Volume 10 Issue 01 Published, February 15, 2006 ISSN 1535-864X DOI: 10.1535/itj.1001.04

Intel® Technology Journal

Converged Communications

Quality Campus VoIP: An Intel® Case Study

Quality Campus VoIP: An Intel® Case Study

Ranjan Sinha, Information Technology, Intel Corporation
Catherine Spence, Information Technology, Intel Corporation
Tim Verrall, Information Technology, Intel Corporation

Index words: VoIP, Campus Voice, QoS, SIP, Host Media Processing, PIMG

ABSTRACT

As IT departments and the communications industry move to a single converged voice and data network solution, the demand on the network to provide specific service levels increases. This demand for higher and differentiated service levels for certain applications has created a need for a reliable and scalable Quality of Service (QoS) model in the network. Prior to the introduction of Voice over IP (VoIP) into a production network, the environment must be equipped to support the reliable transport of voice and a solid end-to-end voice quality plan must be defined and in place.

Typical demands on the network when supporting voice are continuous uptime (99.999%) and consistently low deviations in latency (jitter). As a real-time application, VoIP is far more sensitive to latency, jitter, and packet loss than standard data applications as users have no tolerance for garbled or broken speech. VoIP continues to be compared to traditional telephony in terms of voice quality and reliability as this is the level of service quality to which most users are accustomed.

QoS is a descriptive for defining how IP packets are dealt with through network devices. Mean Opinion Score (MOS) is used to measure the “quality” of a telephone call. QoS on the network helps facilitate a better MOS score for voice on an IP network. QoS is not a standard or protocol, but simply a generic industry term for outlining technologies, standards, and strategies to provide for network quality. In general QoS, to facilitate good voice quality and high MOS, requires that packets carrying real-time voice traffic cannot be delayed and must be prioritized over data traffic, which can better tolerate being slightly delayed. Most jitter in the network is caused by queuing delays associated with momentary or chronic congestion. QoS for voice can help make this queuing delay transparent to the voice packets.

This paper examines a case study based on Intel’s experience of deploying VoIP [2] with voice quality within a campus and converging voice and data on the existing local area network (LAN). The methodology used to enhance end-to-end voice quality includes ensuring bandwidth, enabling QoS and optimizing IP telephony endpoints. This case study demonstrates that even a basic voice quality and QoS strategy can produce high-quality results with minimal infrastructure upgrades.

INTRODUCTION

In this paper we examine a case study based on Intel’s experience of deploying Voice over Internet Protocol (VoIP) with quality voice within a campus where voice and data were converged on the existing LAN using Session Initiation Protocol (SIP). We demonstrate the voice quality methodology used and the benefit it provided. First, we familiarize the reader with the terminology used throughout the paper and provide an overview of the Intel case study. Next, we discuss the network infrastructure elements, voice quality plan specifics, and trial results. Finally, we explore the challenges and solutions and summarize our key learnings.

Background and Terminology

SIP is an Internet Engineering Task Force (IETF) protocol that is used to initiate interactive user sessions with multimedia elements. SIP is specified in IETF Request for Comments [RFC] 2543 [1]. In the Open Systems Interconnection (OSI) communications model, SIP takes place in the Application Layer (Layer 7) and is responsible for establishing, modifying, and terminating the user sessions. In this case study, the user sessions are Internet telephony phone calls.

In a traditional telephone system for a campus, a PBX or Private Branch Exchange is used to provide call switching that is circuit-switched. Functionality typically provided by traditional PBX systems includes local and least-cost call routing, call forwarding, low-density call conferencing, and call detail recording. Most traditional PBX systems use a proprietary digital protocol over a private local area telecommunications network. Cables are run from the PBX to telephone stations in offices and cubicles, often in parallel to data network cables. Voice quality in a traditional voice network is usually a non issue because the end-to-end path is ensured for the duration of the phone call. The dedicated network provides low latency. Jitter is very low since the paths of the signals take the same route throughout the system. Physical faults such as loose connections or cable malfunctions can degrade the quality of the call. However, these impairments are detected by monitoring systems that measure electrical characteristics of the system.

In a VoIP solution for a campus, an IP PBX is used to provide the local telephone system. Because of the greater access to data and the incorporation of open standards, IP PBX systems generally provide the same features as traditional systems with more intelligence and there is greater opportunity to integrate with standard business applications enabling a higher level of automation. Sometimes media processing applications such as voicemail and automated attendant and Interactive Voice Response (IVR) are converged on the IP PBX, making adjunct devices unnecessary. Physically, the IP PBX may be connected to the data LAN or to a private voice LAN. Monitoring call quality for the IP PBX system becomes more important than in a traditional PBX, due to the nature of packetized voice and variability introduced by traditional networking.

MOS is used to measure the *quality* of a telephone call, whether it is a traditional circuit-switched or IP telephony call. MOS is determined by a panel of human listeners in a controlled environment who rate the audio from 1 to 5, with 5 indicating best quality. A MOS of 4 is considered acceptable quality, where 90% of the users are satisfied with the quality of the call.

QoS refers to how IP packets are dealt with through network devices. QoS on the network helps facilitate a better MOS for voice on an IP network. QoS is not a standard or protocol, but simply a generic industry term for outlining technologies, standards, and strategies to provide for network quality. In general, QoS, to facilitate good voice quality and high MOS, requires that packets carrying real-time voice traffic cannot be delayed and must be prioritized over data traffic, which can better tolerate being slightly delayed. Most jitter in the network is caused by queuing delays associated with momentary or chronic congestion. QoS for voice can help make this queuing delay transparent to the voice packets.

Introducing the Intel® VoIP Program

The Intel® VoIP program took place at one Intel site over a six-month period, with over 50 participants who represented a variety of job functions and telephone usage models. The VoIP system was deployed within the enterprise with standard business voice applications, including automated attendant, voice mail, unified messaging, follow-me (call forwarding) features, fax, and remote access.

The IP PBX was built on Intel architecture with additional communication building blocks from Intel. The range of telephone devices used in this trial enabled the demonstration of SIP interoperability of IP endpoints. Additionally, the IP PBX was integrated with Intel's standard Instant Messaging (IM) solution to share basic user presence, along with the ability to launch a phone call.

The project met each of several high-level business goals that were set. Those goals included using and validating some of Intel's VoIP products, giving Intel's IT department the opportunity to introduce a converged communication solution into the enterprise, and demonstrating a phased migration path from a traditional telephone system to a next-generation IP PBX. How the goals were accomplished is outlined below.

- Validate Intel products.

  Through the success of the project, Intel was able to validate that an open standards VoIP solution using SIP can be deployed with Intel products. Intel is maintaining a showcase platform and demo room at the pilot site where interested parties from inside and outside Intel can experience the converged communication solution firsthand.

  The Intel components used in the solution are as follows:

  - The Intel® Xeon® processor is incorporated in the industry-standard server running the IP PBX
  - Intel NetStructure® Host Media Processing (HMP) Software within the IP PBX performs media processing tasks on Intel servers without the use of specialized hardware.
  - Intel NetStructure PBX-IP Media Gateway (PIMG) is an Internet appliance that integrates

---

® Intel, Xeon, and Intel NetStructure are trademarks or registered trademarks of Intel Corporation of its subsidiaries in the United States and other countries.

the proprietary digital telephones with the IP PBX through telephone set emulation.

- Intel® Centrino® mobile technology on laptops runs the IP PBX client software and softphones.

- Introduce converged communications.

  Intel introduced Intel technology into the enterprise for converged communications at three levels. First, applications were converged at the server including base telephone system and media processing applications. Second, voice and data were converged onto the same LAN. Finally, telephone applications were converged at the client including call control and softphone.

- Demonstrate phased migration.

  The production showcase environment demonstrates to customers and others outside Intel how a phased migration from a legacy PBX to IP PBX might proceed. In this case, the IP PBX was deployed behind the traditional PBX. Voice was layered on top of the existing data LAN.

Planning for and monitoring voice quality were critical elements of the trial. Because users are accustomed to the high-quality performance of a traditional telephone system, they have no tolerance for echo, jitter, or delay in VoIP communication. As a real-time application, VoIP requires the same continuous uptime as the system it replaces. High quality is possible, provided careful planning goes into the network design and implementation details.

VOICE INFRASTRUCTURE

In this section we provide an overview of the technical voice infrastructure with an emphasis on characteristics important for voice quality and QoS, which include connecting the IP PBX with the PSTN, gateway integration, and a general discussion of the IP telephony endpoints.

The IP PBX resides behind the traditional PBX, meaning that no changes were needed to provide the connection between the premises and the public telephone network. All calls continue to come into and out of the campus through the existing PBX. The PBX continues to route inbound calls as usual, directly or through the auto-attendant without modification. This configuration enabled quick installation of the new system and removal of the old, with minimal disruption.

To make the connection between the traditional PBX and the IP PBX, a gateway is used to enable translation of TDM circuit-switched calls from the existing PBX to packet-switched calls from the new IP PBX. In this case the direct physical connection to the PBX is through a PRI ISDN T1 interface. The gateway is connected to the LAN through a switch, as shown in Figure 1, making sure that the configured protocols are matched up properly.

Pilot users tested a variety of IP telephony endpoints, including existing legacy phones and IP phones. Most users received IP phones. Legacy phones are connected to the IP PBX using Intel's PIMG as shown in Figure 2. IP phones are connected directly to the LAN.

® Intel and Centrino are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

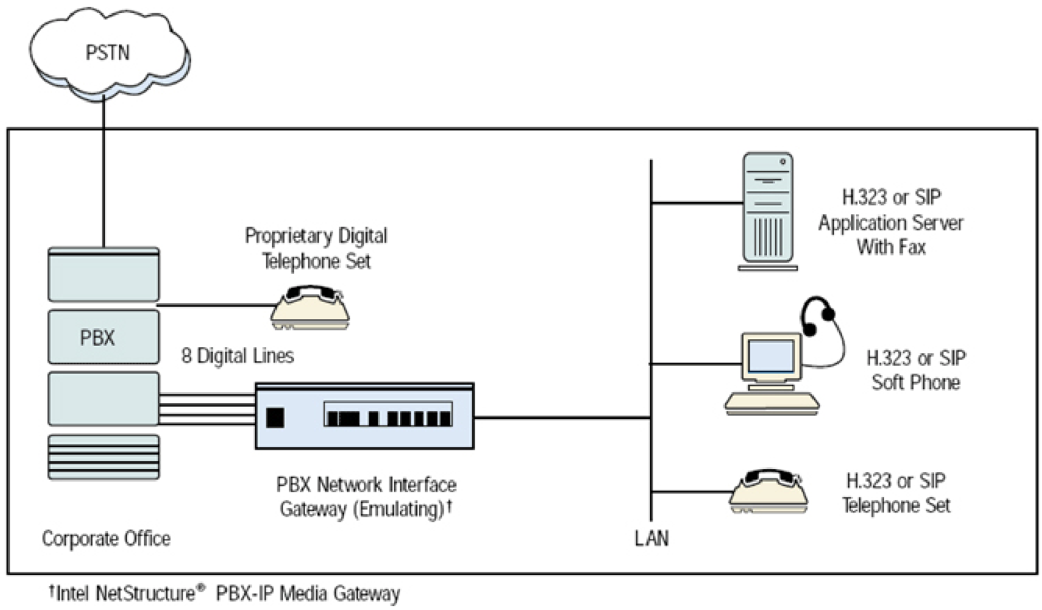

Figure 1: Connecting the IP PBX to the traditional PBX

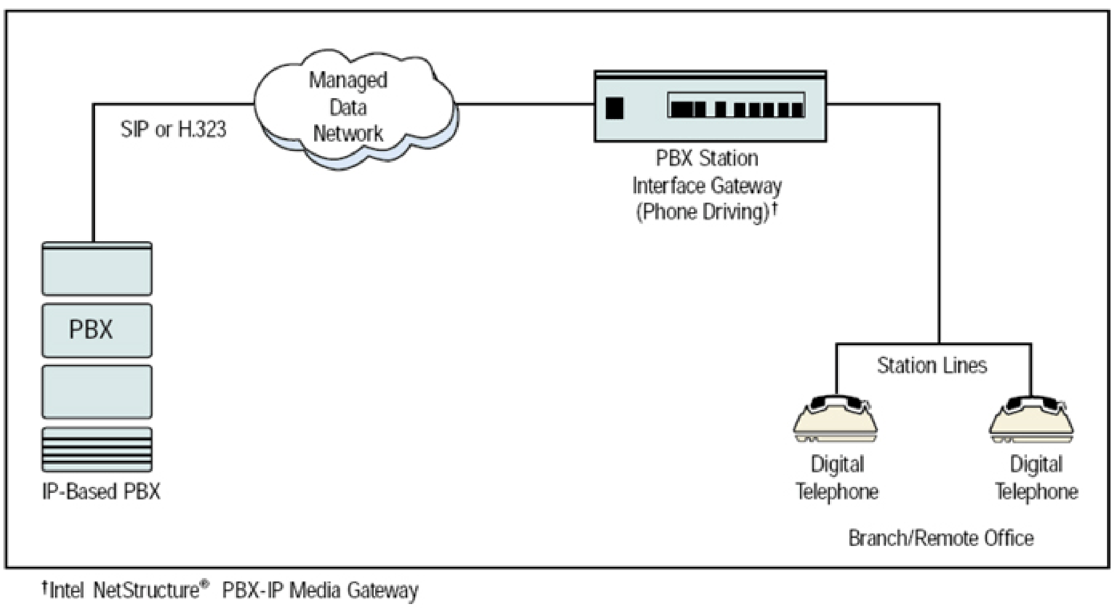

Figure 2: Using PIMG to connect legacy phones

We tested several types of IP-based phones in the pilot, including hardware phones, software phones, and wireless phones. All these were SIP-enabled out of the box. Of users with IP phones, 58% received a hardphone, 28% received a softphone, and 14% received both. Wireless IP phones were tested but not deployed to any live users.

Some hardware IP phones were deployed with Power over Ethernet (PoE). PoE makes for a more streamlined approach since a separate AC adapter for each phone is expensive and takes up valuable outlet space. PoE enables the transmission of electrical power over the same Ethernet cable as voice and data, using industry-standard IEEE 802.3af. For implementation, the access switch had to support the PoE feature, so in some instances the switch was replaced in this trial. An advantage of PoE is that it offers backup power capabilities if power goes out in the facility. Since the PoE cable draws power from the network, power during an outage continues to be supplied via the network's Uninterruptible Power Supply (UPS). This improves the user experience and overall perception of voice quality.

Another important feature of the hardware IP phones with respect to QoS is the ability to do 802.1q virtual LAN (VLAN) trunking and the ability to mark outgoing packets using differentiated services or *DiffServ*. In areas where these phones were on a separate VLAN, the default VLAN was used. In others, 802.1q and DiffServ were used.

Software phones or softphones consist of a software client installed on a PC, connected to a headset with audio and microphone. Good quality audio devices are essential for a good end user experience. The softphones, like hardware IP phones, have the ability to mark the packet with DiffServ allowing those packets to obtain priority on a properly configured network. However, there were other factors that also impacted voice quality. For example, Intel tested wired and wireless headsets from a variety of vendors and found that the selection had a profound impact on call quality and usability.

- **Wireless headsets**–Several types of Bluetooth* wireless headsets were tested and none were able to provide acceptable call quality. Noise and echo were problematic. Further, there were major usability issues concerning power including battery life and awkwardness of turning power on to answer an incoming call.
- **Wired headsets**–Wired headsets come either with their own Digital Signal Processor (DSP) using the PC's USB port, or they use the PC soundcard through the headset and microphone plugs. Users experienced a wide range of quality issues including noise and echo. The more expensive USB headsets with DSP provided the best quality.

VOICE QUALITY PLAN

The plan was very simple and straightforward. The goal was to provide QoS for the voice traffic end-to-end across the IP LAN via a DiffServ model, using Class of Service (CoS) and Differentiated Services Code Point (DSCP) where appropriate. DSCP refers to how the network traffic is marked or *tagged* in the IP layer (Layer 3) to identify it as part of a particular class. Voice packets are marked with DSCP (DiffServ code point) on the end device; this marking is mapped to a CoS tag in the LAN switch. CoS refers to categories of network traffic that are recognized and queued by the network switches. CoS marking or tagging occurs in the Layer 2 header for the 802.1q trunks, and this method is referred to as 802.1p DiffServ, as the overall model defines how these packets are queued and forwarded in the network. Based on the policies set, network switches use CoS and DSCP values to provide high-priority handling for voice traffic.

The life of an audio packet as it is transported over an IP network is given here:

- Human speaks
- Coded with some form of compression
- Inserted into packets with sequence numbers
- Transported onto the network
- Received into network equipment
- Network determines next-hop interface and switches the packet
- Packet is queued for sending to the next hop
- Received in a playout buffer
- Decoded in sequential order
- Played back to human

As packets traverse the network, all elements must recognize which packets are real-time audio and prioritize their handling, yet ensure that no uplink, or any other portion of the network, becomes over utilized by the addition of voice on the network.

Another factor is the coder-decoder or *codec* used to compress and decompress voice transmission. G.711 is a standard for speech codecs that provides toll quality audio at 64 Kbps. Intel elected to use uncompressed G.711 as a codec in this trial since it provides the highest quality available for the system implemented, and bandwidth utilization is not an issue on the LAN.

LAN DESIGN

In this section we provide the architectural details on how the LAN was organized and discuss how we supported the QoS plan. Our goal was to ensure that as little jitter or loss as possible was introduced into the network and that QoS was provided for all voice calls over as much of the network as was possible. Small amounts of jitter could be compensated for in the receive buffers. Small amounts of loss are handled by the codec, although this loss was very little.

The network was designed using Ethernet switches with dedicated full duplex ports for each endpoint device. The LAN design (Figure 3) consists of the core, distribution, and access switches. All the routing is done in the core; the distribution acts as the aggregation point for all the

* Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

access switches out in the communication closets, and the access layer is the connection point for all the users. In this particular design, we focused on four communication closets spread out across the campus. All connectivity between switches is Gigabit Ethernet linked by fiber. The access switches all deliver 100 Mb switched to the desktop and phones.

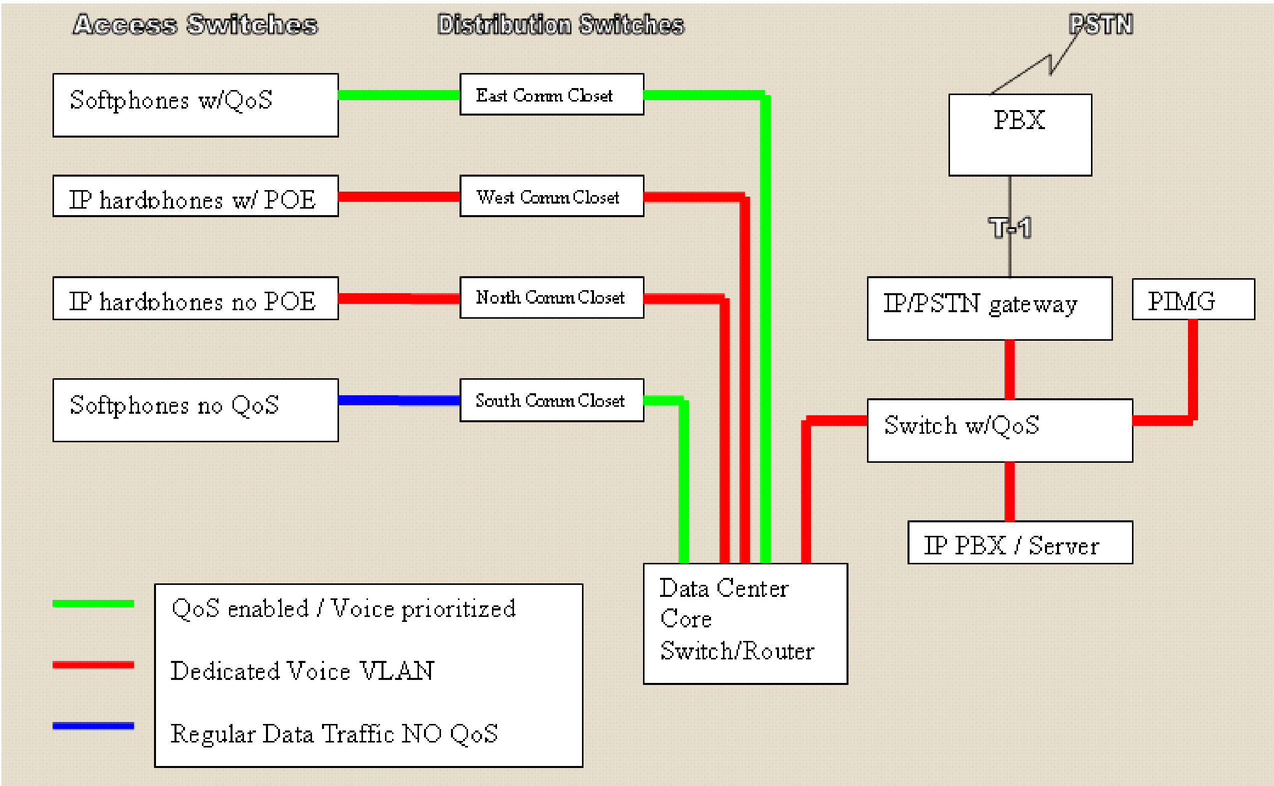

Figure 3: LAN design diagram

A switch located in the data center connects the IP PBX, the Gateway between the IP PBX, and the traditional PBX and the PIMG. All these devices are on the same VLAN, and QoS was enabled for any voice traffic in this segment. This allows for the end-to-end QoS for all the endpoints except for one communications closet.

In the network supporting the pilot user community there is a mixture of PoE and non-PoE access switches in the communication closets. There are four closets: North, South, East, and West, providing access for four different types of use cases identified:

1. Softphones with QoS: All the users in the East communications closet were using softphones on their laptops with headsets. For these users, it was essential to prioritize the voice traffic over their data traffic. This meant enabling QoS on the switches end-to-end from this closet. At the client, we installed the packet scheduler and made sure that the softphone client is marking packets.
2. Hardware IP Phones with PoE: All the users in the West communications closet were using hardware IP phones with PoE. The access switch in the West closet had to be capable of providing the industry-standard 802.3af PoE. The QoS mechanism used in this case was to segment all the hardware IP phones onto a dedicated VLAN using 802.1q and to use priority queuing for voice, to prevent contention with the data traffic.
3. Hardware IP phones without PoE: All the users in the North communications closet were using hardware IP phones but did not have PoE capability. These phones use AC adapters. The QoS mechanism used is the same as that used in the West communications closet including segmented VLANs and priority queuing.
4. Softphones without QoS: The South communications closet did not contain access switches that were capable of prioritizing voice traffic before forwarding it on to the distribution switches. Therefore, we deployed softphones without QoS to any users serviced out of the South closet. This model enabled us to contrast voice quality for users that had some form of QoS mechanism in place end-to-end versus those users that did not have end-to-end QoS.

In addition, we tested Wi-Fi phones in the lab with a dedicated access point so there was no other traffic but voice. The voice quality results were good but we did not perform extensive testing since we had a bigger problem to solve: secure device authentication was not ready in time to deploy users with phones. We decided instead to push Wi-Fi phone deployment to the next phase of the project.

Where the switch was capable of 802.1q and DiffServ, the PCs were also plugged into the switch of the hardware IP phones. This allowed the PC to be set up on a different VLAN than the hardware IP phone. This approach is excellent for a campus with wire runs of one port per cubicle/office or in cases where limiting the number of access switch ports is required to support both voice and data. With this approach the same functionality of using two different network ports is realized at half the real estate.

In summary, the philosophy behind the design was to ensure QoS on a link-by-link basis end-to-end. We enabled DiffServ, which enabled the prioritization of voice traffic over data traffic anywhere that there might have been contention for the network. VLAN separation was also enabled to provide easy differentiation and greater security of the voice network.

The uplinks were all Gigabit and lightly utilized, so we were able to ensure that enough bandwidth for the voice calls was available. Figure 4 shows the portion of the network with DiffServ enabled and how the prioritized packet travels through access, distribution, and the core.

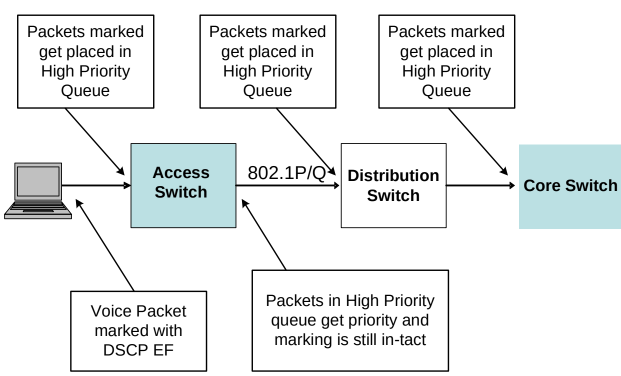

Figure 4: End-to-end packet flow

Figure 5 shows the relationship between the voice quality and various components.

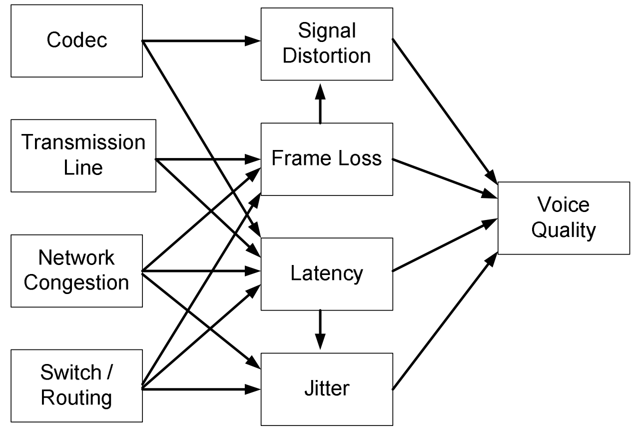

Figure 5: How voice quality maps to components

QoS must be comprehended by all IP telephony end points. All IP endpoints in the pilot were QoS-aware, including all IP phones, softphones, and Wi-Fi IP phones. The following sections describe the QoS considerations we followed for each type of endpoint.

Hardware IP Phones

We used a virtual VLAN segmentation (IEEE 802.1Q) to create a separate broadcast domain. That allowed us to segment voice traffic from network broadcast messages and data traffic. We used a separate VLAN for each supported network access switch. In other words, there was a VLAN for each communications closet.

Softphones

For softphones, we enabled end-to-end QoS in every component in the network. We set up the network and all of its components to distinguish voice and prioritize it above other kinds of traffic. The IEEE standards for QoS are 802.1p and 802.1Q, where a Layer 2 tag is used to specify packet priority. In 802.1p, CoS 5 identifies voice traffic. Using DiffServ, the voice packet is marked with Class EF, which gets priority on a QoS-enabled network. Unlike a hardphone, the softphone resides on a laptop or other computer that is probably running other applications. Like the network, the client environment must recognize and prioritize voice, and the selected softphone must support it.

Non-IP Hardware Phones

Part of the QoS plan included identifying and segmenting the voice traffic on the network from data traffic, and giving voice traffic its own dedicated, logical route. We performed that sort of segmentation only for IP hardphones, not for all voice traffic.

TRIAL VOICE QUALITY RESULTS

This section contains the call quality results experienced in this Intel case study. User reported results are discussed as well as management, reliability, and quality.

User Reported Results

Intel did not perform a formal MOS measurement with users. Feedback was collected through channels such as a weekly health check e-mail (yes/no voting buttons to indicate problems), a detailed qualitative user survey, and interviews conducted by a Human Factors Engineer.

Softphone with QoS users experienced good voice quality all the time, no matter what time of day whether in a conference call, outbound, or inbound call. These users were extremely satisfied with the quality of voice and in fact, with properly tuned audio settings on the laptops and good quality headsets, the voice quality was described as an improvement over their regular phones.

Voice traffic for hardware IP phone users had virtually no contention with the data traffic. These users also reported good quality voice. The make and model of hardware IP phones was important as we deployed a variety from multiple vendors. Users were dissatisfied with certain types of phones, so by trial and error Intel could determine which phones worked best in the environment. Once phones with reported issues were replaced, users described the quality of voice as very good at all times.

Where QoS capability between the access switches and the distribution switches was not available, we had some softphone users. These users expressed satisfaction with the quality of voice most of the time; however, they frequently experienced degraded quality of voice, reportedly several times per day. Often this was explained by the fact that heavy downloads were occurring on the access switches at the time of the calls. Without prioritization for voice traffic, call quality was impacted, even with the high availability of bandwidth in the LAN. This confirmed the necessity of enabling end-to-end QoS.

Users with legacy phones connected to the IP PBX through the PIMG also reported good voice quality. The only area that could be impacted through the voice quality plan was at the data center switch, where voice was prioritized and a separate VLAN was configured. The existing wiring was used between the PIMG and the legacy phones, ensuring dedicated connections.

Management, Reliability, and Quality

The reliability of the LAN was 99.98%. This was achieved through pro-active monitoring. There was no physical redundancy for the access switches. The traffic levels on the network were well understood through the baselining efforts and there were no bottlenecks identified. The utilization on the uplinks from the distribution to the core did not peak above 40% and the average utilization remained below 10%. Tools were in place to monitor the health of the network including visibility into the voice traffic that received priority. All management, monitoring, and alerting is centralized, providing end-to-end visibility into the network. From a management standpoint, the ability to push QoS parameters throughout the network from a central location is important in order to effectively manage a large deployment.

KEY CHALLENGES AND SOLUTIONS

Many challenges were found during this project, many were solved and some are as yet not completely resolved. We examine key challenges in this section including user perceptions of voice quality, softphone QoS issues, and voice quality impact resulting from security and reliability issues.

User Perceptions of Voice Quality

When requesting a user community to report all voice quality issues, you can expect many normally unreported instances to find their way into the report. This presents a challenge. The question is how to separate the real quality issues from the minor issues or those quality issues that have always existed. We did this by utilizing good proactive measurements of the network, evaluating the QoS usage and load along with making regular test calls and qualitatively measuring the quality through user feedback.

Softphone Quality Issues

This is an area where a number of challenges exist and are not all solved at this time. Some of these issues relate more to the quality the user perceives for the voice call than anything technical. The application gaining higher priority of the PC resources is one aspect that still needs to be solved, although many CPU management tools exist. The system should provide QoS for VoIP traffic out of the LAN interface card on the PC itself. The QoS packet scheduler on the PC is a step in the right direction, but a user can place all his/her client traffic in the voice (expedited forwarding) queue and so a centrally managed system is required before the client QoS can really be trusted. There were many areas where we found attractive and simple solutions. For example, it was easy to replace headsets where quality and user acceptance was much better when we used headsets with DSP.

Security and Reliability Issues and QoS Impact

There can be many potential security issues relating to a VoIP design and deployment and some of these played a role in this case study. First there was the question regarding what happens to VoIP if there is a worm or a Denial of Service (DoS) attack on the network. This problem was significantly reduced by providing a dedicated QoS protected VLAN for all the hardware IP phones. This keeps the data traffic separate from the voice traffic and therefore would allow the voice traffic guaranteed access to the network during a DoS attack.

Another area of IT concern is privacy. IP phone calls are much easier to tap and listen in on. If end-to-end encryption is used to protect the content then it makes it more difficult for the network to distinguish and prioritize voice. Also, encryption can introduce latency. One possible solution is to not converge voice on the LAN and instead install a private voice network. However, that approach significantly degrades the business value of converged networks.

CONCLUSION

The Intel VoIP program was a very successful trial on how to introduce VoIP into the enterprise using SIP and converging voice and data on the LAN. As the telephony industry moves from vertically integrated proprietary solutions to horizontal open standards architecture, a key challenge is to provide quality and reliable voice service on a mixed IP network. This project implemented a very basic voice quality and QoS strategy, proving that VoIP can be layered on top of an existing infrastructure with minor upgrades.

In terms of voice quality and QoS, the key learnings from this trial were as follows:

Voice Quality Plan: Intel benefited from having a deliberate plan up front. Enabling QoS leads to better and more predictable voice quality. This is a mandatory item for future deployments of any VoIP implementation in Intel. It is important to address voice quality holistically, from voice endpoint to voice endpoint.

Loading VoIP on a production LAN: Adding voice to an existing data LAN was important because it showed that the voice traffic did not impact the LAN functionality. It enabled the extrapolation of how scaling the pilot impacts the LAN from a capacity planning perspective. In addition, Intel received the benefit of managing only one network–one of the most touted benefits of VoIP. One major issue yet to be resolved is with IT security: end-to-end encryption may be required. If that is not possible, VoIP may reside on a separate voice network, which significantly degrades the business value but may improve voice quality.

Echo is a problem and difficult to diagnose: A considerable amount of time was spent diagnosing and fixing echo problems. Many possible sources/causes were identified, not all of which can be easily fixed. Trying to operate with mixed TDM and VoIP equipment makes it harder. Better industry-standard diagnostic tools would aid this process.

Headset selection is important: The appropriate headset selection is important for good voice quality. Wireless and less expensive wired headsets introduced noise and echo. The more expensive wired DSP headsets provided the best call quality.

G.711 Codec: G.711 is currently recommended for campus calls. Let gateways provide any necessary compression.

Drivers for VoIP include improved user mobility, functionality, presence, and productivity. Cost savings is on a case-by-case basis, with ROI and TCO improving as system costs decrease over time. In this case, the more specific business benefits we identified were as follows:

- Networks were simplified by converging voice and data networks.
- Costs related to data center space and telephone moves, adds, changes, and deletes were reduced.
- New or improved telephony applications were made possible due to VoIP’s open, extensible architecture.
- The stage is set for advanced multimedia applications, services and productivity enhancements through converged communications.
- There is an immediate productivity gain from unified messaging, find/follow me, universal phone number, and presence integration.

These benefits can only be realized if acceptable voice quality is achieved.

ACKNOWLEDGMENTS

Thank you to the other members of the Intel VoIP Program Core Team for their work on this project and input to this paper: Dave Golish, Rick Lansford, Steve Sacker, Matt Santaiti, Rose Sheehy, Shelby Siegel, Geri Sunderlin, Gary Veum, Charlie Weaver, Mike Wilson, Michael Mall, Tracy Little, Roman Stobnicki, Scott Trevor, and Joe Walmach.

REFERENCES

[1] Session initiation protocol – RFC 2543, http://www.ietf.org/rfc/rfc2543.txt*

[2] Deploying VoIP – http://www.extremenetworks.com/libraries/solutions/DeployVoIP_WP.pdf*

AUTHORS’ BIOGRAPHIES

Ranjan K. Sinha joined Intel as part of the Dialogic acquisition in 1999 as the Network/Telecom manager. He is currently responsible for the DMZ Operations at Intel.

He holds an M.S. degree in Computer Science from West Chester University, PA. His e-mail is ranjan.k.sinha at intel.com.

Catherine Spence is an IT computing manager for Intel. She has over 12 years experience with computer telephony and enterprise voice. She earned a graduate certificate in Software Engineering from Harvard University and a B.S. degree in Computers and Engineering from Trinity College, CT. Her e-mail is catherine.spence at intel.com.

Tim Verrall is a staff engineer with Information Technology. He is a lead architect and designer for Wide Area Network, network security, and voice networking at Intel. He has over 15 years of network engineering experience. His e-mail is timothy.verrall at intel.com.

For further information visit:

developer.intel.com/technology/itj/index.htm