The Pitfalls of Scaling VOIP

Net Forecasts – Peter J. Sevcik
BCR Volume 32, Number 3
March 2002

The "experts" either love voice over IP or they dismiss it. They're either completely enamored with how well it works and its market prospects, or they are completely skeptical, claiming that it doesn't work much better than a bad ham radio call and that it is destined to be a small fringe market. There are few balanced views.

In this mad rush to judgment, who is correct? Objective analysis is needed whenever large investments and changes that impact end users are at stake.

Determining Voice Quality

Most of the skepticism is based on the notion that VOIP voice quality is too poor to satisfy users who have grown accustomed to the quality provided by of the old Bell System. There are many ways to judge voice quality, but the most popular is the Mean Opinion Score (MOS).

The MOS tests, which are based on tons of research originally developed at Bell Labs, uses real people who listen to a test sentence through some communications medium. They then rate their ability to understand the sentence by rating the sound as follows: 1=bad; 2=poor; 3=fair; 4=good; 5= excellent. The MOS is the arithmetic mean of all the individual scores.
Running a MOS evaluation is difficult and expensive. The ITU, with the input of many telecommunications research organizations, has developed a mathematical model to calculate MOS under a variety of network conditions. This is documented in the ITU-T specification G.107, "The E-Model, A Computation Model For Use In Transmission Planning." We implemented the E-model to analyze the effects of various network conditions when operating a VOIP call.

We first used the E-model to calculate the MOS of many codecs under ideal conditions in order to better understand the scores. As it turns out, the E-model never returns a score of 5. We found the official attributes (poor, fair, etc.) associated with each score provided very little insight into how to rate a call from a consumer's perspective.

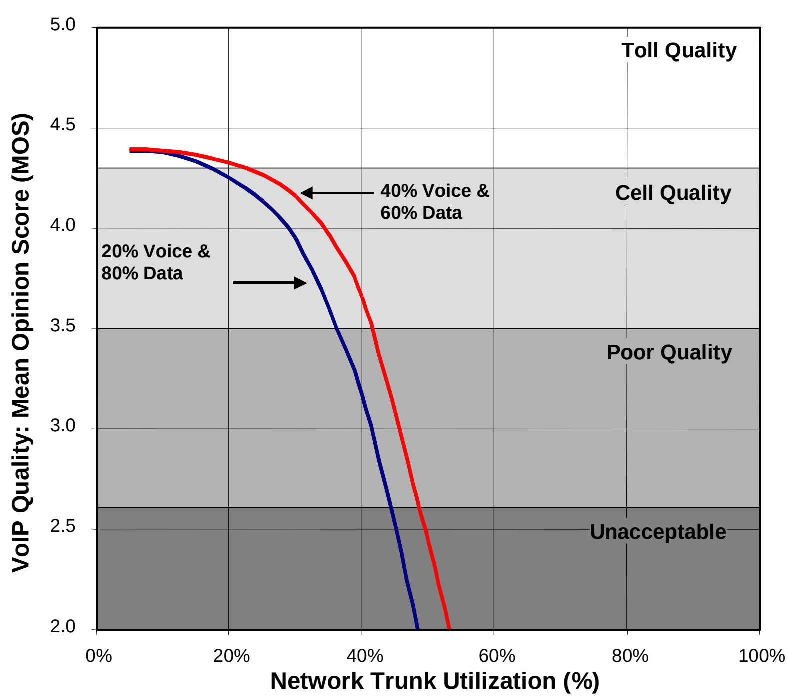

Figure 1 – Voice Quality as a Function of Data Traffic and Utilization

For example, the codecs used in cell phones provided a range of MOS from 3.5 to 4.3. Interestingly, the history of first-, second-, to third-generation cell phone technology exactly follows a progression of MOS scores within that range. We also used the G.107 specification to show that a MOS below 2.6 was unacceptable. This leads to a recasting of the MOS results into four bands of quality as shown in Figure 1.

Analysis Results

We ran the E-model to analyze the results of simulations of data traffic and VOIP traffic on a U.S.-based national network. This represents either a typical Furtune-500 company enterprise network or a next-generation voice carrier. We ran many combinations of traffic mixes, hop counts, trunk speeds, etc.

In general, the model provided expected results, but one of the most interesting is the effect of data traffic upon the VOIP traffic. It is clear that the bursty nature of the data traffic has a serious effect upon the voice quality as seen in Figure 1. If there is only 20 percent voice traffic, the MOS for VOIP falls to cell-phone quality at 17 percent trunk utilization. However, if there is more voice – say, 40 percent voice – and therefore less bursty data traffic, the operator can run the network to 22 percent utilization before the VOIP traffic falls into the cell-quality zone.

The clear conclusion is that the bursty nature of data traffic significantly impacts the VOIP call. There are three ways to avoid the problem: operate at low utilization, operate a voice-only network with no competing data traffic or employ end-to-end QOS to protect the VOIP from the data traffic.

Implications and Recommendations

The first implication of this analysis is that enterprises using VOIP on a corporate WAN simply will not be able to operate the service without a strong QOS technique in place across the entire network. There are many candidates -- 802.1p, COS, DiffServ, MPLS -- but little consensus on how to make them work end-to-end across multiple enterprises and multiple carriers. Enterprises cannot afford separate networks for voice and data, nor can they operate at less than 20 percent utilization during the busy period. The slope of the curves in figure 1 are extremely steep once congestion starts to appear. An enterprise would have to operate at a target utilization of less than 8 percent in order to have the headroom to operate at 16 percent some of the time. A budget of 8 percent utilization in the busy period implies buying twelve-times the bandwidth that will be used – hard to justify to the CFO.

The second implication is that carriers and service providers who operate a VOIP service need a careful strategy for managing user expectations. Some carriers have embraced VOIP to take advantage of cheaper switches, new services, faster implementation and innovative business models.

However, many of these carriers currently operate the VOIP on a separate network from their other networks. Keeping the data and voice on different networks may turn out to be a long-term necessity to supply a high-quality voice service. Of course, national carriers have the resources and economies of scale to operate separate networks. History is simply repeating itself: Networks tend to be engineered for and dedicated to a single purpose.

The third implication is that operating VOIP on the general Internet works but is a risky business proposition. This is what gives the VOIP skeptics their ammunition. Two people can download some software and talk over the Internet. It will work fine some of the time but it can't work fine all of the time to all of the locations served by the Internet.

The real question is why does it even work at all? The answer involves the utilization of the Internet. As I pointed out in a column last year (see *BCR*, January 2001 pp. 10-11), in 2001 Internet bandwidth would outpace demand by 20 to 1, and the gap was growing. In short, the Internet was operating at only 5 percent utilization, or below the target of 8 percent needed to stay within toll quality and well below a safe target of about 18 percent needed to supply cell quality.

But that was then, this is now. The carriers essentially stopped investing in bandwidth during 2001, so general utilization is creeping back up. Voice quality on casual VOIP calls is sure to degrade.

This does not mean that you can't operate high quality VOIP on the Internet. In fact, service providers like iBasis and ITXC do it at a large scale. But they deliver high-quality voice calls because they do not mix voice and data on the access lines, and because they are very careful to give the voice traffic to high-performance international ISPs like Worldcom, Genuity and Cable & Wireless. They constantly monitor the quality of service provided by alternative ISPs and move the traffic to insure good overall quality.

Furthermore, they can take advantage of very good SLAs supplied by a few key ISPs. It appears that that even though Internet service is becoming a commodity, some ISPs really do provide better service in order to get business. Fortunately, all ISPs don't try to maximize utilization in order to offer the lowest possible price. If they did, by now iBasis and ITXC would have run out of choices.

Conclusion

VOIP works well in "clean" environments – networks that are over-provisioned. However, competitive pressures on both enterprises and carriers make such environments harder and harder to find. The industry needs to get end-to-end QOS working before our networks all run out of air. VOIP is the canary in networking's mine.

Peter Sevcik is president of NetForecast in Andover, MA, and is a leading authority on Internet traffic, performance and technology. He has contributed to the design of more than 100 networks, and led the project that divided the Arpanet into multiple networks in 1984, which was the beginning of today's Internet. He can be reached at peter@netforecast.com.

NetForecast measures application profiles, user behavior and network statistics to predict performance, adoption, and market impact of new technologies. The firm has helped leading service providers, enterprises, and vendors navigate the changing competitive landscape of the Internet economy. We supply key technical and market guidance to ensure the success of network-based projects, products and services. www.netforecast.com

NetForecast