


Monitoring and Managing Voice over Internet Protocol (VoIP)

As with most new technologies, Voice over Internet Protocol (VoIP) brings new challenges along with the benefits. The main challenge is VoIP's extreme sensitivity to delay and packet loss compared with other network applications such as web and e-mail services. A basic understanding of VoIP traffic and of the quality metrics provided by VoIP monitoring tools will help you keep your VoIP network running smoothly.



How does VoIP work?

VoIP phones use codecs to translate analog sound streams into digital packets for transmission. On the receiving end, the codec translates the packets back to analog. For two people to converse normally, all of this must happen in as close to real time as possible.

For call setup, most enterprise VoIP solutions include one or more call managers, which are servers that set up calls between VoIP phones, and can also provide gateway connections to the Public Switched Telephone Network (PSTN) for calls outside the VoIP network. Typically, the call initiator contacts the call manager, which then rings the phone being called. Once the receiving party answers, the call manager provides a mechanism for the phones to negotiate codecs and connection parameters.

The connection itself is typically in the form of two full-duplex streams: a Real-time Transport Protocol (RTP) stream that carries the encoded audio, and an RTCP (Real-time Transport Control Protocol) stream to provide communications control. Once the call is set up, the call manager is no longer involved until the teardown phase, when the IP phones inform the call manager the call has been completed so the centralized call queue (a list of what phones are active) can be updated.

What can go wrong with VoIP?


Depending on what components of the network are compromised, users can experience audio quality problems with a call, or they might not be able to connect in the first place if a call manager is affected.

Audio quality problems

Problems with VoIP audio quality are almost always related to network delay, jitter, and packet loss, or some combination of the three. It is common to see them together because they are both related to a general deterioration of network conditions.

In any VoIP deployment, some delay is unavoidable. Codecs take time to encode/decode the audio stream, and even the fastest network medium is not instantaneous.

How VoIP phones send audio streams over a network

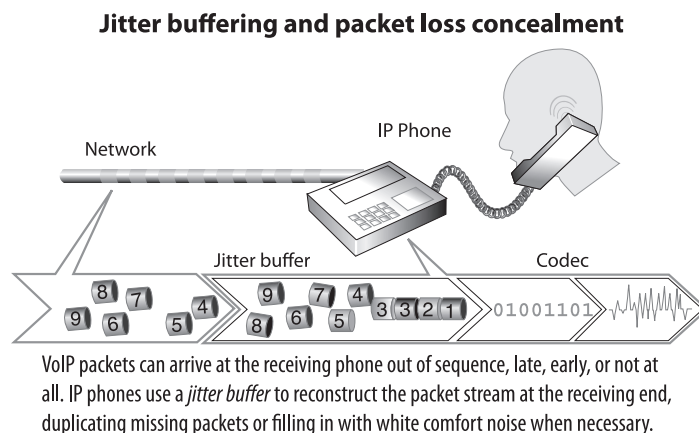


In addition to network delay, the VoIP equipment itself (IP phones, gateways, etc.) subtracts even more processing time from the overall delay budget. The delay budget for reasonable two-way conversations in real time is about 150 milliseconds (one way). When delay exceeds the budget, the callers can get confused about who should be speaking and who should be listening, and begin to talk over and past one another.

Network jitter and delay

Real-time voice communications are sensitive to delay and variation in packet arrival times. Codecs require a steady, dependable stream of packets to provide reasonable playback quality. Packets arriving too early, too late, or out of sequence result in jerky, jumbled playback. This phenomenon is called jitter.

Because no network can guarantee a perfectly steady stream of packets under real-world conditions, VoIP phones use jitter buffers to smooth out the kinks. A jitter buffer is simply a First-In, First Out (FIFO) memory cache that collects the packets as they arrive, forwarding them to the codec evenly spaced and in proper sequence for accurate playback.



While a jitter buffer can successfully mask mild delay and jitter problems, severe jitter can overwhelm the jitter buffer, which results in packet loss (see below). Increasing the size of the jitter buffer can help, but only to a point: A jitter buffer that increases overall round-trip delay to 300 ms will make normal conversation difficult.

Packet loss

As mentioned above, packet loss can be the result of the jitter buffer being overwhelmed. Other reasons include landline media failure and poor wireless signal quality. The latter can be a big problem with VoFi (Voice over WiFi) service. Regardless of the source, VoIP phones and gateways attempt to conceal this type of signal degradation by duplicating packets to fill in the missing data. As with jitter, these techniques can maintain voice quality only to a point.

Packet loss on data networks has long been characterized as a “bursty” phenomenon, which is another way of saying “it never rains, it pours.” Networks tend to either sporadically drop single packets (these periods are called “gaps” in packet loss), or large numbers of contiguous packets in a “burst.” Packet loss concealment techniques typically have no trouble handling packet loss during gap periods; it is the sustained bursts you must watch out for.

Call management problems

If the VoIP call manager (sometimes called the VoIP server) is overwhelmed with requests, or its connection to the network is impaired, call setup delays can reach the point where users abandon calls before they are able to connect to the other party. If IP phones are misconfigured, or their IP connection to the server is impaired, calls remain open in the call queue long after the parties have disconnected.

Managing VoIP quality

You can manage only what you can measure. Managing a VoIP deployment therefore requires some hard numbers beyond subjective user assessments of quality (although these are obviously important as well). Beyond monitoring the network parameters discussed previously in this paper (“What can go wrong with VoIP?”), having an overall quality score such as a Mean Opinion Score or R-factor score can also be a useful VoIP network health index.

VoIP monitoring tools calculate the MOS and R-factor scores using a formula known as the E-model. Using the statistics it has collected from the network, the analyzer calculates how much the various impairment factors (such as codec compression, jitter, delay, and packet loss) would affect the typical user’s perception of call quality.

Choosing between VoIP-specific and all-purpose monitoring tools

There are a number of different options on the market for managing VoIP quality, mainly falling into three categories:

- Dedicated VoIP tools originally developed for the telcomm industry. These tools are great for testing IP phone and gateway designs, but not as good at solving deployment problems on a live network.
- Network protocol analyzers that have added “VoIP Support” by buying technology developed for the telcomm industry and integrating it into their product line.
- All-purpose network monitoring tools that approach VoIP quality management from an IT administrator’s point of view rather than from that of telcomm engineer.

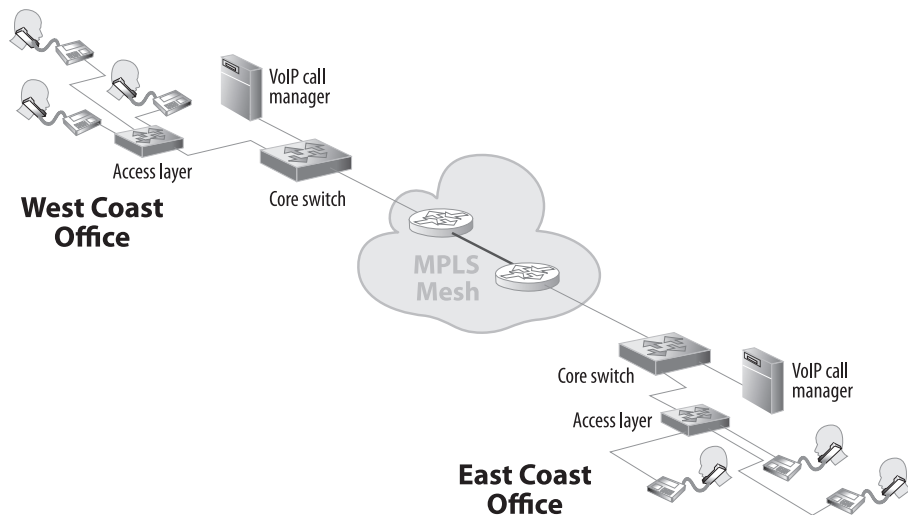
To the IT administrator, managing VoIP quality is just another network task. This makes the third approach (the all purpose network monitoring tool) often the most practical choice. But note that “VoIP support” means more than just decoding the packets of various VoIP protocols; it also means being able to track and display network delay, jitter, and packet loss, and to distill this information into overall quality scores, both per-call and in aggregate.

And to be really useful to the enterprise, the tool should also track, store, and analyze long-term trends. This is so that you can understand what is “normal” VoIP performance, and maintain a database of Call Detail Records (CDRs) from which you can generate reports for management or service providers. The VoIP monitoring tool should also be capable of automatically notifying you when selected statistics indicate a developing problem. On all of these counts, Network Instruments® Observer® meets the requirements.

VoIP points of visibility

In switched environments, where to deploy an analyzer or probe for maximum visibility isn’t necessarily obvious. Complicating matters for VoIP is the fact that each call includes both client-server communications (between IP phones and the call manager during setup and tear-down), and peer-to-peer (the streams of voice data passed between the parties).

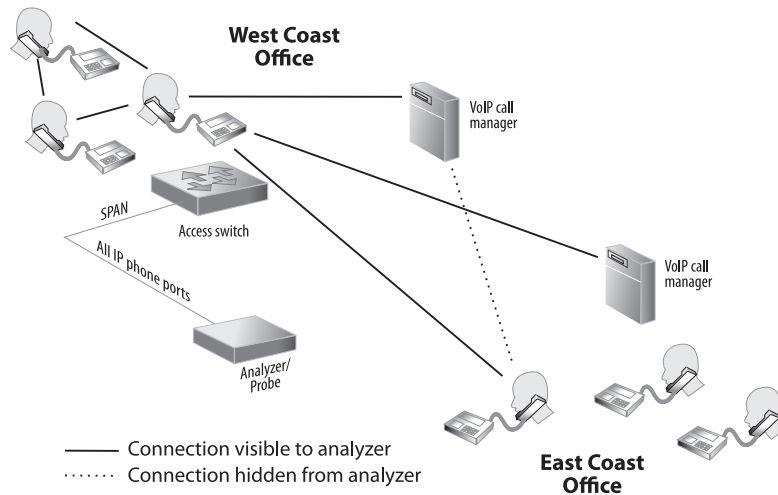
For example, consider the following VoIP network deployment:



Where to place probes on such a network depends on what you want or need to see. If you need access to all local conversations on either coast, including both call setup and actual voice data, use a SPAN session on the access layer switch to mirror VoIP traffic to the analyzer. Assigning all VoIP traffic to a dedicated VLAN makes this fairly straightforward.

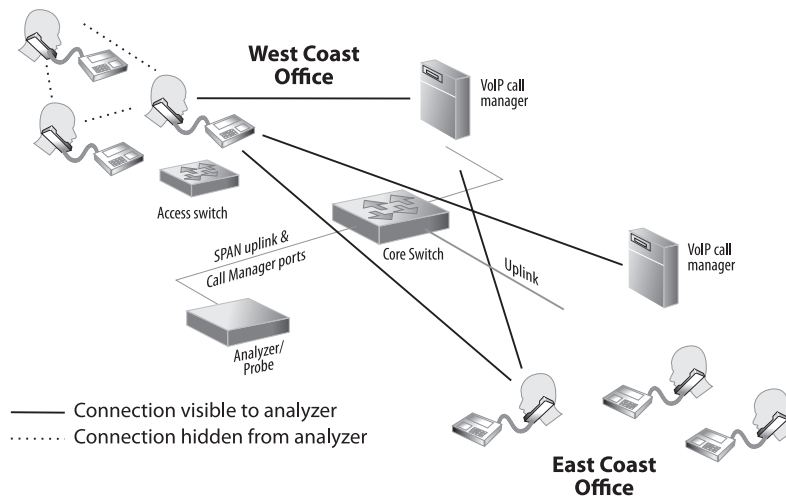
Capturing local IP phone traffic shows:

- Any phone's communications with its local call manager,
- Both sides of the full-duplex connection between local phones talking to each other, and
- Both sides of the full-duplex connection between phones located on opposite coasts.



What you will not be able to see from this probe is any communications between the East Coast and the call manager located on the West Coast, or between the West coast and the call manager located on the East coast.

If you are more interested in a coherent view of calls between the West Coast and East Coast, including all call manager communications, use a SPAN session to mirror both the uplink traffic between the core and MPLS mesh, and all traffic flowing to and from the call manager. This will give you a coherent view of inter-office calls, along with all call manager communications, both local and remote.



With a probe deployed in this manner, you will not be able to see the peer-to-peer voice traffic between local phones. For complete coverage, connect probes to both the core and access layers at each site. Another alternative is to deploy probes at the core 24/7/365, monitoring the access layer with a portable analyzer or software probe only to troubleshoot local call problems as needed.

VoIP network analysis

How can VoIP analysis help manage quality? By closely monitoring the network conditions that affect VoIP, you can begin to address developing infrastructure problems before they result in user complaints or downtime.

Tracking network performance

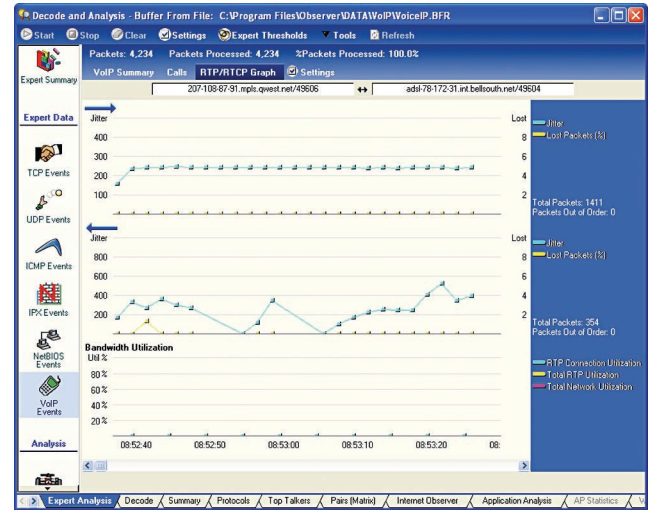
Consider the East Coast/West Coast example described in the previous section. Ken, the administrator responsible for ensuring VoIP quality, has set up a Network Instruments 10/100/1000 Probe Appliance on the core switch to monitor all call manager activity and any VoIP traffic traversing the link. He has configured Observer to send him an e-mail whenever any of the following conditions arise:

- MOS falls to 3.5 or less
- Jitter levels crossing the MPLS mesh exceed 20 ms
- Delay levels crossing the MPLS mesh exceed 80 ms

Any of these are indications that VoIP quality is threatened. Given the topology involved, the most likely source of problems is the MPLS mesh routers, which are under the service provider's control. By digging deeper into the statistics the analyzer provides, you can determine why the MOS is falling, and what is causing jitter, delay, or packet loss.

If jitter is the problem, a good place to start is by comparing jitter levels against bandwidth utilization to see if there is any correlation.

The analysis shown above (taken from Network Instruments' Expert Observer product) shows just such a correlation. Depending on the situation and IT budget, such an observation could mean it is time to invest in more bandwidth, or time to put more controls on employee Internet usage for applications such streaming media and peer-to-peer file sharing unrelated to business.

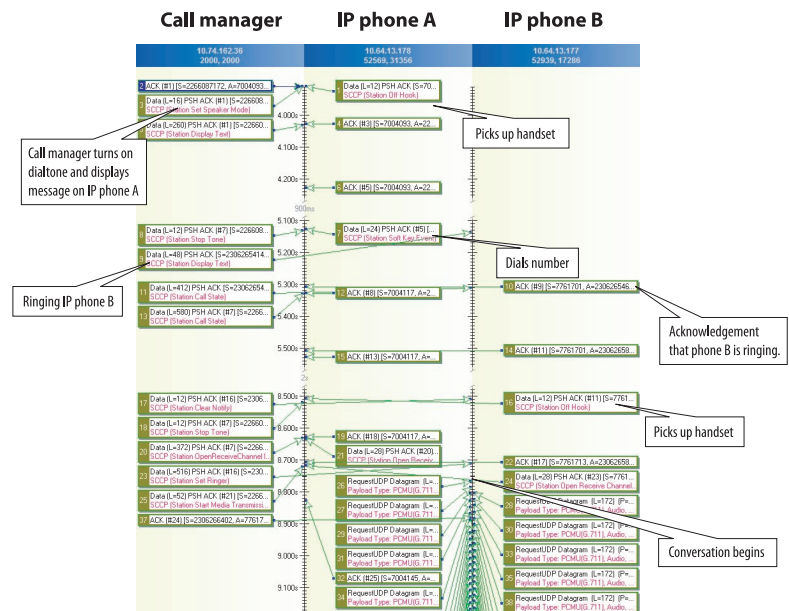


If there isn't an obvious correlation between jitter and bandwidth utilization, the depth of data provided by an all-purpose network analysis and monitoring tool can help you dig deeper for the correct diagnosis. For example, if VoIP traffic across an MPLS mesh is subject to excessive jitter, it could be the result of "route flapping" on the service provider's routers. An analyzer can confirm this and provide documentation that this is the case. Armed with the hard data provided by analysis, you could then contact the service provider so they can address the problem. If delay across the mesh exceeds the contractual obligations of the Service Level Agreement (SLA), the provider may owe your organization some refunds for service failure, in addition to being responsible for fixing the problem.

Troubleshooting connection problems

When a user can't get a dial tone, or there are excessive delays in ringing the other party's phone, examining a graphical display of how the call is progressing between the parties and the call manager can indicate what is going wrong.

Network Instruments' VoIP Expert displays just such a diagram. Simply right-click on any call or connection stream. Because differing protocols dictate differing phone/call manager interactions, some knowledge of the protocol is necessary for detailed troubleshooting. But even if you lack a detailed knowledge of the protocol, the Connection Dynamics display highlights which party isn't responding, or which party is responding slowly.



An example of a Connection Dynamics display showing a VoIP call using the SCCP protocol. It is easy to see how such a diagram is essential to efficiently troubleshoot VoIP connection problems.

Summary of VoIP statistics and quality metrics

The following table summarizes the statistics and quality measurements discussed in this paper, both defining what is measured, and describing its relevance.

VoIP metric	What it measures	How to use the analysis
Jitter	Jitter measures the variability of delay in packet arrival times. In spite of the jitter buffers used to counteract jitter, at excessive levels it can interfere with smooth playback and cause packets to be dropped.	By using triggers to notify you when jitter levels are reaching levels that threaten voice quality, you can examine your routers for problems or contact your service provider and help them solve the problem.
Delay	The amount of time it takes a packet to reach its destination. Whenever packets travel a network, some delay is inevitable. For real-time telephone conversations, there is a one-way "delay budget" of approximately 150 ms.	As with jitter, using automatic notifications to actively manage levels of delay can prevent the problem from escalating to the point where users complain.
Packet loss	The percentage of packets that did not reach their destination.	Sporadic packet loss is usually insignificant. However, sustained bursts (see the next item) can affect quality.
Bursts	Periods characterized by high rates of packet loss. The burst percentage is the percentage of time that the call experienced high-rate packet loss; the burst density is the actual percentage rate of packet loss during bursts.	VoIP phones have no trouble masking a lost packet here and there by duplicating the previous packet or filling longer silences with white noise. But users will notice sustained bursts. If VoIP traffic has been assigned proper QoS and has enough bandwidth, the most likely culprit is media failure.
Gaps	Periods characterized by low rates of packet loss. The gap percentage is the percentage of time that the call experienced low-rate packet loss; the gap density is the actual percentage rate of packet loss during the gaps.	Usually not significant, as packet loss concealment technologies are usually successful in masking the effects of low-level packet loss. Contrast with bursts, described above.
Average call setup/teardown	An average of how long it is taking the call manager to open and close calls.	A spike in these statistics can indicate a problem with the call manager or its connectivity to the network.
Codec	The compression/decompression method that was used for the call.	Different codecs are capable of different levels of quality sound reproduction. Higher compression comes at the cost of lower quality, but may be necessary given the bandwidth available to the call. If it seems as if the codecs in use are using more compression than necessary (or not enough) given the amount of bandwidth available, perhaps the VoIP phones can be reconfigured to use a different codec.
Mean Opinion Score (MOS)	Starting with a theoretical perfect score of 5 (excellent), impairment factors such as codec, delay, jitter, and packet loss are used to calculate how a typical user would rate voice quality.	These are useful as quick overall indicators of VoIP health. If the average MOS falls below 3.5, or the average R-factor falls below 80, it's likely that you have more than a few dissatisfied users. If you see these statistics trending downward, it's time to examine more detailed analysis to determine what is going wrong.
R-factor	Similar to MOS, this scale ranges from 1-100.	

Conclusion

Managing VoIP is similar to managing any other network application; VoIP differs only in its level of sensitivity to network delay. By keeping close tabs on delay, jitter, and packet loss, you can prevent network problems from becoming phone problems. Having a network analyzer that includes sophisticated VoIP analysis in your toolbox will make this essential task much more manageable for the already overworked IT administrator.