



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials Bootcamp Style (Security 401)"
at <http://www.giac.org/registration/gsec>

Securing a University's Bandwidth with PacketShaper

Introduction:

This paper is not limited to universities and could be applied to any network architecture. It is meant to bring attention to the importance of securing any network's bandwidth. This paper will assist the reader in the implementation, installation and configuration of the PacketShaper and the processes that are necessary to apply bandwidth utilization policies. It is important to remember that there is no "one size fits all" solution. I suggest using what is pertinent to your scenario and learn from my mistakes. I am not providing a guaranteed solution or an instructional paper; I am merely providing you with tools, strategies and the technology that I used in securing and providing reliable bandwidth to our institution.

One must also understand that this paper is written with an emphasis on a university network which differs greatly from traditional corporate enterprises. According to Ted Udelson, academic institutions are presented with special and complex challenges which are not faced by commercial or government entities. He further lists the most common threats:

They have difficulty in controlling end users.

The culture cultivates free thinking and "open" access to information.

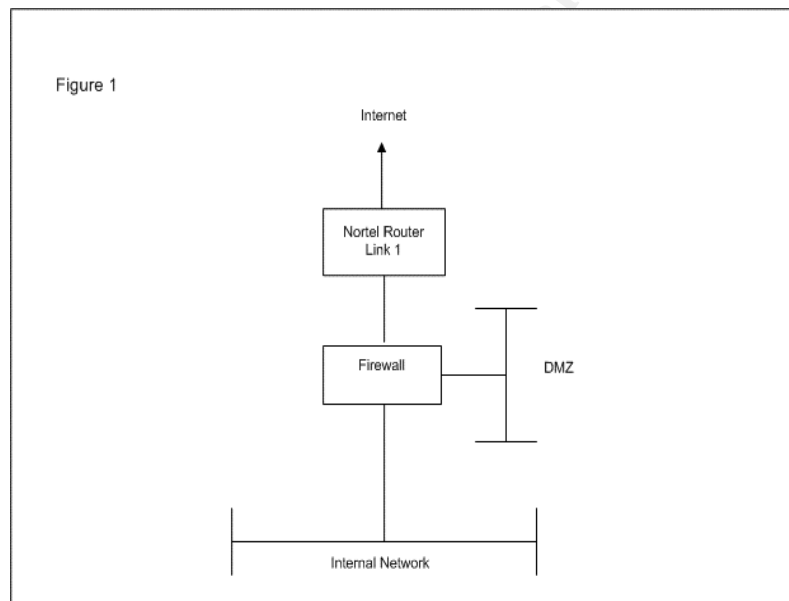
The university serves as a research body, corporation, and Internet service provider. Colleges and universities must analyze each of these functions to determine the proper stance to take with regard to security (Udelson, p. 10).

These points brought up by Mr. Udelson, present a network administrator with many challenging and unique tasks. It is important to first, understand the threats that are specific to your network environment and then develop a solution that will fit best for your specific scenario.

Scenario: Before PacketShaper

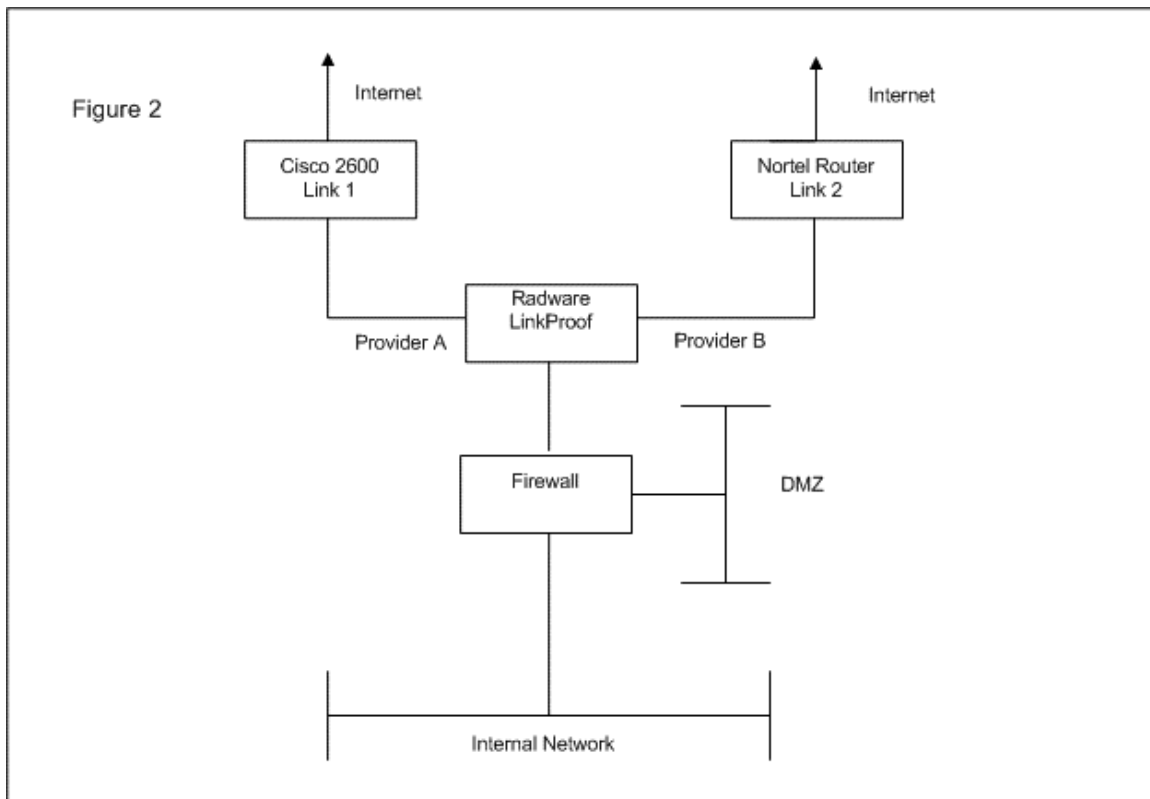
In late 2001, administration had received complaints from several students that the bandwidth that was provided to them was not adequate at times to conduct research. Specifically, students complained that at certain times of the day (a stretch between 10:00pm and 2:00am) internet access would come to a complete halt.

This was brought up to the CIO and the concern was later passed off to me. I conducted some research and monitoring using MRTG tool on our single T1. My report of the utilization of bandwidth showed that the T1 line idled between 80% and 90% utilization on working hours (9-5), and reached 100% during the 10:00pm – 2:00am stretch. **Figure 1** shows the basic public network setup.



My observation was passed along to my CIO and then onto administration. The problem needed to be resolved quickly and thus a very reactive decision was reached. Administration decided that the university should purchase an additional T1. This additional T1 was purchased in early 2002.

The university decided that it would purchase a device called Linkproof by Radware for the integration of both T1 lines. These T1 lines would be setup to provide load balancing, redundancy, and a larger bandwidth capacity. **Figure 2** shows the new design that was created for the integration of the dual T1.



The implementation of an additional T1 and the Radware Linkproof device were to provide the additional bandwidth needed and supply the university with some redundancy. The Linkproof device was able to eliminate

. . . link congestions and bottlenecks from multi-homed networks, for fault tolerant connectivity and continuous availability of web services. By intelligently routing traffic and controlling bandwidth service levels across all Internet links, Linkproof enables effective link utilization, accelerating responsiveness, controlling bandwidth consumption and economically scaling operations. (LinkProof, p. 1)

The additional T1 and Radware Linkproof solution provided the university with larger amount of capacity and offered the university the needed tolerance, but it was not able to monitor internal usage.

Two weeks into the winter semester of 2002, the administration continued to receive complaints of slow internet access. Bandwidth monitoring was conducted once again and during the peak hours for the university (10:00pm to 2:00am) bandwidth readings would burst to the 100% capacity.

My first approach to this situation was to use portions of the “Defense in Depth” strategy and identify the business goals by the administration, faculty, students and the IT Department. Administration wanted a controllable, cost effective and quick solution. Faculty wanted guaranteed bandwidth and the Communications Department wanted designated bandwidth to conduct their streaming video

projects and presentations. Students wanted everything, from peer to peer networks to online gaming and Xbox live gaming. The IT Department wanted a better solution, one that would provide filtering, control and designate bandwidth on a policy based system. The IT Department also needed to be able to implement a VOIP (Voice Over IP) solution with adequate QoS (Quality of Service) in the near future.

It became apparent to the IT department that we could not continue to add T1's, and that we needed to come up with a solution that would be able to measure, monitor, filter and shape the bandwidth traffic. A solution also needed to be backed up by an "Issue-specific Policy". Currently the university had no specific internet utilization policy neither developed nor implemented.

A New Problem:

At around the same time we were beginning to experience constant problems with our firewall. At first we did not know or realize that this problem was part of our lack of bandwidth control and knowledge. The log files would grow at a rate that the OS could not handle. This would cause the firewall to either freeze and hang or the harddrive designated for the log files would fill up and consequently shut down the firewall.

After researching the log files it was determined that the culprit was SMTP traffic initiating from internal clients (specifically students). There were two different options to solve this problem. Allow SMTP to go through the firewall which would propagate SMTP traffic to the outside world, or stop SMTP traffic at the internal core router. Our core router also served as our VLAN manager. We setup an ACL (Access Control List) to not allow student traffic to send SMTP traffic. This solution seemed to work. We began to experience problems with the core router less than a week into the implementation phase. The core router began to crash every 24 hours. Once the router was reloaded some SMTP traffic was still being filtered, but not all. It was agreed that we were going to not filter at the router level, and try to find the culprit students? At this point, I was not able to identify this problem as a miss management of bandwidth.

We decided that we would try to answer the following key questions, **Why? What ? Where? and How?.** **Why** monitor and secure bandwidth? **What** were we going to use to measure and secure bandwidth? **Where** did we need to monitor bandwidth? And **How** would we enforce these solutions?

Understanding the Importance of Securing Bandwidth

Before we can understand **Why** we should secure and manage bandwidth we must define bandwidth. Scientifically speaking,

...bandwidth is the width of the range of frequencies that an electronic signal occupies on a given transmission medium. Any

digital or analog signal has a bandwidth. In digital systems, bandwidth is expressed as data speed in bits per second (bps). In analog systems, bandwidth is expressed in terms of the difference between the highest-frequency signal component and the lowest-frequency signal component. (SearchNetworking.com, p. 1)

Generally speaking we identify bandwidth as the speed in which flow of information is transmitted back and forth within a network or between many networks. Usually the more bandwidth one has the better the flow of information is exchanged. This statement is generally true. We are going to identify some reasons **Why** it is important to secure your network's bandwidth.

The number one reason to secure your bandwidth is cost. Cost can be measured in a many different ways. The most obvious associated cost with bandwidth is your ISP costs. In our scenario, the university was currently using two T1 lines and one point to point WAN link. The total cost of the university bandwidth was about a \$30,000 yearly investment. This investment needed to be monitored, secured and efficiently utilized. Once bandwidth was converted to an investment it became apparent and easier to convince the administration that further studies and policies should be implemented.

Another reason to secure your bandwidth can be performance. We are referring to the overall performance of the university's bandwidth. Bottlenecks, congestions, dropped or lost packets and unnecessary retransmissions are all signs of an ill performing network. Many of these symptoms can be traced back to poorly managed bandwidth. Optimizing performance on a network basically attempts to minimize negative effecting traffic or "less desirable" traffic (P2P, video, sharing) and provide or guarantee the mission-critical applications their needed bandwidth.

Policy may dictate and mandate the need to secure and manage campus bandwidth. Our IT Department had no policies set to limit bandwidth, block "less desirable" traffic or manage bandwidth.

What to use? PacketShaper by Packeteer – A Brief Description

The next question that we needed to answer was, what were we going to use to measure and control bandwidth? We knew that we could setup MRTG tools and measure the overall bandwidth, but it was not going to help us analyze packets, protocols or control bandwidth. After an extensive comparison and research, we decided to use a product by Packeteer called PacketShaper.

PacketShaper is the bandwidth-management solution that brings predictable, efficient performance to applications running over enterprise wide-area networks (WANs) and the Internet. It balances traffic's demands, giving each type of traffic the bandwidth it needs to perform. PacketShaper protects critical traffic, paces bandwidth-greedy traffic, and prevents any single type of traffic from monopolizing resources. It provisions bandwidth to applications, sessions, branch offices, and/or users. (Four Steps Packeteer, p. 3)

PacketShaper was the device that was going to be able to monitor inbound and outbound traffic, as well as analyze and filter. This product would secure our bandwidth and we would be able to set forth "Issue-specific Policies" that could be enforced. Packeteer has produced a simple introductory paper on the PacketShaper product and how to deploy it in your network. It can be found via this URL:

<http://support.packeteer.com/documentation/packetguide/5.2.1/documents/4Steps.pdf>

First Step: "Classify Network Traffic"

This first steps means allowing PacketShaper to identify traffic as it passes through the device. PacketShaper has the ability to identify or classify traffic by applications, protocols, web pages, subnets, users and many more. It has the ability to automatically classify known applications and protocols. Since, new applications are added on a daily basis Packeteer makes new classification features available to customers by introducing new "easy plug in" features. If a vulnerability or application is introduced a new plug in will be offered. After downloading and applying the plug in; PacketShaper is able to automatically classify the new application or vulnerability.

PacketShaper has the ability to manually classify applications, subnets, protocols and other network traffic. As new applications are introduced they become more integrated, more bandwidth intensive and more difficult to classify under one category. PacketShaper has the ability to manually classify these complex applications that may differ from the simple IP scheme and single port applications. Some of the manual classification categories are as follows:

- **Web Classification:** Most of the traffic today resides through HTTP traffic. PacketShaper is able to identify and differentiate HTTP traffic, by direction of traffic, web URL, server based, or host name. This allows for more granularities within the HTTP class.
- **Intricate Port Classification:** PacketShaper is able to classify and analyze difficult traffic that uses multiple ports or conducts in port hopping. Through this same classification it is able to differ classify traffic that may share the same port
- **File-Sharing Protocol:** This category refers to the famous Napster, Kazaa, and Gnutella.

Second Step: “Analyze Behavior”

PacketShaper has the ability to measure the classes of traffic that were previously identified. It will be able to track “...traffic levels, detects network trends, measures response time, and calculates network efficiency” (Four Step Packeteer, p. 5). This period of analysis will help answer many questions regarding the bandwidth traffic of an organization. PacketShaper is managed through a simple web interface. This interface contains many helpful tabs that will be useful to analyze the classified traffic. One of the helpful tabs is the Monitor Tab:


Traffic Class Name	Class Bits	Policy Bits	Current (bps)	1 Min (bps)	Peak (bps)	Guar. Rate Failures	Partition Min/Max	Policy Type (Pri., Guar., Limit)	Top User Analysis
inbound			2.6M	1.6M	3.6M	0	uncommitted-sonst		
Localhost	20346	20346	501k	73.7k	1.2M	0		Priority (0)	
DMZ	123777	123777	0	0	2055	0		Ignore	
P2P			4368	2537	19.8k	0	0-19200		
eDonkey	6	6	0	0	355	0		Rate (0) 0-256	
Gnutella	6281926	NA	4368	2537	18.9k	0			
Kazaa	1483647	1483647	0	4	11.4k	0		Rate (1) 0-256	
Napster	96	NA	0	4	17.8k	0			
Web			1.6M	1.4M	2.9M	0			
HTTP			1.6M	1.4M	2.9M	0			
216_50KredicForum	2748	NA	0	3	1.4M	0			
Remote	0	NA	0	0	0	0			
StudentEmail	19985	NA	628	623	1.6k	0			
Tevemb	12083	NA	0	4	34.5k	0			
Outlook	12487645	12487645	1.6M	1.3M	2.9M	0		Rate (0) 0	
DNS	668888	668888	1003	343	13.2k	0		Priority (0)	
FTP	7642	NA	0	0	28.8k	0	0-240		
Hotline	0	NA	0	0	0	0			
IRC	974	NA	0	0	41.3k	0			
U2TP	1	NA	0	0	9420	0			
LDAP	41	NA	0	0	691	0			
lockd	0	NA	0	0	0	0			
Mirroring	143	NA	37	75	361k	0			
MPEG-Audio	1312	NA	0	18	1.9M	0			
MPEG-Video	8364	NA	0	1	2.6M	0			
MSN-Messenger	41201	NA	218	227	227k	0			

This tab will identify the automatic or manually set classes on the left column, it also will show such columns as Current (bps), 1 Minute (bps), and Peak (bps). This tab will be very helpful in pulling data on desired classes and will become an important gathering tool for controlling bandwidth.

Third Step: “Control Performance”

PacketShaper is able to manage application performance and guarantee a preset amount of bandwidth. PacketShaper controls bandwidth through the usage of partitions. A partition “...creates a virtual separate pipe for a traffic class” (Four Steps Packeteer, p. 5). One is able to set a size for the reserve link, define whether it can expand over the cap and control that growth. Partitions work much like pipes within pipes. Figure 4 shows the relationship of partitions within partitions:

Figure 4:



Picture from Packeteer Website at URL:
<http://support.packeteer.com/documentation/packetguide/5.2.1/documents/4Steps.pdf>, p. 20

There are different kinds of partitions that can be utilized. PacketShaper can use either “hierarchical partitions” or “dynamic partitions”. “Hierarchical partitions” enable one to preset a certain amount of bandwidth within another subset of partitions. For example, one could set 30% of a link designated to HTTP traffic, and then assign different portions of the preset 30% to web servers that utilize HTTP traffic. One could assign half of the 30% to all web servers, quarter to OWA traffic and the remaining to any HTTP traffic. “Dynamic partition”, allows one set partitions on a per-user basis. It allows one to manage a user’s bandwidth allocation across all types of applications.


Step Four: “Report Results”

The reporting capabilities of the PacketShaper allow for a quick visual and comprehensive analysis of the traffic flow. PacketShaper will graph bandwidth based on time, network efficiency, average bandwidth and peak periods. This ability to quickly see what is traversing the network becomes a powerful and helpful tool in reaching your optimal goal of securing desired bandwidth performance.

Where to Use PacketShaper?

Now that we understand what to use to monitor, shape and manage our bandwidth I had to decide where to place this device within our network. The placement of the PacketShaper depended on our needs, desires, budget and the current topology of our network. I will discuss the basic options that we had and the advantages and disadvantages of each placement.

I took a basic and common setup of most university topologies and introduce the possible options of placement. **Figure 5** shows the different options:



Option 1, implements PacketShaper outside the border router. One of the positives to this solution is that you will be able to shape incoming and outgoing packets at this topology level. The other positive is that only external traffic will be shaped all internal traffic will not be accessed or modified. One of the negatives is that internal traffic will not be controlled, or managed. Another negative is that the PacketShaper will need a WAN or T-1 interface which will be more expensive and less flexible.

Option 2, does not require PacketShaper as we are using the router to shape bandwidth. The positives to this solution are that you do not have to buy or manage an additional device. Another positive is that internal traffic is not interfered with or shaped. The negatives to this option are that you are restricted to router based shaping, which is very limited and less effective. The other disadvantage is that you will be taxing the router CPU. Routers are designed to route traffic not to shape it and analyze it.

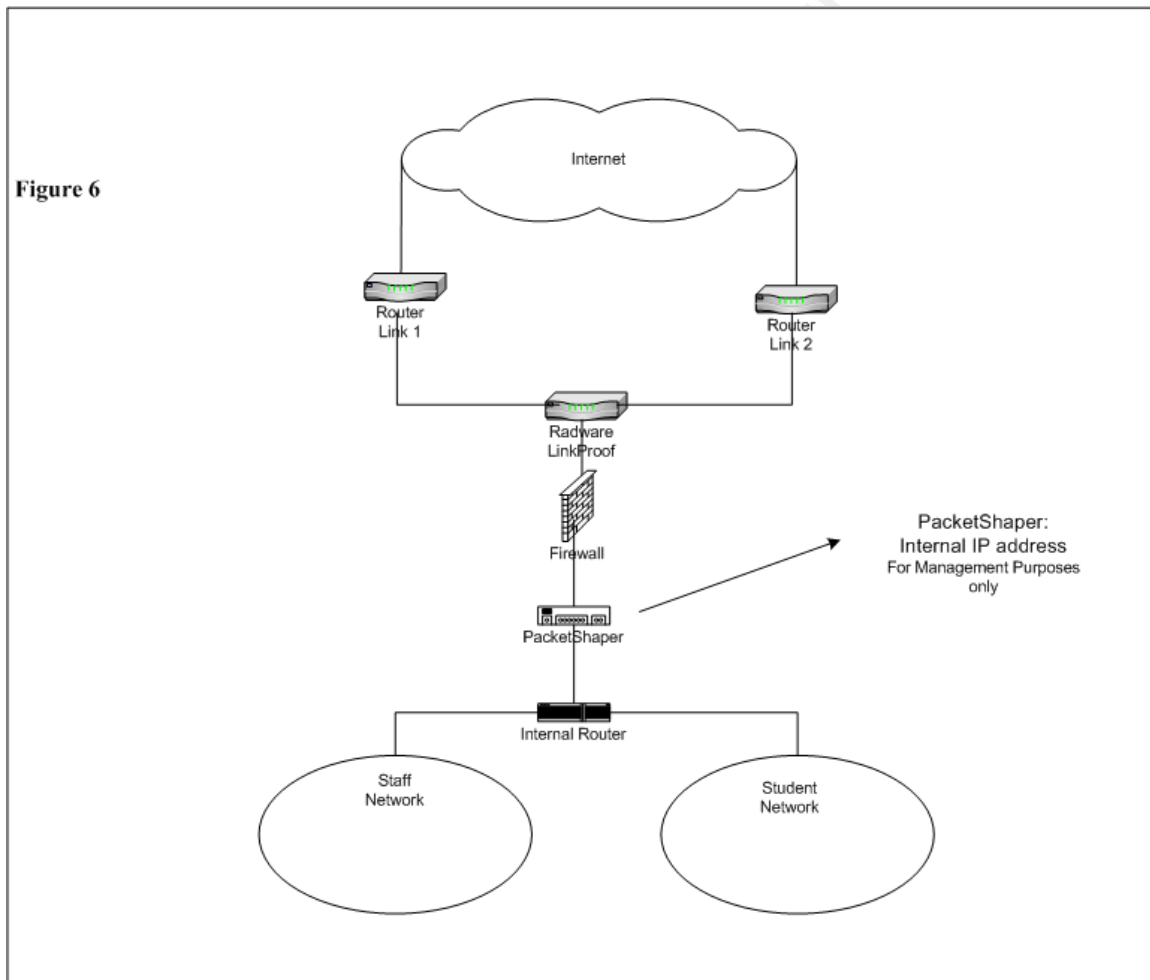
Option 3, implements PacketShaper internally or inside your border router. The positives to this solution are that you can use more flexible and less expensive Ethernet interfaces to manage traffic. Also, this option will allow for partitioning of university's internal network and the use of multiple shapers. Some of the negatives include a greater amount of bandwidth will be managed which may require a more capable and more expensive device. Another negative is that internal traffic will be interfered with and shaped.

Option 3, allows administrators for the most flexibility and manageability of bandwidth.

Now that we understand **Why** there is a need to manage bandwidth? **What** device? And **Where** to place it?, we can start discussing on **How** to use it? For this explanation we will return to the scenario previously mentioned.

Scenario: During Installation

Since our first three questioned have been answered and explored I will move on to the implementation of the PacketShaper and describe what I did to deploy this appliance. **Figure 6** shows where our IT Department decided to install the PacketShaper:




After exploring all of the different options and analyzing the pros and cons, it was decided that in our scenario it was important to be able to manage and shape internal traffic. The best place to do this was to implement the PacketShaper between the firewall and internal router. The PacketShaper has two Ethernet interfaces, one labeled “In” and the other “Out”. They basically describe the flow of traffic. The “In” interface describes traffic flow destined towards the internal

network. While, the “Out” interface describes traffic flow initiated from the internal network destined to the outside world or DMZ. In our deployment of the PacketShaper we will not be able to monitor, shape or manage traffic that does not traverse the PacketShaper. This traffic will include internal peer-to-peer traffic and traffic between internal servers and internal clients.

Configuring PacketShaper

Once I decided where to implement the PacketShaper I needed to figure out how to physically plug the cables and what cables to use. **Figure 7** shows the front end of the PacketShaper:

Figure 7



Picture from Packeteer Website at URL:

http://support.packeteer.com/documentation/packetguide/5.2.1/documents/PacketShaper_Getting_Started_v521.pdf

The RJ 45 interfaces are clearly labeled “Outside” and “Inside”. The types of cable that will be plugged into these interfaces depend on the type of device that you will be plugging into the PacketShaper. In our scenario, I used the firewall and router. Therefore, I will need cross-over cables to plug in to both interfaces. Servers and uplink ports also require cross-over cables, while hubs or switches require straight-over cables. Once, all ports and devices are plugged in correctly one will see traffic begin to flow and normal connectivity will be restored.

After physically connecting the PacketShaper and verifying that traffic is traversing the device I was able to connect to the device and log in. There are three simple ways to connect and configure the PacketShaper:

- Through a direct console connection
- Telnet
- Through a Web Browser

The first time that I connected to the PacketShaper via any of the above mentioned ways I had to use the default IP. This is a factory set IP address that has been assigned to the device. I later changed this IP address to a more meaningful IP address. For the purpose of this paper we are only going to be covering connections via Internet Explorer.

I simply started an Internet Explorer session and typed in the default URL. The first time I connected I was directed straight to the basic configuration or setup page. In this page I was able to modify the following options:

Shaping	ON/OFF
Traffic Discovery	ON/OFF
Easy Configuration	ON/OFF
IP Address	IP for Device
NetMask	Netmask for Device
Gateway	Next hop usually defines outbound flow
Site Router	Optional: Router which Device is plugged into
Domain	Optional: Domains that will be monitored
DNS Server	Name Servers that will be used to resolve host names

Wan Settings:

InBound Rate	Total bandwidth available
OutBound Rate	Total bandwidth available

Lan Settings:

Inside Fast Ethernet NIC Mode	Auto/ 100 Full/ 100 Half/ 10 Full/ 10 Half
Outside Fast Ethernet NIC Mode	Auto/ 100 Full/ 100 Half/ 10 Full/ 10 Half

These are the basic configuration settings for the PacketShaper.

- I made sure to leave the Shaping option on the OFF position, because at this point we are not going to start shaping traffic.
- The Traffic Discovery option should be set to the ON position. This will allow the PacketShaper to begin discovering traffic.
- The Easy Configuration will not be covered in this paper as it is a less flexible option with many limitations; I kept this option set to the OFF position.
- The IP Address option is a management option. Simply select an IP Address that makes sense to your scenario. This depends on the placement of the PacketShaper. In our scenario we decided to place the PacketShaper behind our firewall, so we decided to go with an internal private IP address that made sense with our IP scheme. Remember this IP address option is for management and connection purposes only.
- The Netmask option corresponds to the IP address that you decide to assign to your device set it accordingly.
- The Gateway option will typically refer to the flow of traffic destined to outer networks. In our scenario the internal firewall network interface is the Gateway. Refer to **Figure 6** for a better visual explanation. Typically the Gateway option will represent traffic destined for the outside world or internet.

- Site Router and Domain options are optional settings. Site Router represents a router that will be used to monitor traffic and Domain can be used for FQDN (Fully Qualified Domain Name) or NT domain naming schemes. The DNS server option should be set so the PacketShaper will be able to resolve names to the IP address that it finds. In our scenario I used the NT 2000 internal DNS for both domains and the external DNS servers IP address.

The next set of options are broken down into two separate categories, WAN and LAN. These are supposed to help you gauge the bandwidth that will be used and measured. In our scenario the WAN and LAN settings were used as following:

WAN Settings:

InBound Rate: 3M

OutBound Rate: 3M

LAN Settings:

Inside Fast Ethernet NIC Mode 100 Full Duplex

Outside Fast Ethernet NIC Mode 100 Full Duplex

The WAN setting is used to set a maximum available rate of bandwidth. In our scenario we are currently using dual T1 and therefore our optimal bandwidth rate inbound or outbound is approximately 3.0 Meg. This will help create the pipe that we are going to be using to control bandwidth. If you refer back to **Figure 4** we are creating the outer black pipe which will engulf all of our shaped traffic. The LAN settings are the optimal speed of your internal backbone speed and allow you to specify which kind of duplex mode is being used. If you know for sure the devices that are plugged into the PacketShaper are Full/Half or are 10/100 set it accordingly, if you are not sure you may use the Auto-negotiate option.

Once these settings were configured I selected the apply changes button and the PacketShaper Basic configurations were set.

Other Configuration settings that I would encourage to set are the SECURITY and DATE & TIME Setup Pages. The Security Setup Page will allow you to select a LOOK and TOUCH passwords. The LOOK mode is a read only mode while the TOUCH mode is write mode. Setting the DATE & TIME configuration will help you diagnose problems that are dependent on time and that only occur during specific times.

Variation of the Four Step Deployment Guide

Once I was done configuring and setting up the PacketShaper it was time to start deploying it and let it run on the network.

I decided to follow the Four Step tutorial offered by Packeteer but I also decided to add two important steps to this model. As one can recall the Four Steps were to:

1. Classify
2. Analyze
3. Control
4. Report

The following six steps were created;

1. Classify-Identify and Simplify
2. Analyze
3. Control
4. Report
5. Develop Policies
6. Recognize Unmanaged Traffic

Step One: Classify, Identify and Simplify

In order to analyze traffic I needed to let PacketShaper capture network traffic. Packeteer suggests allowing the device to analyze network traffic for 3 days, but I believed that it would be better to analyze traffic for a full week. By analyzing an entire week, you will be able to capture traffic for all seven days and a more accurate analyzes will be stored.


The first thing that I looked at was the Monitor Tab. This tab showed all discovered traffic and it breaks up the traffic into two categories. The two categories are Inbound and Outbound. Under each category PacketShaper will identify classes of traffic. These classes are well known protocols such as HTTP and known applications like Citrix. I took some time to review and learn what was traveling along our network. The first thing that I did was to place the classes into more descriptive folders. I created a folder by going to the Manage Tab. This tab is similar to the Monitor tab, with all discovered classes on the left most side of the page. On the right side of the page there are some options that I needed to explore. The first button that I looked at was the Class button. This button allows one to create a Class folder. I did the following to add some classes:

Select the Class Button → Then Select the Add Folder option → this brings up a window with an empty field, fill in a descriptive name (P2P) → Select the OK button.

The Manage tab page will now refresh itself and a new P2P Folder will appear under the InBound category. By simply selecting the P2P folder a new configuration page will display on the right side of the page. **Figure 8** shows what the configuration page will look for all classes. The Traffic Classes are shown on the left panel of the web page. On the right panel of the configuration page are the CLASS, PARTITION, and STATISTIC buttons. I will discuss the CLASS button only in this particular section, the other buttons will be discussed later in this paper.

I have already used the CLASS button to create a folder. To move an already existing class into a folder simply select the CLASS button and then select the move option. A new screen will appear. Simply select the desired class “KaZaA” and select the Move Class button. The “KaZaA” application will now be under the P2P folder. I continued to classify and organize our traffic. The more organized and simple you keep your traffic classes the easier it will be to set traffic control settings.

Figure 8



Now that I have described the basics of the Manage tab, I am going to share a simple and useful list that I created and used in organizing our Monitor Tab.

1. Identify critical traffic. For our scenario the following were selected
 - a. HTTP
 - b. SSL
 - c. SMTP
 - d. DMZ traffic
 - e. RDP
2. Identify less desirable traffic. We decided to focus on Peer to Peer Networks and Video protocols
 - a. eDonkey
 - b. Gnutella
 - c. KaZaA
 - d. Napster

- e. MPEG-Audio
 - f. MPEG-Video
 - g. QuickTime
 - h. Real
3. Setup Folders for Steps 1 and 2
 - a. WEB folder created
 - i. HTTP discover traffic within class
 - b. P2P folder created
 - c. SSL discover traffic within class
 - d. DMZ set policy to ignore partitioning
 - e. RDP discover traffic within class
 4. Identify Peak Traffic Classes that are in excess of 500K
 5. Delete everything else and simplify

I setup this simple to follow list to keep our goals and help us organize the Monitor Tab. Number One identifies the mission critical applications and protocols that use our bandwidth. Remember that HTTP will hold web servers, OWA, and other web based applications that will run on HTTP. The SSL class will have secured traffic that uses port 443. The SMTP traffic class will have mail related traffic, and hopefully only mail servers. The DMZ traffic class was identified because we did not want to monitor or have traffic destined to the DMZ or from the DMZ to be shaped. Therefore, we classified this by the DMZ subnets. Remember that a traffic class may be a protocol, application, host, or subnet. The RDP traffic class runs the terminal servers that are used throughout the organization.

Number Two identifies traffic that I knew or expected to be causing bandwidth problems. These are the popular P2P network applications. I also suspected to have many video applications running and consuming most of bandwidth.

Number Three helped us organize the traffic classes into similar groups. I created a folder called WEB. Under this Class Folder I moved HTTP and DNS protocols. Under the HTTP protocol I selected the option to allow Traffic Discovery within Class. This will allow us to discover specific traffic within this class. What I was trying to identify is our web servers and we are looking to discover what or who is using an undesirable amount of HTTP traffic. So under the WEB Folder we should see two traffic classes, HTTP and DNS and under the DNS we will identify our web servers, WebServer1, WebServer2 and so on. We will leave the DNS traffic class alone for now. The next folder that was created was the P2P. This was discussed earlier in this section. We are going to move onto identifying further servers within traffic classes. I enabled the Traffic Discovery within Class option for all SSL, SMTP and RDP traffic. We must make sure that most of traffic using these protocols must be our servers running these protocols. If you see an unknown host consuming or creating most of this traffic make sure to write down the host IP address or Name. We will set partitions or policies to eliminate this unwanted consumption.

Number Four is going help in identifying traffic that has a peak utilization of 500K or more. I picked this number because it represents a considerable amount of consumption. All traffic that has peaked at more than 500K will stay in our Monitoring Tab. The Peak column is visible in **Figure 9**. Notice that all measurements are set by (bps) or bytes per second.

Figure 9

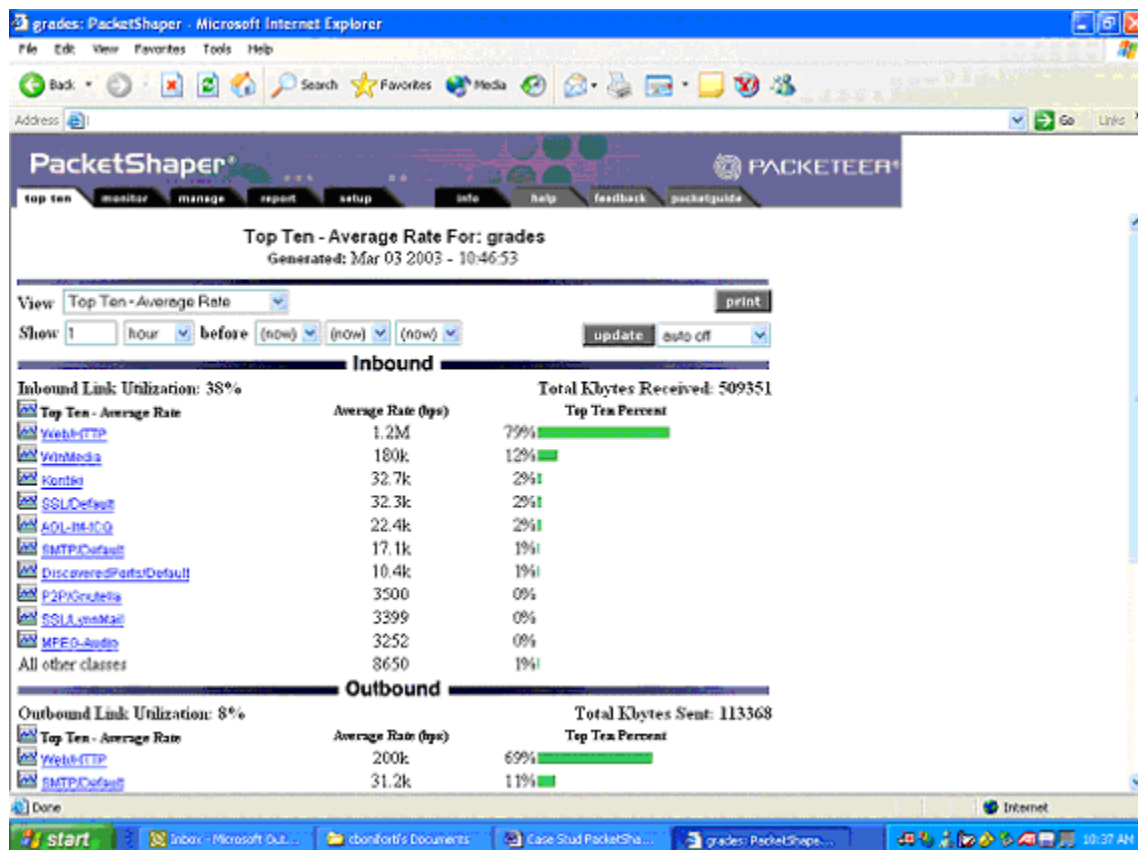
Traffic Class Name	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)	Quar. Rate Failures
Inbound			2.6M	1.6M	3.6M	0
Localhost	20346	20346	301k	73.7k	1.2M	0
DMZ	122777	122777	0	0	2055	0
P2P			4368	2537	19.8k	0
eDonkey	6	6	0	0	355	0
Gnutella	6281926	NA	4368	2537	18.9k	0
KaZaA	1483647	1483647	0	4	11.4k	0
Napster	96	NA	0	4	17.2k	0
Web			1.6M	1.4M	2.9M	0
HTTP			1.6M	1.4M	2.9M	0
218...S/Forum	2748	NA	0	3	1.4M	0
Remote	0	NA	0	0	0	0
StudentEmail	19985	NA	658	623	1.5k	0
Teachb	12083	NA	0	4	34.5k	0
Outlook	12487645	12487645	1.6M	1.3M	2.9M	0
DNS	668888	668888	1003	343	13.2k	0
FTP	7642	NA	0	0	28.8k	0
Hotline	0	NA	0	0	0	0
IRC	974	NA	0	0	41.2k	0
L2TP	1	NA	0	0	9420	0
LDAP	41	NA	0	0	691	0
lockd	0	NA	0	0	0	0
Marimba	143	NA	37	75	261k	0
MPEG-Audio	1312	NA	0	18	1.9M	0
MPEG-Video	8364	NA	0	1	2.6M	0
MSN-Messenger	41201	NA	218	227	227k	0

Number Five will allow us to simplify our objectives. Basically we are going to delete all classes that are not heavily utilized throughout our network. Any class which uses less than 500K.

Step 2: Analyze Behavior

Now that the Monitor tab is clean, I can start analyzing traffic behavior. In order to do this I need to take a look at couple of different tabs. I am going to use the Monitor, and Top Ten. I have already covered the Monitor tab and I am aware that I must look for traffic that peaks over the 500K limit. I identified such traffic and documented. I now need to look at the Top Ten tab. This tab is the first tab available and it is a very useful tab. **Figure 10** shows an example of the Top Ten Tab:

Figure 10



This tab is a quick way to view the utilization of network traffic. It will display utilization on a current status. One has the option to show traffic in measurements of Minutes, Hours, Days, Weeks or Month. This will be helpful and analyzing an entire 8 hour work day, or more importantly non work hours. Also notice that the utilization chart is split into Inbound and Outbound, and that each category has the most used class underneath each category. In order to effectively use this tool, I am going to list a helpful method that I used during the initial analyzes period.

List for Top Ten Usage:

1. Starting with Monday and continuing throughout the week run the Top Ten at beginning of the day set to the current hour. This is the default setting
 - a. Document Inbound –vs- Outbound utilization
 - b. Identify most popular protocol or application
 - c. Review average rates per Class

2. Starting with Monday and continuing throughout the week run the Top Ten during the end of the day set to the current hour.
 - a. Document Inbound –vs- Outbound utilization
 - b. Identify most popular protocol or application
 - c. Review average rates per Class
 - d. Compare difference from Number 1

3. Starting with Monday and continuing throughout the week run the Top Ten during the end of the day set the time option to 3 Hours & 8 Hours
 - a. Document Inbound –vs- Outbound utilization
 - b. Identify most popular protocol or application
 - c. Review average rates per Class
 - d. Document any Class that looks unfamiliar or different then what expected

4. On Friday run an additional Top Ten at the end of the day, set the time option to 1 week.
 - a. Document Inbound –vs- Outbound utilization
 - b. Identify most popular protocol or application
 - c. Review average rates per Class

This simple four step list was used for a one month period. A report was created from the given information and the following information was extracted from the month long report. The Top class for both Inbound and Outbound traffic was HTTP, consuming some 75% of all utilized bandwidth. WinMedia and MPEG-Video were consistently second or third but only consuming some 5-10% of utilized bandwidth. This month long report also provided us with a nice baseline to work from and start working into the next phase of our deployment.

Step 3: Control Behavior

The next step in our successful deployment is controlling behavior. I now know what I need to secure and guarantee HTTP, SSL, RDP, SMTP and DMZ. I also know what may cause unwanted consumption, P2P, WinMedia, and MPEG-Video or Audio. Since I know what I want and do not want I can now begin to control traffic. I am going to use the Manage Tab and develop some partitions and policies. Remember that there are two different types of partitions, hierarchical and dynamic, and that a partition works much like a pipe within a pipe. Refer to **Figure 4** for a visual look at a partition.

I am going to start by reviewing the critical applications or protocols and make sure that they have guaranteed partitions set for them. The critical classes that I picked were:

- HTTP
- SSL

- SMTP
- RDP
- DMZ

During the analysis period I was able to determine that HTTP traffic peaked at some 3.0Megs consuming the entire bandwidth pipe at one time. I also learned that average consumption of bandwidth for a week hovered at around 1.0Megs. I set up the following partition for HTTP traffic. I created a 1.0Meg partition by guaranteeing 1.0Meg at all times and allowing HTTP traffic to burst to 1.5Megs. I will explain about bursts a little later. **Figure 11** shows the options that one will see when creating a partition. Simply go to the Manage Tab, select the desired class in this case HTTP. Then select the Partitions button and one will now see a screen much like this:

Figure 11:



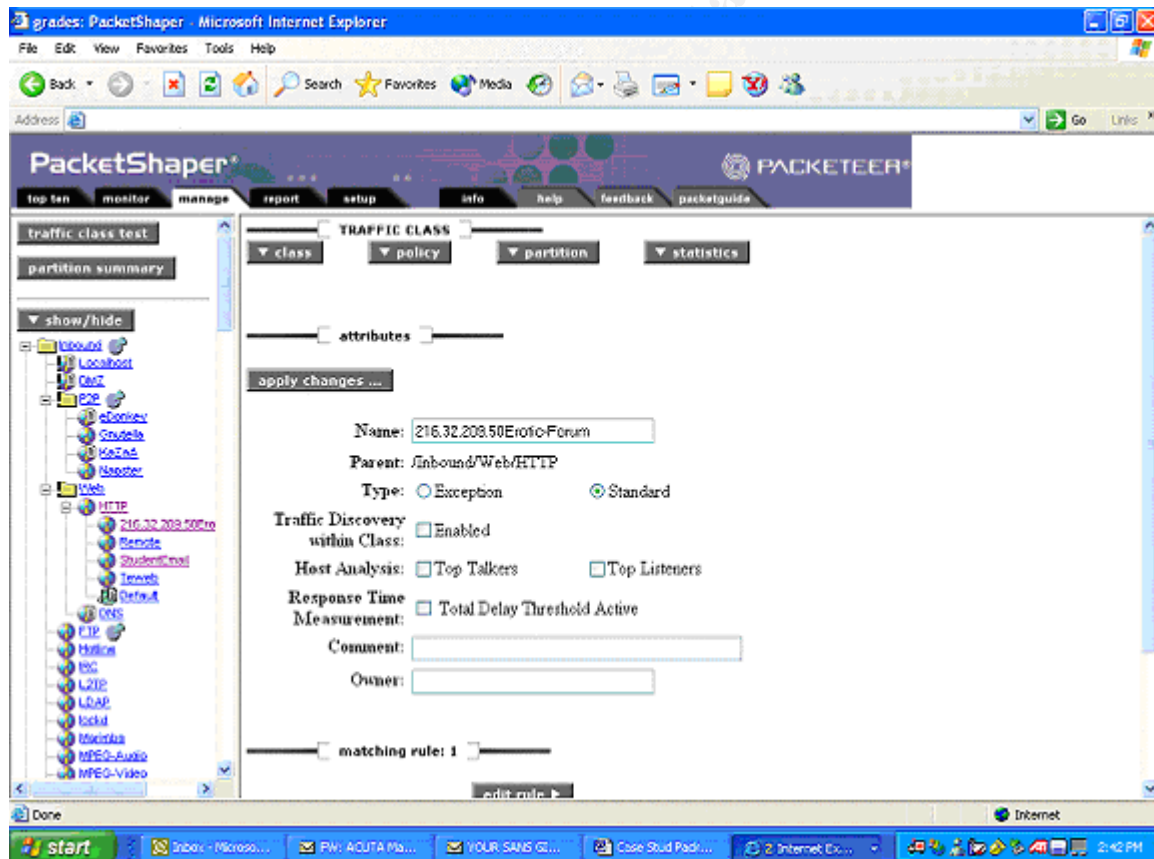
There are two fields: Size and Limit. The number that I applied on the Size field is a guaranteed number that I am going to set for the HTTP class. In my scenario I am going to set this to 1.0Meg. Please remember that all measurements are in bps so I would insert the following number, 1,000,000 bps. The next field is the Limit field, and one can also see that the Burstable option is selected by default. Leave this option alone as we do want to allow burstable usage. The Burstable option is a set amount of bandwidth that is made available for usage, only if it is available. I am setting this amount to 1.5Meg. So, what I

am doing is guaranteeing 1.0Meg to the HTTP class and allowing it to burst or consume 1.5Meg if it is not being taken by any other class. I simply selected the Add Partition button and created a partition for the HTTP class.

The next thing that I needed to do is identify the hosts or servers that are part of the HTTP class. In our scenario we have four web servers that I wanted to identify. I will call them Web 1, Web 2, Web 3, and Web 4. Step one earlier in this section of the paper explained how to classify host or servers. Once the web servers were clearly classified and placed under the proper folder, I was ready to begin setting policies for each web server.

In order to set a policy I needed to select the specific class, in this case Web 1. Once I selected this server I noticed that a new button appeared in the Manage Tab, Policy. **Figure 12** shows what this page will look like:

Figure 12



I then selected the Policy button and the add option. This will bring to another configuration page. This page will contain some additional options. This page will allow for a more specific policy type to be set. These are the options available:

1. Rate

2. Priority
3. Never Admit
4. Ignore
5. Discard

- Rating is used when you want to be able to guarantee a portion of the already partitioned bandwidth. You can specify a guarantee number and a limit to what the burst can be. One must also set a Priority to the burst from 0-7, where 0 is the low and 7 is the highest.
- A Priority is used when one does not care for guaranteeing a certain class a specific amount of bandwidth, but you do want to prioritize its right to bandwidth when comparing the class to other classes. The Priority policy button allows you to distinguish importance of sub-classes.
- The Never Admit option is used when one needs a protocol to be disallowed but still recorded on the utilization chart.
- The Ignore option is used for traffic that one would allow but does not care to shape or monitor.
- Lastly, the Discard options will drop the protocol essentially blocking it, but not recording the action.

I decided to use the Rate option for the web servers. The following table will show how I rated the servers and what bandwidth was set for them.

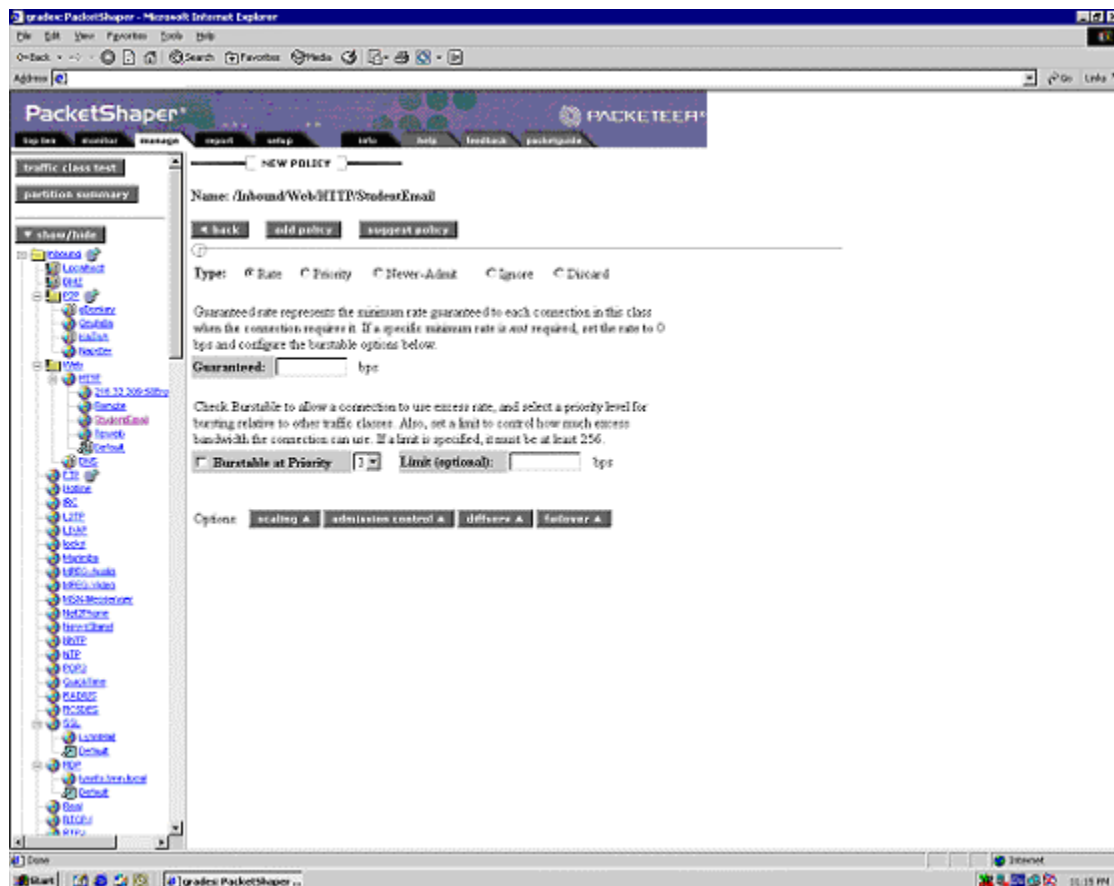
Server	Rating #	Peak	Avg.	Guaranteed	Burstable
Web 1	6	240K	140K	150K	250K
Web 2	3	50K	30K	45K	50K
Web 3	4	100K	40K	45K	100K
Web 4	2	250K	20K	40K	50K
All HTTP	NA			720K	1.5M

- Web 1 was used for Online Registration and it is a highly visible web server.
- Web 2 is a web application server and it is utilized by faculty to enter attendance and other administrative tasks.
- Web 3 is used by the student and it also runs some important web applications.
- Web 4 is used for the front of the terminal server application.

By analyzing and reviewing the importance of each server and their Peak and Avg. consumptions, I was able to come up with a productive rating system. I was also able to come up with some realistic and useful numbers for guaranteed and burstable bandwidth settings.

The next figure will show what the options are once you are in the Policy button option.

Figure 13



I used a similar strategy to setup our SSL services and servers. SMTP was used the same way except that no other burstable bandwidth was set. The reason that I did this was that I thought that no other client or server should be running SMTP traffic. SMTP traffic should only be created by MX servers. The servers were identified with the corresponding rating and proper bandwidth was applied. I set the RDP class much like the HTTP class using the same strategy and techniques.

The Priority method or strategy was used for the P2P folder. This was the Peer-to-Peer group. Since I did not care to reserve any bandwidth for this group; all I needed to do was set a low priority for the entire class. Thus, not disallowing this type of traffic, but merely making sure that it did not hog or consume all available bandwidth. The lowest possible priority was given to this group, "0", and 56K

was guaranteed to this entire class. My thinking was that I would be able to relieve the firewall from blocking these sites, and allow the PacketShaper to control and monitor the traffic. By setting the guaranteed bandwidth to 56K, without any burstable options, users would not be able to consume more than 56K of P2P traffic. By selecting a low or lowest Priority this class of traffic would have to wait until all higher priority was delivered. Allowing us the flexibility of this would make my job easier in creating policies for the proper usage of bandwidth.

The Ignore method or strategy was used for DMZ traffic. I used the Ignore option for all DMZ traffic. Basically, I was able to tell PacketShaper to ignore all traffic destined for or incoming from the DMZ and allow it normal passage. I was able to classify the DMZ via its subnet.

Step 4: Report

This step is a very important visual step. It will demonstrate visually the current and past utilization outputs of the networks bandwidth. I used the Report tab to create reports and useful graphs. I followed the same list that I provided in the Top Ten section and applied it to this section. The Report Tab, was also used to see drastic falls of bandwidth when a partition was created and implemented. This is a very simple tool to use and run. I currently have a network administrator checking and documenting the outputs of pre-configured reports. I have instructed the network administrator to run these reports once a day and twice on Friday. These reports are later reviewed and compared for network trends.

An example of how we used the Report tab was when we found a student consuming 80% of all HTTP traffic. I ran the Report Tab and noticed a drastic jump in Inbound\HTTP traffic. The graph had jumped from 1.2M to 2.3M in a matter of minutes. The 2.3M stayed constant. I immediately went over the Monitor Tab and selected the HTTP class. I allowed PacketShaper to conduct a Top Talker and Top Listener of the class and I received the following outputs. The Top Talker was a site called 216.X.X.X.ExoticForum, obviously a non-educational site. This site was consuming some 800K. On the Top Listener I noticed a user of an internal IP address consuming 800K of HTTP traffic. So, I set a policy to the website and restricted to only use 8K. After refreshing and applying the policy, bandwidth dropped back down to 1.2M. The culprit site and student were no longer registering any bandwidth consumption.

The flexibility and ability of the Reporting showed me that I could use this to help me create, administer and enforce the Issue-Specific policies that I intended to implement.

Step 5: Develop Policies

The next step that I decided that to work on was developing “Issue-Specific Policies”. Through the use of the PacketShaper I could construct and develop policies. The university had un-written policy that the IT Department implemented. More specifically we were blocking or disallowing P2P traffic. I noticed that the following things happened once I implemented the PacketShaper. The firewall was being taxed as it was constantly being asked to allow ports 1214 and other common P2P ports. These requests were being blocked and recorded on the firewall logs. This added more volume to my log files on the already over worked firewall. When we ran the PacketShaper we noticed that there was some noticeable P2P traffic. How was it possible that P2P traffic still existed and traversed our network? I did some research and found out that some P2P applications are able to hop or switch ports once the applications finds out that the default port is being blocked. I also started to think about our router and firewall problem that was occurring because of unwanted SMTP traffic being created by the student population. I thought that I could use PacketShaper to block or control the SMTP traffic that was being created by non mail servers.

I decided to document the partitions that I created and make this part of the “Issue-Specific Policy”. The following table will demonstrate what protocol or applications were targeted and what partitions were reserved:

Class	Peak	Avg	Guaranteed	Burstable	% of Bandwidth Reserved
In\HTTP	3.0M	1.0M	1.0M	1.5M	33%
In\SSL	1.1M	100K	128k	750K	4.3%
Out\RDP	1.3M	300K	512K	1.0M	17%
Out\SMTP	800K	50K	128K	512K	4.3%
All Traffic					42.4%

I was able to set numbers that I could add to our policy by using the resources and data provided by the PacketShaper. This would make my job much easier in developing policies that were being enforced and backed up by numbers. The all traffic column shows that I am leaving an unassigned partition of 42.4% of the 3.0M of bandwidth. This 42.4% could be used to create future partitions if needed.

My main objective in this step was to create an effective and usable Bandwidth Utilization Policy. Before creating and implementing the policy I knew that I had to address the following steps:

1. Identify Risk
2. Communicate Findings
3. Update policy as needed
4. Develop metrics to measure policy

The above mentioned list was gathered from the SANS GIAC training documentation, Chapter in Basic Security Policy.

I decided to follow the “**Issue-Specific Policy**” format:

Bandwidth is utilized by all network devices. Some devices consume more bandwidth than others. This relationship will cause some applications or protocols to not work properly if bandwidth is being consumed in an uncontrollable fashion. Most less-desirable applications will consume majority of bandwidth and hinder critical applications from receiving the required bandwidth.

Solution:

PacketShaper
Bandwidth Utilization Policy

Method:

Define the scope of the Policy—

The mission critical applications need to be defined and reviewed. These mission critical applications will receive guaranteed bandwidth through the use of Partitions. The mission critical applications at this time are:

- HTTP
- SSL
- RDP
- DMZ
- SMTP

Less desirable applications need to be identified and a limit of bandwidth needs to be set. The Less desirable traffic is:

- P2P
- MPEG-Audio and Video
- AOL IM
- WinMedia
- ShoutCast

A Partition with a limit of a bandwidth needs to be set to these applications. A low Priority will be set to these applications. Other traffic needs to be monitored and controlled. Any application consuming more than 500K will be evaluated and a Partition will be assigned. At no time will a single user consume more than 40% of particular class of traffic.

Layer the Defense Strategy—

Partitioning a particular class will act as first layer of defense. Setting a policy for a specific user, subnet, or application will be used to further restrict and control traffic. Thirdly, a Priority rate will be assigned to a specific user, subnet or application to further rate the assignment of allowable bandwidth. Fourthly, a policy of Never-Admit may be assigned. If all of these do not provide desired

results, the user, or application may be blocked at the firewall or action to disallow internet access to the user will be enforced.

Identify Responsibilities--

A network administrator will be responsible for reviewing and creating daily, weekly and monthly reports. The administrator will use the Reporting tab to create such reports. The administrator will conduct bi-weekly meetings and review findings with IT Department. The administrator will also alert the Security Manager any time a red flag appears or when any policy is broken. The Security Manager will implement Layer of Defense strategy to the specific incident, and will review policy and change accordingly. All policy changes will be reviewed and agreed with CIO.

Step 6: Recognize Un-Managed Traffic

Right after creating the Bandwidth Policies, I was alerted by my network administrator that there was a new protocol taking some 550K of bandwidth. The application was the ShoutCast streaming Audio. After doing some research, I realized that this was a streaming audio tool. I wondered how it could consume some 550K. After using the PacketShaper, I was able to identify that this applications was coming or being used by one source. What did not make sense to me was that it was on the outbound traffic category. I went back to the ShoutCast web site and continued to research the application. I noticed that this application could be used as a server; meaning the student was using the application as a radio station and broadcasting using our bandwidth. This prompted us to create a new policy in our User Policy, that no one user could use their computer as a server. Using the Policy implemented I was able to control and throttle the bandwidth used by this user. After, setting the policies on the PacketShaper I noticed that the ShoutCast application no longer registered on the PacketShaper, Monitor Tab.

This is a great example of Un-Managed Traffic that may cause chaos. It became apparent to me that I needed an additional step to my implementation list. This is an on going step and must be constantly reviewed.

Conclusion:

The project that I was responsible for, securing bandwidth has been an ongoing success for our department and organization. I was able to accomplish the two main objectives; control bandwidth and develop a bandwidth utilization policy. The use of the PacketShaper device allowed me to identify the universities critical applications and also identify the culprit bandwidth hogs. Creating the steps and sticking to these steps helped me implement an effective solution. I was able to clearly implement a strategy that we could constantly use to monitor and control bandwidth. The PacketShaper gave me the tools to analyze and collect history data. This data was used to create a controlled performance environment. This process helped me in creating a clear, effective, and measurable policy. I was able to resolve the over consumption of bandwidth.

After implementing the PacketShaper, I was able to reduce overall utilization from 90% on both T1 lines to a constant and controllable 65%. Before the PacketShaper, I would see Peaks or burst into the 100% utilization rate during the Peak hours of 10:00pm and 2:00am. With the PacketShaper I was able to limit these peaks and burst to a comfortable 80%. By these measurements I can confidently say that I was able to meet my expectations and accomplish the administrations goals. I also saved the university from buying an additional T1. The problem with the firewall and router were also solved. I was able to stop SMTP traffic from propagating by applying partitions and controlling SMTP traffic. This allowed me to erase the ACL commands from the router and also reduced the number of SMTP requests to the firewall.

I had to ask myself are there are any vulnerabilities left? The answer to this is yes. I was able to accomplish my goals and desires, but I also understood the PacketShaper's limitations. I was not going to be able to stop internal traffic that did not traverse the appliance. So, traffic from client to client and client to server would not be managed or seen. With my new knowledge gained by the SANS GIAC class I was able to present possible enhancement that would help solve this problem.

I suggested that we should segment our internal network. The student network and faculty network were divided by a router, but no VLAN's were created. I suggest that we should segment and setup more specific VLANs. One of my suggestions is to create two separate VLANs within each network. This would allow us to place the servers and client in separate VLANs and set the gateway to the PacketShaper therefore forcing traffic to traverse the device.

I also suggested in changing our network topology. I suggested that we should add an additional firewall that would create an internal barrier between the two NT networks. This firewall would be implemented between the Faculty network and the internal router. Therefore less restrictive rules could be setup for the Student network and more restrictive rules could be set on the internal faculty firewall.

I feel that I was able to accomplish what was asked for and in addition I was able to supply the university with enhanced security. I was also able to suggest some important needed changes that are not specific to bandwidth but were brought up because of this project. I also feel that without my newly gained knowledge of the SANS GIAC class I would not be able to implement such effective policies, nor add new enhanced security strategies.

Citations:

“LinkProof: Always Online”. 2002

URL: <http://www.radware.com/content/products/lp/default.asp>

”SearchNetworking.com Definitions - powered by whatis.com”, Search Networking.com Website. 2002

URL:

http://searchnetworking.techtarget.com/sDefinition/0,,sid7_gci211634,00.html

“Four Steps to Application Performance Across the Network” Packeteer Website. September 2002

URL:

<http://support.packeteer.com/documentation/packetguide/5.2.1/documents/4Steps.pdf>

Udelson, Ted. “System Security Policy: What it is and Why every Campus Needs One.” Journal of Telecommunications in Higher Education. Fall 2002: Page 6-10

“Welcome to PacketGuide”. Packeteer Website. 2001

URL: <http://support.packeteer.com/documentation/packetguide/5.2/index.cfm>

© SANS Institute 2003, All rights reserved. Author retains full rights.

Upcoming Training

Click Here to
{Get CERTIFIED!}



SANS Prague 2017	Prague, Czech Republic	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Boston 2017	Boston, MA	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS New York City 2017	New York City, NY	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Salt Lake City 2017	Salt Lake City, UT	Aug 14, 2017 - Aug 19, 2017	Live Event
Community SANS Omaha SEC401*	Omaha, NE	Aug 14, 2017 - Aug 19, 2017	Community SANS
Virginia Beach 2017 - SEC401: Security Essentials Bootcamp Style	Virginia Beach, VA	Aug 21, 2017 - Aug 26, 2017	vLive
SANS Chicago 2017	Chicago, IL	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Virginia Beach 2017	Virginia Beach, VA	Aug 21, 2017 - Sep 01, 2017	Live Event
SANS Adelaide 2017	Adelaide, Australia	Aug 21, 2017 - Aug 26, 2017	Live Event
Community SANS Pasadena SEC401 @ NASA	Pasadena, CA	Aug 23, 2017 - Aug 30, 2017	Community SANS
Mentor Session - SEC401	Minneapolis, MN	Aug 29, 2017 - Oct 10, 2017	Mentor
SANS Tampa - Clearwater 2017	Clearwater, FL	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS San Francisco Fall 2017	San Francisco, CA	Sep 05, 2017 - Sep 10, 2017	Live Event
Mentor Session - SEC401	Edmonton, AB	Sep 06, 2017 - Oct 18, 2017	Mentor
SANS Network Security 2017	Las Vegas, NV	Sep 10, 2017 - Sep 17, 2017	Live Event
Mentor Session - SEC401	Ventura, CA	Sep 11, 2017 - Oct 12, 2017	Mentor
Community SANS Albany SEC401	Albany, NY	Sep 11, 2017 - Sep 16, 2017	Community SANS
Community SANS Columbia SEC401	Columbia, MD	Sep 18, 2017 - Sep 23, 2017	Community SANS
Community SANS Dallas SEC401	Dallas, TX	Sep 18, 2017 - Sep 23, 2017	Community SANS
SANS Baltimore Fall 2017	Baltimore, MD	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Copenhagen 2017	Copenhagen, Denmark	Sep 25, 2017 - Sep 30, 2017	Live Event
Community SANS Boise SEC401	Boise, ID	Sep 25, 2017 - Sep 30, 2017	Community SANS
Baltimore Fall 2017 - SEC401: Security Essentials Bootcamp Style	Baltimore, MD	Sep 25, 2017 - Sep 30, 2017	vLive
Community SANS New York SEC401	New York, NY	Sep 25, 2017 - Sep 30, 2017	Community SANS
Rocky Mountain Fall 2017	Denver, CO	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS London September 2017	London, United Kingdom	Sep 25, 2017 - Sep 30, 2017	Live Event
Community SANS Sacramento SEC401	Sacramento, CA	Oct 02, 2017 - Oct 07, 2017	Community SANS
SANS DFIR Prague 2017	Prague, Czech Republic	Oct 02, 2017 - Oct 08, 2017	Live Event
Community SANS Charleston SEC401	Charleston, SC	Oct 02, 2017 - Oct 07, 2017	Community SANS
Mentor Session - SEC401	Arlington, VA	Oct 04, 2017 - Nov 15, 2017	Mentor
SANS Phoenix-Mesa 2017	Mesa, AZ	Oct 09, 2017 - Oct 14, 2017	Live Event