Ditech Networks

| WHITE PAPER |

Voice Quality beyond IP QoS

January 2007

Contents

Voice Quality beyond IP QoS

Executive Summary

Good or excellent voice quality is a key competitive differentiator for voice over IP (VoIP) offerings among service providers, including incumbent local exchange carriers, Internet telephony service providers, multiple system operators, cable companies and wholesale providers. Conversely, poor voice quality becomes a crippling disadvantage for VoIP services that compete with each other, or with the Public Switched Telephone Network (PSTN).

The most competitive VoIP offerings are those that recognize and address the fundamental difference between voice quality and IP quality. Quality of Service (QoS) provisions for packet-based networks all serve to minimize packet loss, delay and jitter. But even in a QoS-enhanced under-subscribed packet network, VoIP voice quality can be unacceptably poor. And under-subscribing to enhance voice quality undermines the primary advantage of VoIP: the cost-effectiveness that is derived from statistical multiplexing gains found only in packet-based networks.

So while packet-level QoS is certainly necessary, it is far from being sufficient.

VoIP networks need something more to improve voice quality beyond just IP QoS. Many similar technologies have been deployed in the PSTN network, which are increasingly becoming IP capable. These impairments may or may not originate in the packet network itself, but every VoIP call to or from the PSTN or a cellular network experiences most of them, and many VoIP-only calls frequently encounter these issues.

These impairments exhibit several common – and annoying – symptoms. Examples include a VoIP user with a choppy or robotic-sounding voice, callers hearing themselves momentarily after speaking, needing to turn up the volume to hear the other party (or being told you cannot be heard), an inability to interrupt the current speaker (no natural "double-talk"), too much noise to carry on a conversation, or the disruptive sensation that the line has gone "dead" for some reason.

The impairments that cause these and other problems include hybrid and acoustic echo, low signal-to-noise ratios, a lack of "comfort" noise, signal levels that are not normalized for consistent volume, use of low bitrate codecs or flawed transcoding between codec types, and poor intelligibility resulting from inevitable packet loss. Individually, these impairments erode customer satisfaction. Collectively, they cause customer churn and make it difficult to attract new, loyal customers.

Achieving "toll quality" VoIP, therefore, requires eliminating each and every impairment to voice communications, just as is done in the PSTN. Fortunately, cost-effective solutions are now available to satisfy just this need.

This white paper is organized into two main sections followed by a brief conclusion. The first section on MOS Expectations with IP QoS covers the fundamentals: how voice quality is measured; user expectations for voice quality; techniques for providing packet-level QoS; and the difference between IP QoS and voice quality. The second section on Maximizing Voice Quality in Packet Networks explains the cause and effect of the five most significant impairments to voice quality in a VoIP network, and describes how the Packet Voice Processor from Ditech Networks addresses all five in a single, purpose-built system. Although the material is somewhat technical owing to the complex nature of the subject, it is written in a way that makes the document suitable for a business management audience.

MOS Expectations with IP QoS

Because quality has always been of paramount importance in the Public Switched Telephone Network (PSTN), substantial research went into assessing voice quality. Mean Opinion Score or MOS (see sidebar) remains the benchmark means for determining voice quality in the PSTN. While the subjective nature of MOS testing works well in the stable and predictable PSTN, the dynamic and "bursty" nature of IP traffic requires a different approach. For this reason, newer, more objective techniques are now available to calculate MOS ratings automatically based on measured values. Of these techniques, the E-Model specified in ITU-T Recommendation G.107 is both the least intrusive and the most cost-effective.

By measuring various parameters of network performance, the E-Model is able to estimate how the "average user" would likely rate a call on the MOS scale (see table).

Mapping E-Model R-Factors to MOS Ratings

User Opinion	R-Factor	MOS
Very Satisfied	90 – 100	4.1 – 5.0
Satisfied	80 – 90	3.7 – 4.1
Some Users Satisfied	70 – 80	3.4 – 3.7
Many Users Dissatisfied	60 – 70	2.9 – 3.4
Nearly All Users Dissatisfied	50 – 60	2.4 – 2.9
Not Recommended	0 – 50	1.0 – 2.4

The R-Factor scale used in the E-Model leverages extensive research that correlates measurable impairments with subjective, human interpretations of voice quality. The specific equation used to calculate the R-Factor takes into account the following: signal-to-noise ratio (SNR) of voice volume to background noise; speech impairments, such as level mismatch and codec distortion; talker and listener echo; and codec tolerance for packet loss. An additional advantage of the E-Model is that a detailed analysis of the R-Factor can be used to help isolate the root cause of persistent problems.

Voice Quality Expectations

Customer expectations regarding voice quality exist at two levels. The first is the PSTN, where the expectation is for "toll quality" voice with a MOS rating of 4.0 or more. And indeed, most PSTN calls achieve an "Excellent" MOS 4.3 rating – about as good as it gets for narrowband speech. The second level has been established by the cellular phone network, where the expectation is considerably lower (around MOS 3.2–3.5). Of course, cellular customers willingly tolerate this "Fair" voice quality for the convenience of mobility.

VoIP networks generally compete with the former, and to be competitive with the PSTN, business and residential VoIP services should strive for a MOS of 4.0 or more. This is easier said than done, naturally, because the PSTN has three design characteristics that together maximize voice quality:

- The use of dedicated time-division multiplexing (TDM) circuits to eliminate contention and minimize latency end-to-end in all connections
- The use of Pulse Code Modulation (PCM), specified in ITU-T Recommendation G.711 at a full 64 kbps (the PSTN's standard DS-0 circuit), for its ability to achieve a MOS rating of 4.4

- Painstaking measures to remove all impairments that degrade voice quality, including eliminating hybrid and acoustic echo, reducing noise or adding "comfort" noise, and normalizing signal levels for consistent volume – all of which serve to maximize intelligibility (These issues are all covered later in the section on Impairments to Voice Quality in a VoIP Network.)

As a result of these and other provisions, circuit-based networks enjoy exceptionally high voice quality. But the PSTN infrastructure is very expensive. Before the advent of VoIP, some carriers attempted to reduce overall costs by taking certain shortcuts, and quality suffered as a result. To remain competitive, one carrier, for example, found it necessary to improve quality and recover marketshare through an extensive advertising campaign promoting its "Pin Drop" quality. The moral of the story is that while good quality has a price tag, poor quality can have a much greater impact on the bottom line.

Mean Opinion Score

Historically, Mean Opinion Score (MOS) was measured subjectively using a team of people to rate the quality of a set of standard sentences recorded with both male and female voices. Each sample would be assigned a rating as follows:

- 5 for Excellent with no perceptible impairments
- 4 for Good with barely perceptible but not annoying impairments
- 3 for Fair with perceptible and slightly annoying impairments
- 2 for Poor with annoying but not objectionable impairments
- 1 for Bad with very annoying and objectionable impairments

The average or mean of these subjective opinions determined the rating or score for the sample; hence the designation as mean opinion score. Because the MOS is an average, the maximum score possible is generally considered to be about 4.5. The PSTN regularly achieves a MOS rating of 4.3, which is close to perfection for narrowband voice quality.

The MOS test methodology, including specific test phrases, along with other voice quality considerations are specified in the ITU-T P Series recommendations on "Telephone transmission quality, telephone installations, local line networks."

QoS in Packet Networks

Packet-based networks, originally built for data communications, are far more cost-effective for offering converged "triple play" voice, video and data services. The lack of dedicated circuits makes these networks very efficient, but the shared bandwidth needs to be managed with special QoS provisions. Four such techniques are available to enhance QoS in packet-based networks.

The first technique is a combination that has been dubbed the "P's and Q's of IP QoS." Class of Service (CoS) specified in the IEEE 802.1p standard is the usual means for providing QoS in IP networks. With CoS, all packets are assigned a priority; VoIP traffic generally receives the highest priority. When congestion occurs, lower priority packets are discarded first. Some carriers also create Virtual LANs (VLANs) specified in the IEEE 802.1q standard to segment VoIP traffic for priority treatment. The combination of CoS prioritization and VLAN segmentation helps to minimize, but cannot eliminate VoIP traffic packet loss, especially in IP networks that experience periods of severe congestion.

The second technique is Differentiated Services, or DiffServ, which is a class-based IP QoS technique specified in IETF RFCs (Requests for Comment) 2474, 2475, 2597, 3140, and 3246. The DiffServ Control Point (DSCP) contained in the IP header is used to control the Per-Hop Behavior (PHB) of routers along the traffic's path end-to-end. For VoIP traffic, the PHB is normally set for Expedited Forwarding (EF) to minimize latency and packet loss.

The third is Integrated Services, or IntServ, which is a flow-based IP QoS technique specified in RFCs 1633, 2211, 2212, and 2215. IntServ is designed for end-to-end provisioning with a Call Admission Control (CAC) feature and utilization of the Resource ReSerVation Protocol (RSVP) specified in RFC 2205. In effect, IntServ reserves bandwidth as a virtual circuit in much the same way as the PSTN reserves bandwidth in a physical circuit. IntServ affords the best possible IP QoS, but the technique fails to scale well in carrier networks.

The fourth and final technique is Multi-Protocol Label Switching. MPLS adds a separate 32-bit header to each packet, which is used to create virtual Label Switched Paths (LSPs) capable of segmenting, prioritizing and expediting traffic. Originally created as a means to implement virtual private networks (VPNs), the MPLS header also has a 3-bit QoS field that can be used for VoIP traffic to minimize packet loss and latency. But as with the IP QoS provisions, VoIP traffic inevitably suffers during periods of severe congestion. The additional overhead is also a problem with the small packet sizes common in VoIP applications.

Packet QoS vs. Voice Quality

The four packet-based QoS technologies can all be configured – separately or in some combination – to prioritize VoIP traffic in a way that minimizes three potential causes of voice quality problems: packet loss, packet delay and packet jitter. But even in an under-subscribed packet network properly-configured with either IP QoS or MPLS, users normally rate voice quality substantially lower than the PSTN.

In a 2004 survey conducted by Keynote Systems, for example, users rated VoIP calls overall at MOS 3.5. What is most disturbing about this finding is that the VoIP calls scored even less than cellular services, which achieved an overall MOS 3.6 rating in this controlled test. Keynote attributed the poor score in large part to the longer delay in the VoIP network – 295 ms vs. 139 ms in the cellular network. (Source: "Voice Over IP's Quality Surprise" article in InformationWeek on 11 July 2005) Other studies report similar results of poor voice quality ratings with VoIP services.

Why is it so difficult to achieve a MOS of 4.0 or more in an IP network? There are three reasons. First is the use of low bitrate encoding in many VoIP networks as a way of minimizing bandwidth demand. But packet loss dramatically affects G.729a encoding (8 kbps), for example, which is only able to achieve a maximum MOS of 3.7 even under the best of network conditions. By contrast, the G.711 encoding (64 kbps) employed in the PSTN can tolerate packet loss quite well.

The second reason is that voice quality inevitably deteriorates as the number of subscribers grows and/or the traffic load increases. Pilot programs and trials seem to work well, and even the First Office Application (FOA) may experience satisfactory results. The underlying problems first become apparent only as take-up rates increase and congestion occurs. Adding capacity mitigates the problem, of course, and sometimes the QoS provisions require a bit of tweaking. But the "bursty" nature of IP traffic makes it inevitable that the network will experience periods of congestion, and hence, diminished voice quality from packet loss, delay and jitter.

The third reason is more often than not the real culprit, and is the result of what remains present in many IP networks compared to the PSTN: impairments to voice quality caused by noise, distortion, echo and mismatched volume levels. Many IP networks simply lack the provisions to eliminate these impairments by design. Either it was assumed (incorrectly) that eliminating such impairments is no longer necessary, or it is assumed (also incorrectly) that their elimination in the PSTN is sufficient.

It is enlightening to consider the problem another way: With VoIP, achieving consistent circuit-like performance is the best a packet-based network could ever achieve. Which means packet-switching can at best only simulate the physical circuits employed end-to-end in the PSTN. Using high bitrate encoding and under-subscribing helps, but these also undermine the cost-effectiveness of using packet-based networks. So while IP QoS or MPLS is necessary for minimizing packet loss, delay and jitter, neither is sufficient for achieving an MOS of 4.0 or more. "Good" or "Excellent" voice quality is possible only in VoIP networks where each and every impairment to voice quality is successfully removed, just as it is in the PSTN.

Maximizing Voice Quality in Packet Networks

The previous section explained that the PSTN has been able to achieve consistently high MOS ratings by eliminating all impairments to voice quality. This section describes why and where these impairments exist end-to-end in VoIP networks, and how they can be eliminated using the Packet Voice Processor™ platform from Ditech Networks.

Impairments to Voice Quality in a VoIP Network

The diagram depicts a typical VoIP network end-to-end. Each network segment – from the subscriber premises through the core and beyond with peering arrangements – is assumed to offer some form of QoS.

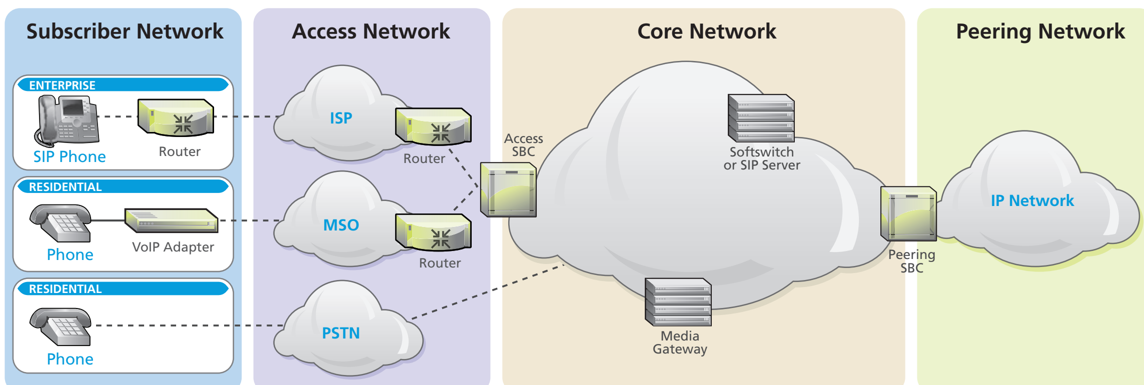

Figure 1 : : Impairments to Voice Quality Occur in the Subscriber, the Access, and the Peering Network Segments

The subscriber network introduces three separate impairments:

- Acoustic Echo – Acoustic echo is caused by poor acoustic isolation between the microphone and speaker in user devices, including handsets, speakerphones and IP softphones. Acoustic echo becomes more problematic with VoIP-induced packet delay, making the echo more noticeable. The problem is quite common in VoIP networks where an estimated 10-15% of all calls suffer from this impairment, potentially in quite annoying ways.
- Ambient or Background Noise – Noise is present anywhere people that live or work. Noise in the form of static, hum, crosstalk and "popping" sounds can also be introduced by various VoIP and packet processing systems. When the noise level is sufficiently high, voice intelligibility and quality suffer substantially. In most VoIP networks, the undesirable noise is simply incorporated right into the encoded packets along with the voice signal. Interestingly enough, eliminating all noise can also be a problem by causing users to perceive the line as going "dead" and prompting the familiar question: "Are you still there?" The perception of a "dead" line is why some codec standards contain Comfort Noise Generation (CNG) or a similar feature. Of course, comfort noise must counter-balance ambient background noise to provide the proper customer experience.
- Audio Level Mismatch – The volume levels of calls between two VoIP endpoints are often unbalanced with one side of the call higher or lower than the other. This can be caused, for example, by devices from different manufacturers or differing microphone sensitivity levels between various phone equipment. Users may be able to compensate for this impairment somewhat by adjusting volume settings, but such inconveniences shouldn't become a constant when delivering higher quality services.

The access network introduces two additional impairments:

- Hybrid Echo – Hybrid or line echo is an electrical signal reflection that occurs at the two-wire to four-wire conversion in the analog tail circuit at the edge of the PSTN. Although hybrid echo is not generated in a pure VoIP network (this does not include Analog Terminal Adapter or ATA connected customer premise devices that do have a hybrid at the POTS phone to VoIP connection), the majority of VoIP calls continue to originate or terminate to the PSTN or a cellular network, or via the PSTN into a cellular network. The additional transmission delays encountered in an IP network worsen the effect of hybrid echo to the point that it can become unbearable. And while some media gateways now include a hybrid echo cancellation feature, this typically addresses only the relatively short hybrid echo tail circuit and not the much longer end-to-end voice path.
- Codec Distortion – Encoding and decoding voice calls using low bitrate VoIP codecs reduces the sharpness of speech and can lead to poor voice intelligibility. With an increasing number of IP borders and multiple transcoding points across a network, providers must compensate for this impairment otherwise there can be severe negative consequences on user satisfaction.

Despite the use of IP QoS or MPLS, there is inevitably some packet loss, delay and jitter introduced in all IP network segments. These packet-related problems exacerbate voice quality impairments in several ways. Most codec standards have some provision to accommodate for packet loss. These Packet Loss Concealment (PLC) or Packet Loss Recovery (PLR) algorithms do indeed help – up to a point. For example, the PLC algorithm in G.711 codecs can recover reasonably well from a packet loss of up to 5%. But with very low bitrate codecs, even with their built-in PLC or PLR algorithms, packet loss rates of less than 1% are normally perceptible, and often annoying.

Jitter or variations in delay increase distortion by making voice sound "garbled" and therefore, less intelligible. For this reason, most VoIP systems incorporate dejitter buffering to eliminate the variations. But these jitter buffers work by imposing additional delay to remove these variations!

Voice quality problems are often caused or exacerbated by packet delay. Packet delay in IP networks is considerable compared to the PSTN, being caused by the accumulation of delays in encoding/decoding, packetization, serialization, output queuing, jitter buffering and, finally, transmissions across the wide area network (WAN). Echo – both acoustic and hybrid – becomes noticeable beyond 25 ms. It is interesting to note that this 25 ms of latency is the same for both a round-trip transcontinental PSTN landline call and the rate speech resonates to the ear through the skull in a natural human "sidetone" of sorts. As the amount of delay increases, voice quality worsens from a "hollow" sound to a truly irritating echo. A one-way delay of more than 150 ms (300 ms round-trip) interrupts the conversational flow to the point where it becomes difficult if not impossible to carry on a normal conversation.

Individually these impairments lessen voice quality, perhaps substantially. Collectively, these impairments can make VoIP voice quality unacceptably poor.

The Packet Voice Processor from Ditech Networks

Ditech's carrier-class Packet Voice Processor is an integrated system purpose-built to remove all voice quality impairments in any IP or MPLS network infrastructure. The Packet Voice Processor platform with both Voice Quality Assurance (VQA™) and Packet Quality Assurance™ (PQA™) technologies enables VoIP carriers to significantly improve voice quality end-to-end in VoIP networks. VQA features improve voice clarity on every call by providing noise reduction, enhanced voice intelligibility, voice level control, acoustic echo control and hybrid echo cancellation. PQA features address additional voice issues specific to VoIP, such as packet loss, delay and jitter to dramatically improve overall call quality and clarity.

The Ditech Packet Voice Processor integrates six separate technologies that eliminate all impairments to achieve PSTN-like voice quality in any packet network:

- Acoustic Echo Control (AEC) solves the acoustic echo problem found in most VoIP networks. Ditech's bidirectional AEC feature suppresses a wide range of echo variances using algorithms based on talker energy levels and Weighted Acoustic Echo Path Loss (WAEPL) to effectively suppress acoustic echo and significantly improve voice quality.
- Adaptive Noise Cancellation (ANC) technology includes a high-precision noise reduction algorithm that removes the noise components of a call without reducing talker volume. This feature effectively suppresses ambient or background noise by up to 21 dB to improve the quality of a call, resulting in greater customer satisfaction and enhanced call clarity.
- Automatic Level Control (ALC) technology dynamically detects level imbalances and automatically adds up to 15 dB of gain or attenuates as much as needed to bring both sides of the call to the same, specified volume. ALC also prevents clipping and codec distortions, and can compensate for background noise by improving the signal-to-noise ratio.
- Hybrid Echo Cancellation (HEC) completely eliminates hybrid echo bidirectionally and end-to-end for VoIP calls that traverse a PSTN hybrid. Ditech is recognized as the industry leader with its advanced HEC technology that cancels echo with an Echo Return Loss (ERL) of up to 0 dB, compensates for network tail delays of up to 278 ms, and affords fast, stable convergence in less than 50 ms.

Beyond MOS 4.0 for VoIP Voice Quality

The ultimate goal of VoIP is to exceed the voice quality provided by the PSTN. With equal or better voice quality, the "triple play" of converged voice, video and data services will enable IP networks to completely replace the TDM-based infrastructure.

Two industry advances hold the promise to make this possible. The first is High-Definition (HD) VoIP specified in the ITU-T G.722 series of codec standards. By employing a high fidelity frequency range of 50-7000 Hz, HD VoIP has the ability to produce voice quality superior to that of the PSTN's G.711, which utilizes only a narrowband frequency range of 300-3400 Hz. HD VoIP might be used, for example, in "telepresence" applications that employ video and audio conferencing, normally with white-boarding and other real-time collaboration tools. These high-definition telepresence services – already being used by early adopters – give participants an experience equivalent to being in the same room together – virtually, of course.

The second advance involves Dynamic Voice Quality. DVQ makes it possible to apply different codecs and VQA algorithms based on available bandwidth to maximize the infrastructure's revenue potential. With this capability, carriers will be able to adjust service levels in real-time to accommodate customer requirements. A premium telepresence session, for example, could be allocated additional bandwidth made available by temporarily employing low bitrate codecs for standard services.

- Intelligent Packet Restoration (IPR) is an advanced feature that reconstructs missing packets within a VoIP packet stream. With the capability to analyze and process the received voice packets, the IPR feature uses a predictive speech model to reconstruct a missing packet's voice payload and rebuild it upon packet play out. This offers an enhanced ability to support high-quality VoIP, even in congested IP networks experiencing substantial packet loss.
- Enhanced Voice Intelligibility (EVI) improves the quality of speech by rebalancing the spectral signature and boosting the critical speech formants (soft sounds) to provide increased clarity and speech recognition. EVI is particularly adept at enhancing the intelligibility of VoIP services that utilize low bitrate codecs with compensation for ambient noise, and without distorting or amplifying the signal.

The Ditech Xi™ solution is a powerful voice quality monitoring tool for the Packet Voice Processor that measures and reports a comprehensive set of real-time voice quality statistics for a VoIP network. Unlike other voice quality monitoring solutions that require intrusive test calls or predict MOS based only on IP problems, Ditech's Xi non-intrusively monitors every call passing through the network. Information such as speech and noise levels, echo delay, return loss, packet loss, and jitter are all collected in real-time for every call. MOS ratings are calculated for both Listening Quality (LQ) and Conversational Quality (CQ) using the E-Model standard. The visibility afforded by Xi enables service providers to proactively identify and address voice quality issues. Finally, the ability to retrieve and analyze records of call quality statistics allows customer issues to be addressed readily through simple queries.

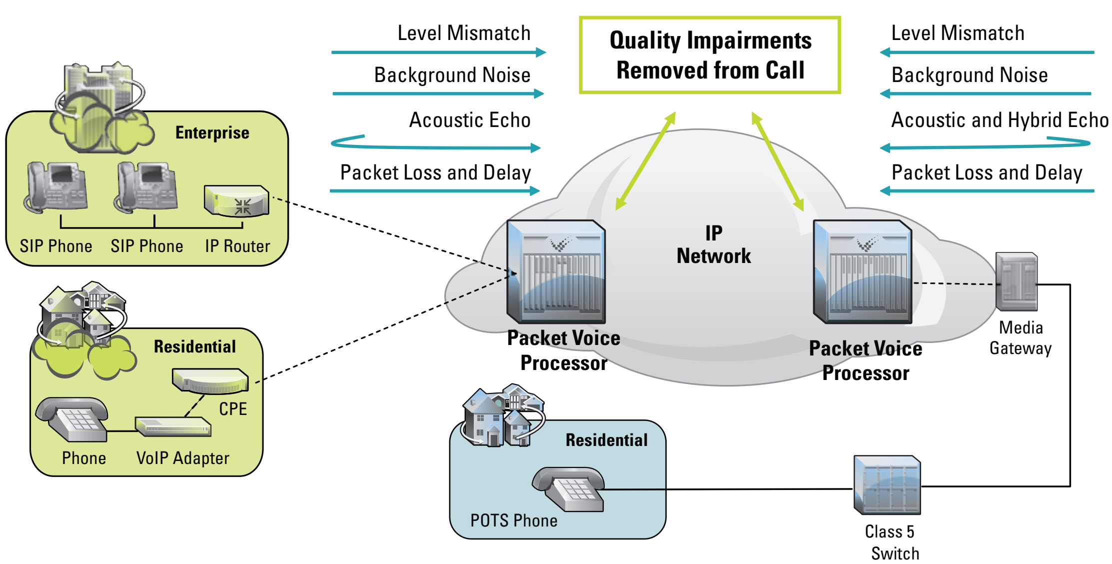

Figure 2 : : The Packet Voice Processor from Ditech Networks Removes all Impairments to Voice Quality

The Cost-effectiveness of Exceptional Voice Quality

Adding a Packet Voice Processor to a VoIP infrastructure pays for itself in numerous ways. PSTN-like voice quality increases customer satisfaction, which reduces costly churn. Acquiring new customers is expensive. And because customers are obviously conversing routinely with many others via their VoIP service, referrals from satisfied users can become a powerful marketing tool that reduces customer acquisition costs. In addition, consistently "Good" or "Excellent" voice quality (around MOS 4.0) minimizes calls to the Help Desk, which previously could do little or nothing to actually help the caller.

A single carrier-class Packet Voice Processor can support up 13,440 simultaneous calls; three units can be configured together seamlessly in a single rack to support over 40,000 calls. The system is conveniently deployed at the boundaries of the core network at both the access and peering network demarcation points. The built-in Hybrid Echo Canceller and Acoustic Echo Control feature with their long echo tails, for example, together eliminate the need for separate echo cancellation systems. And the ability to transcode among more codec types than any other solution available allows this critical capability to be consolidated on a single network location for maximum benefit across all backbone VoIP equipment. (Additional information about the transcoding capabilities of the Packet Voice Processor can be found in a separate white paper titled "Cost-effective Transcoding for Carriers" on Ditech's website.)

Finally, the ability to deliver satisfactory voice quality using low bitrate codecs enables carriers to get the most from the entire IP infrastructure investment. This also eliminates the wasteful but sometimes necessary practice of under-subscribing the network, as well as helps to postpone potentially costly and disruptive upgrades to the network. Indeed, many carriers are able to cost-justify the relatively modest investment in a Packet Voice Processor based on this capability alone.

Conclusion

Unless and until the impairments to voice quality are all successfully removed, no amount of IP QoS provisioning can ever hope to achieve the PSTN's MOS rating of 4.0 or more. And as VoIP services become more widely available, voice quality will emerge as the major competitive differentiator among providers.

Although some planning is required to improve voice quality in any IP or MPLS network, the implementation effort is itself made substantially easier – and far less expensive – with the availability of purpose-built voice processing systems like the Packet Voice Processor from Ditech Networks. Some additional education on quality issues may also be required for business management and the technical staff. Such training is often beneficial for two reasons. First, with a full appreciation for voice quality's impact on the bottom line, management is generally quite willing to support taking the steps necessary to make improvements. And second, being armed with a thorough understanding of voice quality issues prepares the engineering and operations staff to do exactly what needs to be done to eliminate any and all impairments.

To receive additional training and planning assistance on improving voice quality in your VoIP network, please contact Ditech Networks by calling 1-650-623-1300 (1-800-234-0884 in the US and Canada) or emailing us at ditech@ditechnetworks.com.

About Ditech Networks

Ditech Networks supplies voice processing equipment for telecommunication networks around the world. Ditech Networks' technology solutions include voice, media processing, SIP, and security delivered on carrier-grade scalable platforms to enhance the delivery of communications services over mobile, Voice over IP, and wireline networks. Ditech Networks' customers are premier network operators including Verizon, AT&T, Orascom Telecom, and others that collectively serve more than 150 million subscribers.

Corporate Headquarters

US Sales – West
825 East Middlefield Road
Mountain View, CA 94043
1-800-234-0884 (Toll Free)
1-650-623-1300 (Direct)
1-650-564-9599 (Fax)
NA-Sales@ditechnetworks.com

US Sales – East and Central
8360 Greensboro Drive
McLean, VA 22102
1-703-909-1747 (Phone)
1-866-758-2862 (Fax)
NA-Sales@ditechnetworks.com

Canada Sales Office
2275 Lakeshore Blvd West, Suite 500
Toronto, ON M8V 3Y3
Canada
NA-Sales@ditechnetworks.com

Mexico Sales Office
Torcuato Tasso 245-6o. Piso
Col. Polanco
México, D.F. C.P. 11570
México
+52 (55) 5254-5422 (Main)
+52 (55) 5350-8679 (Sales)
Mexcio-Sales@ditechnetworks.com

Brazil Sales Office
Alameda Araguaia, 933
Alphaville
Barueri, SP 06455-000
Brazil
+55 (11) 4208-6266 (Main)
+55 (21) 3521-5543 (Sales)
Brazil-Sales@ditechnetworks.com
South-America-Sales@ditechnetworks.com

Middle East/Africa Sales Office
21 El-Fawakeh St., 3rd Floor
Dokki, Giza 12311
Egypt
+20 2 336-5100 (Phone)
+20 2 761-8964 (Fax)
MEA-Sales @ditechnetworks.com

South Asia Sales Office
No. 8/11, Sarvapriya Vihar
New Delhi – 110016
India
+91 9810372555
South-Asia-Sales@ditechnetworks.com

South-East Asia Sales Office
Lippo Plaza 3rd Floor
JI Jend. Sudirman Kav. 25
Jakarta 12920
Indonesia
+62 (21) 5291 3780 (Phone)
+62 (21) 522 1977 (Fax)
SE-Asia-Sales@ditechnetworks.com

China Sales Office
Unit 3010, No. 500
Xiangyang South Road
Xuhui District,
Shanghai 200031
PRC
+86 (21) 5456 0305 (Phone)
China-Sales@ditechnetworks.com

Spain Sales Office
Torres Quevedo, 1 – (P.T.M.)
28760 Tres Cantos – Madrid
Spain
+34 (91) 803 74 44 (Main)
+34 (91) 829 26 90 (Sales)
Spain-Sales@ditechnetworks.com
Europe-Sales@ditechnetworks.com