

Technical Overview

PROTECTING VIDEO OVER IP TRAFFIC



*Using Packeteer's PacketShaper to protect
Video over IP traffic across the WAN*

July 2002

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



©2002 Packeteer, Inc. All rights reserved. Packeteer, the Packeteer logo, AppCelera, AppVantage, PacketShaper and PolicyCenter are trademarks or registered trademarks of Packeteer, Inc. in the United States and other countries. All other company trademarks are the property of their respective owners. No part of this document may be reproduced, photocopied, stored on a retrieval system, transmitted, or translated into another language without the express written consent of Packeteer, Inc.

This paper examines a significant problem associated with IP convergence and multi-service networks — how to protect IP videoconferencing traffic from competing voice and data traffic. In particular, this paper explains how the unique capabilities of Packeteer's PacketShaper[®], an application performance solution, resolves this problem.

Background

Today's network users have come to expect similar levels of quality and experience as they do of traditional real-time communication systems, such as POTS* and broadcast television. They have little tolerance for quality degradation in video over IP environments and far less tolerance for unpredictable data transactions running over the same IP networks. For this reason, ensuring effective quality of service (QoS) is extremely important when implementing a converged network in which video and data share the same IP circuit. Without effective QoS, bursty data traffic can easily disrupt video traffic, causing glitches, cutouts, frame loss, and other problems that jeopardize a user's satisfaction with this technology.

Because of its real-time nature, IP-based voice and video are particularly sensitive to jitter (variation), latency (delay) and packet loss (bandwidth). Packet loss rates of two to three percent or jitter of 200ms can render an IP video signal unusable. A good QoS system must protect video traffic from these potentially harmful effects.

This involves controlling traffic across the wide-area network, particularly at the most bandwidth-constrained points of the network — WAN links. These links require policy control for multiple reasons: to protect the video traffic so that it does not suffer packet loss, high latency, or jitter from queue delays, and so that excessive video traffic will not disrupt the performance of other mission-critical applications. TCP is very aggressive and bursty, so even a small number of competing users accessing file servers, surfing the web, running print jobs, or sending email can interfere with the performance of video traffic.

Packeteer's TCP Rate Control technology eliminates network congestion by preventing bursty TCP flows from consuming excessive amounts of bandwidth and router queues from causing delay, jitter, and packet loss. Packeteer's UDP Rate Control and Admissions Control technologies also ensure that critical data traffic on the network is not disrupted by video traffic. These QoS technologies, combined with sophisticated Layer 7 classification and performance analysis, enable Packeteer's PacketShaper[®] to provide an advanced QoS platform that ensures efficient, reliable performance of critical applications over the WAN and Internet.

PacketShaper Methodology

Implementing PacketShaper in a video over IP environment involves three steps:

1. Assessing network requirements through traffic classification and performance analysis
2. Implementing policies
3. Monitoring results

All three of these elements are equally important.

Assessing Network Requirements

Traditional routers and queuing systems use static port numbers to detect applications. Consequently, these products are incapable of tracking the protocol when it splits and jumps to randomly selected port numbers. When this happens, network administrators lose any ability to monitor and control that traffic. PacketShaper, by contrast, peers deeply into each flow at Layer 7 and determines its identity, regardless of what port number it uses. This is a critical capability for supporting H.323 installations, because after all, you can't control what you can't classify.

PacketShaper has an extremely robust, application-aware classification technology, allowing it to identify applications at Layer 7. In the case of H.323, this is extremely important because the protocol does not use "well-known" TCP or UDP ports. It starts on a well-known port number for call setup negotiations but quickly splits into different channels such as RTP (data channel) and RTCP (control channel) which each have different bandwidth and policy requirements. Not only are these channels different, but they use randomly selected port numbers.

PacketShaper also has Automatic Traffic Discovery, enabling it to immediately detect and characterize H.323 traffic on your network with no operator intervention. This makes setup and getting started simple and quite convenient.

H.323 Elements	Protocol
Call Setup Protocols	Q.931, H.323 Gatekeeper, etc
Video & Voice Channels	RTP
Control Channels	RTCP
Collaboration Channels	T.120
System Control	RSVP, Proprietary Management
Some of the different protocols that may be used in a single H.323 call. PacketShaper detects all of them with its Layer 7 classification capabilities.	

Another key benefit of PacketShaper is its application monitoring capabilities. PacketShaper automatically records more than 50 different statistics about each application on the network, ranging from network utilization to network health to application response time. All data is stored on board.

Implementing Policy

PacketShaper delivers a comprehensive set of traffic control mechanisms for managing traffic and enforcing policy. Policy provisioning is very flexible, providing the opportunity to set minimum (CIR) and maximum (EIR) bandwidth allocations. These can apply:

- Per flow
- Per user
- Per application
- Per group of users or applications
- Any combination of the above

Policies are hierarchical and can be nested so that, for instance, you could limit the total amount of video traffic on the network to 1,000k but guarantee each individual RTP stream at least 200k. Flows that cannot be serviced are handled by the admissions control mechanism and can be dropped, rejected, or put on hold until bandwidth becomes available. With PacketShaper, all policies share their unused bandwidth with other lower-priority traffic so that bandwidth does not go unused. Admissions control can limit the maximum utilization of video traffic on the network either by number of active sessions or amount of active bandwidth. New sessions exceeding the limit will be dropped or rejected.

Packeteer offers a comprehensive set of traffic control technologies, including:

- TCP Rate Control
- UDP Rate Control
- Per-flow queuing
- Class-aggregate queuing
- MTU-size management
- Admissions Control

Monitor Results

Once you've implemented desired policies, you can monitor the network to validate their success. Set revised policies as needed.

Improving Video over IP

Packeteer's traffic management technologies are useful in a number of ways in a video over IP environment:

Measurement

- Characterize bandwidth requirements of different equipment and codecs
- Monitor overall bandwidth usage of video as a percentage of the network, usage per stream, etc.
- Trend analysis to chart growth of video traffic and conduct capacity planning for bandwidth upgrades
- Monitor top users of video and how much bandwidth they are consuming
- Monitor network latency and round-trip times to assess delay

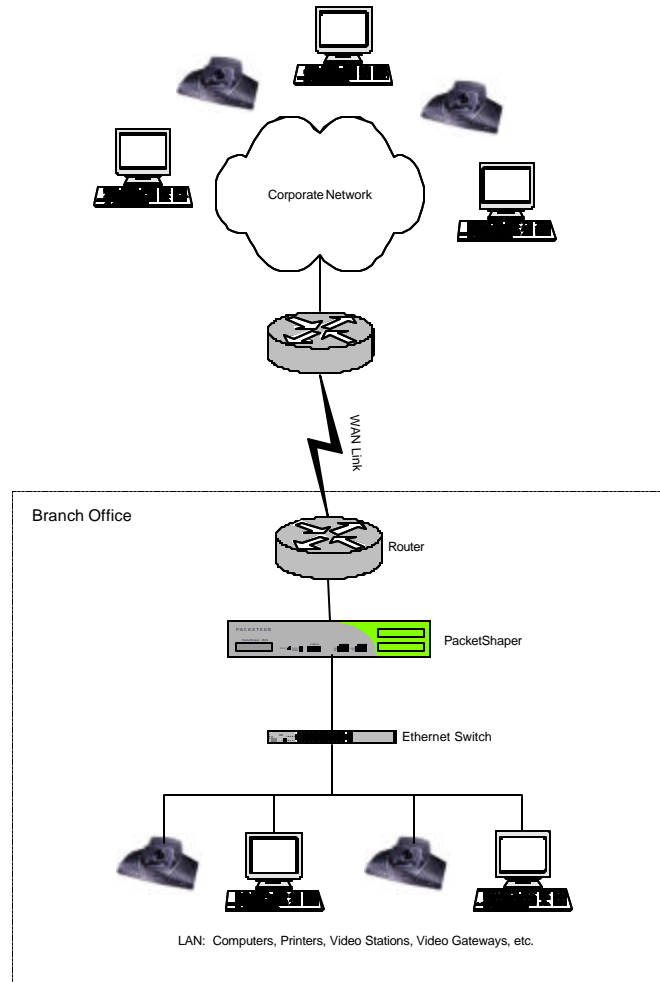
Management

- Set guaranteed minimum bandwidth to protect video traffic from bursty data traffic
- Set maximum bandwidth or session limits to prevent video traffic from consuming excessive resources and squeezing out important data traffic
- Set alarms to report back to central management console when traffic conditions or performance levels, such as latency, exceed specific thresholds
- Execute synthetic transactions to continuously measure network latency, availability, and response time, and send alerts when they go out-of-spec*
- For smaller (<1Mb) links, set a maximum packet size (MTU) on TCP traffic to assure that it integrates well with video streams and does not create excessive jitter
- UDP Rate Control and latency bounds to help neutralize jitter in RTP flows

* The Synthetic Transactions feature is only available on AppVantage models

Deploying PacketShaper

PacketShaper is deployed at branch offices, next to WAN-link routers on the Ethernet side. From this vantage point, its Layer 7 visibility provides organizations with thorough visibility into traffic heading to and from the central site and between other branch sites. This visibility provides the intelligence needed to utilize PacketShaper's policy control features for managing bandwidth allocation, latency, and the efficiency of each application based on its relative business importance. Like a switch, PacketShaper fits transparently in the traffic path and does not add a hop or alter routing configurations. Because it resides on the Ethernet side of the router, PacketShaper is transport-agnostic. In other words, it integrates seamlessly with frame-relay, point-to-point, ATM, DSL, wireless, and most other WAN transport mechanisms.



Do I need a PacketShaper on both sides of the link?

In most cases, the answer is no. Traditional queuing-based QoS systems can only effectively control traffic in the outbound direction. Therefore, you must use two per link, one on each side, to manage traffic. However, Packeteer's unique TCP Rate Control technology manages inbound and outbound TCP traffic, requiring only one PacketShaper unit per WAN link. In the case of UDP, video over IP traffic is not bursty, and thus the PacketShaper's predictive scheduler can accurately measure and predict bandwidth needs, even for inbound traffic, in order to schedule TCP traffic around it.

A second PacketShaper on the other side of the link may be required in cases where there is significant bursty or excess UDP traffic on the link that needs to be limited or eliminated. Normally, this case only occurs when sites have excessive video over IP traffic and some flows need to be denied with bandwidth-based admissions control.

Analyzing Video Traffic

To examine the impact of competing TCP traffic on video over IP, we ran a series of tests. For the video traffic, we used a Polycom iPower 900 with 768k video selected, running across a T1 line between San Francisco and Vancouver. The competing traffic was TCP from a traffic generator simulating a number of users surfing the web, downloading large email attachments, and accessing file servers.

Without QoS




Figure 1. Bandwidth consumption of video over IP traffic in the face of competing web surfing traffic. No traffic shaping is present.

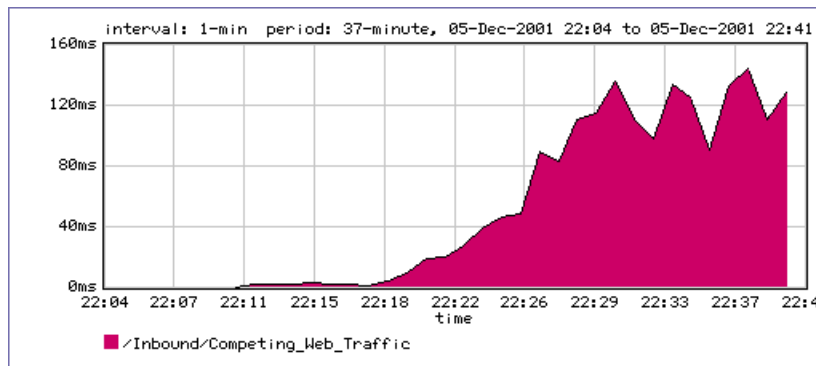


Figure 2. Additional network latency added by congestion.

Test Developments:

- 22:04** You can see that video traffic (red) accesses the bandwidth it needs (~812k) because there is no competing traffic (22:04 – 22:11)
- 22:11** Competing traffic is introduced. The volume of competing traffic builds from one user at 22:11 to 15 users at 22:40.
- 22:17** With just two users active and a third starting up, the TCP traffic has already expanded to fill all of the available remaining space on the T1 link. Total link

consumption of the video plus the competing traffic now consumes 1.544 Mbps.

- 22:19** With five users active, the video traffic is still getting all of the bandwidth it needs. You can see on Figure 2 that latency added by congestion (above and beyond the basic link transit time) is now starting to increase rapidly. Glitches are beginning to show on the screen, likely because of increased jitter caused by the rise in latency.
- 22:20** As the sixth user starts up, the video is no longer able to compete and begins losing packets. Its bandwidth share drops to 790k (2 percent loss) and the video starts to get glitchy.
- 22:23** At eight users, the video is only managing to hold about 730k (15 percent loss) and is no longer usable at all (frame rate = 1 fps). Performance degrades rapidly from here as additional users come online. Congestion-induced latency soars.

Summary: Without QoS, it took no more than six active users downloading files from a file server or the web to destroy the performance of even a robust 768k video stream. Smaller video streams would have even less margin for error and be more sensitive to these congestion effects.

It is not a big stretch to imagine 10, 15, or more users simultaneously accessing a T1 line that serves a typical 50- or 100-person branch office, especially in the morning when employees arrive, sync up their laptops, download email, log on to systems or view and print summary reports from the night before. Or consider a big event that would prompt many people to use the network at the same time, such as breaking news, a stock market move, a press announcement by the company or its competitors, a significant sporting event, or someone emailing a large Powerpoint file to a dozen colleagues before an upcoming meeting. Any of these scenarios are common.

With QoS From PacketShaper

We ran the same test again, but this time with the PacketShaper deployed on the link to protect the video traffic. A partition was configured on the PacketShaper to give the video (RTP) traffic a guaranteed bandwidth of 820k. Auto-policy provisioning was turned on to allow the PacketShaper to automatically assign suitable policies to competing traffic.

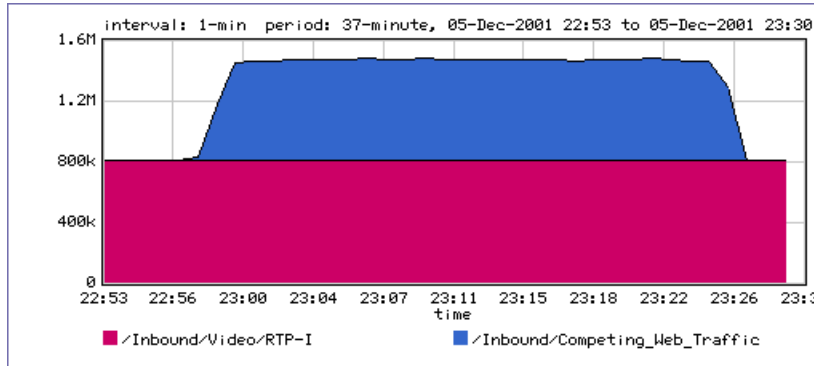


Figure 3. Video’s bandwidth consumption in the face of a competing traffic load equivalent to 20 users downloading files.

The competing traffic load was set at 20 simultaneous users for the entire test. Note that the video maintained all of the bandwidth it required for quality performance, even in the face of a heavy load and a full pipe. In fact, as shown in Figure 4 below, packet loss rates for video were actually lower during this test than they were with no competing traffic.

	Total Packets	Loss	% Loss
No Shaping, No Competing Traffic	20,000	47	.23
Shaping On, Heavy Competing Traffic	230,000	162	.07

Figure 4. Packet loss with no other competing traffic on the link was actually higher than with shaping on and 20 users generating competing traffic.

Summary: With PacketShaper QoS, video over IP traffic is protected even in the face of a heavy, competing traffic load. In fact, PacketShaper’s UDP Rate Control appears to neutralize video jitter and achieve lower loss rate than unshaped video without competing traffic.



©2002 Packeteer, Inc. All rights reserved. Packeteer, the Packeteer logo, AppCelera, AppVantage, PacketShaper and PolicyCenter are trademarks or registered trademarks of Packeteer, Inc. in the United States and other countries. All other company trademarks are the property of their respective owners. No part of this document may be reproduced, photocopied, stored on a retrieval system, transmitted, or translated into another language without the express written consent of Packeteer, Inc.