Satellite links pose challenges in setting up onboard VoIP

While defining and monitoring the quality of service provided by an onboard VoIP system is essential to extracting the maximum possible benefit from the technology, measuring IP performance over a satellite link is no easy task

By Guy Adams*

The SatManage network correlator module shows network events relating to Mean Opinion Score and jitter at 15 minute intervals over a 24h period

Before the prevalence of high bandwidth shared-use circuits, maritime vessels typically used narrowband satellite links for critical voice communications. Dedicated voice channels via low earth orbit or Inmarsat satellites can be very expensive, especially when call volumes are more than very occasional. Voice over IP (VoIP) therefore represents a highly attractive solution for maritime voice communication needs.

Typically, VoIP takes the form of an enterprise system with special handsets on the vessel and a digital switch on shore to link to the conventional voice network. However, there are now other legitimate VoIP solutions that run directly from users' laptops or desktops, typically for crew welfare purposes, such as Skype and Windows Messenger. In recent times, growth of full broadband access, not only to very large vessels, but to medium size commercial tonnage and even smaller fishing vessels and mid-size yachts, has made VoIP a cost effective solution for most.

Managing traditional voice circuits was fairly straightforward, given the usually lower latencies and the fact that circuits were generally un-contended. However, the issues of running voice, and particularly VoIP, over a dedicated geostationary satellite circuit, as is the case for VSAT, need to be much more carefully considered.

With shipping operators seeking to reduce costs, 'luxuries', such as the dedicated voice circuits, are becoming less and less common. Very large vessels may have two satellite antennas, but these are generally to provide for redundancy or multi-tenancy arrangements rather than dedicated circuits for high priority traffic.

Therefore, even in these large scale deployments, voice, even mission critical voice, must share bandwidth with all other traffic – right down to low priority but high bandwidth traffic generated by crew and by automatic software updates initiated by Microsoft and other applications. While most vessels will have an alternative voice circuit for emergency use, such as an Iridium phone or similar, this cannot be used regularly, simply because the primary broadband link is congested or there are voice quality issues.

There are therefore three key issues that need to be addressed:

- how these circuits are specified and contracted
- how service providers classify and prioritise voice (and, for that matter, other high priority traffic with real-time requirements)
- how this is monitored.

It is interesting that the proliferation of mission critical and high priority voice has in no way triggered a change in the way clients specify circuit requirements. Shipping companies typically undertake significant exercises to analyse and understand the nature of the data traffic that they will need to transport, and go to enormous lengths to specify exactly how much bandwidth they should acquire, how they will test this, what circuit availability should be and exactly how this will be measured.

Yet many neglect to specify basic circuit level performance parameters, ie, not just circuit availability and whether adequate bandwidth is available, but whether the circuit is running at the performance needed for applications to work, or more importantly, is the relevant part of the circuit running at the performance needed for the application to work?

Take for example a VoIP call requiring, say, 12 kbps of bandwidth, running over a 1 Mbps/256 kbps circuit: for the purposes of this call, only the performance (latency, jitter, packet loss) of the 12 kbps is important. In this scenario the 'keep it simple' adage does not apply.

While specifying no performance criteria for a circuit is inadequate, insisting that the entire circuit perform to the level required for VoIP is also shortsighted since whatever anyone says, maintaining this level of performance on the entire circuit is considerably more expensive, and ultimately the end user will pay for the extra cost. The ideal philosophy can therefore be paraphrased as: 'ask for exactly what you need, but no more' (see table).

The method used by satellite equipment to

SAMPLE SPECIFICATION FOR A 1 MBPS (DOWN)/512 KBPS (UP) CIRCUIT OVER A 5 MINUTE PERIOD

	Max average latency	Max average jitter	Max average packet loss
Voice – up to 64 kbps	700ms	40ms	1%
Citrix – up to 256kbps	800ms	60 ms	1%
General service – remaining bandwidth	900ms	100ms	1%

classify and prioritise different types of traffic is now very sophisticated, generally using any combination of IP address, port, protocol type and several service flags. In effect, this creates multiple virtual circuits within the primary circuit.

Each of these virtual circuits is dynamic; sometimes they will not exist (if there are no voice calls, there will be no voice service) and when they do exist their size will be variable up to the limit defined in the specification. What should vary is the monitoring of each of these virtual circuits according to the parameters defined in the circuit specification.

The matter is further complicated, however, by the fact that many of the metrics to which casual reference is made, such as jitter and packet loss, are not rigorously or well defined terms. There is no widely accepted definition that specifies exactly how raw data should be measured and from this how such metrics should be calculated.

Probably the closest are the IP Performance Metrics** (IPPM) proposed by the Internet Engineering Task Force (IETF) but even these are based on a number of assumptions, which cannot be applied in the case of satellite networks. For example, it is assumed that both ends of the link are fully available for detailed monitoring, that cost is no issue and there are no limits on the bandwidth between both monitoring points. The two main issues are:

- measuring circuit performance metrics one-way. Since all of the monitoring capabilities are on the hub side, it is only really possible to measure these metrics since round trip metrics, ie, measures of latency, jitter and packet loss, are a sum of the respective downstream and upstream metrics
- measuring circuit performance based on actual end user data crossing the network is problematic.

Take as a simple example measuring one-way packet loss (which happens to be an extremely important metric for assessing VoIP quality) of actual user VoIP traffic. How can this be measured?

It turns out the only way is to have two devices, one on either end, with extremely accurate time synchronisation, which communicate with each other to determine how many packets entered the link at one end and compare this to the number that exited the other, over a given time period.

Even this is plagued with problems, particularly over a high latency connection, since several packets are likely to be in transit across the network at the end of the period but as they have not arrived, they will be marked as missing. This is one example, but there are many others, and almost all of them converge to a single monitoring premise: a user may not be able to monitor it directly, but he can infer it.

This means that rather than being able to measure packet loss, latency and jitter of actual network traffic, it is necessary to inject artificial traffic, measure the performance of this, and use this to infer the performance of the actual traffic. Based on over a decade experience of monitoring complex satellite networks, this method has proven to be entirely satisfactory, and if done correctly will show up all the significant issues on the circuit.

In response to these issues, SatManage has developed a specific set of tools for monitoring quality of service (QoS). These tools allow calculation of the base metrics described above as well as advanced VoIP metrics such as R-value and MOS (Mean Opinion Score), the de facto standard for measuring actual user perception of voice quality. As with all SatManage tools, these have been developed specifically to work on satellite networks. ***MEC***

***** This article was authored by Guy Adams, chief technical officer at Parallel Ltd.***

***** To learn more about the Internet Engineering Task Force's proposed performance metrics, visit: www.ietf.org/html.charters/ippm-charter.html***

Software compresses call costs

Communications technology specialist, TriaGnoSys, has launched what it claims to be the 'world's most efficient satcoms compression' software. Telephone calls made using conventional voice over IP (VoIP), says the company, can use as much as 96 kbps bandwidth, while calls placed with GSM can consume up to 30 kbps.

TriaGnoSys' new software, VoCeM, is said to reduce necessary bandwidth to only 6 kbps, therefore significantly reducing throughput over the satcoms link and leading to much lower call costs for the end user. Reducing the bandwidth requirement also increases the number of simultaneous calls that a satcoms link can handle at any moment in time. As such, the payload usage could increase between five or 10 times, which for cruise ships or ferries, could result in markedly higher revenue generation from passenger calls.

The efficiency gains are achieved by employing state-of-the-art transcoding and compression techniques, combined with innovative kilobit transmission technology. The VoCeM software has been designed to be installed on existing GSM, 3G/UMTS and VoIP communications servers so no hardware upgrades are involved.

Managing director of TriaGnoSys, Axel Jahn, says: "The promise of onboard and remote communications services has always been that the experience will be the same as standard terrestrial services. The technology to make that happen has been available for some time, but at a cost. The introduction of VoCeM, bringing down call costs, means the promise is now a significant step closer to becoming reality."

Remote management leads to greater uptime

Parallel has welcomed Uplogix, a specialist provider of secure remote management (SRM) solutions, as its latest technology partner.

The combination of SRM technology from Uplogix deployed at remote sites with SatManage in the network operations centre (NOC) is said to give satellite operators 'unprecedented visibility' into their remote networks. Uplogix appliances can be remotely deployed to monitor and automate routine administration of both satellite and IP network devices common in today's hybrid networks.

In the event of a network outage, Uplogix connects to the NOC over an out-of-band link, ensuring transfer of device information to SatManage and providing a channel for remote troubleshooting. The co-location of the Uplogix appliances makes it possible to gather richer information more frequently about the deployed devices without increasing the burden on network traffic or the devices themselves.

Parallel's managing director, James Dell, says that pairing SatManage in the NOC with Uplogix at the network's edge "increases uptime, improves SLAs, and reduces support costs to hybrid networks."

SatManage integrates with several vendor technologies and where networks have Uplogix SRM solutions, SatManage takes Uplogix data feeds to highlight abnormal events on the network as well as allow at-a-glance analysis of network performance. This can show problems developing over time and even highlight wasted bandwidth and resources. Because problems can be spotted in advance, pre-emptive action can be taken on-site by Uplogix to eliminate negative impact on the business.

"The combination of Uplogix secure remote management with SatManage brings additional value to satellite customers, who are often in the most extreme locations with the most challenging architectures," comments Tom Goldman, chief executive officer at Uplogix.