secure VPN connections to provide server access to FTP as well. This method tunneling around the firewall is confirmed in the simulations shown later in the report. The server only required the parameter of supported services to be set to "All". Lastly, each workstation has its attributes set to utilize the previously created 'VPN User' profile.

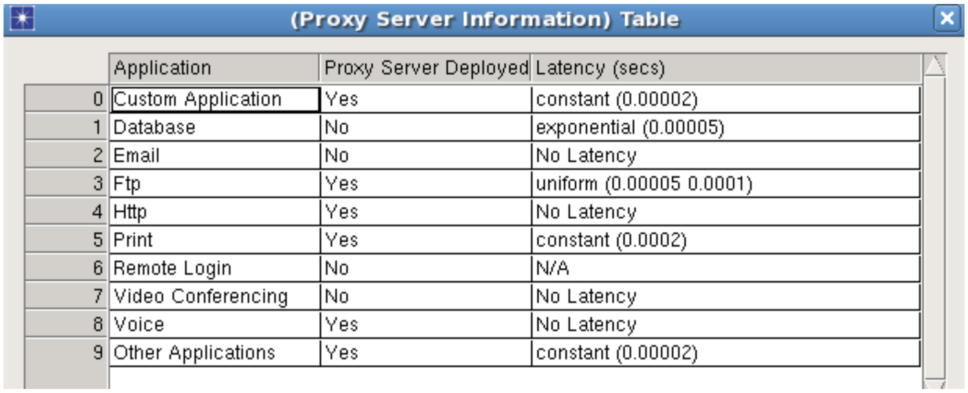

(Proxy Server Information) Table

	Application	Proxy Server Deployed	Latency (secs)
0	Custom Application	Yes	constant (0.00002)
1	Database	No	exponential (0.00005)
2	Email	No	No Latency
3	Ftp	Yes	uniform (0.00005 0.0001)
4	Http	Yes	No Latency
5	Print	Yes	constant (0.0002)
6	Remote Login	No	N/A
7	Video Conferencing	No	No Latency
8	Voice	Yes	No Latency
9	Other Applications	Yes	constant (0.00002)

Figure 7: Firewall Proxy Settings

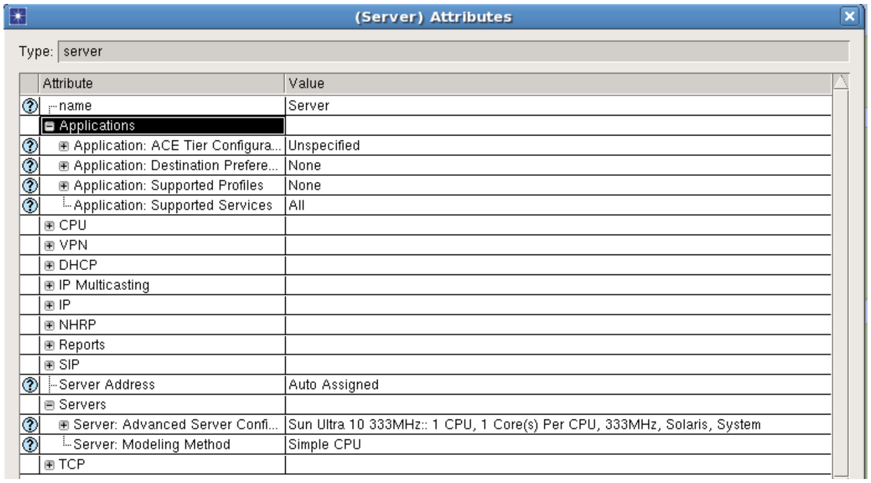

(Server) Attributes

Type: server

Attribute	Value
name	Server
⊟ Applications	
⊞ Application: ACE Tier Configura...	Unspecified
⊞ Application: Destination Prefere...	None
⊞ Application: Supported Profiles	None
Application: Supported Services	All
⊞ CPU	
⊞ VPN	
⊞ DHCP	
⊞ IP Multicasting	
⊞ IP	
⊞ NHRP	
⊞ Reports	
⊞ SIP	
Server Address	Auto Assigned
⊟ Servers	
⊞ Server: Advanced Server Confi...	Sun Ultra 10 333MHz:: 1 CPU, 1 Core(s) Per CPU, 333MHz, Solaris, System
Server: Modeling Method	Simple CPU
⊞ TCP	

Figure 8: Server Attributes

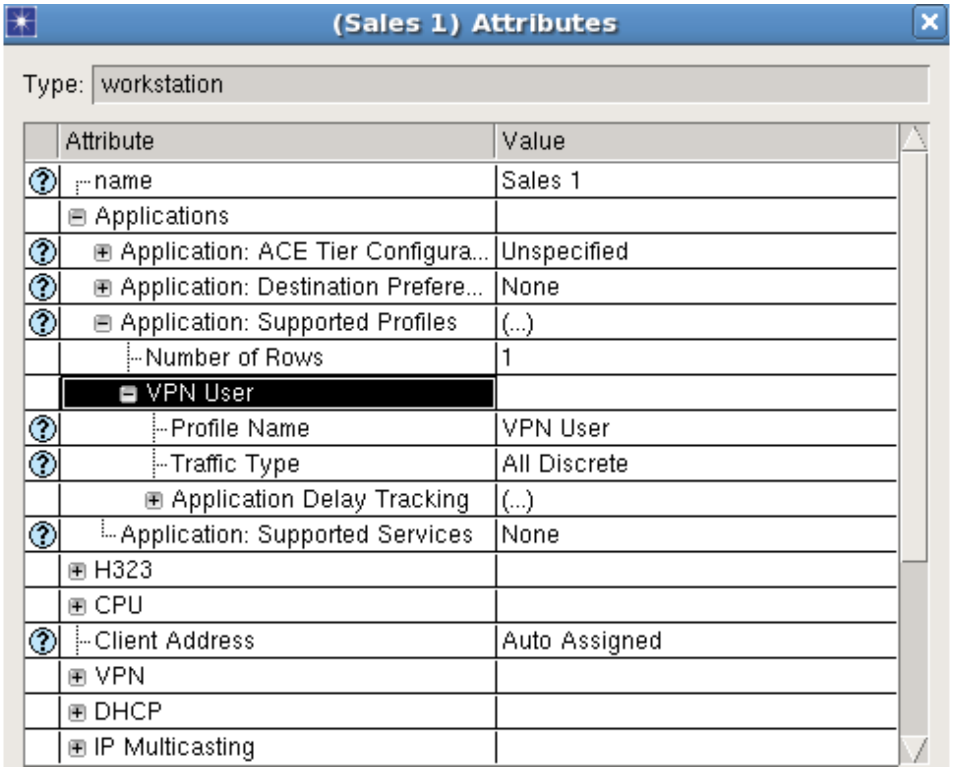

(Sales 1) Attributes

Type: workstation

Attribute	Value
name	Sales 1
⊟ Applications	
⊞ Application: ACE Tier Configura...	Unspecified
⊞ Application: Destination Prefere...	None
⊟ Application: Supported Profiles	(...)
Number of Rows	1
⊟ VPN User	
Profile Name	VPN User
Traffic Type	All Discrete
⊞ Application Delay Tracking	(...)
Application: Supported Services	None
⊞ H323	
⊞ CPU	
Client Address	Auto Assigned
⊞ VPN	
⊞ DHCP	
⊞ IP Multicasting	

Figure 9: Workstation Attributes

3. Discussion

The statistics selected in the DES will be shown in this section. In simulating both scenarios simultaneously, the profile was set to send traffic for the first 10 minutes while the simulations were run for 30 minutes to allow the values to reach a steady state if possible. The settings shown in section 2.1 were made for the profile to continue application use for the duration of the entire simulation. This result will be shown as well. OPNET ran the simulations in roughly 5 to 10 seconds based on the scenario.

3.1 Global Received Traffic

The following two figures depict the global received traffic in the network. As Remote-Login requires constant connection and therefore has the highest spike in packets transferred. Email and FTP access do not require a continuous connection in the sense that the user will receive an email and no further traffic is required for that event, so the traffic is fairly low when compared. Furthermore, the expanded network's traffic has a peak value of 0.7 packets/sec versus 0.14 packets/sec for the baseline network. This gives a ratio of about 5:1, which agrees with the topology of the expansion scenario having five times as many workstations.

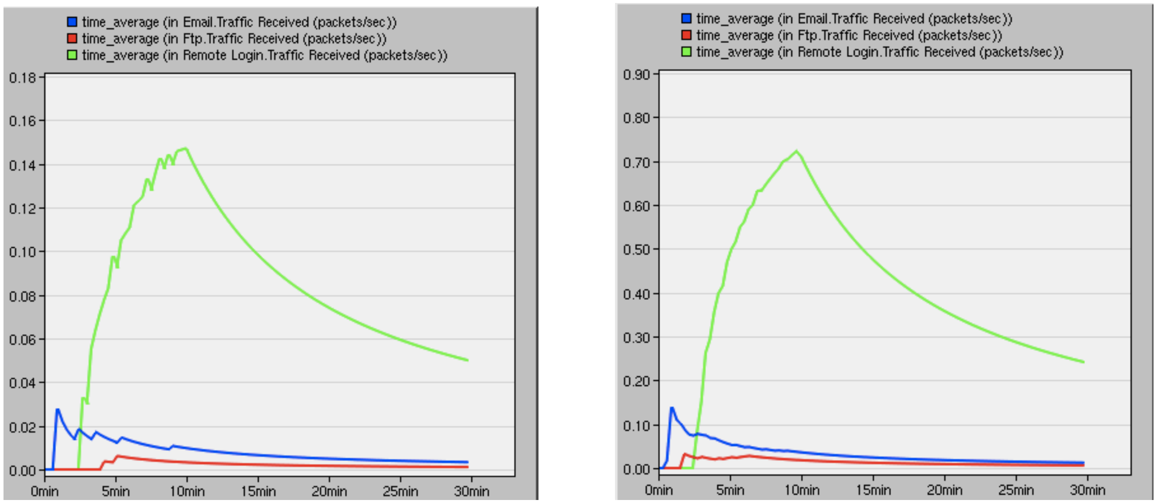

Figure 10: Global Received Traffic (Packets/sec): Baseline (Left) / Expanded (Right)

3.2 Server Processing Time

The graph in Figure 11 illustrates the server's task processing time for both scenarios. As you can see, our expanded network (red line) causes the server's task processing time to peak at 0.016 seconds at 2 minutes, whereas at the same point in time, the baseline network (blue line) is merely 0.006 seconds. This makes sense because as the number of workstations increase, so does the number of requests, which in turn puts more load on the server and slows it down, resulting in longer processing times for each task. As we move forward in time, the task processing times decrease until the 10-minute mark. This shows that the server has dealt with the initial surge of requests and is steadily making its way through the rest. Beyond the 10-minute mark, the task processing times stay at a constant level of around 0.0043. This is due to the fact that we have defined our profile in such a way that all workstations stop communicating with the server after 10 minutes. On a side note, a weird anomaly occurs around the 3-4 minute mark for the baseline network. For some reason, the task processing time decreases, reaching a minimum that's even lower than the constant level beyond 10 minutes. We are not sure why this happens and would require future comparisons with topology changes to see if the same effect is observed.

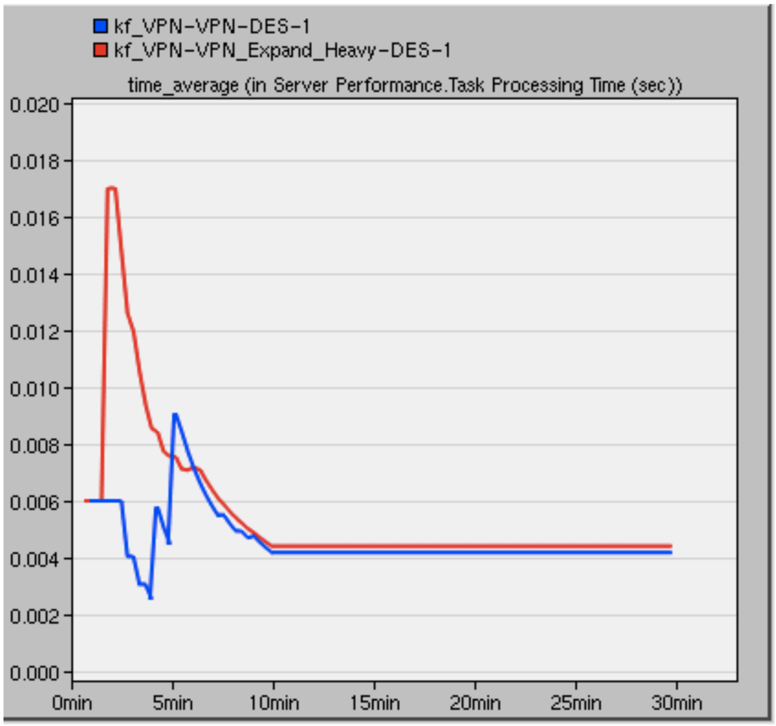

Figure 11: Server Processing Time

3.3 IP VPN Tunnel Delay

As mentioned earlier, the goal of this project was to observe the performance of VPN on a baseline and expanded network. To get a clear observation of this, we looked at the IP VPN Tunnel delay from Router D to A and Router D to B. As seen in Figure 12 (left), although there is a small increase in overall delay, the tunnel's performance over the course of the simulation is fairly similar for both scenarios. This shows that under our network topology, VPN maintains its quality of service regardless of the increase in workstations as well as the server's task processing delay. However, this result is still based on the applications being used for only 10 minutes out of the simulation time. Figure 12 (right) is the result of the profile configuration settings shown in section 2.3 where the workstations are active for 8 hours or, in other words, an entire workday. In this case, the tunnel delay is not time-averaged so as to show the variations based on application use, but the overall performance is similar to the 30-minute simulation. In addition, the 8-hour simulation shows the tunnel managing the increase in traffic better as the variations are smaller, when compared to the baseline network.

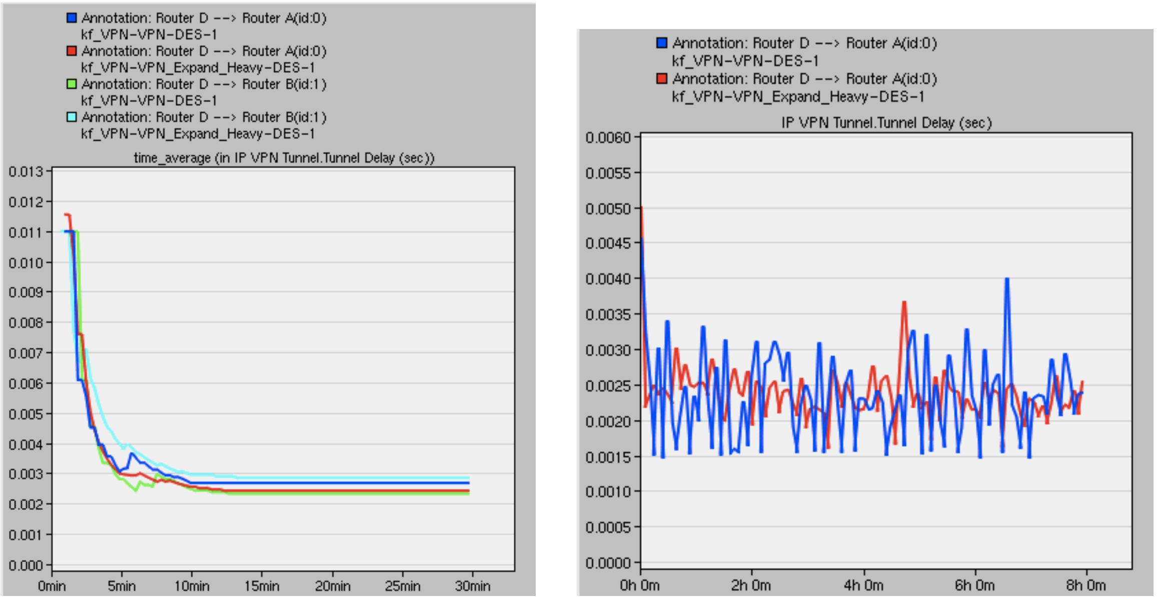

Figure 12: IP VPN Tunnel Delay: 10 minutes (Left) / 8 hours (Right)

3.4 Difficulties

Initially, the project included video-conferencing as an application choice since it is so widely used in office environments. While implementing the connection, it was seen that only the workstations were sending data but nothing was being received in terms of video-conferencing communication. Due to time constraint, and requirement for implementing other applications, this application was eliminated from the statistic collection. Secondly, the expanded network's star topology around Router A and B was chosen due to state machine inequalities when connecting a third router to the Internet cloud. The error was traced to a state-machine fault occurring in any new router that was connected to the Internet cloud in addition to A and B. The OPNET 16.0 administrator was later notified of this occurring in the project file, but our team continued with the network topology seen in this report so as to meet the deadlines and deliver some results. Lastly, it was important to understand what statistics were actually relevant in getting a better understanding of what VPN can deliver in a network. This led to trial of various simulation strategies and statistic(s) collection until the selection was narrowed down to those discussed in this report. This included experimenting with firewall parameters to help understand the effect VPN has on the proxy server.

3.5 Future Work

For future work, we should add more applications to our profile so we can model a more realistic scenario and obtain a more complete summary of VPN's effects. In particular, we should try to implement video conferencing in such a way that it can both send and receive data to and from the server. This is an important step forward because video conferencing is such a widely used tool in the office environment. We should also try to model different types of network topologies, such as tree, ring, and mesh. As we mentioned in the previous section, OPNET 16.0 would not allow us to use more than two routers to construct the clients' network. If we are able to resolve this problem in the future, then it would definitely be interesting to see whether the network's topology affects the quality of service. Another aspect that we can improve is our applications' traffic models. For our project, we let email start at 60 seconds, ftp at 120 seconds, and remote login at 180 seconds, with all of them having normal distributions. In reality, clients would not be requesting these applications in such a manner. Some clients may only check email for a few minutes, while others may use remote login for a few hours to work on projects from home. Lastly, to expand on the idea of investigating the effects of VPN, we can try to implement Multiprotocol Label Switching (MPLS) VPN for the same office network.

4. Conclusion

By modeling two VPN networks (one with two workstations and the other with ten) and analyzing the server and tunnel performances, we can come to a few conclusions. Firstly, the IP VPN tunnel delay results have shown us that the number of clients and the amount of traffic does not significantly affect the VPN tunnel delay. Both the 30 minute simulation and the 8 hour simulation showed no obvious difference between the delays. Secondly, the servers' performance results showed that the server was being affected the most by the increase in workstations. The server in the expanded network experiences a huge increase in task processing time during the first few minutes of the simulation. Since the VPN tunnel delay stays constant, it makes perfect sense that the server would be the one to experience the most change. Taking everything into consideration, it seems that site-to-site IP VPN is a fairly good method of transferring data over the Internet because it provides an adequate level of security and the delay is unaffected by traffic or number of clients.

5. References

[1] A. Zaballos, G. Corral, I. Serra, J. Abella, "Testing Network Security Using OPNET," OPNETWORK.2003, Washington DC (United States), August 2003.

[2] E. Aboelela, "Firewalls and VPN: Network Security and Virtual Private Networks," Computer Networks: A Systems Approach - Network Simulation Experiments Manual, 5ed, March 2011.

[3] H. Bourdoucen, A. Al Naamany, A. Al Kalbani, "Impact of Implementing VPN to Secure Wireless LAN," International Journal of Computer and Information Engineering, 2009.

[4] R. Malhotra, R. Narula, "Techno-Evaluation and Empirical Study of Virtual Private Networks Using Simulations," Journal of Computing, Volume 3, Issue 7, July 2011.

[5] Matutes Mestre, Manuel, "Simulation-Based Analysis of VPN Technologies," 2003.