THE BROADBAND SERVICE OPTIMIZATION HANDBOOK

Allot communications

# CHAPTER 8

## VoIP: A Major Opportunity

While VoIP has existed for over a decade, its recent rise in popularity has made it a major opportunity for service providers and a must-have offering, especially in light of the need to increase revenues. VoIP is taking off: According to a recent Infonetics Research report, worldwide residential and SOHO VoIP subscribers are forecast to double between 2005 and 2006 to 47 million.

Many of the major broadband providers have included it in their service portfolios, and several free or low-cost VoIP alternatives are options for anyone with a broadband connection. Meanwhile, the telcos that have provided traditional phone service over the public switched telephone network (PSTN) are scrambling to build a VoIP option into their service packages.

If you are planning to get on the VoIP bandwagon or even if you already have it in your service portfolio, you should be aware of the challenges of providing VoIP to your subscribers. This chapter explains how using the three basic intelligent IP service optimization steps can help overcome those problems. But first, we will take a look at the challenges of VoIP service provisioning and maintenance.

### VoIP Challenges the Basic Nature of IP Networks

Circuit-switched voice networks have set the standard over the years for high-quality and reliable voice communications. While VoIP can offer significant savings and flexibility – such as its ability to be integrated with other applications for multimedia sessions – its main challenge lies in providing the quality and service expected by subscribers accustomed to circuit-switched voice communications. The IP network environment was originally designed for data applications that were fairly tolerant of delays. Unlike e-mail, file transfers, Web surfing and other data-centric applications, VoIP is extremely delay-sensitive and is also affected by other network metrics such as jitter and packet loss.

VoIP has small bandwidth requirements — thankfully — but the bandwidth that a voice session does consume must be consistently available. Its vulnerability to ***network congestion*** and ***malicious traffic*** makes VoIP a prime candidate for degraded user quality of experience (QoE) and, as a result, increased customer churn. The quality of VoIP is affected by all factors that can degrade voice communication and result in garbled, delayed, or dropped sound. These factors include:

- ***Latency*** — The length of the delay between the transmission and reception of a voice packet
- ***Jitter*** — The variation in the amount of latency. Jitter may occur because of network congestion, improper queuing, and configuration errors, for example.
- ***Packet loss*** — The dropping of packets due to router congestion, full queues, or packet corruption. Note, though, that a few dropped packets are more desirable than packet retransmissions, which prolong latency and result in degraded conversations.

These factors derive from the network congestion and malicious traffic potentially faced by all service providers. Intelligent IP service optimization can address both situations and minimize the three main VoIP-degradation factors.

### Intelligently Optimizing VoIP

We saw in Chapter 4 how intelligent IP service optimization can prioritize traffic through policy-setting based on a variety of variables, such as service tier, application type, and subscriber category. With the unique needs of VoIP, service optimization becomes critical to managing network congestion. You can guarantee performance and reliability by using service optimization to "nail up" the minimum amount of bandwidth needed by a given VoIP session on a per-call basis and to set a limit on the number of simultaneous calls that can be made based on total bandwidth reserved for VoIP. This capability, often referred to as ***call admission control***, prevents too many calls from degrading the quality of the calls already in progress by stealing bandwidth from each and leaving all underserved. An application in your voice gateway can be used to program a response that a caller might get if “all lines are busy” — such as a busy signal or a recording to “please try your call later.”

We also learned in Chapter 5 that intelligent IP service optimization can be used to identify malicious traffic (in particular, any that originates from subscriber contacts) through deep-packet inspection (DPI) *before* denial of service (DoS) attacks can cripple the network. For VoIP traffic, whose sensitivity to latency is paramount, the slightest disruption of service will mean failure for any subscriber using VoIP at the time of the attack.

The key to delivering reliable and consistent VoIP service that ensures high QoE among all subscribers is to use a refinement of the three basic steps of intelligent IP service optimization we have explored throughout this handbook:

1. ***Establish a baseline*** to know what VoIP protocols run on your network
2. ***Use policies to set priorities*** for the various types of VoIP traffic likely to appear on your network
3. ***Monitor traffic in real-time*** to protect VoIP from network vulnerabilities

Let’s look at how each of these tasks is accomplished.

### Which VoIP Traffic at What Times?

In addition to VoIP packets generated by subscribers to your own VoIP service, other VoIP packets from alternative applications and services will likely find their way onto your network. Even if the VoIP service is not yours, you will be blamed if the VoIP service is not performing as expected on your network. So you must include *all* VoIP traffic in your baseline determination. (See Figure 1.)

**Figure 1. VoIP Protocols on the Network**

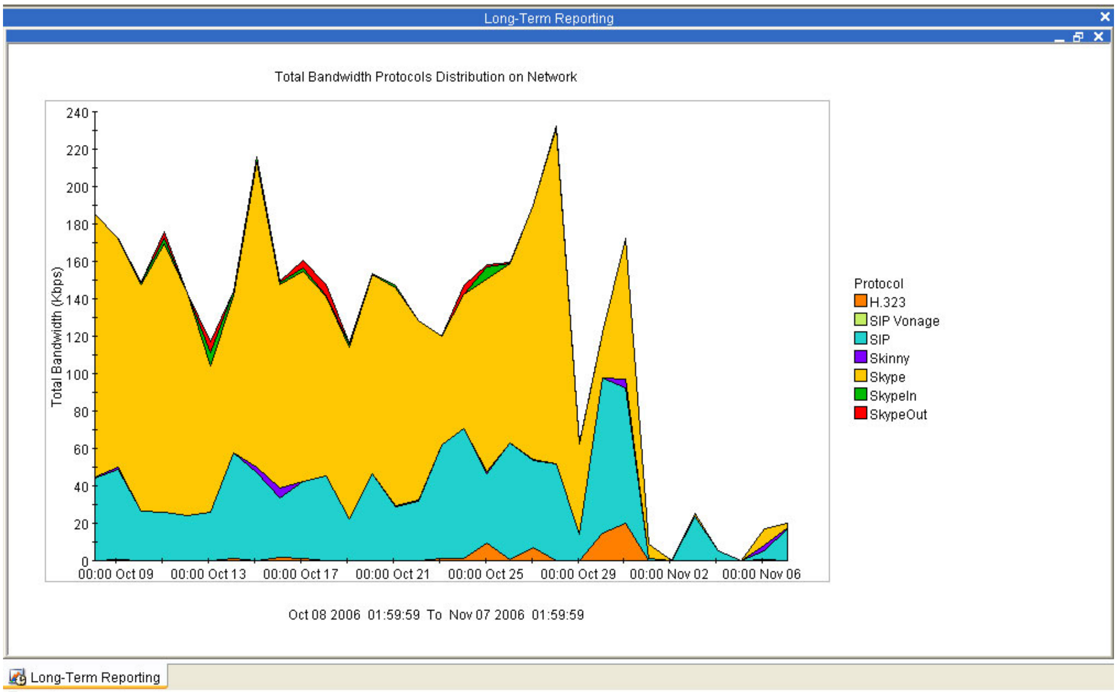

**The figure above shows the two most prevalent VoIP protocols on the network to be Skype and Session Initiation Protocol (SIP). Vonage and other VoIP services use SIP.**

---

You can use this visibility to assess how much VoIP is on your network as well as segment subscribers into usage groups that have different priorities. For example, the highest priority setting might be for subscribers to your own VoIP service offering. You can also use this knowledge to target the VoIP users of alternative services with special marketing campaigns aimed at getting them to switch to your service with the promise of higher-quality VoIP.

Lastly, you can use this baseline to determine how your network managers should manage the traffic on a temporal, day-to-day, hour-by-hour basis through an analysis of how many users are using what VoIP protocols at what times. The heaviest users — the "top talkers" — will shake out as a group pretty quickly, and they can be assigned priorities (and possibly migrated to a higher-tier service) and be put into a specialized subscriber group. High congestion periods can be singled out for special management strategies for those times.

### Give VoIP Traffic the Priority It Needs

Because latency, jitter, and packet loss derive largely from traffic congestion, VoIP traffic must be identified, categorized, and assigned policies that give it its required quality-of-service (QoS) resources. You can use intelligent IP service optimization to set policies that manage traffic on a hierarchical basis that addresses both application and subscriber management issues. (See Figure 2.)

**Figure 2. QoS Assignment for VoIP**

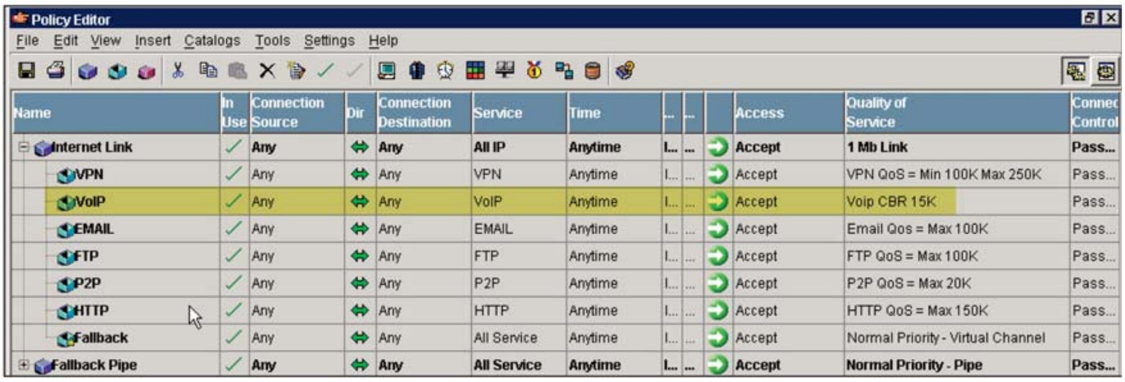

| Name | In Use | Connection Source | Dir | Connection Destination | Service | Time | ... | ... | | Access | Quality of Service | Connec Control |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| ⊟ Internet Link | ✓ | Any | ⇔ | Any | All IP | Anytime | I... | ... | ➔ | Accept | 1 Mb Link | Pass... |
| VPN | ✓ | Any | ⇔ | Any | VPN | Anytime | I... | ... | ➔ | Accept | VPN QoS = Min 100K Max 250K | Pass... |
| VoIP | ✓ | Any | ⇔ | Any | VoIP | Anytime | I... | ... | ➔ | Accept | Voip CBR 15K | Pass... |
| EMAIL | ✓ | Any | ⇔ | Any | EMAIL | Anytime | I... | ... | ➔ | Accept | Email Qos = Max 100K | Pass... |
| FTP | ✓ | Any | ⇔ | Any | FTP | Anytime | I... | ... | ➔ | Accept | FTP QoS = Max 100K | Pass... |
| P2P | ✓ | Any | ⇔ | Any | P2P | Anytime | I... | ... | ➔ | Accept | P2P QoS = Max 20K | Pass... |
| HTTP | ✓ | Any | ⇔ | Any | HTTP | Anytime | I... | ... | ➔ | Accept | HTTP QoS = Max 150K | Pass... |
| Fallback | ✓ | Any | ⇔ | Any | All Service | Anytime | I... | ... | ➔ | Accept | Normal Priority - Virtual Channel | Pass... |
| ⊞ Fallback Pipe | ✓ | Any | ⇔ | Any | All Service | Anytime | I... | ... | ➔ | Accept | Normal Priority - Pipe | Pass... |

**VoIP traffic has been set up with a QoS policy that ensures a constant bit rate (CBR) and 15 Kbps of bandwidth. This will guarantee that the required bandwidth for VoIP traffic is consistently available.**

___________________________________________________________________

Because the bandwidth demands of VoIP are minimal compared to those of gamers and P2P applications, you can assign VoIP traffic very high priorities on a per-subscriber basis using the granularity that DPI allows. In addition to employing call admission control, you can prioritize one VoIP service (say, your brand) over another and still prioritize *all* VoIP traffic over any or all other traffic on your network. By identifying individual subscribers and their packets, you can associate specialized policies down to the smallest granular level on your network, thereby gaining complete control over the behavior of VoIP traffic.

### Protect VoIP Traffic at All Times

With more and more businesses using VoIP, service providers must avoid disruptions in service as never before. To this end, intelligent IP service optimization can help providers identify malicious traffic on the network.

For example, a distributed DoS (DDoS) attack involves sending an anomalous amount of SYN-flood packets. This action will signal red flags in an intelligent network that a DDoS attack has begun and then initiate immediate corrective action, including blocking any worm traffic. Intelligent networks can also use signature recognition to detect other suspicious traffic flows.

This dual approach to traffic protection — both behavioral monitoring of traffic and DPI for identifying, classifying, and prioritizing services — is important to VoIP applications. This is because the packet inspection alone can introduce latency by queuing, buffering, examining, and then re-queuing and transmitting. Behavioral identification as a first line of defense mitigates this latency problem.

There are other aspects to protecting the integrity of VoIP. For example, the User Datagram Protocol (UDP), rather than TCP, is nearly always used as the underlying transport protocol for VoIP traffic. The reason is that UDP reduces latency by not requiring acknowledgments between communicating devices and not retransmitting dropped packets. These functions, which ensure integrity in data communications, add too much delay to a VoIP session for a conversation to be intelligible.

It’s important to note here that some service optimization systems use QoS techniques that apply only to TCP and do not address UDP. If you wish to offer VoIP services, though, it is important to use a system that can apply effective QoS actions to UDP streams, as well.

### Chapter Summary

You can use intelligent IP service optimization to protect the integrity of VoIP traffic on your network, including both your branded VoIP offerings and other VoIP services. You can do this in some service optimization systems by prioritizing VoIP traffic on a per-flow basis.

In addition, using call admission control ensures QoE for calls in progress, and applying QoS techniques designed for UDP-based traffic reduces VoIP latency. As an indirect result of reduced latency, you minimize VoIP's other nemeses: jitter and packet loss.

Finally, intelligent IP service optimization protects the integrity of VoIP services by managing traffic congestion and recognizing anomalous signatures or protocol behavior. Once anomalous traffic is recognized, it can be blocked to prevent DDoS attacks that could disrupt your voice service.

###