"How To" Series

AT&T

Critical Steps for a Successful VoIP Deployment

INTRODUCTION

Enterprises are increasingly deploying Voice over IP (VoIP) delivered via their network-based Virtual Private Networks (VPNs). They expect that voice and data network convergence will help reduce management complexity, drive down operating costs and enable communications to be more efficient. By integrating voice and data on a single IP-based network, businesses can leverage existing LAN (Local Area Network) and WAN (Wide Area Network) investments to create open, scalable VoIP applications, customized to their unique organizational requirements. The convergence of voice and data can be an enabler for new IP telephony features, such as "locate me" and IP soft-phones. However, deploying a real-time application like VoIP across a network designed for data presents unique challenges for network managers. This paper focuses on how to successfully deploy VoIP in a network-based IP communications environment.

Factors Impacting the Quality of VoIP Calls

Users expect the quality of their VoIP calls to be similar to calls delivered on traditional voice networks. This presents a significant engineering challenge, as the IP network must account for the impact of loss of packets, delay and jitter in transmitting IP telephony calls. Packet loss can occur for a variety of reasons, including link failure, Ethernet line errors and high levels of congestion that lead to buffer overflows in routers. Delay and jitter result from encoding, decoding and packetizing voice samples. A combination of these factors can cause voice distortion that makes it difficult to completely understand the remote person.

Importance of a Solid Infrastructure

To mitigate these intermittent quality issues, enterprises need an appropriate WAN infrastructure designed to support differentiated application treatment. MPLS (Multiprotocol Label Switching)-based VPN services provide network managers the control and agility to mark and classify voice and data traffic. This enables them to mitigate the effects of link congestion at the edges of the network and bottlenecks in the IP network. This classification process uses Class of Service (CoS) queues to differentiate the end-to-end treatment of applications, giving real-time applications, such as voice, priority over non-mission critical applications, such as e-mail. As compared to tandem routed hub-and-spoke traffic, MPLS provides an any-to-any network topology that may reduce overall end-to-end latency. This is delivered with the same level of security that traditional Frame Relay services offer. Service providers that offer MPLS-enabled network-based VPNs are key partners in the success of a VoIP deployment.

Quality of Service Vital to Voice Traffic

Ensuring Quality of Service (QoS) is the biggest technical challenge in transitioning to a voice over IP environment, for both the LAN and the WAN. Unlike typical data applications, voice packets are very sensitive to variations in delay and high packet latency. Therefore, networks carrying VoIP must be designed and configured properly to ensure that real-time packets are delivered as consistently and efficiently as possible.

> Ensuring Quality of Service (QoS) is the biggest technical challenge in transitioning to a voice over IP environment.

After establishing a solid network infrastructure and defining QoS and CoS requirements, a holistic end-to-end approach is needed. The critical steps outlined below are a "how to" guide to a successful deployment.

Critical Steps

1. Assess Your Network

A network assessment determines if an enterprise's network is prepared to support both traditional data and VoIP applications. The network manager needs to understand what data applications are currently running on the LAN and WAN, at what times, and whether they are performing adequately. Application types must be differentiated in order to gain a full understanding of individual performance requirements. A variety of tools can be used for the assessment, including data capture software, network management applications and traffic probes. In addition, the voice traffic requirements will need to be derived. For each location where VoIP will be implemented, determine the peak call volume that the WAN will need to support.

2. Select VoIP Codec

VoIP codecs are coders/decoders on the enterprise LAN and WAN that convert analog signals and encode voice packets for transmission across IP networks. There are a variety of codecs supported by different CPE vendors. The most important factor to consider in selecting a codec is the amount of bandwidth consumed per call. G.711 and G.729 are popular voice codecs used in enterprise VoIP deployments; G.711 offers superior audio quality, compared to voice streams encoded using G.729.

G.711 is often used in LAN environments, where bandwidth is not a concern, and requires approximately 80kbps, including overhead. G.729 is often used in WAN environments, where bandwidth is limited, and requires approximately 30kbps, including overhead.

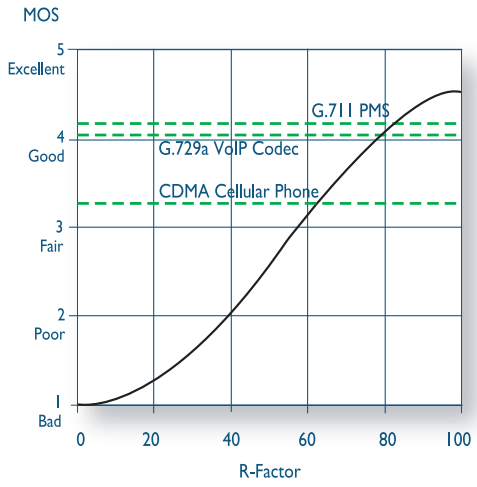

Obtaining high Mean Opinion Scores and R-Factor measurements are critical in ensuring voice quality.

3. Obtain Metrics

End-to-end VoIP simulations should be completed, including measurements of delay, packet loss and jitter. These metrics are often combined in an overall estimate of voice quality such as R-Factor or Mean Opinion Score (MOS). In voice communications, the R-factor provides a subjective numerical measurement [1(worst) to 100(best)] of the quality of human speech when transmitted over a communication network. MOS ranges from 1(worst) to 5(best), and are obtained through a grading process by trained listeners (though some tools will provide an expected MOS based on the measured quality metrics). There are many tools available for conducting this type of testing. The VoIP simulation testing should be based on the

peak call volumes identified for each location. The results of these tests will provide information on whether the existing WAN ports at each location are sized appropriately.

4. Review Results

The results from the previous steps help to determine the answers to the following questions:

- Do the current applications perform adequately across the existing LAN and WAN?
- Do the delay/jitter/packet loss metrics, over the LAN and WAN platforms, meet the thresholds required to support VoIP?
- Does the expected VoIP quality (R-factor and MOS) meet acceptable metrics for the proposed environment?
- Is the WAN at each location sized appropriately to support data and VoIP peak call volume traffic?

If the answers to the above questions are yes, continue with step 5. If not, then it's important to engage the service provider and/or CPE vendor to assess problems and make recommendations.

5. Assign the Right Class of Service

Assigning the right CoS is essential to ensure real-time voice applications do not suffer when contending with data applications simultaneously. An MPLS-enabled service provider, who offers QoS/CoS, can assist you in establishing traffic queues and allocating bandwidth appropriately. Traffic queues become extremely important when congestion occurs in the network.

AT&T's MPLS-enabled VPN network supports CoS traffic queues to differentiate applications: Real-Time, Bursty High, Bursty Low and Best Effort. VoIP applications are classified in the Real-Time class, which gets assigned to a priority queue that gives absolute priority to voice packets over the other classes. This is a "strictly" sized queue and therefore is policed inbound, at the edge of the MPLS network. The bandwidth allocation for this class should be based on the number of busy hour simultaneous calls the network is expected to support. A failure to size this queue properly could result in the inability to place additional VoIP calls or degradation of overall VoIP call quality. Critical data applications – those that run the business or more latency sensitive applications – are classified in the preferred data queues (Bursty High/Bursty Low). All non-critical data or unknown applications should be placed in the default Best Effort queue.

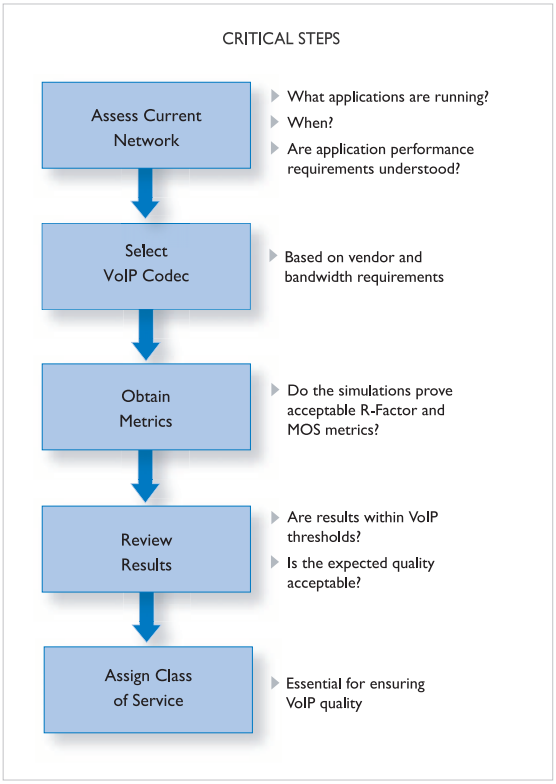

Additional Considerations in Deploying VoIP

Optimizing WAN Bandwidth

Deploying VoIP on low-speed WAN connections presents additional challenges. Even with voice in the priority queue, it is possible for a data packet to be in transmission when a voice packet arrives. On low-speed WAN connections equal to or less than 768 kbps, the amount of time it takes to insert a single large data packet can be enough delay to disrupt VoIP call quality. Fragmentation of large data packets using Multilink Point-to-Point Protocol (MLPPP) can address this challenge and keep data insertion delay to a reasonable level. Compression of headers in the voice packets using compressed Real-Time Protocol (cRTP) can

reduce the bandwidth used on low speed links, further conserving network resources.

Develop a LAN QoS Strategy

LANs have an abundant amount of bandwidth (10Mbps, 100Mbps) and can handle the bursty nature of traditional data applications and real-time characteristics of VoIP simultaneously. The presence of LAN hubs, however, can cause excessive collisions as LAN traffic rates increase. This leads to increases in packet loss and jitter which impact overall call quality. When possible:

- Remove hub devices from the LAN and replace them with 100Mbps LAN switches with high-speed uplink ports.
- LAN switches and router interfaces should be paired together for full-duplex. Duplex mismatches can cause high rates of corrupt packets and degrade VoIP call quality.
- A LAN QoS strategy should be implemented to ensure network-based Trojan and virus attacks do not affect VoIP performance. Many LAN switch vendors support QoS mechanisms at layer 2 (802.1P).

Understand Network Address Translation

Network Address Translation (NAT) is typically implemented in firewalls and routers to "hide" internal IP addresses from the outside world. There are two components to any VoIP call: the call signaling channel (i.e., H.323, SIP) and the media stream. The call signaling channel carries IP addressing and control information and is used to establish uni-directional media streams between two VoIP end-points (one media stream from the source calling party to the destination and one media stream from the destination called party to the source). Within the signaling channel, the destination IP addresses are used to tell each VoIP end-point where to send the media steam.

If the source calling party is behind a NAT device, calls will be established but the source calling party cannot hear the destination party because the destination party does not have the correct IP address to send the media stream to. This problem occurs because the VoIP end-points are unaware that NAT is being performed and the devices performing NAT are unaware that the VoIP signaling data packets needs to be modified. Understanding the inherent problems that NAT brings to VoIP deployments is important. There are solutions available to solve these problems, including the use of vendor-specific IP PBX, firewall, and router-based mechanisms.

Summary

This paper outlines the factors that need to be taken into consideration for VoIP deployments. Using the process and guidelines described in this paper will enable a business to effectively assess the current state of their network and successfully transition to a converged voice and data environment.

About the Authors – AT&T Network Design and Consulting Division

VoIP is one of many complex applications supported by the consultants and analysts of AT&T's Network Design and Consulting Division. Based on years of experience supporting thousands of customers, this team leverages state-of-the-art design methodologies and tools to help customers meet their business requirements through best-in-class network solutions.

For more information, visit AT&T's *Networking Exchange* at www.att.com/networkingexchange/voip.

AT&T
The world's networking company®