Route Optimization

Making The Promise of VoIP a Reality

An Internap White Paper

Table of Contents

Executive Summary

Integrating voice and data onto a single network – known as “network convergence” – has unprecedented momentum in the business world today. Companies are attracted to this model not only because the investment and maintenance of separate networks for data and voice communications is cumbersome, but it is very costly as well. The Internet, with its inherent benefits of reach, adoption and cost efficiency, has become the *de facto* standard for this convergence. However, its success is largely dependent on the successful migration of voice applications (i.e., telephony, fax, video conferencing, etc.) onto the Internet, which was not originally designed for voice communications.

To ensure a seamless transition to network convergence, enterprises are turning to Internap. Internap helps customers to achieve the service-level quality and performance guarantees they require in order to migrate their business-critical voice applications to the Internet with confidence. This white paper will discuss the promise of Voice over Internet Protocol (VoIP), its benefits, its most salient hurdles, and how to leverage route-optimization solutions to make the public Internet predictable and reliable for real-time voice communications.

Introduction

The promise of VoIP is very compelling. Opportunistic enterprises can save money along with increasing revenues and customer satisfaction by implementing VoIP-based applications like IP call center and web services. But with these benefits comes concerns about quality of service.

Obtaining adequate levels of network performance, reliability and predictability is the No. 1 challenge for businesses that have deployed or are considering VoIP. Since the Internet was not designed for voice applications; it was originally designed for data applications - like email – it has little or no sensitivity to transmission delays. These data applications are known as “bursty” and do not require dedicated transmission paths through the network. Real-time voice traffic, on the other hand, is extremely sensitive to delay. Packets that do not arrive in time are useless for real-time communication because they do not create a complete and seamless conversation.

Compounding this challenge is the fact that the public Internet lacks accountability. It is comprised of disparate networks and service providers and is not managed by a single, centralized operator to coordinate the flow and quality of interactions. It is for this reason that many enterprises have historically opted for expensive private networks where they could have more control over service quality and maintain tighter security over their voice traffic.

Today, companies can leverage new solutions that use route optimization to alleviate these challenges and make the Internet mimic the reliability of private voice networks and other real-time voice communications.

The Promise of VoIP – What Is Fueling This Growth?

The principal driver behind new VoIP implementations is cost reduction. By migrating voice platforms to the Internet, enterprises can significantly reduce long distance costs. However, equally significant to reducing the Total Cost of Ownership (TCO) of a VoIP investment is the ability to lower the administrative and maintenance costs associated with phone systems. Companies using VoIP today derive a lower TCO by reaping savings in one or all of three general areas: toll bypass, bandwidth efficiencies and operational efficiency.

Lowering VoIP Total Cost of Ownership

Toll Bypass

The first area of savings is the cost of telephone charges associated with the Public Switched Telephone Network (PSTN) versus a VoIP network. A phone call on the PSTN is priced based on distance and duration, whereas a call using a VoIP network avoids or bypasses the PSTN toll charges entirely. This is the origin of the term “toll bypass.” This toll bypass effectively eliminates any and all distinctions among local, long distance and international calling so enterprises save in every calling category. The table below details the difference in how calls are handled on a PSTN versus a VoIP network.

PSTN vs. VoIP: How Are Phone Calls Processed?

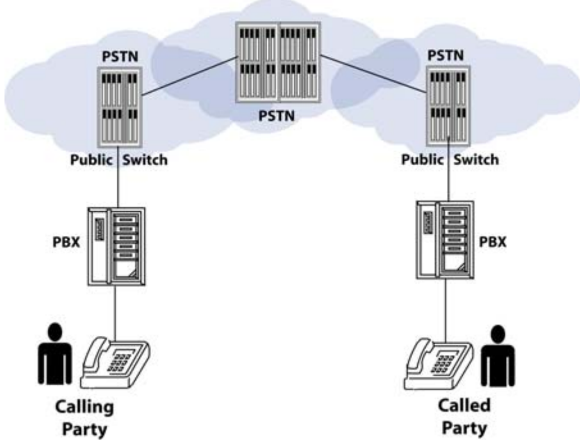

Figure 1: Traditional Circuit-Switched Network

On the Public Switched Telephone Network (PSTN), calls between two parties are set up by a series of private and public switches. The resulting fixed communications link is dedicated for the duration of the call. When an individual makes a phone call over a circuit-switched network, a connection is made between a company's PBX and the local telephone company, also known as the PSTN. Depending on the destination, this connection might extend to the national or international exchange before reaching another local exchange, where it will be passed on to the PBX and the person who receives the call. This end-to-end link, established by a series of public and private switches, is 100% dedicated on a single, per-call basis and cannot be shared or used for another function as long as the call is in progress. For this reason, these dedicated circuits cannot be shared and the carrier bills the call on a time and distance rate.

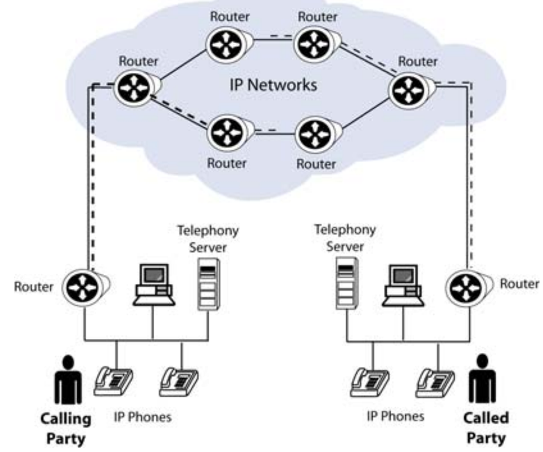

Figure 2: VoIP Networks

The Internet does not use switches to link calling parties. Instead, the analog voice signal is digitized by an Internet Protocol (IP) and broken up into thousands of small data packets by a router - the VoIP equivalent to a switch. These data packets are sent, or routed, over the public Internet to their destination, enabling calls to bypass the PSTN entirely.

In addition, migrating other voice services that are usually deployed over dedicated private lines can secure cost savings, as these voice services are billed on a flat-rate basis. These are private lines between offices where call volume is high enough to justify expensive dedicated circuits, such as Tie Lines, Ringdown and Hoot 'n' Holler circuits. Successfully migrating these voice applications onto the less expensive VoIP network will also reduce TCO.

Major Insurance Company Migrates Hoot 'n' Holler to VoIP

One large insurance company's hoot 'n' holler system linked analysts and brokers from 300 branches so the latest news and changes in market conditions could be broadcast across the company. By deploying modular access routers, it was able to migrate their entire hoot network to VoIP, saving $24,000 per month in dedicated circuit charges. Even after buying the new routers, the payback period was less than one year.

Increased Bandwidth Efficiencies

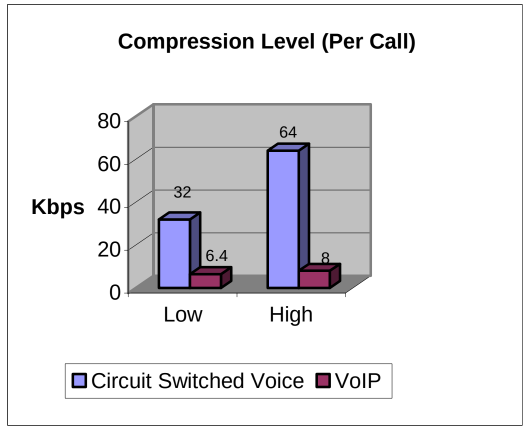

A second factor that reduces costs is the ability to better utilize existing bandwidth and mitigate the degree to which enterprises need to overprovision their network. Since dedicated bandwidth is not necessary to complete each call on a VoIP network, bandwidth can be shared and leveraged for other applications simultaneously. These applications could be other phone calls, faxes, email and video, among others, enabling companies to use bandwidth more efficiently. Because VoIP typically utilizes 85%-90%* less bandwidth than traditional circuit-switched voice communications, it can further reduce TCO. VoIP is compressed - 64 kbps supports eight simultaneous calls on a PBX instead of only one.

A good example explaining VoIP's substantial bandwidth efficiency over PSTN is a typical two-person telephone conversation. Half of the interaction is actually silence - gaps when no one is speaking. So half the dedicated capacity of circuit switched time division multiplexing (TDM) networks goes unused during a typical toll call. Whereas PSTNs are based on TDM technology, in which communication capacity is continuously dedicated to the user - even when the person is not speaking. VoIP, on the other hand, ignores the silence that is inherent in voice communications. Since VoIP capacity is not dedicated, the natural gaps that occur in conversation do not result in unused bandwidth. Given this, using a VoIP network, enterprises can buy only the bandwidth they need without having to oversubscribe.

Greater Operational Flexibility

Lastly, the third component to reduce TCO is the convergence of voice and data traffic onto a single network, eliminating infrastructure and maintenance redundancies while increasing network flexibility. IP technology simplifies interoperability between separate voice and data systems, providing enterprises with the flexibility, extensibility and customization they need while reducing operational costs. In traditional circuit-switched voice networks, the transport, call control and application layers typically are grouped into single proprietary systems. In IP networks, these layers are represented by individual components that can be integrated or substituted as necessary to fit the needs of the overall system, allowing the system, applications and services to be more dynamically designed and managed.

For example, IP phones have physical locations with regular extension numbers, but these locations and numbers are mapped to IP addresses, enabling numbers to be reallocated to another location simply by re-mapping the IP address. This means that VoIP devices can be managed remotely - with most changes taking just seconds to make - without dispatching a technician to the site. This feature of VoIP significantly reduces the costs associated with Moves, Adds, Changes, and Disconnects (MACDs). In addition, because small sites require only a router/gateway, LAN switch and IP phones, bringing a new site on line can be accomplished in hours rather than days, as is typical with circuit-switched technology.

The VoIP Transport Challenge

VoIP Performance Requirements

As discussed, the promise of cost saving makes VoIP an attractive application to implement, but with the lack of dedicated capacity comes the challenge for a VoIP network that meets rigorous quality of service (QoS) requirements to ensure that all calls are of toll quality.

* Wyss,Balz, Voice over Internet Protocol, Microsoft Windows CE White Paper.

Unlike traditional PSTNs, the Internet was not originally designed for real-time voice communications. It was designed for “bursty” applications that are not sensitive to delay and do not require dedicated transmission paths through the network (such as email, or viewing web pages over a corporate Intranet). For this traffic, as long as data packets arrive within a reasonable amount of time, the needs of the end user and the application are satisfied. Voice traffic, however, is extremely sensitive to delay. Packets cannot be used for real-time communication if they do not arrive in time to be part of the conversation. If you have ever suffered through a “choppy” long distance call, or felt like you and another caller were talking over one another while making a mobile call, then you know how frustrating this can be. The bottom line is that voice packets traveling over the Internet cannot be used for real-time communication if they never reach their destination or do not arrive in time to be in synch with the conversation.

To meet the VoIP quality of service requirements, enterprises need to be aware of how the key IP performance metrics – packet loss, latency and jitter – impact VoIP.

Packet loss, the most important component of a good route for a VoIP call, is the percentage of transmitted packets that never reach the intended destination. Packet loss of just 1-2% can affect service, resulting in confusion and frustration for VoIP users. This is because VoIP technology utilizes User Datagram Protocol (UDP) rather than Transmission Control Protocol (TCP) to route packets.[†] TCP is a connection-oriented protocol that automatically re-transmits lost packets in the event that they fail to be delivered, which is ideal for data transmissions. UDP, however, is not connection-oriented and does not re-transmit lost packets. In a real-time communications environment, old voice packets are of no use to callers and only cause confusion if they arrive after the conversation has progressed beyond the point where the packet fits. When UDP packets are lost, callers hear clipping, or short losses of conversation.

The second most important component is jitter. Jitter causes packets to arrive at their destination in uneven patterns, which degrades the call and causes inconsistent voice quality. Depending on the type of VoIP codec equipment the enterprise uses, VoIP can be extremely jitter-sensitive. The codec will deal with jitter by buffering the call, which creates an audible delay. When the delay goes above 100 milliseconds (ms), it is noticeable and begins to interfere with the conversation. If jitter exceeds the levels for which codecs can buffer, the call will begin to clip and may be dropped. Although thresholds vary by equipment manufacturer, jitter should be less than 15 ms.

The last metric is latency. Although thresholds for latency are much larger than for packet loss, latency varies significantly for domestic and international calls. Also, there is a lower expectation on the user’s part when it comes to the delay on international calls. Ideally, latency should be below 200 ms round trip internationally and less than 100 ms round trip domestically. When latency is more than 100 ms, listeners hear a slight pause in the conversation which may be acceptable for international calls but not for domestic. When the delay is 250 ms, conversations become very stilted and experience many awkward silences and accidental interruptions. The biggest contribution to latency internationally is the use of satellite links. A single satellite hop will add up to 500 ms of delay to a round trip connection making the best paths for international calls undersea fiber connections.

[†]Note that some VoIP equipment manufacturers are beginning to use TCP, relying on network service providers to reduce packet loss.

The Internet Phenomenon

As discussed in the previous section, packet loss, latency and jitter are the network components that need to be measured in order to manage VoIP traffic over the Internet and provide a predictable quality of service. Therefore the network solution for enterprises looking to deploy VoIP needs to takes these metrics into consideration when routing traffic over the Internet.

Network Service Providers(NSPs) guarantee these metrics will be met - but only for packets on their network. This is due to the fact that NSPs have peering relationships with other NSPs and exchange packets, regardless of the application, using Border Gateway Protocol (BGP). With BGP, routing decisions are not made based on packet loss, latency or jitter metrics – critical components to voice grade communication – but rather on information that is irrelevant to VoIP quality, like number of network hops between destinations. This routing technology makes it impossible for the NSP to route VoIP traffic based on path performance. In other words, when paths degrade due to latency, packet loss or jitter along the route, BGP is unable to recognize that there is a problem and will continue to send traffic along the path, rerouting only in the event of a hard failure or an administrative policy change.

To overcome this limitation, many businesses attempt to access more direct paths to target destinations by multi-homing. But since the largest NSP controls only 28% of destination addresses, many enterprises need to multi-home with three or more providers, which can result in significant management and network maintenance costs. This is counter to the central promise of VoIP – cost savings. The fact remains that VoIP traffic will still not be routed based on performance (i.e., an email packet will be routed the same as a VoIP packet).

Internap Solutions

To enable migration with confidence, Internap built unique enterprise solutions integrating route optimization technology that provides the stable, predicable, high-performance environment needed to support the sensitivities of VoIP. The combination of technology, infrastructure and unique Service Level Agreements (SLAs) enable Internap to provide its customers with something private and public relationships cannot offer: a differentiated quality of service.

For guaranteed timely delivery of VoIP traffic, the Internap® network-based route optimization solution, called Performance IP, is one way to ensure that voice communication is performing at its highest level. Featuring guaranteed 100 percent operational uptime and the lowest packet loss and latency in the industry, Performance IP uses data lines of various connection speeds to connect a company to an Internap Private Network Access Point (P-NAP®). Internap-patented route optimization technology residing in the P-NAPs accesses the global routing tables being advertised by all of the backbones directly connected to each P-NAP. Rather than relying on BGP to attempt to select the best route based on non-performance-related metrics, Internap intelligent route optimization technology uses empirical measurements to determine current, real-time Internet performance characteristics spanning the eight largest NSP backbones. This real-time observation of traffic and path characteristics enables Internap to route via optimal path selection based on latency and packet loss performance metrics, avoiding potential problems before VoIP traffic is affected. For example, with route optimization technology, Internap can automatically avoid high-latency satellite hops when undersea fiber connections are available to a destination.

Internap Network-Based Solutions

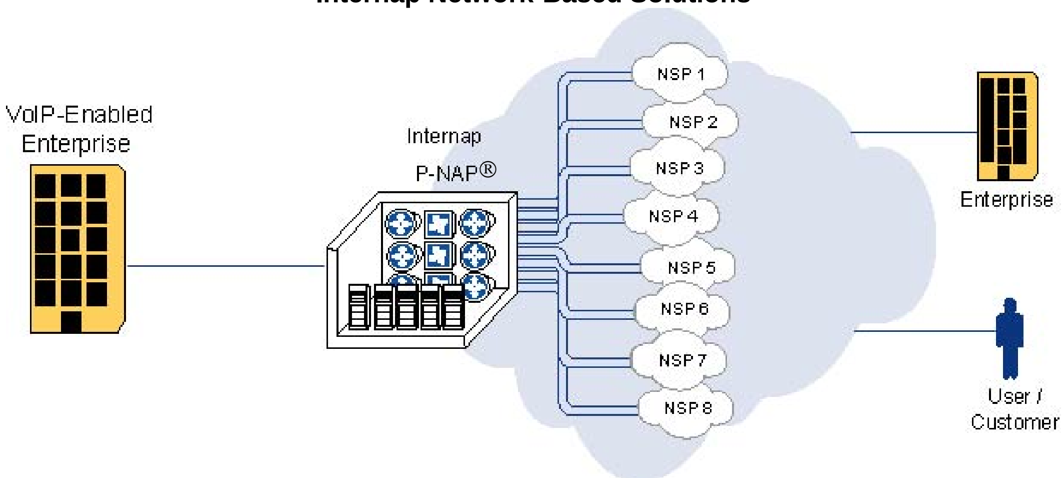

Internap world-class SLAs cover VoIP traffic running over all connected backbones in its multi-carrier network and include several service availability metrics that demonstrate its commitment to being the best performing platform in the industry. Internap guarantees levels of reliability and performance required for VoIP, including 45 ms latency, less than 0.3% packet loss and less than 0.5 ms jitter.

Route Optimization for Other Applications

For applications other than performance-sensitive applications like VoIP or user-based routing, Internap offers a patented-hardware route optimization solution that can direct applications to bandwidth that is consistent with performance requirements and economic value. This solution, called Flow Control Platform™ (FCP), is a client premises-based route-control solution that optimizes traffic flow for performance and cost.

The FCP solution is a network appliance deployed at the enterprise network edge that offers flexible hardware options to optimize IP traffic. By intelligently routing traffic across multiple ISP links to conform to company performance and cost policies, the FCP solution enables control of the network by optimizing bandwidth costs by 20-50%, while reducing latency by an average of 35%.

Internap Certified for VoIP

The International Telecommunication Union (ITU) has developed a numerical representation of voice quality called Mean Opinion Score (MOS). MOS ratings are measured on a scale of 1 (bad) to 5 (excellent), and are derived from research conducted at Bell Labs in which people were asked to document their perceptions of VoIP calls.

NEC Business Network Solutions, Inc. recently conducted an independent VoIP readiness assessment, in which they tested the overall QoS over Internap's IP overlay network between Internap facilities in Dallas, Chicago, Atlanta, New York City, and Seattle. Internap **far exceeded the ITU's performance levels for packet loss, latency, and jitter**, registering a 4.10 MOS score, placing it significantly above both the "acceptable" (3.60) and "good" (4.03) levels.

Using the two solutions (network- and premise-based) together not only enables the most appropriate performing path no matter what the application, but also provides the ultimate in performance and quality of service.

.

Summary

Though the Internet was not initially built for routing sensitive applications such as VoIP, Internap route optimization technology enables VoIP traffic will utilize the best performing path, ensuring the highest level of quality of service. With this, enterprises can reap the cost savings and efficiencies associated with network convergence and deploy a VoIP network with confidence. The promise of VoIP can now become a reality for the enterprise.