

Lab 3: WAN Capacity Planning for an Application Deployment

Objective

In this lab, OPNET's IT Guru Academic Edition modeling software is used to perform a capacity planning study. IT Guru predicts the impact a new application deployment will have on network resources, and the impact the network will have on the application performance.

You will learn the benefits of simulating applications on the network infrastructure prior to their deployment. You will understand how a new application affects the performance of other applications and what important decisions related to network infrastructure upgrades are needed to improve application performance.

Lab Project

Grow Money is a finance company with over \$20 billion in assets, and offices throughout the US. Their plan is to deploy a new application called Ubiquitous, which will disseminate their strategy, over the course of the day, to all offices, so that executives in various locations can take their daily trading positions in line with the centralized strategy. Due to huge financial stakes, every user must be able to access every Ubiquitous transaction in less than 5 minutes. As one of Grow Money's network engineers, you have been assigned the task to predict, before Ubiquitous' deployment, the impact by answering:

- (1) How will the network's background load affect Ubiquitous' response time as experienced by the end user? , and
- (2) What effect will the expected Ubiquitous' load have on the network?

Overview

In this lab, there are 3 scenarios. The first scenario studies the network and application performance prior to Ubiquitous' deployment. In the second scenario, we deploy Ubiquitous, study its response time, and determine the impact it has on the network and other applications. The third scenario focuses on evaluating an option of a WAN upgrade to improve the performance.

Lab Instructions

Note: Bold marked words indicate menu/clickable options.

You should be familiar with OPNET software and its environment. We will skip detailed instructions and focus more on results.

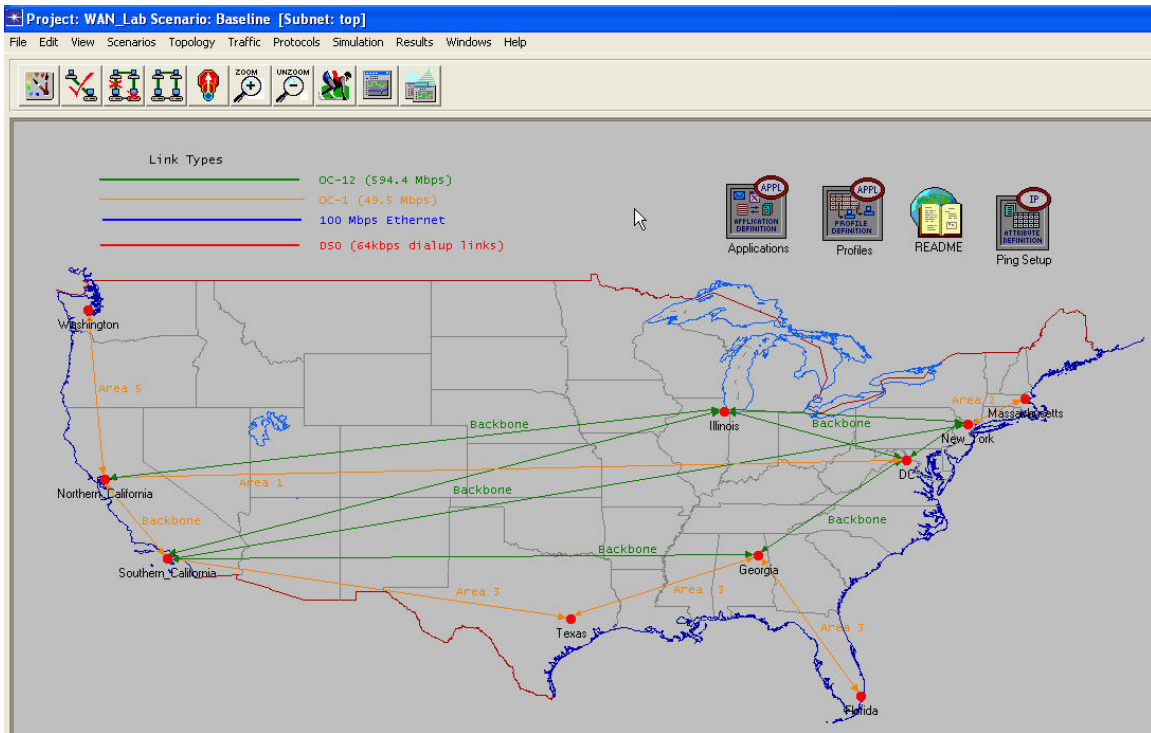
Part 1: Understand the Model

1. Start IT Guru Academic Edition 9.1.
2. Select **File** -> **Open...**

3. Scroll to the project named **WAN_Lab**, select it and click **OK**. If you do not see this project in the list, it means you have not downloaded and installed the labs properly. For complete instructions, follow the steps on this page:

http://opnet.com/services/university/Install_Lab.html


4. The first scenario is named “Baseline”. You can verify the project and scenario names by observing the title bar.



5. Double-click on the README file icon to learn more about the scenario.



Return to the topology by clicking

6. You may wish to double-click the  icon that marks every region in this view to drill down into the inter-city and/or intra-city infrastructure. Return to the topology

by clicking .

Part 2: Configure and Run the Simulation

The goal is to evaluate network performance by simulating 50 minutes of a typical day. This is done by running a high-fidelity Discrete Event Simulation.

1. Click on the **Configure/Run Simulation** toolbar button .

2. Make sure the Simulation **Duration** is set to **50** minutes.
3. Click **Run**. Monitor the progress bar as the simulation proceeds. Ignore the messages, if any, that state "Recoverable Error". This simulation may take many minutes to complete. Be patient.
4. When the simulation completes, Click **Close**.

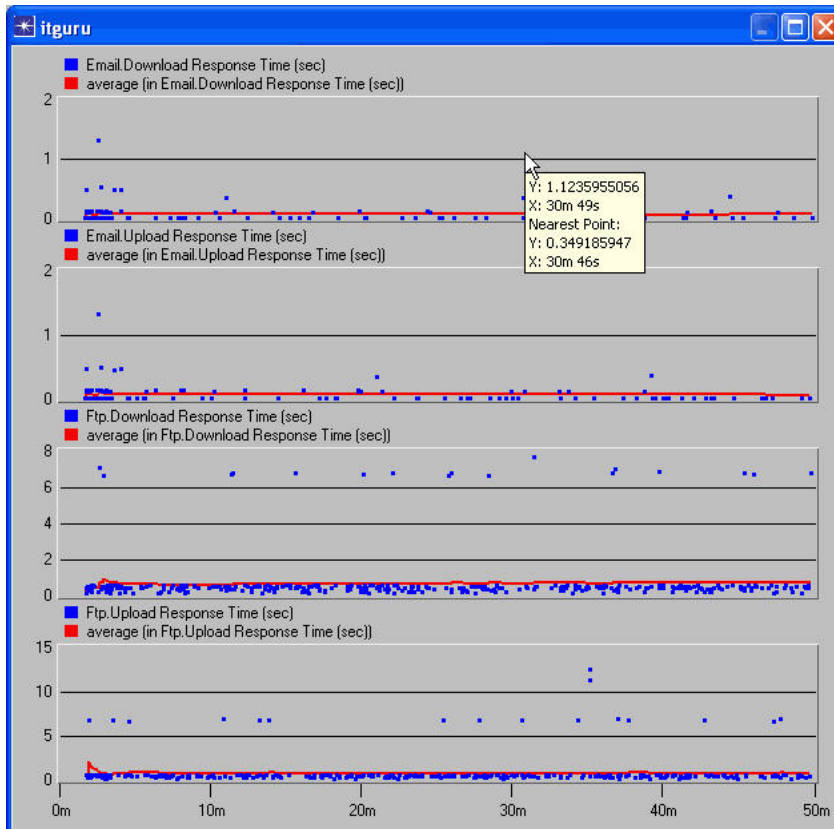
Part 3: View Results

The goal here is to observe the results of (a) application response times for HTTP, FTP, and Email traffic, (b) packet delay and packet jitter values for Voice traffic, (c) top 10 link utilization values, and (d) top 10 queuing delay values.

1. Click on **Results->View Results**.

View Application Response Times:

2. Expand **Global Statistics -> HTTP**. Repeat for **Email, FTP, and Voice** applications.
3. For both Email and FTP, select **Download Response Time (sec)** and **Upload Response Time (sec)**. The selected graphs will be plotted on the right-side of the frame. On the lower-right side, make sure the settings are **Stacked Statistics, This Scenario** and **As Is**. Click **Show**. A window is created with the graphs containing the raw data for the Email and FTP application response times.
4. In the View Results Window, click the **As Is** drop-down menu and select **average**. This will convert the raw data into an average curve. Click **Add** and then click on the top-most graph of the window that you created in step 3. Now you have both the raw data and average graphs for Email and FTP application response times.



Note: If you did not follow the underlined instructions above, you may have mismatched raw data and average graphs, so close the graph, choose **Delete** and go through steps 3 and 4 again

You may place the resulting window anywhere on your desktop by dragging it by its header bar.

5. Click **Unselect** to clear the current selections.

6. For HTTP, select **Page Response Time (sec)** and plot the raw and average data on one graph using steps 3-5 above.

7. For Voice, select **Packet End-to-End Delay (sec)** and **Packet Delay Variation**. Packet delay variation is synonymous with packet jitter. Plot the raw and average data on one graph using steps 3-5 above.

What can you conclude looking at voice packet delay values?

Answer: Voice packets experience an average delay of around 15-20 msec. It is recommended that voice packets have no more than a 50 msec delay to maintain a high quality of conversation. In this simulation, voice application quality is quite good.

View Top 10 Link Utilizations and Queuing Delays:

8. Click **Results -> Find Top Results**. Expand **Link Statistics -> point-to-point** and select **Utilization** and click **Find Top Results** button. You will see link utilization values in a table sorted in descending order.

What is the most utilized link? Also, write its utilization value.

Answer: _____

9. If you wish to graph these values, click on **Graphs Stacked** drop-down menu and select **Graphs Overlaid**. Click **Graph**.

10. Go back to **Select Statistics** Window and click **Close**.

11. Repeat steps 8-10 above but choose **queuing delays** instead of **Utilization**.

Which link has the maximum queuing delay? Also, write the maximum delay value.

Answer: _____


12. Close the **Top Objects: point-to-point.utilization** and **Top Objects: point-to-point.queuing delay** window.

Part 4: Simulate Ubiquitous' Deployment

1. Click on **Scenarios ->Switch Scenarios**.

2. Choose scenario **Ubiquitous_Deployment**.



3. Read the "README". Click the  icon to return to the topology view. Verify that Ubiquitous application is deployed by double-clicking **DC** node. Notice the Ubiquitous server is deployed on the same physical box where the Email server resides. Clients running Ubiquitous application will communicate with this server. Configuring custom applications is beyond the scope of this manual.

4. To verify the application deployment, double-click on **Southern_California** node and then, right-click on **Malibu_LAN** object. Select **Edit Attributes**. Select **Supported Profiles** and click the (...), select **Edit** and observe that **Ubi_profile** is configured. 2 clients within this LAN will be acting as clients that run the Ubiquitous application. Click **Cancel** on all the windows to return to the scenario.

5. The traffic from other clients accessing Ubiquitous has been abstracted into IP Traffic Flows. The blue dotted lines pointing towards **Malibu_LAN** and

Irvine_LAN are examples of such **IP Traffic Flows**. Select a traffic flow, right-click on it, and select **Edit Attributes**. Click on **Traffic (bits/second)** and select **Select**. You will see the variable traffic profile that the IP flow represents. Click **Cancel** on all the windows to return to the scenario.

6. Run the simulation for the **Ubiquitous_Deployment** scenario. Follow the steps outlined in **Part 2: Configure and Run the Simulation** section.

7. Once the simulation completes, view Ubiquitous' performance by clicking **Results -> View Results**. Select **Custom Application** and choose **Application Response Time (sec) <Ubi_profile / Ubiquitous>**. Place both the raw and average data on one graph. Refer to steps 3 and 4 in **Part 3: View Results** section. It shows that Ubiquitous response time is between 2-5 minutes, which is acceptable.

8. Does Ubiquitous deployment impact other applications? We will now compare results with the previous scenario to study the impact Ubiquitous has had on other applications by graphing response times for Email, HTTP, FTP and Voice.

9. Click **Results -> Compare Results**.

10. Expand **Global Statistics -> HTTP**. Repeat for **Email, FTP, and Voice** applications. At the lower right corner, make sure the settings are **Overlaid Statistics, All Scenarios, and As Is**.

11. For Voice, select **Packet End-to-End Delay (sec)**. Click **Show**. Click the **As Is** drop-down menu and select **average**. Click **Show** again. You now have two graphs (one raw data and one average) comparing the voice performance across the two scenarios. What do you observe?

Answer: The voice quality has degraded due to Ubiquitous deployment as seen by increased packet delay.

12. Repeat step 11 for **Packet Delay Variation (jitter)**. Note the jitter values increase with the simulation time indicating it has not reached steady state when the simulation finished. Increase in jitter values will lead to unpredictable voice quality.

Note: Licensing for IT Guru Academic Edition 9.1 renders a limit on the number of simulation events it can complete. Ideally we want to run simulation until it reaches steady state but due to this limitation we set the Simulation Run Time to 50 minutes. The enterprise version of IT Guru does not have such limitation.

13. Repeat step 11 for **HTTP** choosing **Page Response Time (seconds)**. Choose **average** values from here on for graph purposes.

14. Repeat step 11 for **FTP** choosing either **Download** or **Upload Response Time (seconds)**. Notice the trend is similar.

15. Lastly, repeat step 11 for **Email** choosing either **Download** or **Upload Response Time (seconds)**. Once again, notice the trend is similar.

Ubiquitous' deployment has adversely impacted the performance of other applications. Applications can impact each other based on what priority they have on the network. Ubiquitous was configured to be a best-effort service (on par with other applications), so the performance degradation points to a potential over-utilization of link bandwidth.

It is essential to know about such problems before an actual application deployment. OPNET's advanced modeling software helps predict such issues before they occur in actual deployment.

16. Click **Results -> Find Top Results**. Expand **Link Statistics -> point-to-point** and select **Utilization** and click **Find Top Results** button. Notice there are several links that reach 100% maximum utilization over the course of simulation. When a link gets utilized completely, the incoming packets get queued up thereby increasing the queuing delay. Click **Close**.

List all links that reach utilization values of 100%.

Answer: _____ (use multiple lines if necessary)

17. Right-click on the link between **Southern_CA** and **New_York**. Select **Edit Attributes** and note the model – it says **PPP_SONET_OC12**.

We will upgrade this link and some other 100% utilized links in the next scenario to see if the above problem caused by Ubiquitous goes away.

18. Repeat step 16 but choose **queuing delays** instead of **Utilization**. Note the top 10 values. You may wish to compare these values after you complete the third scenario.

Which link has the maximum queuing delay? Also, write the maximum delay value.


Answer: _____

Note: In general, there may be several ways to address the types of problem faced in our second scenario – the solution we offer in the third scenario is just one of the ways.

Part 5: Simulate WAN Upgrades.

1. Click on **Scenarios ->Switch Scenarios**.

2. Choose scenario **WAN_Upgrade**.

3. Read the **README** file by double-clicking the icon. Click the  icon to return to the topology view.

4. Note the upgraded link between **Southern_CA** and **New_York**. Right-click on the link and select **Edit Attributes**. Note the model is **PPP_SONET_OC24**.

Note: The bandwidth utilization % values have been adjusted because it is input as a percentage of the link bandwidth. Since we doubled the bandwidth from PPP_SONET_OC12 to PPP_SONET_OC24, we reduced the background utilization by a factor of 2 to keep background load the same across both the Ubiquitous_Deployment and WAN_Upgrade scenarios. You don't have to make any changes in the lab files – they have been made for you. Bandwidth Utilization is an attribute of a link should you desire to take a look at values.

To note the other upgraded links, explore each city by double-clicking the icon starting with **Southern_CA**. Upgraded links will be marked.

The full list of upgraded links is as follows:

- a. Southern_California to NewYork.
- b. Texas to Georgia.
- c. Los_Angeles to Malibu (double click on Southern_California node to view)
- d. Los_Angeles to Irvine
- e. Irvine to Irvine_LAN

5. Run the simulation and compare the results by doing steps 6 – 15 outlined in **Part 4: Simulate Ubiquitous Deployment** section.

What do you observe?

Answer: After doing the WAN upgrade, the response times of other application have returned to levels seen in the baseline scenario. Ubiquitous Application Response Times have not changed much and stay within acceptable levels.

6. Repeat steps 16 & 18 (outlined in **Part 4: Simulate Ubiquitous Deployment** section).

7. Compare the values you get for this scenario. Click **Results -> Compare Results**. Expand **Object Statistics -> Southern_California <-> NewYork [0]**. Further, expand **point-to-point** and select **queuing delay -->**. Click **Show** to see

how queuing delays varied across scenarios. Select **Unselect** in the **Compare Results** window, select **queuing delay <--** and click **Show**.

What do you observe?

Answer: We also observe that no link has reached 100% utilization, and hence effect of queuing delays on application response times has reduced.

We do note that towards the end, there is very small jump in the jitter value. This suggests that (a) we have not necessarily reached the steady state, and (b) this network may not be completely ready for Ubiquitous deployment. Perhaps this warrants a further study using various QoS schemes based on the principle of traffic engineering.

Note: We did not upgrade the 64 kbps links between users in Miami, Boston, Seattle and their routers. Their utilization continues to be 100%. The impact of those links are reflected in the higher-than-median values of FTP, Email and HTTP response times that are observed. You may wish to duplicate a scenario by clicking **Scenarios -> Duplicate Scenario** to study the effects of upgrading the 64 kbps links as well.

CONCLUSION: We have successfully sized our network by understanding the impact of Ubiquitous' deployment on existing applications' response times and on network capacity. Likewise, we understand the impact of the WAN link upgrade required to bring Voice application response times to acceptable limits without changing Ubiquitous' response times too much. We acknowledge the study may warrant a further QoS study in light of a small jump in jitter values towards the end of the simulation.