

SPEAKEASY
QUALITY OF SERVICE:
VQ TECHNOLOGY

August 2005

TABLE OF CONTENTS

Overview 1

Maintaining VoIP Quality of Service..... 1

The Building Blocks of Voice Quality (VQ) Technology..... 1

Differentiated Services Infrastructure 2

How VoIP works for SOHO (Small Office/Home Office) 2

 Network Internet routing protocols 2

 An Analogy of Per-Hop Behaviors 3

Business VoIP—more complex, same quality of service 4

Summary 5

OVERVIEW

Speakeasy's business VoIP service provides an integrated voice and data solution that delivers voice and data over a single broadband connection. With Voice over Internet Protocol (VoIP), voice is digitized and transmitted in Internet Protocol (IP) packets rather than using traditional circuit-switched Time Division Multiplexed (TDM) resources of the public switched telephone network (PSTN).

As a further advance, Speakeasy's Voice Quality (VQ) technology prioritizes voice packets over data. Because Speakeasy has a private, dedicated IP network, it provides dependability and performance, which allows Speakeasy to maintain a voice Quality of Service (QoS) guarantee from the desktop to the PSTN in a manner not possible for most providers.

This paper addresses issues related to Quality of Service and Speakeasy's mechanism for handling each related issue.

MAINTAINING VOIP QUALITY OF SERVICE

As business managers think about cutting costs and gaining efficiencies with Hosted VoIP, IT managers think about quality and reliability of service. One of the biggest components in choosing a reliable Hosted VoIP vendor is determining how they will maintain Quality of Service (QoS).

Speakeasy's Voice Quality (VQ) technology solves the problems of keeping voice service over the Internet reliable. By prioritizing VoIP traffic at all the constraint points of WAN links between Speakeasy and customers, Speakeasy is able to provide a high quality and reliable service. This is done by smoothing traffic at the traditional bottlenecks: Toward the customer WAN and through use of router prioritization. Application Layer Gateways (ALG) are used in a similar way to manage the traffic flow from the customer to the network.

Speakeasy addresses the issue of customer WAN circuit congestion with a variety of techniques. Outbound voice traffic from the customer is protected with the use of an Application Layer Gateway (ALG) that is capable of ensuring that voice traffic receives priority through the LAN router. Additionally, Speakeasy edge routers are configured to prioritize voice traffic more favorably than data traffic when crossing the WAN link to our customers, frequently, the largest and most common bottleneck. Finally, the Speakeasy Network directly connects to all the voice network equipment through private interconnects so no voice traffic destined for the PSTN is passed over a third party network or any public peering points. This ensures consistent quality voice service.

THE BUILDING BLOCKS OF VOICE QUALITY (VQ) TECHNOLOGY

Historically in IP networks, there have been two major fields of thought when it comes to maintaining Quality of Service: Differentiated Services (DiffServ) and Integrated Services (IntServ). Speakeasy's Voice Quality technology is based upon the DiffServ model rather than the IntServ model. The primary difference between the two is how they deal with the problem of resource allocation.

IntServ very closely resembles the way the Public Switched Telephone Network (PSTN) works. An attempt is made to reserve capacity from end-to-end and then to deliver service if that reservation is successful—your call either goes through or you get a fast-busy signal. Speakeasy ruled this out as a viable option for two reasons: signaling and the economics involved.

For IntServ to essentially 'guarantee' that resources will be available from end-to-end, there needs to be a way for the set-up of the call to occur end-to-end. The Resource Reservation Protocol (RSVP) was created to address this problem. In an IP network without end-to-end signaling such as that which RSVP provides, maintaining the 'guarantee' that IntServ requires is difficult. While access networks such as Speakeasy could deploy RSVP, it would also have to extend the protocol all the way through partner networks who provide the IP to PSTN gateways. Today, such cooperative examples of interprovider signalling for purposes of traffic prioritization are nearly nonexistent. Additionally, IntServ essentially recreates a circuit switched network on top of a packet switched network, which removes a lot of the

statistical multiplexing benefits of the packet infrastructure. This overlay quickly erodes many of the benefits of a converged network.

The primary advantage of DiffServ is that the signaling protocol is not required. DiffServ does not attempt to solve the problem from end-to-end, but rather on a per-hop basis. Speakeasy considers the DiffServ model to represent a better vehicle to implement converged services across a single broadband connection, as the statistical multiplexing characteristics of the packet network are not eroded

DIFFERENTIATED SERVICES INFRASTRUCTURE

Fundamentally, there are two main concepts in any sort of a Differentiated Services infrastructure:

- Classification
- Scheduling

Classification takes place on the ingress to the domain and ultimately is only an attempt to identify and mark the different forms of traffic so that it can be mapped into the correct per-hop behavior.

Scheduling is where the resource allocation actually takes place; i.e., "I am going to give more buffer capacity to per-hop behavior 'A'." That is, Speakeasy is going to ensure that it will attempt to service the queue for per-hop behavior 'B' at a higher frequency than 'C' or 'D', etc.

HOW VOIP WORKS FOR SOHO (SMALL OFFICE/HOME OFFICE)

Before any sort of per-hop behavior can be applied to a given packet, it must be classified. For Speakeasy Home VoIP customers, both the core routers and Analog Terminal Adapter (ATA) perform classification. Speakeasy's core routers classify traffic coming from its voice gateways and mark the header in such a way so that the voice per-hop behavior can be queued appropriately through Speakeasy's network. The customer's ATA also performs classification. Any voice traffic that is converted to IP is considered to be in the voice per-hop behavior.

Once classified as a packet carrying voice data, Speakeasy 'protects' the voice signal from regular data flows. The customer's ATA shapes the traffic that is sent across the DSL line before it hits Speakeasy's network. It is important to note that ALL traffic must be passed through the ATA for it to have the desired effect of scheduling the voice traffic before data traffic.

The other area where scheduling comes into play before a packet is sent down a client's DSL line is on the Speakeasy Edge Router itself. Also called an access router, the edge router sits at the periphery of a network, in contrast to a core router that is in the middle of a network. Edge routers and core routers are somewhat relative terms, but may also indicate that different vendors or models of equipment are being used. All are routers, but of different size and capacity. They may have been built from Application Specific Integrated Circuits (ASICs), which are optimized for different tasks. The Edge/Access router is where the customer aggregation takes place.

The scheduler on the Speakeasy Edge Router in the default configuration divides all network traffic equally. This includes buffer capacity as well as how frequently a packet can be transmitted through the interface. For example, if we have an OC3c with 1000 customers bound to it, then each customer will be serviced 1/1000th of the time and have 1/1000th of the buffer pool available for his/her use. In this case, everyone is treated with classic best effort forwarding.

However, with Speakeasy VQ, we have defined four per-hop behaviors, so that the customer now has four differentiated behaviors. These four behaviors are:

- Network Internet routing protocols
- Voice—Packets from Speakeasy gateways
- OnNet—Packets from Speakeasy address space
- OffNet—Packets from the Internet

For every client that is VQ enabled, there are four different per-hop behaviors, each of which is modified from the default case for maintaining different service levels. For the purpose of preserving a quality VoIP conversation, the most critical are Voice, the valued resource to be protected, and OffNet, the large volume of traffic competing for resources. An earlier reference was made that should 1000 users be attached to a given aggregation circuit, at any given time, each would have access to 1/1000th of the resources. Should one of those users subscribe to Speakeasy's voice service, then the scheduling frequency is no longer equal, but looks like this:

	Non-voice user	Voice user
Voice	-	256
Data	1	1

The numbers reflect the frequency at which the interface scheduler will look in each of the queues to see if there is any data waiting to transmit. As such, VQ is essentially Speakeasy's way of 'stacking the deck' in favor of voice packets. Assuming 1000 users per backhaul, the resource allocation looks like this:

Any data user	1/1255
VQ user 1 (Data)	1/255
VQ user 1 (Voice)	256/1255

This means that each time VQ user 1's data queue is checked, his/her voice queue will be checked 256 times. The important thing to note here is that the voice traffic will not be sent until the queue is checked, and the voice queue is checked more frequently than any of the data queues.

An Analogy of Per-Hop Behaviors

Imagine that you are student going on a school field-trip from Seattle to our nation's capitol. Because of the distance involved, flying is the obvious choice rather than driving. At school (a non-DiffServ enabled network), everyone is equal, so when it comes time to board the bus for the airport everybody lines up and files in one by one. Getting off the bus and heading into the main terminal is the same—the bus empties in the same order as it filled, First In, First Out. Of course, now that you are not on school property (a DiffServ enabled network), things are slightly different for some travelers.

Several have frequent flier miles to get an upgrade to business class—a different behavior aggregate that triggers an entirely different per-hop behavior. At the ticket counter, people flying business class are in a MUCH smaller queue (lower latency). Here, the ratio of airline agents to assist business class travelers vs. those flying economy is much more favorable (scheduler weighting). In the economy section, the lines are long and there are many directional ropes to traverse (larger buffer capacity).

The loud speaker announces, *"Last call for boarding flight DCA_1776 at gate J4."*

At this point, it is quite possible that those with business class tickets are already in the gate area, or perhaps even on the plane as they were processed through the lines much more quickly. Others, who are flying with an economy ticket, may miss their flights altogether due to having to wait in queue for so long (packet loss, similar to old world order TTL expiration). In this case, those with economy tickets have no choice other than to look for a later flight (re-transmit).

In Speakeasy's view, queues at the airport are very similar to the DiffServ model, particularly when considering the elasticity of resource allocation. Travelers may have noticed that the helpful airline attendants working with business-class customers will assist those in the economy if they are otherwise unoccupied. The same can happen in DiffServ where resources from one per-hop behavior can be "borrowed" from another if there is a momentary surplus.

BUSINESS VOIP—MORE COMPLEX, SAME QUALITY OF SERVICE

The principles that apply to Small Office/Home Office (SOHO) VoIP hold true for business VoIP; however, another scenario exists. It is possible for the IP data from one phone handset to go to another directly without passing through a voice gateway. As a result, multiple classification possibilities exist. Additionally, there are times where a packet's per-hop behavior will change as it passes through the system.

From a DiffServ perspective, calls from our customers to and from the PSTN are handled in an identical fashion in both the Business and SOHO VoIP products. At the customer's location, there is still a need for a packet shaping device to assist in protecting the customer's voice traffic. As before, this device is primarily there to prioritize the voice packets that are sent to Speakeasy. However, the device used for business voice also has some common features of IntServ functionality. The device deployed at the edge of a customer's LAN is an Application Layer Gateway (ALG), called an Edgemark. The ALG is able to perform call admission control as a way to help ensure adequate resources exist end-to-end to sustain calls. If enough resources are available to sustain the voice call end-to-end, then the call goes through. If not, the user gets a fast-busy signal. The ALG does not implement this exactly, as the end-to-end Resource Reservation Protocol (RSVP) needed as the basis for IntServ was never widely adopted.

The ALG does, however, know how many active calls are in progress. Remember that all of the phones register themselves with the ALG, and the ALG is in their forwarding path. The ALG is also configured with a specific number of calls the provisioned network capacity should be able to sustain. With these two pieces of information (number of active calls in progress and number of calls configured), Speakeasy can determine if a call has a chance of going through. In the event that a new call pushes the call capacity over the limit, then the ALG has the ability to disrupt the signaling. In this situation, the caller gets a busy signal of sorts, but the quality of the existing calls that are being forwarded through the ALG is not degraded.

It would be ideal if the network never dropped a packet and the customer only wanted to use their circuit for voice traffic. However, customers usually want to converge their infrastructure to run both voice and data over the same links. The ALG does not have the ability to control non-VoIP based traffic inbound to the customer's location: It cannot fully prioritize traffic entering the customer's LAN in a fully converged voice and data environment. This is where the Speakeasy edge router comes into play.

In the table below, there are three possible sources/destinations from which the VoIP traffic, or any traffic for that matter, can flow. As an aid, here are the acronym expansion and definitions used in the table:

- PSTN: Public Switched Telephone Network (voice)
- **OnNet**: An IP address that is within Speakeasy's address blocks
- **OffNet**: An IP address that is NOT within Speakeasy's address blocks

Sender	Receiver		
	PSTN	OnNet	OffNet
PSTN	N/A	Voice	N/A
OnNet	OffNet	OnNet	OffNet
OffNet	N/A	OffNet	N/A

Under the receiver section of the table, we see what the actual scheduling behavior of the traffic is. Note: there are three receiver options.

Below is a table which displays Speakeasy's relative scheduler weights:

	Non-Voice User	Voice User
Voice	-	256
OnNet	-	64
OffNet	1	1

As before, the numbers represent the relative frequency that a particular queue is serviced.

In the business voice scenarios, the network of classification is:

- Speakeasy Core Routers, which classify traffic coming from the Voice Gateways and mark the header in such a way so that the voice per-hop behavior can be acted upon through Speakeasy's network. [PSTN -> OnNet]
- The ALG classifies voice traffic based upon registration [Voice -> Speakeasy across the backhaul]
- The Speakeasy Edge Router rewrites the IP precedence bits based upon the source and destination of the packet header. If both are within Speakeasy's Classless InterDomain Routing (CIDR) space, then the packet is given the behavioral aggregate for OnNet. Otherwise, it is considered OffNet.

Scheduling looks like:

- The ALG sends voice traffic ahead of data traffic [Voice -> Speakeasy across the backhaul]
- The Speakeasy Edge Router maps all packets based upon their behavior aggregate into the appropriate traffic class and, by extension, queue.

SUMMARY

Built for dynamic businesses, Speakeasy's business VoIP offers an integrated voice and data solution that lowers operating costs, increases business efficiency, and guarantees world-class reliability while delivering crystal clear call quality. Speakeasy's hosted system eliminates the capital expenditures and long-term leasing requirements of traditional PBX systems, resulting in significant cost savings. And with Speakeasy, voice quality and service reliability are never compromised. This solution truly is the technology of the future.

VoIP is a lot more complicated service to maintain than casual web browsing. Before deciding on a managed VoIP solution for your business, be certain that your provider takes your VoIP service as seriously as you do, as Speakeasy does.

Author: Brad Volz, Senior Network Engineer, Speakeasy, Inc.