

PipelinePub.com

May, 2006 | Volume 2, Issue 10

Getting Through... Clearly IMS Call Session Control and Quality of Service

By Barb Lancaster

The basic idea of IMS is to take advantage of the power, flexibility, and cost reductions associated with "IP-everywhere" and add to it some of the disciplines associated with "traditional" telecoms, providing, for example, Quality of Service, access control, network security, and billing for calls and traffic.

This all seems like good news for service providers, especially the traditional carriers, as they strive to contain new-generation competition, while at the same time migrating to their own all-IP networks as quickly as possible.

An interesting outcome of the implementation of IMS is that the business of telecommunications may evolve to look more like the old PSTN than like the Internet. (This is not necessarily a criticism: to move forward it's sometimes helpful to learn from the past.) To illustrate this, we shall look at just one of the featured benefits of IMS – QoS.

The ability to deliver QoS through session management is an important feature of IMS. By making the transport layer aware of call sessions, then it is possible to deliver guaranteed QoS, at least by one measure. The decision about what to measure becomes very, very important.

Today, the discussion of QoS tends to resolve around minimizing jitter and latency in voice and video streams. Jitter can cause sound/picture distortion and loss of information. Too much latency can cause conversation to become difficult and unnatural. However it is possible to define QoS in many other ways. For a user accessing a database, response time may be important and data error rates will certainly be of interest. For a customer watching a movie, latency may be of no importance (because the video stream can be buffered for a few seconds without damaging the overall enjoyment of a two-hour movie) but any residual jitter that affects sound and picture quality will be a concern.

In the traditional PSTN environment, QoS focused on call completion success rates, not voice quality and not latency. Latency is naturally low in a circuit-switched environment. Voice quality is uniformly not-too-bad, and difficult to measure objectively anyway.

In the circuit-switched world, the number of calls connected is limited by network capacity. Once a trunk is completely loaded with voice calls, no new calls can be accepted and so those new calls fail. PSTN carriers measure that failure rate and use it as a measure of QoS. In a situation of rapid customer growth (which the PSTN network has experienced for most of its life), failed calls were a consistent cause of customer complaint. Measuring call failure therefore provided a reasonable proxy measure of customer satisfaction. It also provided data that could be used to pinpoint priorities for growing network capacity.

By comparison, in the unmanaged world of IP telephony it is possible to initiate new call sessions even when the route between the end points is already heavily loaded. The result can be that voice quality and latency on the new call is less than satisfactory. The problem is compounded because it is not just the new session that experiences poor sound quality: the initiation of the new session will itself cause deterioration of the quality on all the other sessions already in progress.

This sort of cumulative problem is just not consistent with the notion of "carrier class". With currently little or no control over this service quality problem, it is no wonder that service providers want to find a way to do it better.

Most attempts to manage QoS are at the packet level, and are based largely on prioritization of traffic. But giving precedence to the packets that make up vulnerable services (such as voice) and relegating less sensitive traffic to "best efforts" may mean significant delays and lost packets in a heavily loaded network.

There are several problems with prioritization at the packet level. One is that the originator of the packet can decide which priority to apply, which may not be the same priority that the carrier wants to apply. But a bigger concern is: what do you do when you have more than enough "top priority" traffic to fill the pipe? Which packets do you leave behind? When random packets are dropped, they need to be re-transmitted, adding even more traffic, increasing the level of congestion. Prioritization at the packet level is clearly not capable of coping with severe congestion.

One possible solution in such circumstances is to manage the traffic at the session level, not at the packet level. This is the task of the Call Session Control Function in IMS. What that means in practice is that once a session is set up, every packet will get through. As the network approaches overload levels, new sessions will not be permitted. The session controller will simply not pass on the session initiation request.

The result of this intervention is to redirect the focus of QoS for voice calls. Instead of being concerned primarily about the voice quality delivered, we can instead look at the proportion of attempted sessions that are successfully initiated. By avoiding network overload, we can replace the assumption that all sessions are successfully set up, with the assumption that all successful calls will receive acceptable voice quality.

From the customers' perspective this changes the experience. Every successful call will sound good, but a few call attempts may be unsuccessful, and the customer will have to try again later. By comparison, with packet-level QoS the customer will find that almost all calls are connected "successfully", but on some proportion of connected calls, network congestion will cause sound quality to be less than satisfactory.



Business Operations Architects



Which do customers prefer? I'm not sure anyone has asked them recently, at least not in a large-scale study. However, based on my unscientific observations I think the answer is, predictably: it depends. Business people in general, like good quality audio, but if the phone call is the one that's going to clinch a deal seconds before the deadline, even a CEO with high standards will put up with some distortion, provided the message is understood. On the other hand, people who rely on cheap and cheerful VoIP to

make international calls will generally be happy with lo-fi connections as long as the price is low. But if one call in three becomes so distorted that the call is useless, then those calls effectively fail anyway. Getting a busy tone might be less frustrating.

Note that if there is sufficient capacity to carry all sessions demanded by customers, it doesn't really matter whether we adopt a packet-based or session-based approach. The calls will mostly be fine. So the QoS debate resolves to what measurement is more useful: call failure rates, or some assessment of distortion and customer acceptability. From the service providers' perspective, call completion statistics provide hard information to assist capacity planning. Subjective opinions on voice call quality are perhaps not so helpful.

The QoS feature of IMS is therefore of fundamental importance. Carriers need session control to manage and deliver QoS in a way that is remarkably similar to the way it happens in the traditional PSTN network. IMS session control provides carriers with exactly the same useful parameter (call connection success rates) that they have used for years.

Are there alternatives to the IMS approach to quality management? Possibly. Throwing capacity at the problem before there is a problem is one way. Reducing bandwidth requirements of video and audio sessions would also help. It is also possible to improve the ability of end devices to measure audio/video performance and terminate the call quickly if a good session can't be created. We can certainly expect to see continued developments in edge-device intelligence. It may be that ultimately cheap enough and good enough will be sufficient for commoditized voice/video sessions. No fancy network intelligence may be needed to deliver those benefits.

Meantime service providers need to control, measure, and manage to deliver services as profitably as possible. IMS, for the present, provides the network intelligence to do that.