Application Note

Voice over IP Call Quality Analysis using the OptiView™ Integrated Network Analyzer

In today's competitive environment, many organizations are turning to Voice over Internet Protocol (VoIP) as a way to cut costs, boost productivity, and increase their competitive advantage. In fact, more than 70 percent of organizations now use VoIP in some capacity.

However, VoIP systems present some special IT challenges. First, because end users expect IP phones to work as reliably as land lines, VoIP systems must meet exceptionally high quality and performance standards. Second, even if the system works fine initially, later network changes could affect call quality – or a growing VoIP system could affect other mission-critical applications. So to minimize problems, you need to verify (before deployment) that your infrastructure can support VoIP; thoroughly examine all system elements during deployment; and manage your VoIP system proactively after deployment, including ongoing monitoring, troubleshooting, and planning for future growth. To meet this challenge, Fluke Networks offers a comprehensive VoIP lifecycle management solution that provides end-to-end management data, giving IT managers and frontline support personnel complete vision into the network from edge to core.

Voice quality

What is voice quality? Voice quality is usually referred to as the subjective quality of a conversational voice call. But, when voice quality metrics are quoted, they generally relate to a listening test which may not take into account factors such as delay. Typical voice quality metrics are mean opinion score (MOS), perpetual speech quality measure (PSQM) or perceptual evaluation of speech quality (PESQ), or R factor.

Factors that affect voice quality can generally be grouped into three areas, network impairments, CODEC related impairments and environmental factors. Impairments due to the network have a reasonably well characterized effect on quality. Packets containing speech frames can be lost in transit due to buffer overflow, or can be discarded at the receiver if they arrive excessively late. Packet loss rate is affected by the size of the receiver jitter buffer – a larger buffer will result in fewer discarded packets, but may increase overall delay. The transit time of a packet through the network will vary for a number of reasons. This effect is usually removed by the jitter buffer at the receiver where incoming packets are delayed, re-ordered and forwarded to the CODEC. However, this process, while removing the jitter, may increase packet loss and overall delay.

Low levels of delay, say under 100ms, are not noticeable during a voice call. Larger delays cause the conversation to become disjointed. Delay occurs in initial packetization, voice compression, transmission through the network and the buffering and decoding process.

CODECS are used to convert the analog voice signal to digital and back. The G711 CODEC provides the best voice quality since it does no compression, therefore introducing the least delay, and is not so sensitive to packet loss. Other CODECS, like G729 and G723, consume less bandwidth by performing compression, but reduce clarity, introduce distortion and delay and make the voice quality extremely sensitive to lost packets.

Voice activity detection, or silence suppression (where packets that contain silence are not transmitted), can sometimes cause clipping at the beginning of talk bursts. While external background noise is outside of the VoIP system, it can affect the user's perception of quality.

Voice call ratings

The "recency" effect reflects how a listener would remember call quality. In tests, a 15-second burst of noise was inserted and moved from beginning to middle and then to the end of a 60-second call. When the noise burst occurred at the start of the call users reported a mean opinion score (MOS) of 3.82. When the noise burst occurred at the end of the call, users reported a MOS of 3.18. The resulting difference in MOS of 0.64 shows a 20% quality rating effect depending on where the noise burst occurred. This recency effect is even

OptiView™ Series III Integrated Network Analyzer

The only portable, integrated analyzer with enterprise-wide vision.

more significant when considering the typical range for MOS is 2.5 to 4.0, resulting in 40% impact over that range. This recency effect is believed due to the tendency for people to remember the most recent events and the way auditory memory functions, which typically decays the recollection over a 30-second interval. So, the effects of recency should also be considered in the model to arrive at an accurate end-user-perceived call quality assessment.

Voice quality is determined by the analyzer and displayed as a MOS or as R-factor. The most widely used is MOS, which is an average rating by individuals who judge the quality and is rated on a scale of 1 to 5.

More recently, a new analytical technique was introduced to calculate MOS scores. This technique is known as R-factor and is based on an analysis of packet loss, delay, jitter and an estimate of user perception.

	R-factor	MOS
Very satisfied	90 - 100	4.34 - 5.00
Satisfied	80 - 90	4.03 - 4.34
Some users satisfied	70 - 80	3.60 - 4.03
Many users satisfied	60 - 70	3.10 - 3.60
Nearly all users dissatisfied	50 - 60	2.58 - 3.10
Unacceptable	0 - 50	1.00 - 2.58

Table 1

Table 1 maps user satisfaction with the voice call against the R-factor and the MOS score. “Toll quality” sound is generally associated with a MOS score of at least 4 and is equivalent of using a conventional land line. G.711 starts with an intrinsic MOS value of 4.4 or an R-Factor of 94.3. So, we would never see a MOS score greater than 4.4 while the G.729, which performs significant compression, has an intrinsic MOS value of 4.1 or an R-Factor of 84.3.

Analyzer placement

Where does the analyzer need to be placed to capture VoIP traffic? In order to capture the VoIP traffic, the analyzer must be in the path of the packets. This can be accomplished by spanning or mirroring the traffic to replicate the data to another switch port. While this is an inexpensive solution, it can introduce new timing variables, making the voice quality measurements less precise, it will also mask any physical layer issues since the switch will not forward errored frames.

Another method of connecting the analyzer is to use an aggregating tap. This tap can be inserted between the switch and the phone and will not introduce any new timing variables. The tap will combine a full duplex data stream into a single stream which can be analyzed by a single port analyzer. The aggregating tap also allows an OptiView™ Integrated Network Analyzer to transmit packets on to the network so the active discovery, SNMP analysis and remote control features can be used when connected in this configuration.

Once the VoIP traffic has been captured, what do we need to do in order to analyze the call quality? While real-time protocol (RTP) provides the end-to-end delivery services, real-time control protocol (RTCP) provides information on the transmission and reception quality of the data carried by RTP.

RTCP provides data to the sender regarding the quality of RTP sessions, continuously calculates inter-arrival jitter and generates performance “reports” on delay, jitter and packet loss. Since this adds additional traffic to the network, the control traffic should be limited to a small and known fraction of the session bandwidth: small so the primary function of the transport protocol to carry data is not impaired; known so the control traffic can be included in the bandwidth specification given to a resource reservation protocol and so that each participant can independently calculate its share. It is suggested that the fraction of the session bandwidth allocated to RTCP be fixed at 5%. While the value of this and other constants in the RTCP interval calculation are not critical, all participants in the session must use the same values so the same interval will be calculated. Therefore, these constants should be fixed for a particular profile.

However, rather than looking through all the RTCP decodes, it is much easier to use the OptiView VoIP Analysis Option.

The OptiView VoIP Option recognizes and decodes all major VoIP protocols:

- The H.323 suite of protocols specified by the ITU including Q.931, RAS, H.245 and T.120.
- Session initiation protocol (SIP), a signaling protocol for Internet conferencing, telephony, event notification and instant messaging. The protocol initiates call setup, routing, authentication and other feature messages to endpoints within an IP domain.

- Skinny client control protocol (SCCP) is the proprietary signaling and communications protocol in Cisco's Architecture for Voice, Video and Integrated Data (AVVID).
- Media gateway control protocol is used for controlling telephony gateways from core agents.

The option also recognizes and decodes all major CODEC protocols used for VoIP.

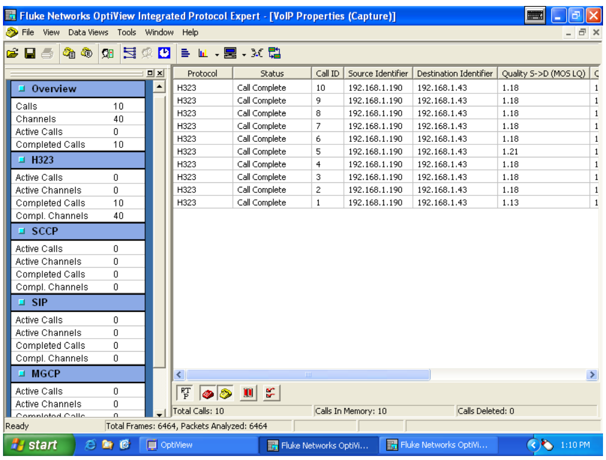

Figure 1.

The VoIP properties screen, shown in figure 1, displays all the calls captured and can be configured to show everything that is critical to assessing the quality of the call, including jitter, packet loss, estimated packets discarded, user and network R-factor.

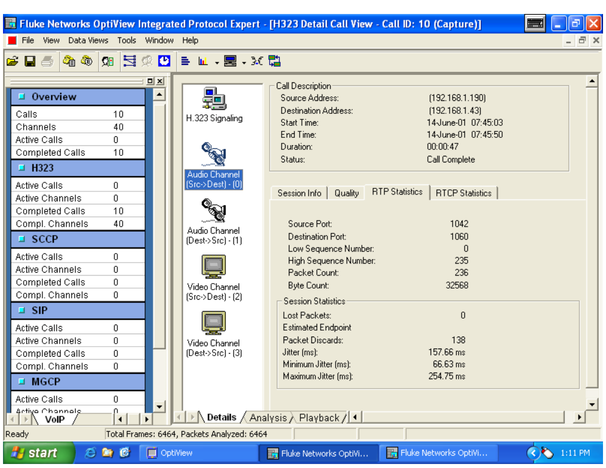

Figure 2.

By highlighting the signaling data icon in the left side pane (shown in Figure 2), the analyzer shows the source and destination names, IP address, port number, call manager and call reference number. Selecting the Channel Detail icon provides quality related information about the individual call. By clicking on the RTP statistics tab, we will get detailed information calculated from the RTP packets captured by the analyzer and we can also view RTCP Statistics.

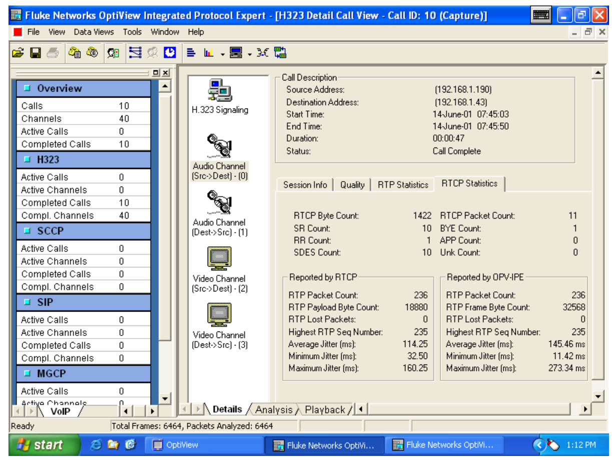

Figure 3.

The VoIP option provides two different views of call jitter and dropped packets, one calculated by the analyzer and the other extracted from the RTCP packets. These two values will probably not match exactly, and, in some cases, may differ substantially. It depends primarily on the point in the network where the analyzer captured the information. If the analyzer captured the packets close to an endpoint, then the jitter values for the far endpoint will nearly match the RTCP reported jitter value. However, if the analyzer captures the traffic towards the mid point between the end points, then the calculated and RTCP value may differ significantly. This may indicate network impairment somewhere between the analyzer and the end point.

The VoIP endpoints are aware of, and report impairments for their receive side only – they have no knowledge of the quality of their transmitted traffic as it is received at the far end.

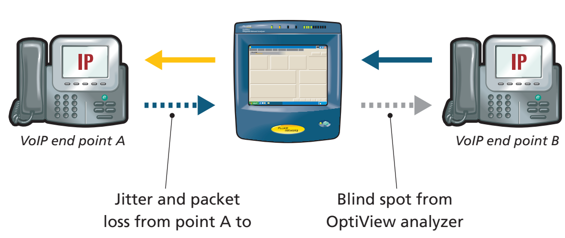

Figure 4.

The OptiView analyzer can see both RTP flows and calculate jitter and packet loss for the full duplex communication. So in the example shown in Figure 4, the OptiView analyzer can see jitter and packet loss caused by the network between endpoint A and the analyzer, but is blind to what happens between the analyzer and endpoint B. The converse is also true: endpoint B to the analyzer, but no visibility from the analyzer to endpoint A. So, if high jitter and packet loss are detected, then there are issues with two of the four network segments, but there is no visibility into the other two segments without some end-to-end feedback mechanism between the endpoints. Enter RTCP.

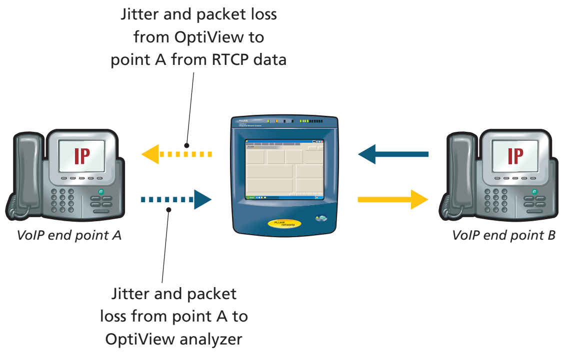

Figure 5.

In this case, we can diagnose and isolate an impairment by using the combination of RTP calculations performed by the analyzer and data from RTCP decodes. Say, for example, the analyzer indicates probable reception problems at endpoint B due to impairments in the RTP flow between endpoint A and the analyzer itself. Unlike the previous example, endpoint A is sending RTCP reports to endpoint B and endpoint B is sending to endpoint A. These RTCP reports are decoded by the OptiView analyzer and can be compared with the calculated values for jitter and packet loss for each direction of the traffic. A significant difference in these values will indicate the direction of the problem in relation to the analyzer.

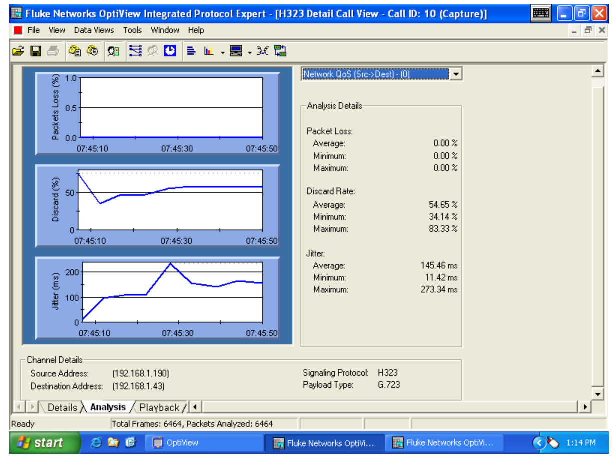

Figure 6.

By selecting the Analysis tab, we can see the phones involved in the voice exchange, the signaling protocol and the CODEC being used – in this example G.723. Packet loss, discard rate and jitter statistics in both tabular and graphical views for the duration of the call are also displayed and the information can be viewed for each direction of the call.

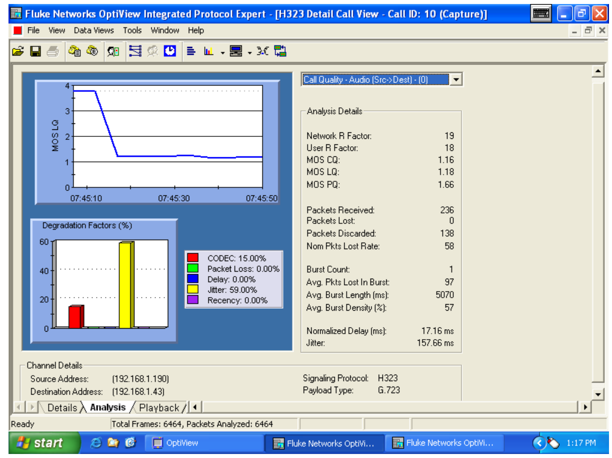

Figure 7.

Individual call quality analysis can also be selected from the drop down menu. In Figure 7, we can see the network and user R-factor score, together with the MOS for call quality (CQ), listener quality (LQ) and PESQ. The call quality incorporates factors such as delay and recency, which affect conversational quality and implicitly affect listening quality. The listener quality does not contemplate the impairments that cause conversational quality problems and can be compared to subjective MOS scores. The MOS PESQ score maps the R-factors to the ITU-T standard P.862 end-to-end measure of voice quality using a test signal. We can also see the factors contributing to degradation – in this case, 56% is attributed to jitter and 15% to CODEC impairments.

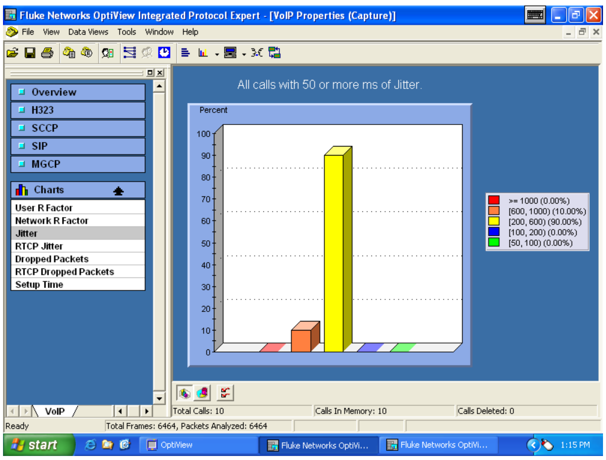

Figure 8.

The analyzer can display parameters for all the calls we have captured. In this case, the jitter measurements for 10 calls. Here, 90% of the calls are experiencing between 200 and 600 milliseconds of jitter. As a general rule, jitter less than 75 ms is considered good, 125 ms acceptable and over 200 ms is poor. Both the network and user R-factor can also be displayed and indicate the quality of the call. The network R factor is calculated based on the physical equipment impairments and the user R factor adds recency and delay to the network R factor in an attempt to add "perceived" annoyances that a user may experience during a call.

Conclusion

In this short application note, we have looked at the primary characteristics of networks that can impair the quality of a VoIP call such as jitter, delay and packet loss, and explained how the OptiView Integrated Network Analyzer with the VoIP option can be used to measure these values. We also discussed the point at which the network parameters affect measured values, how jitter measured close to a device sending RTCP reports will be different to those measured near the receiver of the reports and how the OptiView analyzer can be used to pinpoint the source of a problem – something which is often impossible based solely on the information obtained from the endpoints.

We have also seen some of the broad range of critical measurements obtained, post capture from the trace file, where OptiView provides extensive detail about the network and call characteristics that determine the quality of service VoIP end users experience.

For more information on VoIP, visit the Fluke Networks' VoIP Lifecycle Management web page at **www.flukenetworks.com/voip** to download special reports, white papers, technical documentation application notes and view webcasts.

For more information about the OptiView Integrated Analyzer, visit **www.flukenetworks.com/optiview**.

Printed in U.S.A. 1/2007 2674050 A-ENG-N Rev B