Technical Product Overview

Four Steps to Application Performance Across the Network

With Packeteer's PacketShaper®

November 2001

Packeteer, Inc.
10495 N. De Anza Blvd.
Cupertino, CA 95014
408.873.4400
info@packeteer.com
www.packeteer.com

PACKETEER

Table of Contents

Four Steps to Application Performance 3
What Is PacketShaper? 3
Ease of Deployment and Use 4
Easy Configure Option 4
A Four-Step Snapshot 5
Step One: Classifying Network Traffic 6
Classification, a Step Above 6
Applications and Protocols — Automatically 8
Step Two: Analyzing Traffic 9
Utilization Analysis 9
Performance Analysis 13
Raw Metrics 17
Step Three: Controlling Traffic 19
Partitioning Bandwidth 19
Partition Description 19
Variations on the Partition Theme 20
Per-Session Rate Policies 21
Other Policies 22
Rate-Control Features 22
Universal Translator 23
Detecting and Avoiding Attacks 24
Step Four: Generating Reports 25
Following a Typical Investigation 27
Recommended Environment 28
For More Information 28

Four Steps to Application Performance

Sluggish mission-critical applications are bad for business. Unfortunately, non-critical or less urgent applications tend to dominate when applications battle for bandwidth on congested WAN access links. Large email attachments or high-capacity file transfers consume more than their share of bandwidth, while Oracle, SAP, PeopleSoft, Citrix, and other critical applications struggle to get any.

Do any of these problems sound familiar to you?

- A high-speed user downloads a large file and SAP performance lags.
- An employee synchronizes his laptop with the message server and clogs the branch office's WAN link for 15 minutes.
- Music enthusiasts' MP3 downloads cause more urgent, interactive applications to struggle.
- A stock market plunge prompts employees to check their online portfolios en masse, bringing all business-critical applications to their knees.
- Repeated bandwidth upgrades fail to address performance problems but do increase costs substantially.

Today's enterprises require performance, predictability, and consistency from their networks and the applications that traverse them. And that's precisely what PacketShaper® from Packeteer delivers.

This paper describes a process to avoid problems like those described in the examples above and serves as a technical product overview for the PacketShaper product line. The paper is intended for system and network administrators in organizations that manage their own applications and network performance.

What Is 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 keeps critical traffic moving at an appropriate pace through bandwidth bottlenecks. Less urgent traffic still moves steadily but uses a smaller slice of available bandwidth. With PacketShaper, no single type of traffic monopolizes the link.

PacketShaper is an ideal solution for any site on the network where a speed disparity causes performance bottlenecks. In recent years, the advent of fast Ethernet and gigabit Ethernet has reduced network congestion on the LAN. Simultaneously, the deployment of fiber infrastructure in the WAN backbone has reduced contention in that portion of the network. However, the bridge between the two, the WAN access link, has remained the slow, weak link in the chain. WAN access-link capacity is still constrained, expensive, and difficult to upgrade.

In enterprise networks that are overwhelmed by increasing amounts of traffic, unmanaged congestion at the WAN-access link undermines application performance, resulting in impaired employee and company productivity. Network managers spend increasing portions of the budget on bandwidth upgrades in attempts to solve the performance problems, only to find that the problems persist.

PacketShaper's four-step approach to safeguarding application performance and maximizing return on network capacity provides insight into and control over congested WAN access links. PacketShaper discovers and classifies applications, analyzes their performance, enforces policy-based bandwidth allocation, and generates reports on the results. With PacketShaper, you can control application performance to suit business priorities and get the highest value from your existing network infrastructure.

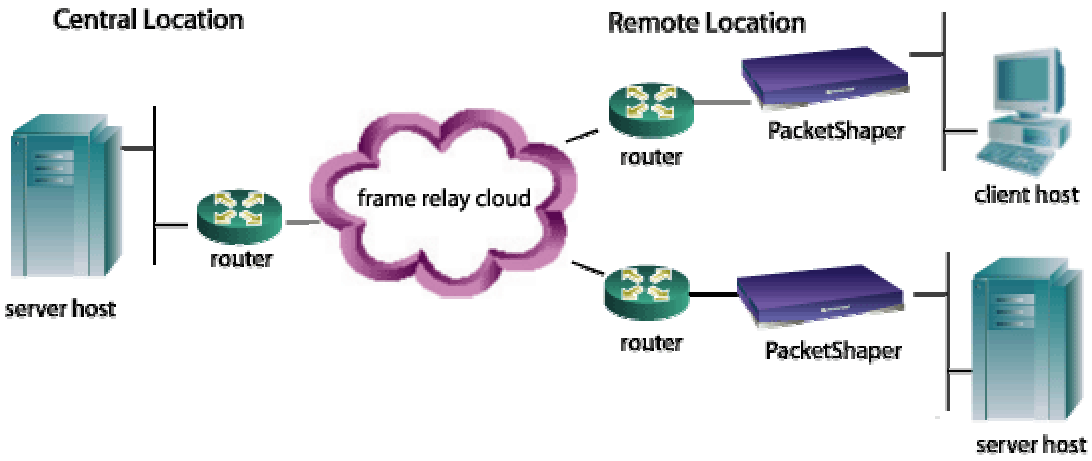

PacketShaper sits behind WAN-link routers and alleviates the performance degradation that results from bandwidth bottlenecks. Branch offices and remote sites benefit most from PacketShaper; it makes efficient use of a limited-capacity link, increases throughput, and protects critical applications.

Ease of Deployment and Use

PacketShaper installation is easy and consists of plugging in two cables and entering address and access information on a web-based setup page. PacketShaper integrates cleanly with existing network infrastructure, imposing no changes on router configuration, topologies, desktops, or servers. It also integrates smoothly with central, third-party management platforms and reporting tools such as HP OpenView®, HP Policy Xpert™, Micromuse NETCOOL™, InfoVista™, Microsoft Excel®, and others. In addition, Packeteer's PolicyCenter™ provides the convenience of centralized management in large deployments.

A web-based user interface offers access to PacketShaper from any desktop with a web browser. A command-line interface offers fast, detailed control from a Telnet session. You choose the level of security required to examine and alter PacketShaper's configuration and measurement data.

Easy Configure Option

PacketShaper has an Easy Configure option that helps new users spot trouble and fix it – quickly and without a big learning curve. First, PacketShaper automatically calculates the top 10 traffic types – those generating the most traffic. You assign an urgency category, such as

MissionCritical, to each of your top 10 applications, telling PacketShaper how to manage your traffic. Instead of creating bandwidth-allocation rules yourself, PacketShaper does it for you.

With Easy Configure, you can easily and *automatically* identify applications on the network, spot the primary contributors to congestion, define and enforce bandwidth-allocation rules, and fix the performance problems.

A Four-Step Snapshot

Each step in PacketShaper's four-step bandwidth-management strategy is an integral part of managing application performance. The steps are introduced below and explained in more detail in the remainder of the paper.

One: Classify Network Traffic

It's hard to protect an application if you can't differentiate it from other types of traffic. PacketShaper detects and identifies thousands of types of traffic. You can isolate traffic associated with applications, protocols, subnets, web pages, users, and more.

Two: Analyze Behavior

How is limited bandwidth consumed? Why do critical applications move so slowly? PacketShaper provides detailed analysis of network and application behavior. It tracks traffic levels, detects network trends, measures response time, and calculates network efficiency.

Three: Control Performance

Protecting the performance of critical applications is a matter of precisely allocating bandwidth according to business requirements and the needs of the applications themselves. PacketShaper enables the control of all types of traffic: steady rates for voice or video streams; immediate passage for small, delay-sensitive traffic such as Telnet; and a balance of consistent access and a bandwidth limit for applications such as Microsoft Exchange that are both bandwidth-hungry and critically important.

Four: Report Results

Comprehensive reports, graphs, and tables provide easy insight into historical performance, load, efficiency, and application-based service-level compliance. If performance is not suitable, control policies can be modified to bring response times within range and keep new traffic and applications from unintended impact. Reports offer proof that applications continue to perform as desired, even during network growth.

More examples and details about PacketShaper's classification, analysis, and reporting techniques are in the Packeteer paper titled "Gaining Visibility into Application Performance."

Step One: Classifying Network Traffic

Identifying and categorizing the types of network traffic that compete for limited bandwidth is the first step toward solving performance problems. Packeteer calls this process c*lassification*. Rich traffic classification is crucial—bandwidth controls are useful only if you can apply them to the precise traffic you have in mind. In addition, administrators are usually surprised to see the diversity of their network applications.

You can classify traffic by application, protocol, port number, URL or wildcard, host name, LDAP host lists, Diffserv setting, MPLS labels, IP precedence bits, IP or MAC address, subnet, travel direction (inbound/outbound), source, destination, host speed range, Mime type, web browser, Oracle database, Citrix published application, Citrix ICA priority tagging, VLAN varieties, and more. PacketShaper builds a hierarchical traffic-classification tree, inserting an entry for every distinct traffic type it observes. Each traffic category is called a *traffic class*.

While most products can differentiate traffic based on layers two through four of the standard OSI networking model, PacketShaper classifies traffic based on layers two through seven, telling you precisely which applications are in use.

More specific traffic identification yields better results. Relying on TCP or other low-level protocols, or port numbers such as the popular HTTP port 80, to classify traffic precludes the discovery of real traffic trends. Moving up on the OSI networking model and differentiating one application's traffic from another's enables application behavior to be distinguished as well. For example, both Telnet and FTP use TCP as their transport protocol, but FTP is much more likely than Telnet to cause congestion. Distinguishing between the two makes it possible to treat the two types of traffic differently.

Classification, a Step Above

Today, the number of applications is rising, and many of them are extremely bandwidth-intensive. Limited bandwidth and shrinking or flat-line budgets are placing a further strain on networks. The growing complexities associated with network traffic make sophisticated classification techniques a necessity. The simple IP-address or static-port schemes that routers use fall short. PacketShaper detects dynamic port assignments, tracks transactions with migrating port assignments, differentiates among different applications using the same port, and uses layer-seven indicators to identify applications.

Web Classification

Sometimes it seems that two-thirds of network traffic is web traffic — from web browsing to web-based clients for mission-critical applications to XML-based e-commerce. But not all web traffic requires or deserves the same treatment. Web traffic can vary in urgency, sensitivity to latency, and performance requirements. PacketShaper can differentiate between different types of web traffic so that you can manage each with an appropriate strategy.

For example, PacketShaper can use travel direction, server location, and/or URLs to distinguish employee browsing from your own (potentially profitable) customers' shopping. It can separate critical HTTP XML from recreational HTTP MPEGs, preventing someone's enthusiasm for

Madonna or Metallica from interfering with e-business. Some very thin web-based clients, such as Oracle's WebForms™ and web-based S390/AS400 host access, can look like normal HTTP traffic, but PacketShaper can classify them separately.

Citrix and Oracle Classification

PacketShaper can automatically isolate each published application running within a Citrix environment. For example, it can distinguish Microsoft Word from PeopleSoft, and it can distinguish interactive PeopleSoft traffic from its print traffic.

PacketShaper differentiates many types of Oracle traffic, allowing you to tailor your analysis or management strategies. For example, PacketShaper can separate Oracle applications accessing the "`sales`" database from those accessing the "`accounting`" database. PacketShaper identifies Oracle8i/9i and netv2 protocols using a multi-threaded server as well as the older Oracle 7 (or before) and netv1 protocols using a dedicated server.

Intricate Port Classification

When multiple applications use the same port, it is usually difficult to apply different management or analysis strategies because they appear as one application. Not so for PacketShaper. For example, PacketShaper can distinguish TN3270 and TN5250 sessions from other Telnet sessions even though all use well-known port 23.

In addition, an application that hops from port to port can be a challenge identify, because it looks like different applications. But, as before, not so for PacketShaper. For example, AOL instant messaging and passive FTP can both hop ports frequently, but PacketShaper tracks them both throughout their journeys.

File-Sharing Protocols

Although Napster traffic comes and goes with political and legal battles, there are now a whole genre of peer-2-peer applications that facilitate file sharing. These favorites of music lovers have taken a heavy toll on network performance. Schools, businesses, and other organizations have no desire to issue oppressive and controversial mandates regarding unsanctioned use of the network. However, they do want to maintain control over their networks and uphold acceptable performance for mission-critical applications.

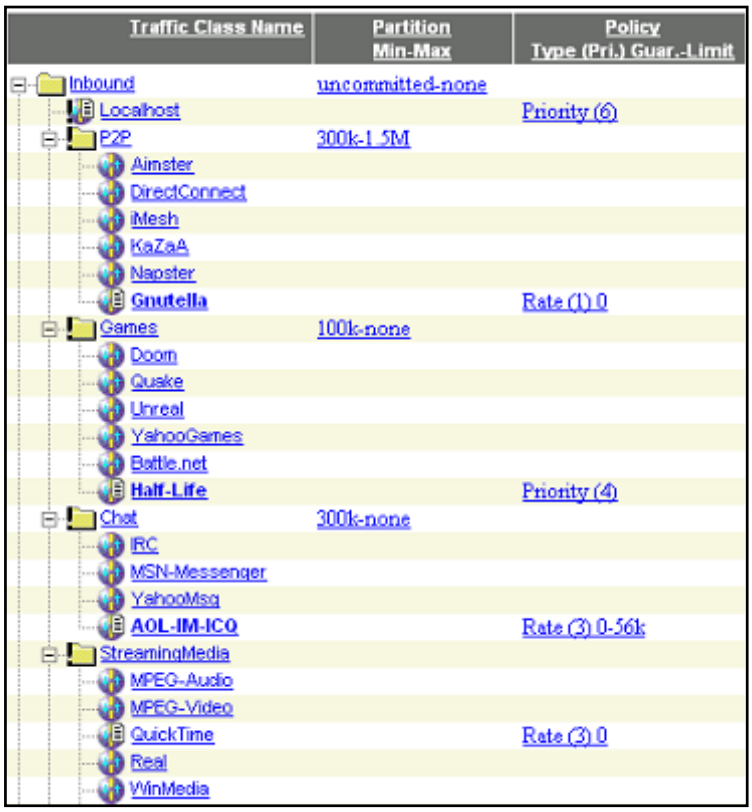

PacketShaper automatically identifies a large number of these file-sharing applications, enabling you to control their behavior (with control features, coming later). These include Aimster,

AudioGalaxy, DirectConnect, eDonkey2000, Gnutella (BearShare, Furi,Gnotella, Gnucleus, Gnut, Gnewtellium, Hagelslag, LimeWire, Mactella, Newtella, Phex, ToadNode), Groove, Hotline, Imesh, KaZaA, Napster (amster, BeNapster, BitchX, TekNap, crapster, gnap, gnapster, gnome-napster, hackster, iNapster, jnap, Knapster, Lopster, MacStar, MyNapster, nap, NapAmp, TkNap, Riscster, Shuban, snap, webnap, XmNap, AudioGnome, Rapigator, potlight, StaticNap, Swaptor, WinMX, macster, Rapster, PMNapster QNX, phaster), Scour, Tripnosis, and, believe it or not, even more popular vehicles.

Classification in Heterogeneous Networks

PacketShaper can also identify and track hosts by their DNS name, even if DHCP (Dynamic Host Configuration Protocol) changes the host's IP address frequently. You can use LDAP (Lightweight Directory Access Protocol) host lists to isolate any traffic associated with any host in the list.

PacketShaper fits well within VLAN environments, both both ISL and 802.1p/q VLANS, as it classifies and perpetuates traffic division based on virtual LAN. PacketShaper can classify (and mark, but that comes later) traffic based on IP COS/TOS bits, Diffserv settings, and MPLS labels, allowing traffic types to have uniform end-to-end treatment by multi-vendor devices in heterogeneous WANs.

Applications and Protocols — Automatically

Packeteer continually adds to the list of protocols and applications PacketShaper classifies automatically. A recent list includes the following services:

Client/Server
CORBA
Folding@Home
FIX (Finance)
Java Rmt Mthd
MATIP (Airline)
MeetingMaker
NetIQ AppMngr
OpenConnect JCP
SunRPC (dyn port)

ERP
Baan
JavaClient
JD Edwards
Oracle (7,8,9i)
SAP

Internet
ActiveX
FTP, Passive FTP
Gopher
IP, IPIP, UDP, TCP
IPv6
IRC
Mime type
NNTP
SSHTCP
SSL
TFTP
UUCP
URL
Web browser

Database
FileMaker Pro
MS SQL
Oracle 7/8i
Progress

Directory Services
CRS
DHCP
DNS
DPA
Finger
Ident
Kerberos
LDAP
RADIUS
TACACS
WINS
whois

E-mail, Collaboration
Biff
cc:MAIL
IMAP
LotusNotes MSSQ
Microsoft DCOM (MS Exchange)
Novell GroupWise
POP3
Kerberos
SMTP

File Server
AFS
CVSup
Lockd
NetBIOS-IP
NFS
Novell NetWare5

Games
Asheron's Call
Battle.net
Diablo II
Doom
EverQuest
Kali
Half-Life
MSN Zone
Quake I, II, & III
Tribes I,II
Unreal
Yahoo! Games

Host Access
ATSTCP
Attachmate SHARESUDP
Persoft Persona
SMTBF
TN3270
TN5250

Legacy LAN and Non-IP
AFP
AppleTalk
DECnet
IPX
FNA
LAT
NetBEUI
MOP-DL/RC
PPPoE
SNA

Messaging
AOL Instant Msging
ICQ Chat
MSN Messenger
Yahoo! Messenger
Internet Relay Chat

Misc
Time Server
Date-Time

Multi-Media
Multi-cast NetShow
NetMeeting
QuickTime
RTP
Real Audio
Streamworks
RTSP
MPEG
ST2
SHOUTcast
WebEx
WindowsMedia

Music P2P
Aimster
AudioGalaxy
DirectConntect
eDonkey2000
Gnutella
Groove
Hotline
Imesh
KaZaA
Napster
Scour
Tripnosis

Network Management
Cisco Discovery
ICMP by packet type
Microsoft SMS
NTP
RSVP
SNMP
SYSLOG

Print
LPR
IPP
TN5250p
TN3287

Push
Backweb
EntryPoint
Marimba
PointCast

Routing
AURP
BGP
CBT
DRP
EGP
EIGRP
IGMP
IGP
MPLS (+tag, +app)
OSPF
PIM
RARP
RIP
Spanning Tree
VLAN (802.1q/p)

Security Protocol
DLS
DPA
GRE
IPSEC
ISAKMP/IKE key exchange
L2TP
PPTP
SOCKS Proxy

Session
REXEC
rlogin
rsh
Telnet
Timbuktu
VNC
Xwindows

Thin Client or Server Based
Citrix Published Apps and VideoFrame
RDP/Terminal Server

Voice over IP
Clarent
CUSeeMe
Dialpad
H.323
I-Phone
MCK Commun.
Micom VIP
RTP
RTCP
T.120
VDOPhone

Step Two: Analyzing Traffic

While discovery reveals *what* is on the network, analysis shows *how* it behaves, helping strategies for performance management, server balancing, topology planning, and capacity planning. Several types of analysis are summarized in the next sections.

Utilization Analysis

Utilization analysis determines how bandwidth is used and answers questions such as:

- How is link capacity divided?
- Which traffic, which hosts, and which web destinations are the most popular?
- What is my DCOM (Microsoft Exchange) traffic's current bandwidth rate? Its peak last month?
- How much of my Oracle throughput is wasted on retransmissions?
- What portion of my link services non-critical traffic? Unsanctioned traffic?
- How much bandwidth did a given user or group of users consume?

PacketShaper tracks average and peak traffic levels, identifies top users and applications, evaluates network efficiency, and more. PacketShaper gives you an automatic breakdown of usage statistics for each traffic class and, if you want it, even for each user.

For a comprehensive list of PacketShaper graphs, consult PacketGuide's Reference section and look for Graphs.

Monitoring Table

PacketShaper keeps several groups of usage statistics. The simplest group is a collection of peak and current rates for every class in your traffic tree and provides a handy glimpse into traffic behavior for all classes on one convenient screen.

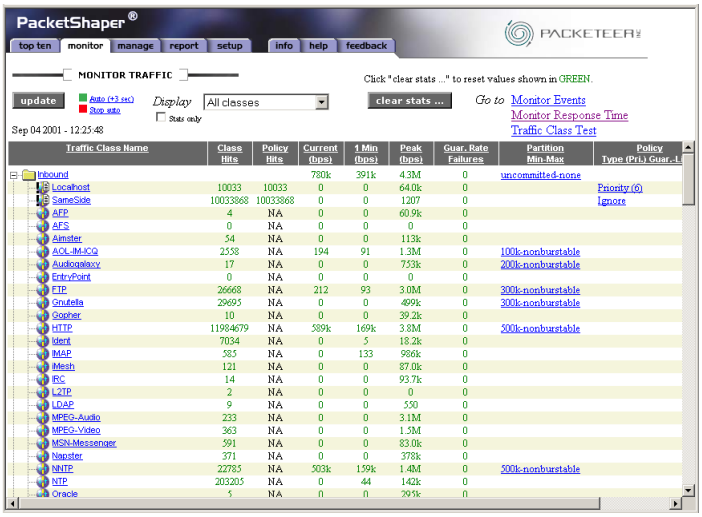

The Tops

PacketShaper offers a variety of mechanisms to track top consumers of bandwidth. There are lists, tables, pie charts, line graphs, and bar charts. And you can evaluate the top applications, users, groups of users, web sites, and other possibilities. Remember the flexibility in criteria for a traffic class.

A pie chart gives an intuitive picture of your traffic classes that generate the most network traffic.

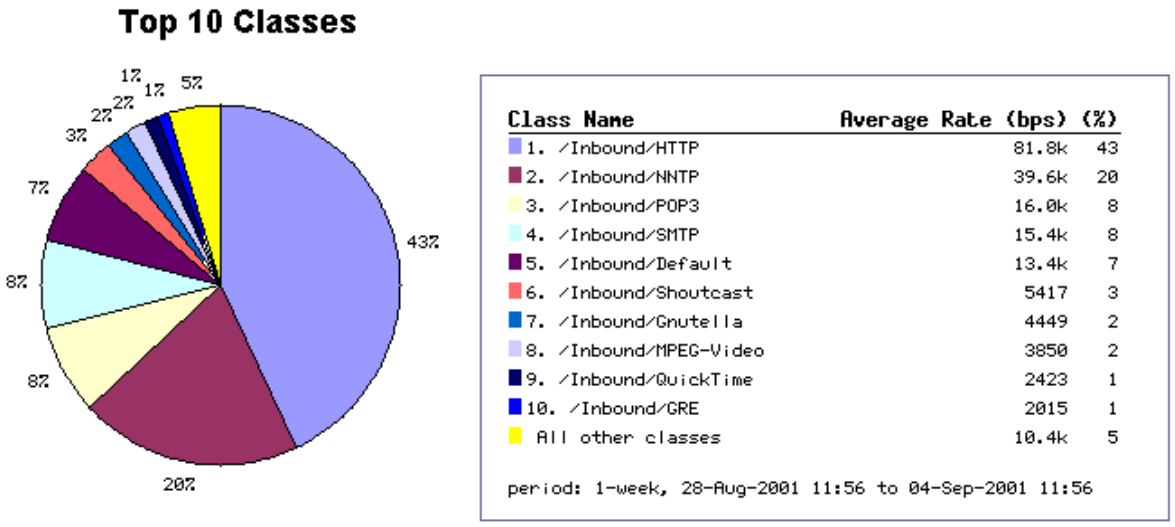

The Top Ten list ranks the top traffic classes with the percentage of bandwidth consumed for each.

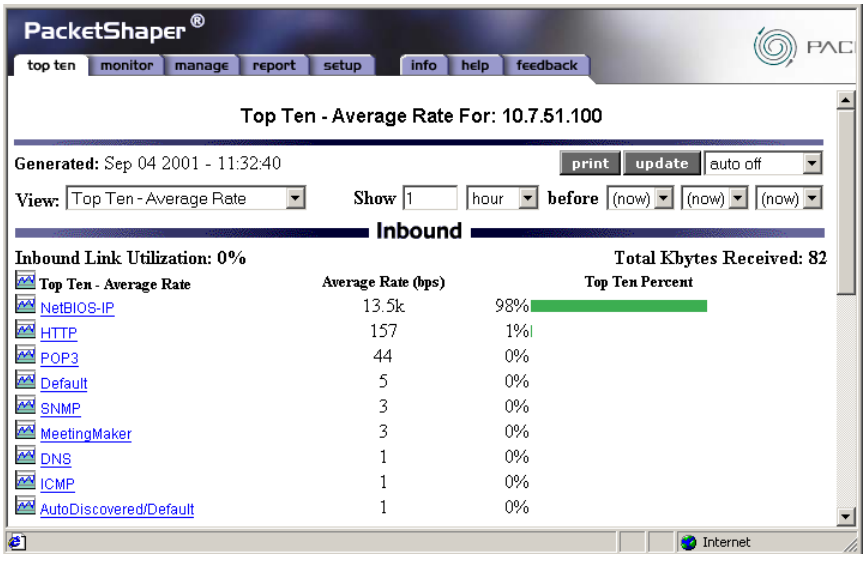

PacketShaper's Top Talkers/Top Listeners feature lets you delve into a traffic class to explore its heaviest contributors. Top Talkers displays the top 10 generators of a class' traffic, and Top Listeners displays the top 10 recipients of a class' traffic. For example, if Top Talkers is turned on for HTTP traffic, you'll find out which webpages are the most popular. And if Top Listeners is turned on for RealAudio, you'll find out who's using the network to listen to music.

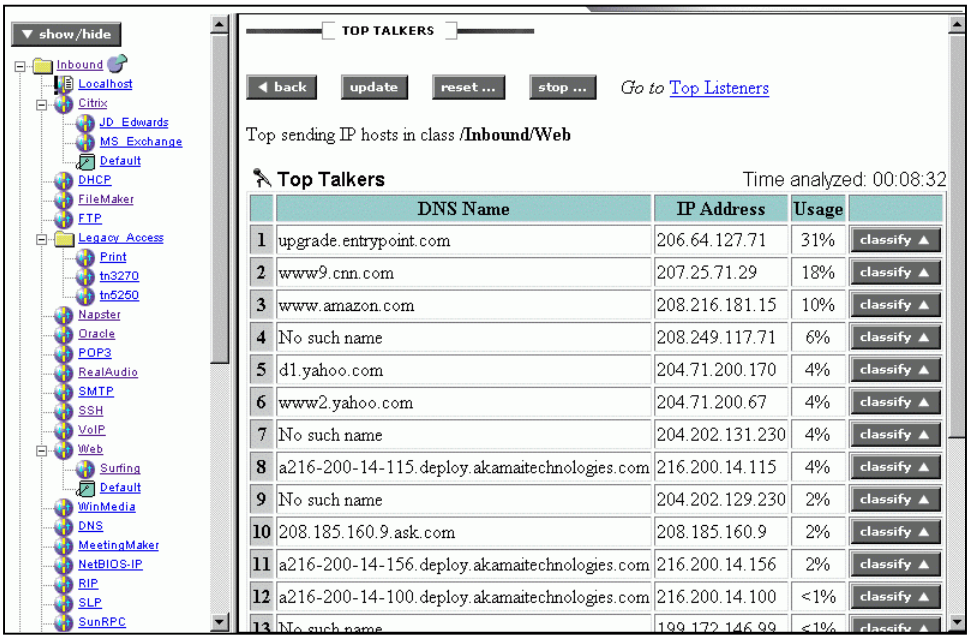

Another PacketShaper feature, called *host accounting,* tracks historical usage levels for each IP address and offers statistics summed for each user, host list, or subnet. If you import this data into a tool such as MS Excel, you can sort the results, yielding top users, or even a fully ranked list of users from top to bottom. Using these features, you could, for example, retrieve usage figures for each department in a company — this much by Marketing, that much by Accounting, and so on.

Utilization over Time

Customers frequently have an artificial sense of security if they examine graphs of average network usage over sizeable chunks of time. When PacketShaper adds *peak* usage and more frequent intervals, graphs can highlight a hidden contention problem. An average-rate line might mislead someone into thinking that usage never approaches link capacity. A peak-rate line tells the real story — frequent spikes that use the entire link.

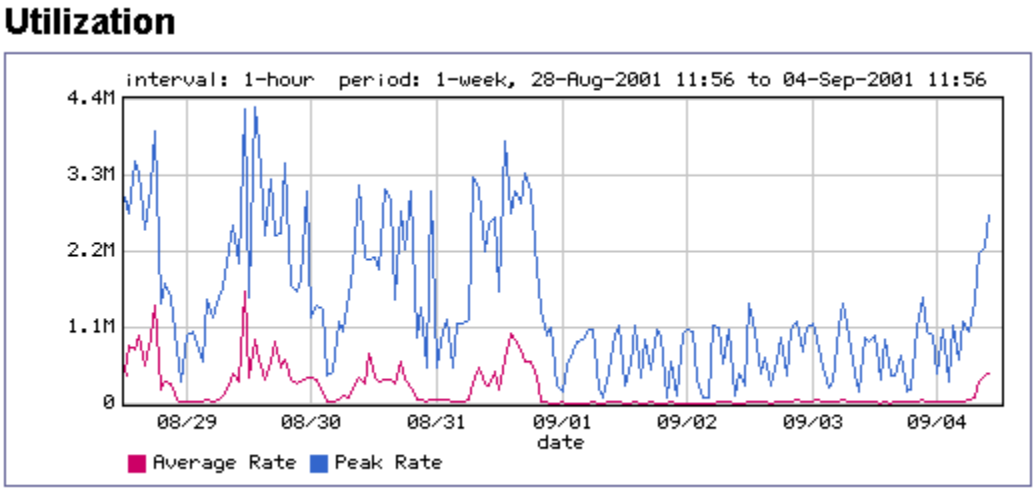

Peak Rate

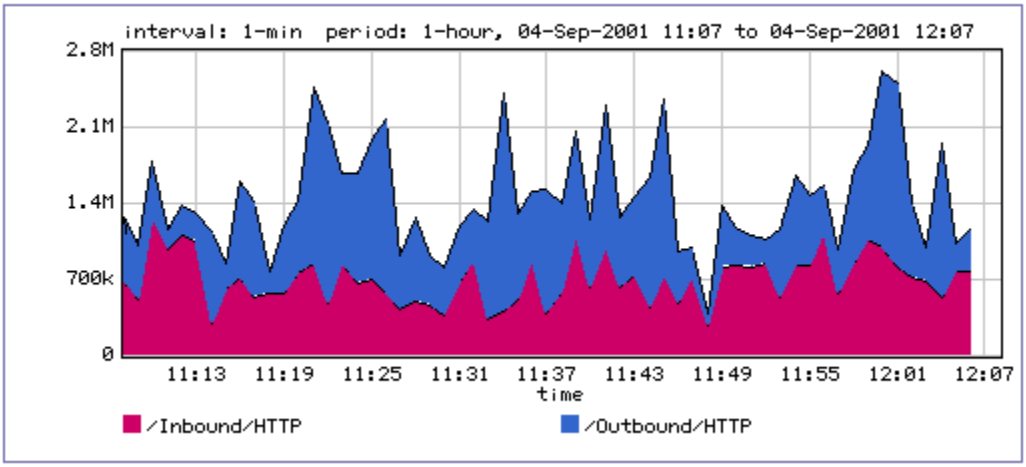

Diagnostic Aids

Some of PacketShaper ISP's graphs aid the diagnostic process. Among them:

TCP Health gives you a comprehensive picture of TCP connections for a link, partition, or traffic class. It compares the number of TCP connections that were started, aborted, and ignored or refused by the server.

TCP Health

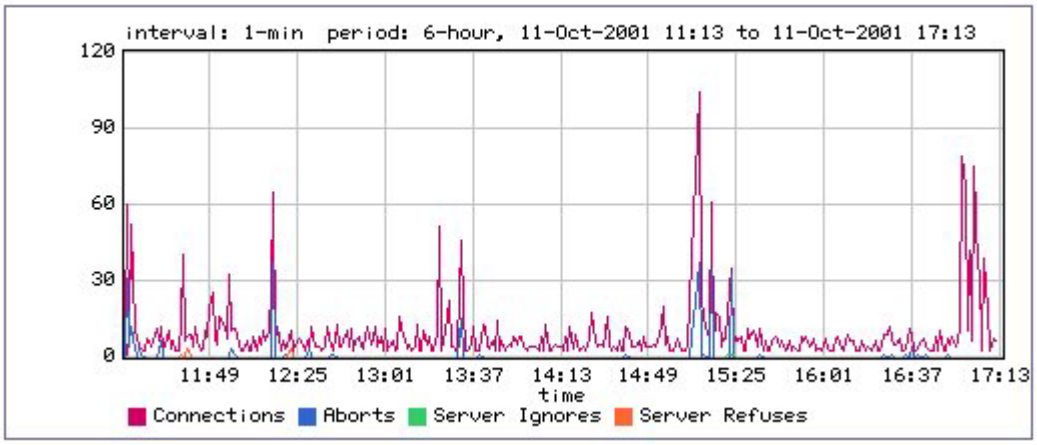

The Network Efficiency graph was designed to expose the hidden cost of retransmissions. The graph shows the percentage of bandwidth wasted by retransmissions. Other technologies might give a count of retransmitted packets, but PacketShaper knows the packets' sizes, a requirement for calculating bandwidth percentages. You can track the current retransmission rate or explore its history. You can focus on the traffic that is of interest: your link as a whole, an application, a protocol, a subnet, a user, a server, or a web destination. For example, you might note that although your network efficiency is 95 percent for your network as a whole, it plunges to 55 percent for your company's web page.

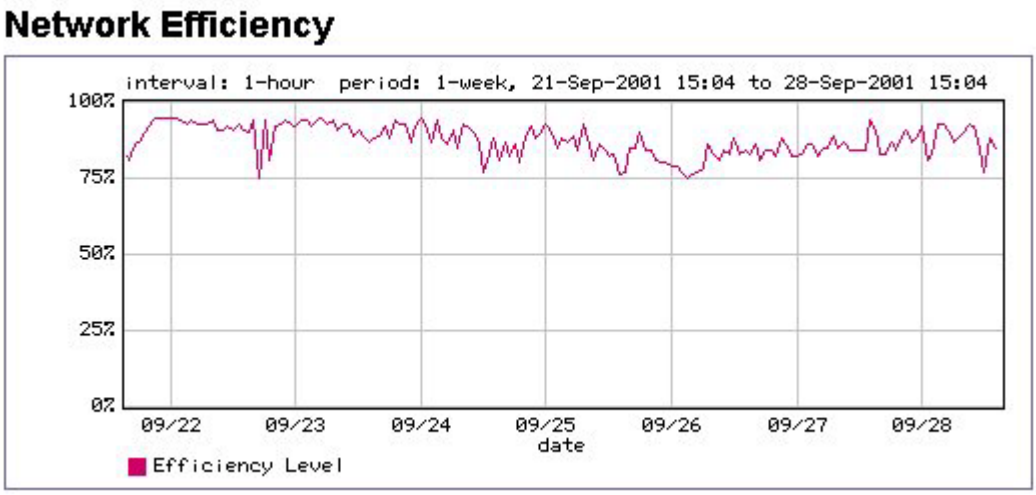

The Connection Retransmissions graph can help focus your investigation into retransmissions, comparing the retransmitted and tossed rates.

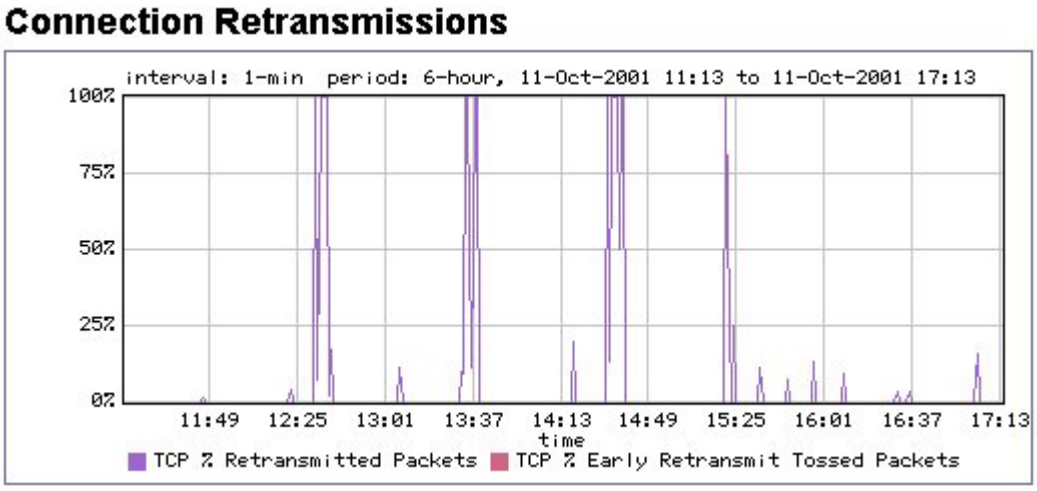

Performance Analysis

PacketShaper's response-time management facility (RTM) provides performance statistics, threshold monitoring, high-level problem indicators, and performance graphs. It quantifies what has traditionally been subjective, anecdotal information. Response-time measurement has many advantages. Response times allow you to recognize performance problems before they impact business. Concrete figures enable performance comparisons to assess the results of configuration changes. With a mechanism to compare actual and promised performance, service-level agreements (SLAs) become meaningful.

RTM offers the following features:

- Tracks delay statistics for flexible, user-defined traffic categories. For example, you can measure response times for applications such as Oracle, applications running over Citrix MetaFrame, individual hosts, subnets, and the web pages of your choice.
- Breaks each response-time measurement into network delay (time spent in transit) and server delay (time the server used to process the request).

- Identifies the users and servers with the slowest performance, called Worst Clients and Worst Servers.
- Allows users to set acceptability standards and track whether performance adheres to the standards. You can set the speed that divides good response from bad (900 ms, for example). And you can set the percentage of transactions that should be within your performance goal (95 percent, for example).
- Offers current and historical performance data in intuitive tables and graphs, in a MIB (management information base), or as raw data. SNMP management tools, such as HP OpenView, and third-party reporting tools, such as InfoVista, integrate smoothly.

PacketShaper's position in the corporate network—monitoring and controlling all the traffic that passes—gives it a unique opportunity to provide accurate response-time measurements at a very low cost. Because it already classifies every packet, PacketShaper can easily calculate the time traffic spends travelling between a client and a server and the time used by the server itself. Rather than collecting response data, PacketShaper notes response times as traffic passes. PacketShaper generates accurate performance measurements at a low operational cost and with low network overhead.

A few of PacketShaper's screens for performance analysis are in the next couple of pages.

The Monitor Response Time window displays all performance-related statistics for each measured class on one screen.

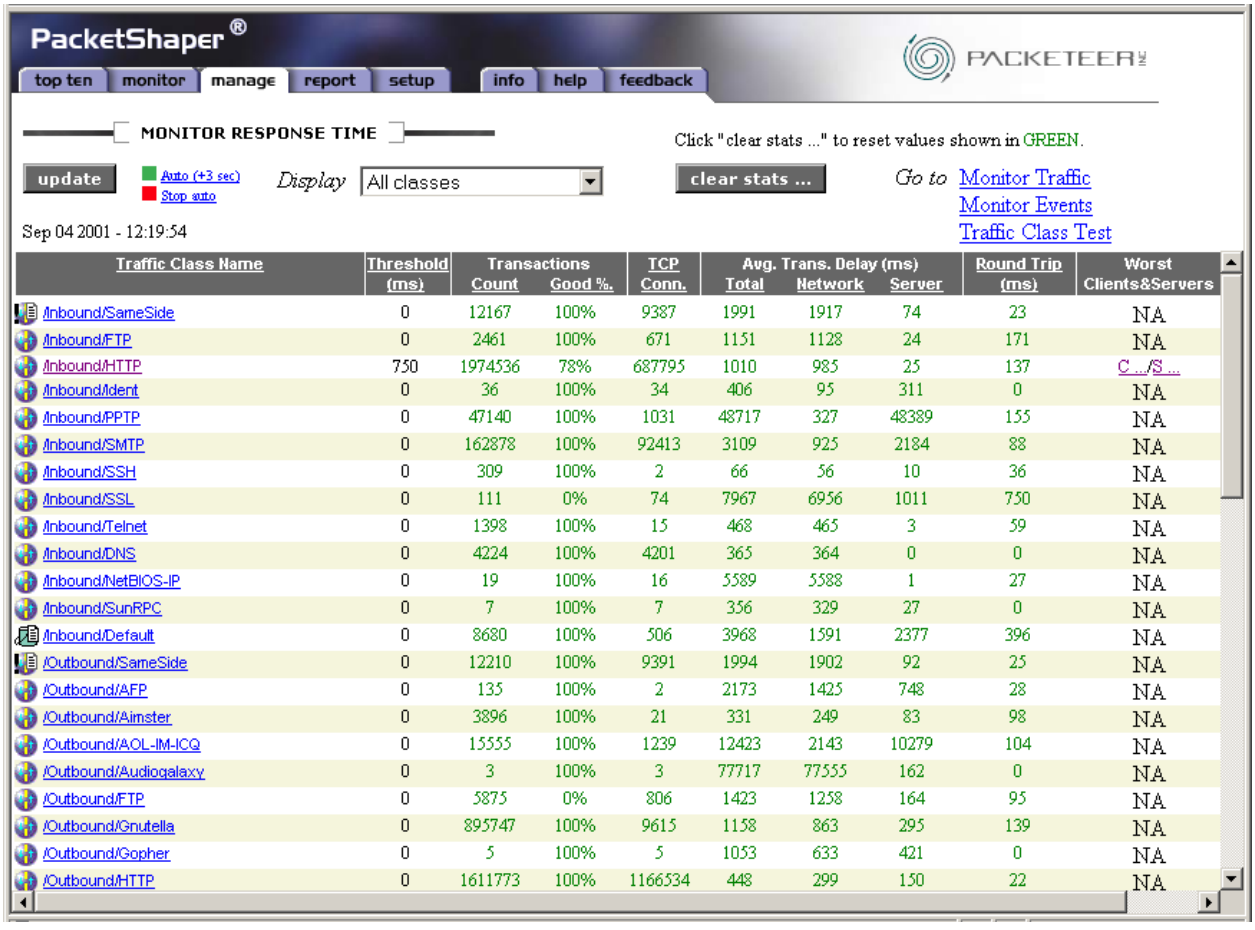

Traffic Class Name	Threshold (ms)	Transactions Count	Transactions Good %	TCP Conn.	Avg. Trans. Delay (ms) Total	Avg. Trans. Delay (ms) Network	Avg. Trans. Delay (ms) Server	Round Trip (ms)	Worst Clients&Servers
/Inbound/SameSide	0	12167	100%	9387	1991	1917	74	23	NA
/Inbound/FTP	0	2461	100%	671	1151	1128	24	171	NA
/Inbound/HTTP	750	1974536	78%	687795	1010	985	25	137	C .../S ...
/Inbound/Ident	0	36	100%	34	406	95	311	0	NA
/Inbound/PPTP	0	47140	100%	1031	48717	327	48389	155	NA
/Inbound/SMTP	0	162878	100%	92413	3109	925	2184	88	NA
/Inbound/SSH	0	309	100%	2	66	56	10	36	NA
/Inbound/SSL	0	111	0%	74	7967	6956	1011	750	NA
/Inbound/Telnet	0	1398	100%	15	468	465	3	59	NA
/Inbound/DNS	0	4224	100%	4201	365	364	0	0	NA
/Inbound/NetBIOS-IP	0	19	100%	16	5589	5588	1	27	NA
/Inbound/SunRPC	0	7	100%	7	356	329	27	0	NA
/Inbound/Default	0	8680	100%	506	3968	1591	2377	396	NA
/Outbound/SameSide	0	12210	100%	9391	1994	1902	92	25	NA
/Outbound/AFP	0	135	100%	2	2173	1425	748	28	NA
/Outbound/Aimster	0	3896	100%	21	331	249	83	98	NA
/Outbound/AOL-IM-ICQ	0	15555	100%	1239	12423	2143	10279	104	NA
/Outbound/Audiogalaxy	0	3	100%	3	77717	77555	162	0	NA
/Outbound/FTP	0	5875	0%	806	1423	1258	164	95	NA
/Outbound/Gnutella	0	895747	100%	9615	1158	863	295	139	NA
/Outbound/Gopher	0	5	100%	5	1053	633	421	0	NA
/Outbound/HTTP	0	1611773	100%	1166534	448	299	150	22	NA

Another window makes all performance-related metrics and graphs for one individual class available. This is also the window where you set performance acceptability standards.

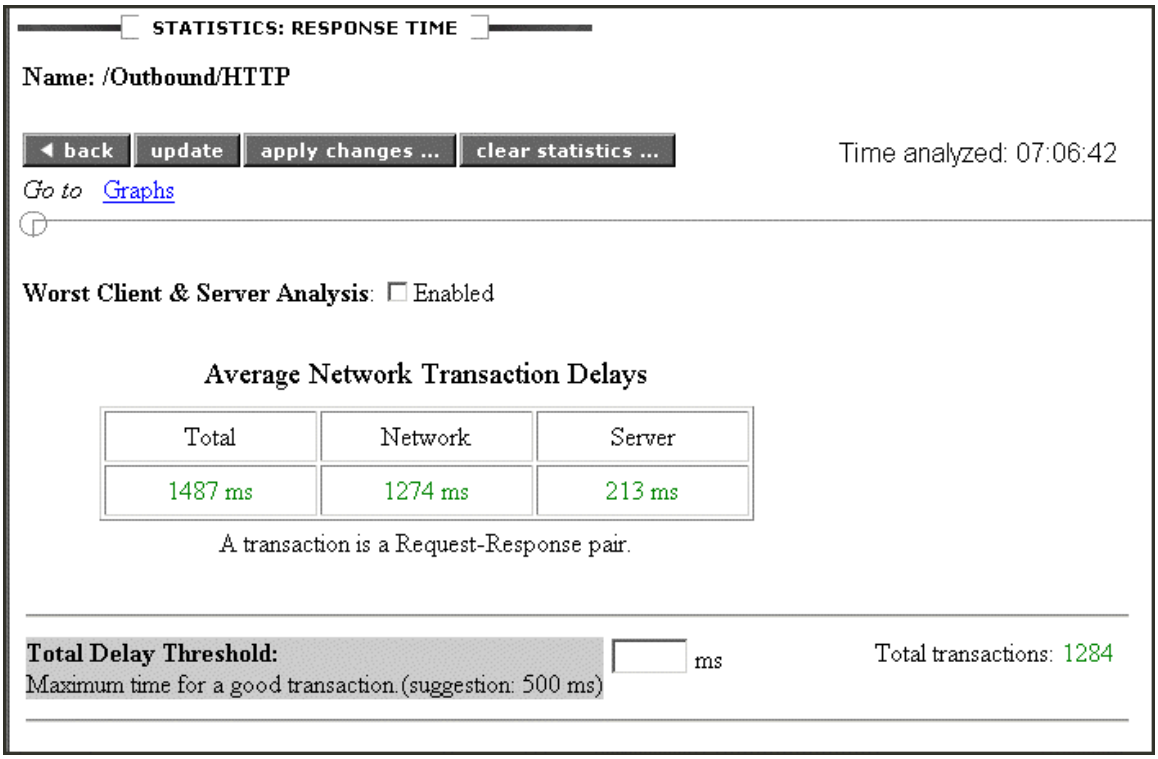

The Transaction Delay graph overlays a history of the network, server, and total delays so that you can review response time history and correlate slow response to a slow network or slow servers.

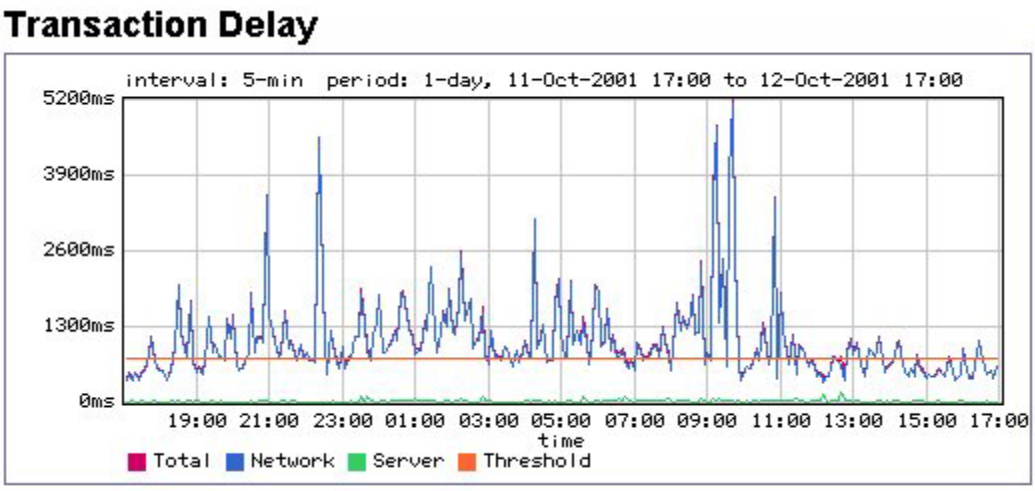

A variety of graphs depict network, server, and total delays over time for one or more traffic classes. You can view actual delays (elapsed time for a transaction), packet round-trip times (elapsed time for just one packet to traverse the network), or normalized packet times (elapsed time for a transaction of a standardized size).

Server Transaction Delay

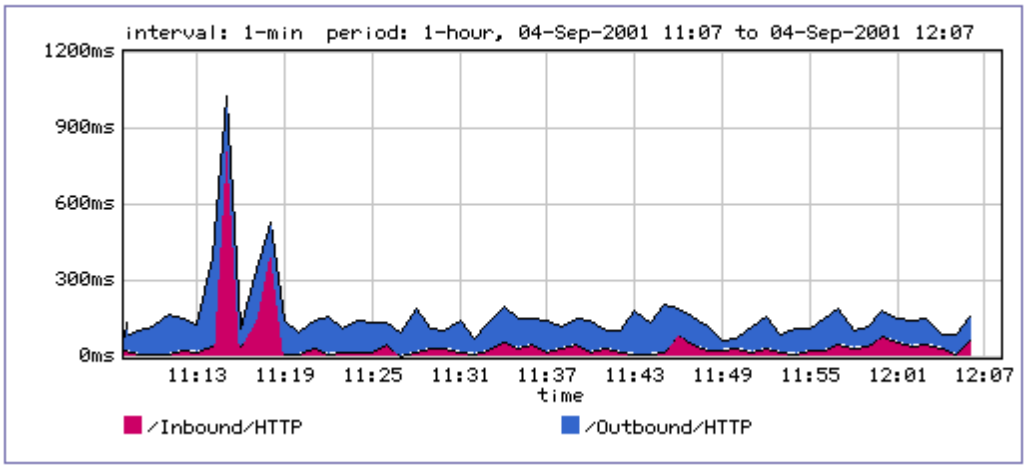

Packet Round-Trip Time

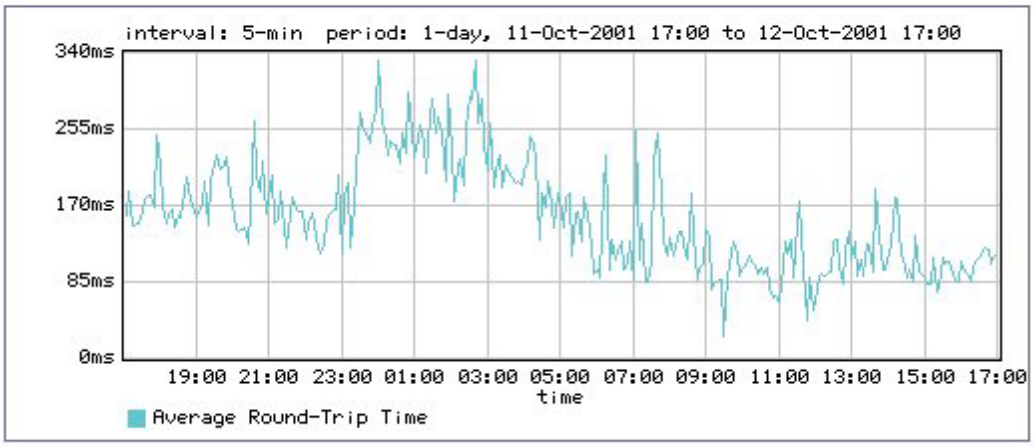

Normalized Network Delay

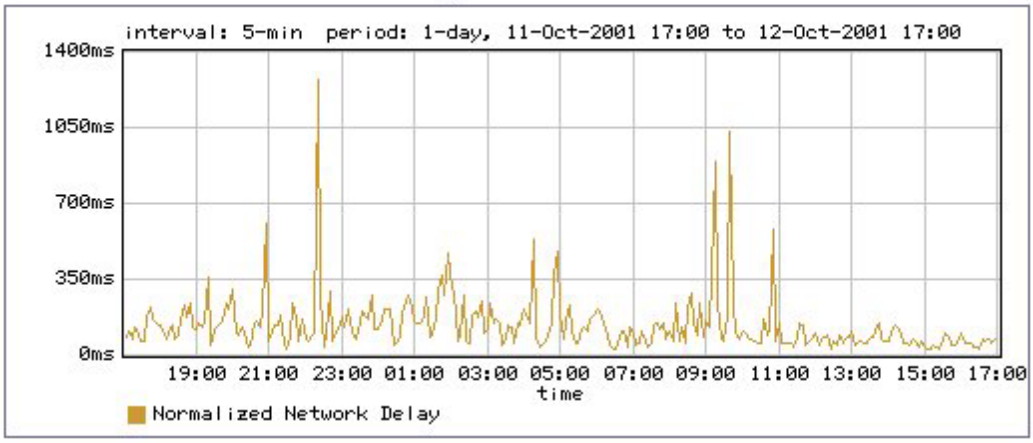

And finally, the Service Level Compliance graph tells you how well you've been doing in meeting the performance standards you set for an application.

Service Level Compliance

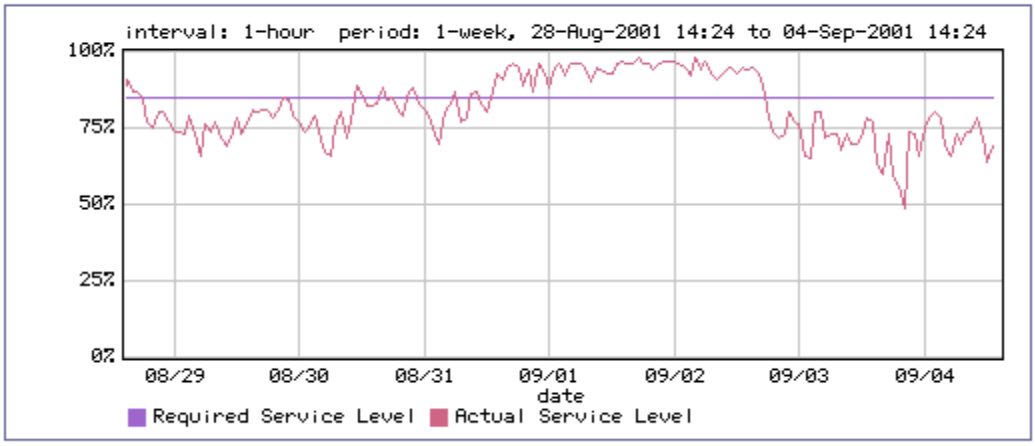

Raw Metrics

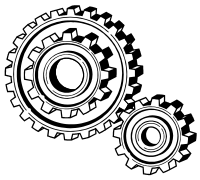

PacketShaper's rich set of metrics can be viewed with PacketShaper's own tables and graphs, or they can be retrieved for use by other tools. Metrics can be extracted using a variety of APIs (application programming interfaces) or protocols and incorporated into databases and other reporting tools.

Administrators can extract data using HTML, XML, and CGI APIs, or the PacketShaper user interface. In addition, SNMP requests, SNMP traps, and POP3 email traps all work for gathering either synchronous or asynchronous data. Extracted data can be saved in a variety of formats including SML, CSV, TSV, and ASCII.

PacketShaper's HTML and XML APIs provide database connectivity. Most modern database packages, such as Oracle and ODBC, provide data extraction agents based on these standards.

Third-party reporting packages, standard web reporting tools, and the reporting functions within HP OpenView all generate reports based on PacketShaper data.

Some of PacketShaper's measurements that might be of interest to you include:

- Throughput in units of bytes, packets, transactions, connections
- Byte throughput for any traffic class: counts, averages, and peaks
- Throughput counts for any IP address, host list, subnet
- Counts and percentages of TCP connections that were denied by a policy, denied because of resource contention, ignored by servers, aborted by users, refused by servers
- Counts and percentages of retransmitted, received, tossed, dropped, and good TCP packets
- Number of HTTP response messages with 2xx success codes, 3xx redirection codes, 4xx client error codes, and 5xx server error codes.
- Largest number of simultaneous TCP connections
- Connection-speed and packet-size histogram data for profiling users

- Histograms, medians, and averages for components of transaction response time: network delay, server delay, total delay, round-trip time, and normalized network delay
- Counts and percentages of transactions that satisfied (or did not satisfy) performance requirements
- Time intervals within service-level compliance
- Time intervals that a service was unavailable
- Top applications, URLs, users; worst performing clients and servers
- Number of users per dynamic partition, using dynamic partitions, and denied access to dynamic partitions
- Counts of traffic flows that were blocked after exceeding a configurable flow limit (suspected DoS attack involvement)
- Numbers of software licenses allowed and in use

Step Three: Controlling Traffic

Abundant data, elastic protocols that swell to use any available bandwidth, speed-conversion bottlenecks, and new, popular, and bandwidth-hungry applications seem to conspire against application performance. Mission-critical applications suffer while less important traffic dominates a link. Unfortunately, most network-management tools stop after identifying performance problems. But that's not enough. PacketShaper's third step — control — gives the power to solve those problems.

PacketShaper guarantees critical interactive traffic the bandwidth it requires while containing less urgent traffic. Each traffic class maps to a specific bandwidth-allocation policy ensuring that each type of traffic receives an appropriate slice.

Partitioning Bandwidth

Suppose a telecommunications company posts a new software release to be downloaded by interested parties. Although these file transfers are critically important, they are not time sensitive. When too many users grab the file, interactive applications grind to a halt — including urgent applications such as those needed for sales processing and manufacturing. If performance is unmanaged, or even monitored but not controlled, the company's network and application performance does not support the company's business goals.

What's needed? After all, the file transfers are necessary, just not at the expense of urgent applications. The company needs to be able to divide its link into multiple, unequal portions. Downloads should have some bandwidth but should be capped. And urgent applications need to be protected. PacketShaper's *partitions* do just that.

Partition Description

A *partition* creates a virtual separate pipe for a traffic class. You specify the size of the reserved link, designate whether it can expand, and optionally cap its growth. Partitions function similarly to frame-relay PVCs, but with the added important benefit of sharing their unused excess bandwidth with other traffic.

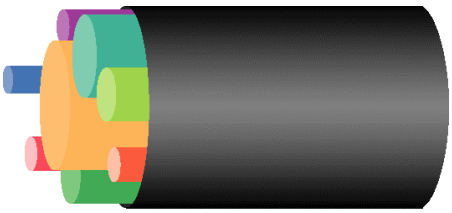

With PacketShaper, the administrator at the telecommunications firm can create an FTP traffic class and assign it to its own partition. When someone initiates a file transfer, bandwidth is available, no matter what other traffic is present. But the transfers are contained to a maximum amount of bandwidth. It's unlikely that transfers taking three minutes instead of two would impact users.

When there's less FTP demand, other traffic can lay claim to the FTP partition's unused bandwidth. Similarly, another partition can reserve an appropriate amount of bandwidth for sales and/or manufacturing applications. If unused, it goes to others. But if needed, no amount of FTP, web surfing, or large emails can infringe on its bandwidth reservation.

A partition is also an appropriate solution for managing Microsoft Exchange traffic. You might define a partition to ensure that during periods of network congestion, Exchange always receives

at least 25 percent of the available bandwidth and never uses more than 50 percent. The partition safeguards Exchange traffic while at the same time keeps it from consuming the whole link.

When does a partition make sense? It's for situations where you don't need to be concerned about each individual flow or session of traffic, but just with an aggregate total. For example, the partition for Exchange did nothing to prevent one Exchange user from impacting another.

Variations on the Partition Theme

Two variations on the partition theme are of particular use: hierarchical partitions and dynamic partitions.

Hierarchical partitions are embedded in larger, parent partitions. They enable you to carve a large bandwidth allotment into managed subsets. For example, you could reserve 40 percent of your link capacity for applications running over Citrix, and then reserve portions of that 40 percent for each application running over Citrix – perhaps half for PeopleSoft and a quarter each for Great Plains and Sales Logix.

Dynamic partitions are per-user partitions that manage each user's bandwidth allocation across one or more applications. They're useful for situations when you care more about equitable bandwidth allocation than about how it's put to use.

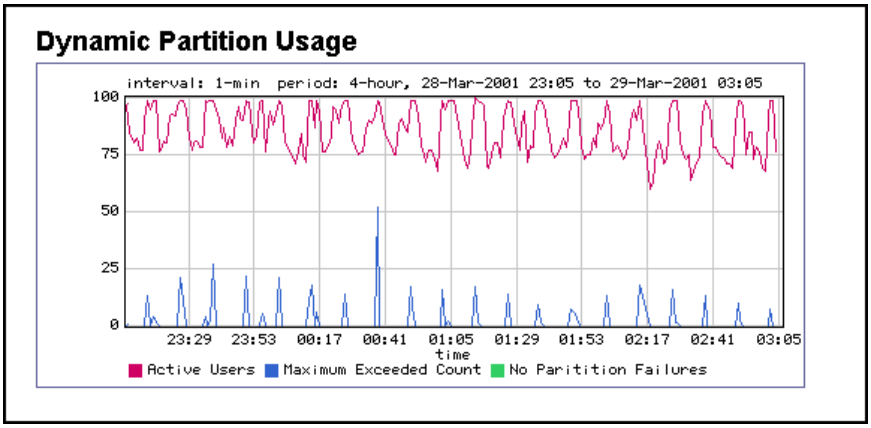

Dynamic subpartitions are created on the fly as users initiate traffic of a given class. When the maximum number of subpartitions is reached, an inactive slot is released for each new active user. They greatly simplify administrative overhead and allow over-subscription. As always, PacketShaper lends any unused bandwidth to others in need.

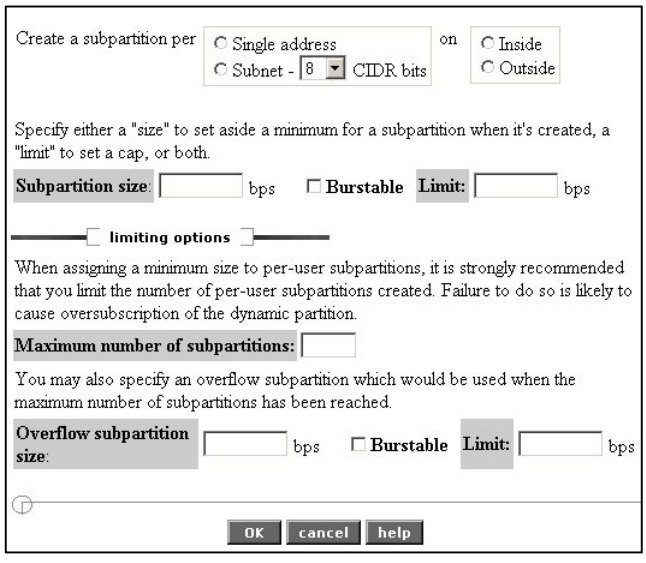

In addition, these same subpartitions can be created for a group of users within an IP address range.

For example, a university can give each dormitory student a minimum of 20 Kbps and a maximum of 60 Kbps to use in any way he/she wishes. And a company can protect and/or cap bandwidth for distinct departments (accounting, human resources, marketing, and so on).

Per-Session Rate Policies

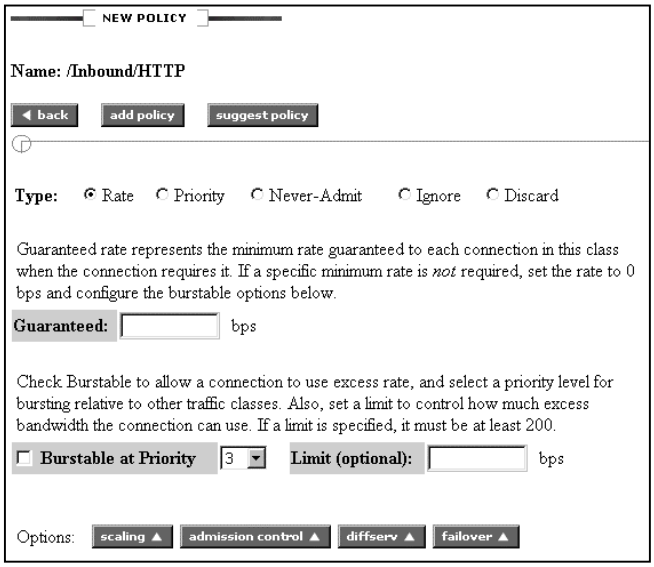

VoIP (Voice over IP) can be a convenient and cost-saving option, but not unless it delivers good service consistently. When delay-sensitive voice traffic traverses congested WAN links on a shared network, the result can be delay, jitter, packet loss, and poor reception. Each flow requires a guaranteed minimum rate or the service is unusable. After all, a video or voice stream that randomly speeds up and slows down as packets arrive in clumps is not likely to attain wide commercial acceptance.

PacketShaper's *rate policies* can deliver a minimum rate for each individual session of a traffic class, allow that session prioritized access to excess bandwidth, and set a limit on the total bandwidth it can use. A policy can keep greedy traffic in line or can protect latency-sensitive sessions. As with partitions, any unused bandwidth is automatically lent to other applications. VoIP traffic needs a per-session guarantee to prevent annoying jitter.

Other applications such as distance learning, NetMeeting, QuickTime, Real Audio, Streamworks, SHOUTcast, WebEx, WindowsMedia, WebEx, or streaming media of any sort would all be appropriate for rate policies with per-session minimums to secure good performance. Print traffic, emails with large attachments, and file transfers are all examples of traffic that need policies with a bandwidth limit at a lower priority than that for critical traffic.

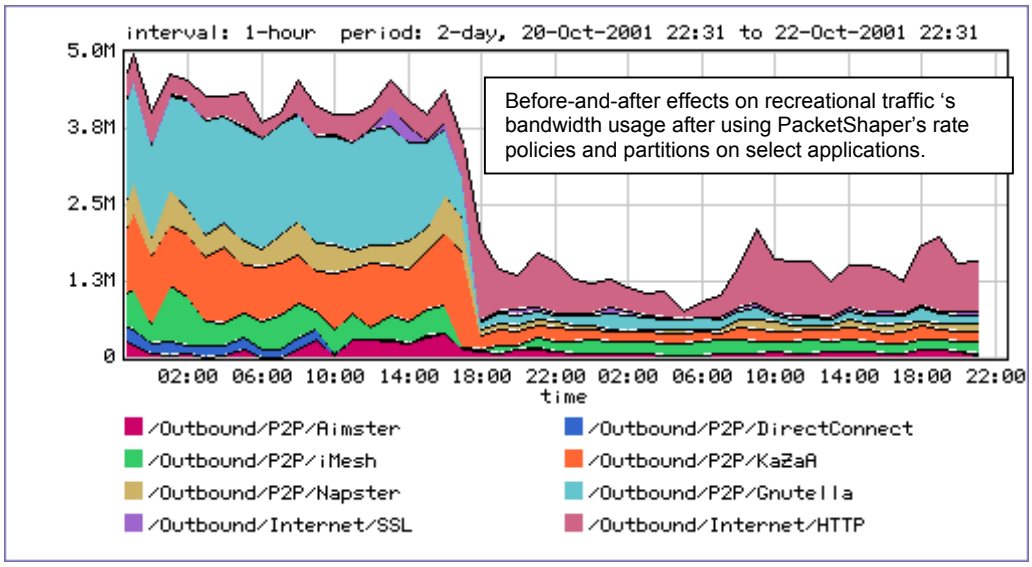

Before-and-after effects on recreational traffic 's bandwidth usage after using PacketShaper's rate policies and partitions on select applications.

The telecommunications firm in the partition example might want to go one step further to ensure that file transfers don't create problems. Suppose someone who is equipped with a T3 initiates a file transfer. Assuming the partition is in place and doing its job, that user could dominate the entire FTP partition, leaving other potential FTP users without resources. Because a partition applies only to the aggregate total of a traffic class, the individual users would still be in a free-for-all. A policy could fix this problem. A policy that caps each FTP session at 100 Kbps, or any appropriate amount, would keep the downloads equitable.

Other Policies

PacketShaper offers a variety of other types of policies, including:

- *Priority* policies allocate bandwidth based on a priority, 0 to 7. There are no minimum bandwidth settings and no bandwidth limits. Priority policies are frequently appropriate for small, non-bursty, latency-sensitive traffic such as Telnet.
- *Discard* policies intentionally block traffic. They are appropriate if you want to deny access to certain traffic unconditionally. The packets are simply tossed and no feedback is sent back to the sender.
- *Never-Admit* policies are similar to discard policies except that the policy informs the sender of the block. For example, you can redirect a web request to an alternate URL that explains the service denial.
- *Ignore* policies simply pass traffic on, not applying any bandwidth-allocation control at all.

Rate-Control Features

Although partitions and policies require explicit action from the network administrator, another control mechanism, *TCP Rate Control*, operates behind the scenes, optimizing a limited-capacity link. PacketShaper's TCP Rate Control overcomes TCP's shortcomings, proactively preventing congestion on both inbound and outbound traffic. TCP Rate Control paces traffic, telling the end stations to slow down or speed up. It's no use sending packets any faster if they will be accepted only at a particular rate once they arrive. Rather than discarding packets from a congested queue, TCP Rate Control paces packets to prevent congestion. It forces a smooth, even flow rate that maximizes throughput.

Contrary to TCP Rate Control, queuing-based bandwidth-management products wait for queues to form and congestion to occur, and then reorder and discard packets. Queuing-based solutions do not proactively control the rate at which traffic enters the wide-area network at the other edge. More importantly, queuing-based solutions are not bi-directional and do not control the rate at which traffic travels into a LAN from a WAN, where there is no queue.

TCP rate control measures network latency, forecasts packet-arrival times, adjusts window sizes accordingly, and meters acknowledgements to ensure just-in-time delivery of the transmissions.

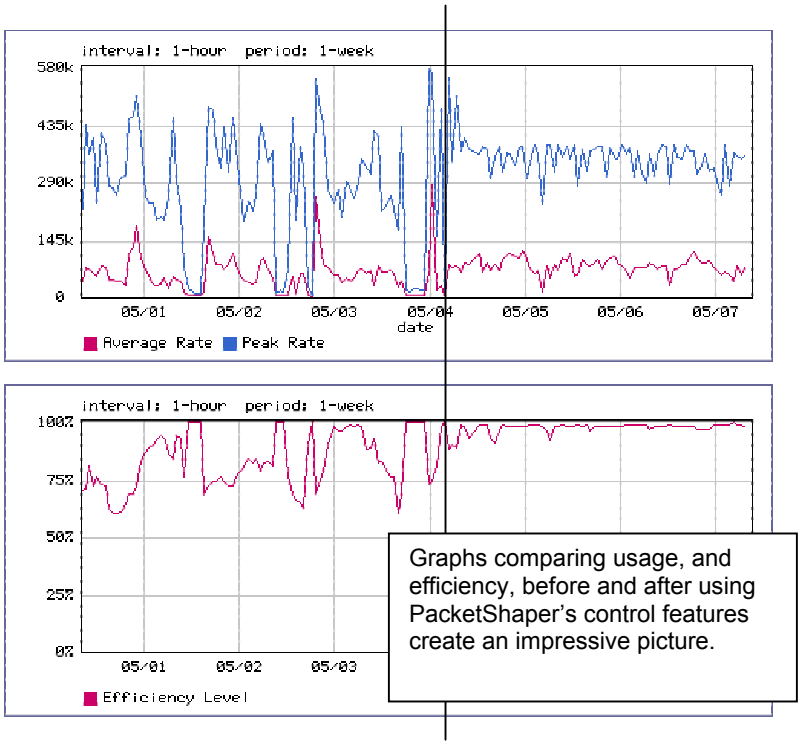

Graphs comparing usage, and efficiency, before and after using PacketShaper's control features create an impressive picture.

Imagine putting fine sand through a straw or small pipe. Sand passes through the straw evenly and quickly. Now imagine putting chunky gravel through the same straw. The gravel gets stuck and arrives in clumps. PacketShaper conditions traffic so that it becomes more like sand than gravel. These smoothly controlled connections are much less likely to incur packet loss, and, more importantly, the end user experiences consistent service.

Comparison Table: Queues and TCP Rate Control		
	Queuing	**TCP Rate Control**
Efficiency	• Tosses packets (RED, WRED) • Induces packet loss (tail-end drops) • Generates retransmissions (timeouts)	• Doesn't form queues • Transfers data more efficiently (more throughput, less time) • Reduces packet loss and retransmissions
Precision	• Limited traffic classification • No bits-per-second control • No per-session or per-user control	• Flexible, application-layer, extensive traffic classification • Explicit bits-per-second control • Rate-based QoS for individual application, sessions, users, and more
Full Duplex	• Outbound control, but no inbound control	• Inbound and outbound control
Proactive	• Reactive • Congestion has already occurred if queues form	• Proactive • Prevents congestion *before* it occurs

Universal Translator

Packet marking is a growing trend to ensure speedy treatment across the WAN and across heterogeneous network devices. First, CoS/ToS (class and type of service bits) were incorporated into IP. Then, Diffserv became the newer marking protocol for uniform quality of service, essentially the same as ToS bits, just more of them. And more recently, MPLS has emerged as the newest standard, integrating the ability to specify a network path with class of service for consistent QoS.

PacketShaper can classify, mark, and remark traffic based on IP COS/TOS bits, Diffserv settings, and MPLS labels, allowing traffic types to have uniform end-to-end treatment by multi-vendor devices in heterogeneous WANs. In attending to marking and remarking, PacketShaper acts as a type of universal translator, detecting intentions in one protocol and perpetuating those intentions with a different protocol as it forwards the packets.

For example, in an organization that has an MPLS core network but non-MPLS LANs, PacketShaper can assign an MPLS label to an outbound business-critical application that ensures a speedy path through the MPLS core. PacketShaper adds its layer-7 application awareness to MPLS capabilities and eases bandwidth bottlenecks, adding significant performance gains to those possible with MPLS alone.

Detecting and Avoiding Attacks

Although PacketShaper is not a firewall, it can help detect and avoid DoS (denial of service) nightmares. Recent DoS attacks against popular websites have raised consciousness about vulnerability.

These insidious attacks employ a variety of mechanisms to wreak havoc. For example, flood-type attacks initiate a large number of illegitimate connections that consume bandwidth and overwhelm receiving hosts.

PacketShaper employs a variety of methods to deal with the attacks. It can limit the number of connections from or to any host. Or limit the amount of ICMP traffic (a frequent attack vehicle that normally contributes just a small percentage of traffic). Or limit the number of flows in one application or traffic class. Or detect and block CodeRed, Nimda, and similar worms. Or block traffic that's only pretending to come from a trusted source.

Some ideas for extra protection against DoS involvement include:

- Limit the rate of new flows to or from a unique host.
- Place a limit on the number of concurrent flows for a traffic class.
- Limit ICMP to a maximum of 5 percent of the link size.
- Block traffic carrying the telltale signs of current worms.

Step Four: Generating Reports

PacketShaper's comprehensive reports, graphs, and tables provide clear insight into historical performance, load, and efficiency. Tools to support reporting functions are frequently similar to the tools for traffic analysis, but the motivations are different. Measuring historical data and formatting results into reports is usually associated with substantiating a strategy or purchase, evaluating service-level compliance, or searching for historical trends.

Examples of questions that can be answered by generating the appropriate report include:

- What was the most bandwidth that FTP consumed last week?
- Is my JD Edwards partition size overly generous or not generous enough?
- How much money did I squander on connectivity that was used for retransmissions? What percentage of my link was wasted on retransmissions?
- Is my SAP response time within its two-second goal?
- Even if I optimize my WAN link with an effective bandwidth-management strategy, when will I *have* to upgrade capacity?
- Were there sufficient dynamic partitions for all interested users?
- What caused slow response yesterday—a slow network or a sluggish server?

With PacketShaper, you can view one of the preconfigured reports, or you can define and create your own reports using any of PacketShaper's stored metrics.

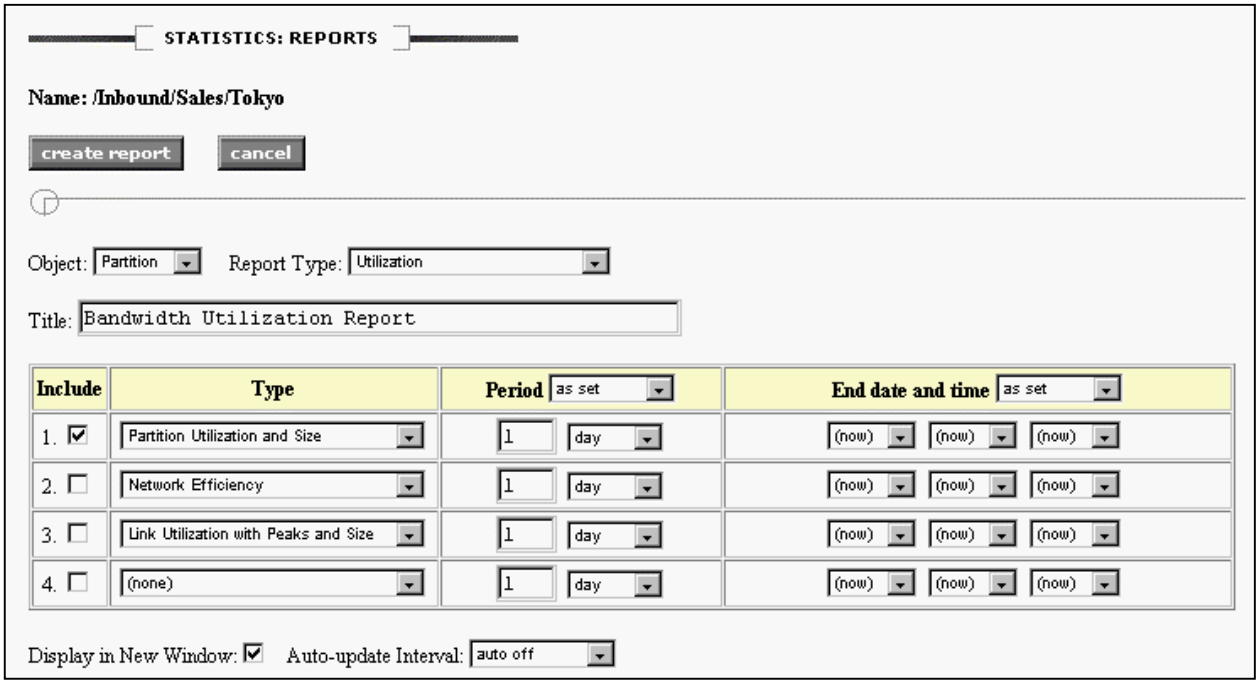

PacketShaper's pre-configured graphs include:

- Bytes Transmitted
- Class Utilization and Class Utilization with Peaks
- Connection Retransmissions
- Dynamic Partition Usage
- Guaranteed Rate Failures
- Link Utilization and Size or Peak and Size
- Network Delay and Network Delay Distribution
- Network Efficiency
- Normalized Network Delay
- Packet Size Distribution
- Packets Transmitted
- Partition Utilization with Size and with Peak and Size
- Packet Round Trip Time
- Service Level Compliance
- Server Delay and Server Delay Distribution
- TCP Connections Initiated
- TCP Health
- Top-10 Partitions, Top-10 Classes
- Transaction Delay and Transaction Delay Distribution

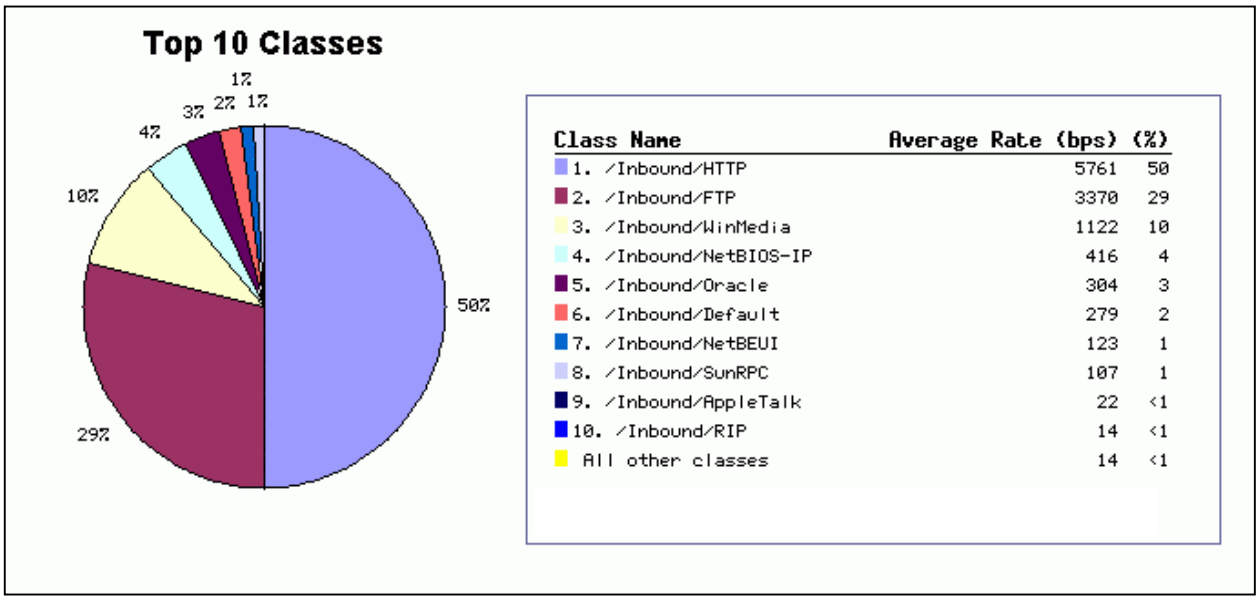

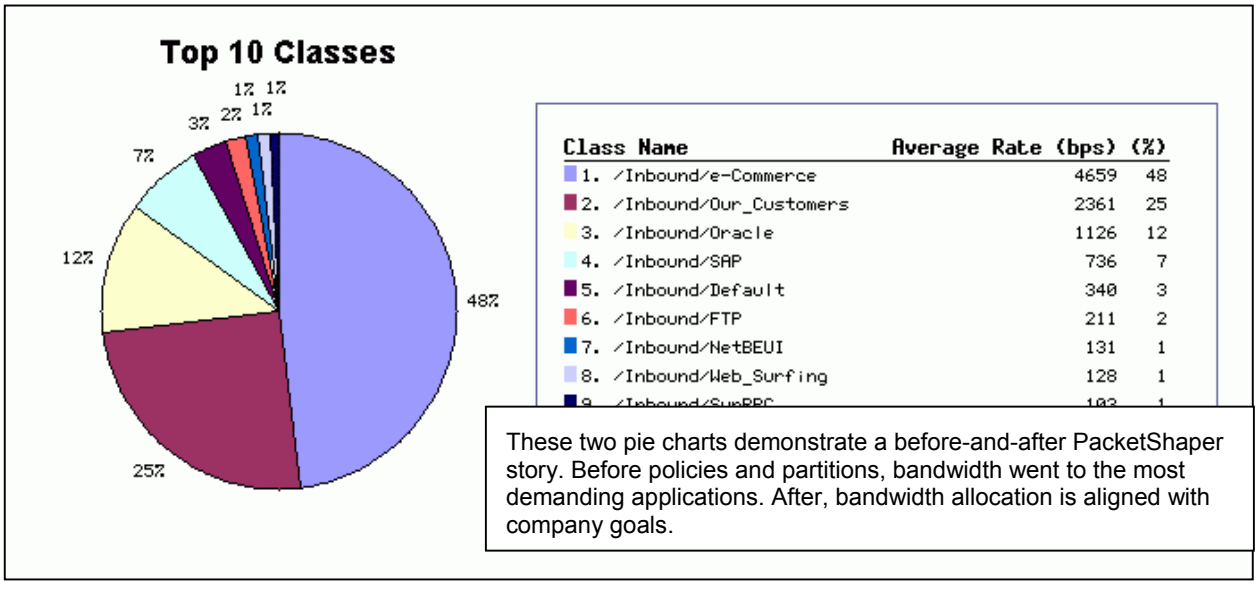

These two pie charts demonstrate a before-and-after PacketShaper story. Before policies and partitions, bandwidth went to the most demanding applications. After, bandwidth allocation is aligned with company goals.

Following a Typical Investigation

Suppose you're a network administrator in a large, worldwide corporation. You have Intranet servers in several different locations. You've defined a traffic class called *OurSites* to monitor Intranet traffic. OurSites' definition includes all HTTP traffic to any Intranet server. As you're perusing the Response Time Summary page, you notice that OurSites' summary says only 60 percent of its transactions have acceptable performance. Users are noticing it too; a few phone calls have started. What's next?

1. You view OurSites' response-time page and notice that the percentage of acceptable transactions has declined to 50.
2. Examining the time-series graph of delay times, you notice that although the server delay has increased only a small amount, the network delay took a sudden and dramatic increase yesterday.

   This tells you that although the servers might be a small bit overloaded, the network path to the servers is experiencing more major problems.
3. The Service Level Compliance graph, showing a time series of performance acceptability, confirms that performance crossed the line to unacceptable starting yesterday.
4. Next, you look at the Worst Clients list. Nothing stands out to you. If all the worst clients were, for example, on the first floor of the headquarters building, that could have helped narrow your search. But no such luck.
5. The Worst Servers list shows something interesting. One Intranet server shows only 17 percent good transactions. The others show percentages in the high 80s or better.

   The problem is becoming clear. The network path leading to one Intranet server has a problem.
6. You get another viewpoint on the servers by displaying the TCP Health graph on OurSites.

   TCP Health compares the number of TCP connections that were started, aborted, refused by the server, and not serviced by the server. It shows a high number of connections were started with an unusually high rate of aborted connections. The numbers of connections that were refused or ignored are fairly low.

   The statistics make sense with rest of your investigations. Users are attempting to access the Intranet (number of connections), but getting frustrated by slow response and clicking the Stop button on their browsers (high rate of aborted connections). But once the servers get the requests, they're serviced.
7. You consult usage statistics in PacketShaper's *Monitor Traffic* page. FTP shows an increase. FTP's Top Talkers list shows a very active FTP server. Switching to a graph of FTP's usage, you see that FTP has been bursting to 80 percent of the link the past two days. It seems that a particularly large and popular file was posted two days ago. The affected Intranet and FTP servers are located together and share a link.
8. You solve the problem by creating a PacketShaper rate policy that allows FTP traffic to burst only if there are no other consumers. Otherwise, it's contained to a slower usage rate. Then you create a rate policy for OurSites that safeguards its performance.
9. OurSites' performance returns to normal. Users are happy. You take a well-earned break.

Recommended Environment

PacketShaper adds value in almost any network with the possible exception of an isolated LAN with more than ample bandwidth and no WAN connection. But the environments where PacketShapers have been especially successful include:

- Networks that support expensive, business-critical applications (SAP installations, for example) together with bursty, less urgent traffic such as web surfing.
- IP networks that carry critical legacy applications with TN3270 or TN5250 protocols or with web-enabled host access.
- Thin-client environments where small interactive packets are very delay sensitive but get overrun by large, bandwidth-hungry applications.
- Businesses or organizations that provide streaming audio or video (VoIP or distance-learning applications, for example).
- Environments where building facilities (HVAC, for example) share the network with other applications.
- Branch offices that must transact business with a corporate headquarters over the WAN.
- Environments using messaging and collaboration applications (Microsoft Exchange or Lotus Notes, for example) where a performance balance is needed between messaging and other critical applications.
- Organizations whose networks support recreational use in addition to officially sanctioned applications. For example, at universities, students use the university network to download music files.

For More Information

If you'd like more information about PacketShaper or to order a PacketShaper, consult Packeteer's web site or call (408) 873-4400 or (800) 697-2253.