PACKETEER®

Successful Voice, Video and Data Convergence Over the WAN

Intelligent LifeCycle Series

Contents

Successful Voice, Video and Data Convergence

Introduction

Voice over IP (VoIP), or IP Telephony (IPT), is fast winning acceptance, with over half of distributed enterprises planning to complete site-to-site VoIP deployment in the next three years (according to a 2006 Forrester Research study). The benefits of a converged network seem obvious: not having to create staff and run two different networks for voice and data, running just one cable to desktops instead of two and a reduction in phone tariffs are all good for business. Simple things like being able to transfer your own extension number while visiting a branch office or just being able to have the same phone extension at multiple locations can also add up in terms of cost savings and convenience.

More interestingly, the converged network introduces a new paradigm in inter-office communications. Real-time communications, such as escalating an Instant Message (IM) to a voice call mid-message or integrating presence data with video and "traditional" communications like e-mail, are rapidly being recognized and embraced as a true competitive advantage. Integrating call-control capabilities with Microsoft's Live Communications Server (LCS), and Office Communicator present further scope for major efficiencies.

In practical terms, most organizations will not replace their depreciated PBX (Private Branch Exchange) with traditional equipment but with IP-PBXs. Similarly, their annual network subscription will be replaced, if it hasn't been already, with IP-based MPLS (Multiprotocol Label Switching) services promoted by their service providers.

So the time to implement Voice and Video on a converged network is now. Significant savings and flexible new ways to work lay ahead once your services and infrastructure converge. What's the shortest, most painless route to fully taking advantage of these benefits?

Quality has emerged as the major obstacle to deployment success as network congestion and competition for bandwidth interfere with voice and video traffic. Many barriers stand between your ability to fulfill voice, video and data convergence's promise of flexible ways to work and significant savings:

- Poor visibility into performance of Voice and Video Calls
- Managing jitter, latency and packet loss, often created when bulk recreational traffic or files crowd out voice or when voice queues are oversubscribed
- Differentiating between sanctioned and unsanctioned VoIP, video or IM traffic
- Inability to increase capacity for convergence without buying additional bandwidth

All these challenges exist against a backdrop of a user community that expects the same level of service and the same levels of quality that they enjoyed with the Plain Old Telephone Service, and a successful implementation must meet or exceed that expectation while supporting increased usage of rich converged inter-office communications.

Application Types	IP Telephony / VoIP Video Conference
Problems	Poor call quality Slow video setup and freeze frames Troubleshooting tools are not voice savvy Jitter and loss due to congestion Oversubscribed CoS in MPLS or aggregate queues
Causes	Congestion—other applications interfere Oversubscription—too many calls in a class of service Lack of capacity—more traffic and ITP sites consume bandwidth

The Intelligent LifeCycle Approach

Packeteer's Intelligent LifeCycle describes a proven methodology for assessing, troubleshooting, and managing application performance across a wide variety of environments. This process calls for four distinct activities that help customers focus their efforts on providing fast, effective solutions rather than temporary fixes for application performance issues. The four steps are: **Assess**, **Provision**, **Accelerate**, and **Extend**.

The Intelligent LifeCycle methodology for voice and video entails assessing the readiness of your network and identifying the different types of voice and video traffic your business utilizes. Then, using provisioning, you are able to reserve a portion of the network to support all of your traffic, protecting individual voice and video users from each other, and protecting business critical applications when voice and video demand swells.

Let's see how this methodology can help you fulfill the promise of voice, video and data on a converged network, while providing consistently high quality of service for users.

Assess: Determine the Right Tools for the Job

Identify what applications are running on the network, what approaches to take to resolve issues, and continuously monitor performance.

Are You Ready?

Assessing the readiness of your network is key to a successful deployment. You need to ask yourself if the same WAN bandwidth today is able to support your efforts tomorrow. If not, can you borrow bandwidth from less business—or time-critical applications or will you have to add capacity? And what will be the impact to your existing applications? How well will voice and video perform on your WAN?

You need to ensure your existing networks and the service levels established for data traffic are capable of supporting the additional demands placed on them by voice and video. Packeteer gives you Layer 7 Plus visibility into the applications on your network and the bandwidth they're consuming. Packeteer's unique Layer 7 Plus visibility technology automatically discovers virtually all types of voice and video traffic automatically, including Skinny, SIP, RTP and RTCP, H.323, MGCP, H.248/Megaco, T.120, as well as Skype, CU-SeeMe and other external voice services.

By discovering applications, tracking network utilization and monitoring performance, you can scope for bandwidth allocation. This will help you determine if you need additional bandwidth.

Measure and Model Usage

Voice and video clients typically use UDP streams. They have different flows and protocols for initiation, control, and data flows. For example, H.323 starts a conversation on one port (using H.323 protocol), jumps to another port (for Q.931), and eventually splits up into a data flow (RTP) and control flow (RTCP).

You will want to see how that type of initiation and teardown and data and flow control affects bandwidth, so with the Packeteer system in place, start and stop a few voice calls or video-conferences to generate some traffic. Leave one of the sessions running at least 10 minutes to collect a substantial amount of measurement data.

Packeteer measures both data payload and header overhead, so requirements will be slightly higher than the nominal bandwidth rate of the protocol. In general, you should expect to see an increase ranging from 10 to 20 percent. Thus, a 768 Kbps stream might take 820 Kbps when headers are included.

Typically, if a manufacturer claims that its voice flow requires 8 Kbps, it will actually need 17 to 21 Kbps due to additional overhead and forward error correction. In addition, it is best to overstate the bandwidth needs of UDP traffic by 15 to 20 percent. If in doubt, try 25 Kbps as your guaranteed rate for voice flows.

Determine an appropriate minimum amount of bandwidth you want for your voice and video traffic, even during times of contention. In addition, determine the maximum amount of bandwidth you want all voice and video to be able to access, even during times of little or no contention. These numbers will be the minimum and maximum sizes for your voice and video “allocation”.

Determine how many concurrent voice and video users you want to support at a minimum. For a rough estimate of bandwidth needs, multiply that number by the amount of bandwidth you observed being used during your one prolonged test session.

- For voice, for example, you might want to support 10 concurrent sessions at 25 Kbps each, for a total of about 250 Kbps in each direction.
- For video, if you’re using a 384 codec, for example, you might want to support two concurrent sessions of 420 Kbps each (remember, headers and control information are included).

Provision: Find and Fix

Create network resource policies to align network resources with the business.

If left unchecked, the additional bandwidth required for voice and video traffic on a converged network will undoubtedly impact the performance of business applications whose data shares that bandwidth. With Packeteer, however, you can shape traffic to wring the maximum value from existing bandwidth without harming the performance of the applications most critical to your business.

In order to do so, first determine what applications are running on your network with Packeteer’s Layer 7 Plus classification and analysis technologies. You can then group applications together in a way that makes sense for converged networks by isolating:

- Applications that are time critical, including business applications—like SAP or Oracle—and voice and video
- Bandwidth-hungry applications (sometimes called “bursty”) like e-mail, file transfers, backup, database sync and distributed storage
- Recreational or malicious traffic

Before deciding that new voice and video traffic automatically merits the purchase of additional bandwidth, it makes sense to prioritize your applications. Use your application list and Packeteer control and application quality-of-service (QoS) technologies to shape traffic and create room for your latency-sensitive business applications as well as the new voice and video traffic.

Protect if Voice and Video Demand Swells

Typically, bandwidth is limited (often by cost) on your WAN, and when a link is bandwidth-constrained, you need to impose QoS rules that give voice and video priority. This is the only way to ensure smooth end-to-end delivery insulated from the impact of “bursty” data traffic, and guarantee service levels for voice calls or video sessions.

Use the Packeteer system to set a rate policy on each of your voice and video data classes (RTP or others) to accomplish several goals:

- To indicate the relative importance of your voice and video traffic so that the Packeteer system knows how to distribute excess bandwidth
- To insulate voice and video users from each other to gain the benefits of rate control and in particular:
  - Reduce retransmissions that waste bandwidth and introduce latency
  - Reduce the appropriate per-session bandwidth to ensure quality calls and voice sessions

Use a guaranteed rate of the minimum bits per second that are required for acceptable voice call or video session quality. Make your policy burstable with a relatively high priority. In doing so, you need to also respect the service levels already in place or desired for other business-critical applications, so use firm but fair policies. Avoid simply prioritizing voice and video (and also see “MPLS Integration” below).

Identify your other critical applications and targeted service levels, and set a bandwidth minimum for that class (for example, a minimum of 20 percent of the link, burstable to 50 percent). To guarantee bandwidth on a per-session basis, create policies to support those requirements.

Treat Sanctioned and Unsanctioned Voice and Video Differently

Recreational voice and video traffic, such as that offered by Skype or instant messenger services, is becoming a greater concern as the technology is rapidly assimilated into our daily lives. Bandwidth intensive recreational voice and video traffic can tie up even the highest bandwidth links. If your business relies on sanctioned IP telephony, you need to differentiate between this and unsanctioned voice or video traffic.

The Packeteer system gives you the ability to differentiate between sanctioned and unsanctioned voice and video traffic by looking beyond standard quality of service (QoS) and letting you shape network traffic to suit your needs. By monitoring the network in this way, you’ll have the ability to view—and block—traffic that can potentially disrupt business and customer critical applications. Policies can be applied to protect critical voice and video traffic and smooth disruptive, bandwidth-intensive or bursty traffic from the other applications running across the network. Recreational traffic can be contained and malicious traffic can be blocked.

Recognize MPLS Limitations – and Integrate to Solve Them

If you are using an MPLS-based IP Subscription, you will clearly want to assign voice and video with the real-time priority class of services. But beware—there are limitations in the queuing and packet-based QoS features inherent in MPLS.

For example, if you have provisioned a class of service for four calls, each packet of each of these four calls will be given the highest priority. In an oversubscription situation—too many calls in a class of service—the network subscription starts having to process more of that class of traffic than it can support. So when calls 5, 6 and 7 occur, all packets have to be downgraded to the lower class of service category. All calls suffer a reduction in quality.

Recognizing this limitation, the Packeteer system allows you to integrate with MPLS. Packeteer QoS works at the IP flow level to control traffic, and then marks them with the proper DiffServ Code Point (or VLAN tag) to get them into the proper service class across your MPLS network connection.

In this way, real-time voice and video traffic in need of immediate and priority bandwidth is tagged with the correct MPLS class—but only after Packeteer has managed the traffic at the flow level to assure that oversubscriptions do not occur.

If you are using MPLS to route voice and video and application traffic on different paths, assign an appropriate Diffserv value to your voice and video and application classes. This assumes that the Label Edge Router will be doing MPLS tagging based on Diffserv or IP ToS tags. You can use the Packeteer system to assign the MPLS labels (instead of the LER).

When special traffic is assigned an MPLS label to facilitate prescribed performance, you can configure the Packeteer system to classify the traffic separately. Then, once you have traffic classes for your MPLS-enabled traffic of interest, you can apply analysis techniques or assign partitions or policies to control performance.

Accelerate: Take it to the Next Level

Apply technologies to enhance performance and capacity

Especially high on the list of potential problems for convergence projects is making room on the network for the increased traffic from video and voice sessions going between sites. Network convergence has many proven benefits; however, it does increase pressure on the network with additional traffic.

An all-too-common response is to throw bandwidth at the problem or invest in a compression point-product. Unfortunately, this is not only an expensive reaction to the unwelcome traffic, it also tends to be ineffective. Since many recreational applications—be they voice, video or data—are opportunistic, they simply ramp to fill any bandwidth that becomes available.

Without control in place, it's a waste of money to add bandwidth or invest in a compression product because nothing has been done to ensure the additional capacity supports business-critical applications like Oracle or SAP, as well as the new sanctioned video and voice traffic.

Warning—Accelerators and Voice

Clearly, protocol acceleration is not applicable to the real-time protocols that carry voice and video. In addition, there is a risk of creating problems for voice or video traffic if the protocol acceleration technology you use is not application-aware and does not allow you to protect voice and video traffic as we discussed in the Provision section above.

These charts illustrate the problems with applying acceleration without control.

Acceleration can have a negative impact on the network characteristics that cause problems for voice, including latency, jitter and loss.

Look what happens to the network when acceleration is applied without control:

- Latency increases 50-100%
- Jitter jumps to 140 ms
- Loss virtually doubles

Taken from PacketShaper 1700 VOIP Summary Report (8.2.0) Based on 100ms latency (round trip) with 20 calls and background HTTP (20 flows retrieving 100K files)

VoIP Latency

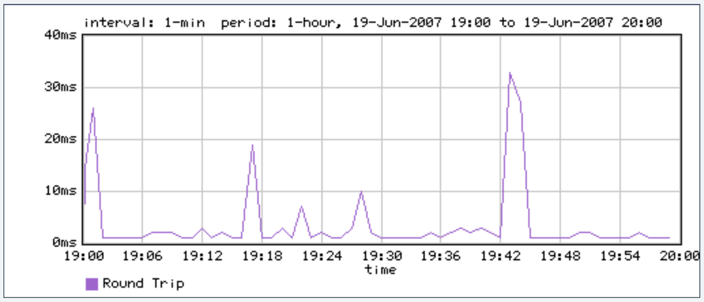

VoIP Jitter

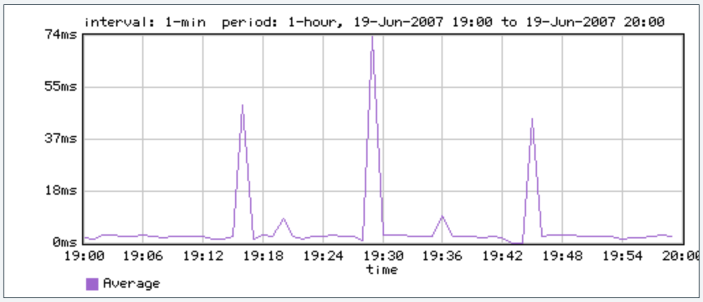

VoIP Loss

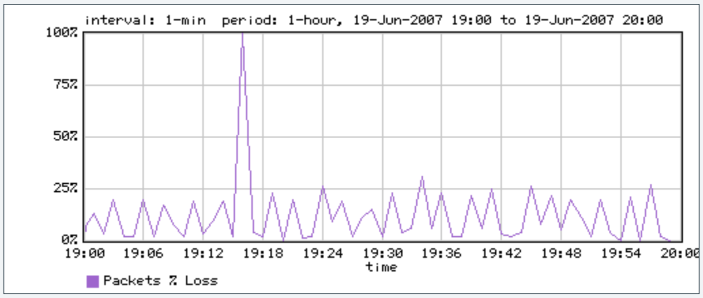

Packeteer Intelligent Acceleration

Packeteer Intelligent Acceleration combines compression with control, giving you the best of both worlds. Packeteer compression uses advanced compression algorithms to provide low-latency, high-impact compression to increase WAN capacity and create room for your convergence project, while the application QoS feature provides the control needed to ensure that the applications most critical to your business get the bandwidth they need to perform at their peak.

Real-time protocol traffic (like the payload in voice and video traffic) is already very highly compressed by the codec. But there is scope for streamlining, as the packets overall contain additional overhead: the RTP header, UDP header and IP headers. Compression techniques can reduce the headers on voice and video packets, reducing the total traffic by as much as 30 percent. Further techniques, such as concatenation and further compression of combined packet headers can further reduce the amount of data actually sent across the WAN.

Packeteer's unique RTP compression packing technology super-compresses the headers of those traffic types and packs them together across many different calls—even with other compressed traffic—significantly saving bandwidth even as call volume increases.

Extend: Everyone, Everywhere

Create an intelligent overlay that extends and adapts current infrastructure to new and emerging issues

At this point we've put some powerful solutions into place to provide a consistently high level of service for voice and video, ensure that other key business applications perform as expected and maximize the performance overall for a converged network. Having controlled traffic and managed the performance of key applications, we can now put in place procedures and ongoing mechanisms to maintain service levels for all business critical applications, be they voice, video or data-centric.

Measure Quality Metrics to Apply SLAs to Voice and Video

Being able to measure and report on quality metrics will help you establish guarantees of quality, in the form of Service Level Agreements (SLAs), and get in front of trends or increases in usage. In this way, you can provision in advance of problems arising and even get the most out of your MPLS subscription.

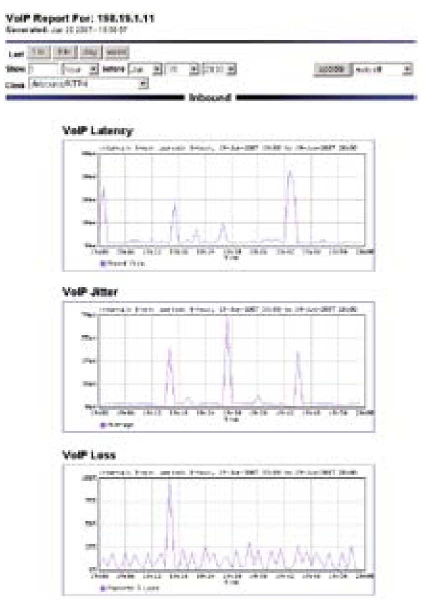

Quality metrics for converged networks measure jitter, loss and delay for RTP traffic—voice and video. The Packeteer system reports in these terms helping you monitor call quality for IP Telephony and Video conferencing. Also, because the classes reported can be tied to particular MPLS classes of service, you can monitor performance of real time classes for MPLS WAN Service.

Leverage the monitoring and reporting tools to look at voice and video traffic data and get an idea of bandwidth trends. Observe the measurements for current rate, one-minute average, and peak rate.

Create utilization graphs covering the time period you left the one session running for the large VoIP data flows (RTP or others).

The Packeteer system provides pre-configured SLA reporting for voice and video

Establish service level reporting to set thresholds and measure compliance with the SLAs thus put in place. Track the service levels and amend your policies accordingly.

When traffic reaches a threshold, Packeteer Application SLA monitoring provides a clear and present view of warning and critical situations.

Set performance thresholds against a known set of quality metrics, and configure the Packeteer system to proactively send alerts or react immediately—for example, to reduce bandwidth for applications that are not time- or business-critical when thresholds are reached.

Further Reference

For more information, see:

- Packeteer 'The Intelligent LifeCycle for Networked Applications' Whitepaper—
  http://www.packeteer.com/resources/prod-sol/Intelligent_LifeCycle_Introduction.pdf
- Packeteer 'Best Practices for Application Traffic Management' Site—
  http://www.packeteer.com/support/BestPractices
- Manage Voice and Video Sessions
  http://support.packeteer.com/documentation/packetguide/8.0/solutions/app-control/manage-voip.htm
- Control Streaming Media
  http://support.packeteer.com/documentation/packetguide/8.0/solutions/app-control/streaming-media.htm
- Enhance MPLS QoS
  http://support.packeteer.com/documentation/packetguide/8.0/solutions/app-control/enhance-mpls-qos.htm

Packeteer, Inc.
10201 N. De Anza Blvd.
Cupertino, CA 95014
United States

T + 1.408.873.4400
F + 1.408.873.4410
www.packeteer.com