Radio World ENGINEERING EXTRA

REPRINTED FROM FEBRUARY 22, 2006 WWW.RWONLINE.COM

WHITE PAPER

by Bob Band

A Packetized Look at Audio Over IP

Exploring Two Approaches of IP Audio Transport and The Challenges of Packet Audio Distribution

Bob Band is the channel manager for Harris Networking Solutions.

Today's radio broadcasters are dependent upon their contribution and distribution networks to deliver broadcasts to listeners.

Contribution networks bring audio sources from multiple locations into one or more studios. Distribution networks are typically point-to-multipoint, and carry program audio to remote transmitters and translators. These networks handle the signals that ultimately feed the transmitters; as such, they must reliably and seamlessly pass high-fidelity audio for transmission. They also allow for remote control, signal monitoring and data transmission.

To date, fixed-bandwidth digital services such as T1, E1 and X.21 have been the systems of choice, as they provide reliable communications over dedicated circuits. With the advent of packet networking, and in particular high-speed Internet Protocol, the need to use expensive fixed-bandwidth services is being challenged.

IP services offer scalable bandwidth, are widely accepted and available and are considerably less expensive than dedicated, fixed-bandwidth lines. However, meeting the mission-critical needs of broadcasters can be a challenge. The reasons have to do with the way packet networks operate.

FIXED-BANDWIDTH VS. PACKET NETWORKS

In a fixed-bandwidth network such as E1 running at 2 Mbps, you get a "nailed-up" circuit dedicated to your use alone, carrying a constant stream of information. Picture a conveyor belt moving non-stop, onto which you can place your user data (audio programs or other); the belt keeps moving at a steady rate, whether you've put anything onto it or not.

The reliability on such a circuit is extremely high, but it comes at a price. You're paying for the full 2 Mbps of transport, 24 hours a day, 365 days a year, regardless of how much or how little you actually use. Also, such circuits are essentially point-to-point, which means you may need multiple E1s if, for example, you are distributing a program from one studio to multiple transmitters, or connecting several regional studios to a main production facility.

Packet networks such as IP, by contrast, are designed to be shared resources. Anyone with access to the network can put data onto it.

IP networks typically operate in one of two modes: Transmission Control Protocol or User Datagram Protocol. Essentially, TCP is designed for point-to-point bi-directional communication, while UDP is more suitable for unidirectional links, either point-to-point or point-to-multipoint broadcast.

Regardless of which mode, or protocol, is used, the information is broken up into small segments and placed into packets, a process knows as packetization. Along with user data, each packet contains several pieces of housekeeping data such as the destination address, a sequential packet ID and, when implemented, a mechanism for priority tagging.

For example, a packet containing bits of real-time audio could be tagged as urgent, while a packet containing bits from a data file transfer would be tagged as lower priority. Once the data is packetized, the packets are given to the network for delivery.

> **Packet loss can occur in several ways ... if the jitter is so high that a packet arrives too late to be inserted back into its proper place in the jitter buffer before playout, then it is essentially lost.**

In the network, the packets are sent through to their destination via one or more switches and routers. Each device looks at each packet, makes a decision on the best route available at that moment in time, and sends the packet on its way. Gateways, which allow packets to move between networks using different protocols, can easily become sources of short-term congestion, as they must handle all packets moving between the networks involved.

For the audio to be successfully decoded at the far end, it is essential that all the packets reach their destination and they are played out in the same sequence as they originated. For this to happen, we have to address two inherent challenges presented by IP networks: jitter and packet loss.

Jitter is a condition where packets arrive at irregular intervals, or even out of sequence altogether. It occurs because there is no way to pre-determine the exact route any given packet will take through the network. As a packet reaches each switch, that device sends it on to the next one based on constantly-updated information about the congestion conditions on the various possible paths. Thus, sequential packets may be sent along different routes.

As the routes vary, so too will arrival times. Differences in arrival time translate into packet jitter, and as packet jitter increases, packets may begin to arrive out of order. To cope with this, a well-designed IP audio transport system should have a receive jitter buffer deep enough to hold the incoming audio packets until any delayed ones have arrived, so they can be played out in proper sequence.

Packet loss can occur in several ways. One relates to the foregoing discussion of jitter. If the jitter is so high that a packet arrives too late to be inserted back in its proper place in the jitter buffer before playout, then it is essentially lost.

Another is that IP networks are designed such that if a switch along the way becomes congested, packets can be dropped as needed to reduce the congestion.

There are several methods available to deal with packet loss. One is priority tagging, which allows the switch to differentiate between high- and low-priority packets. As congestion increases, low-priority packets are dropped first.

Second, it is important to have an adequate Service Level Agreement (SLA) with one's network provider. An SLA is a service contract in which the provider agrees to provide a certain level of guaranteed throughput across their network for your traffic.

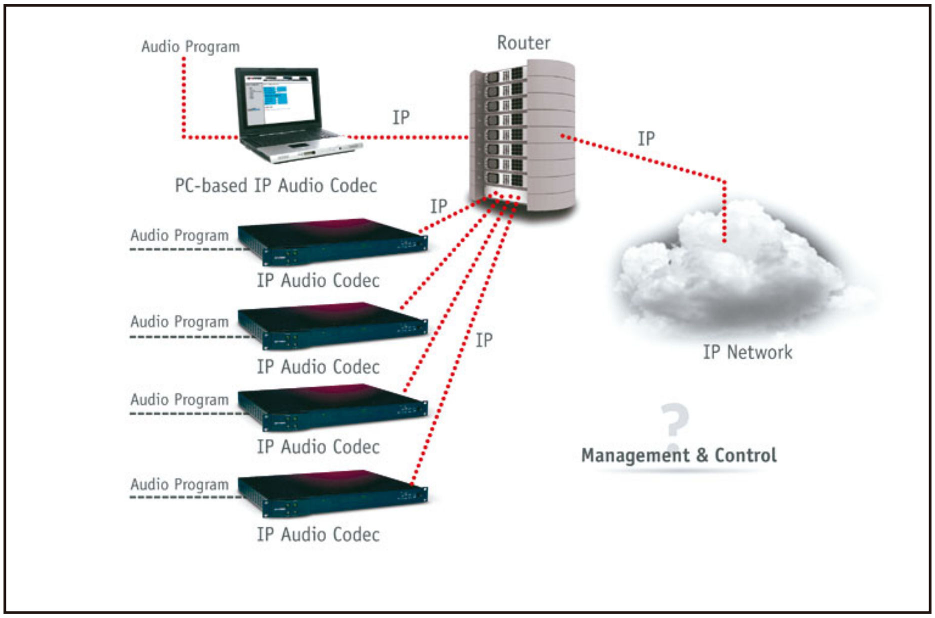

Traditional approach to transporting audio over IP networks

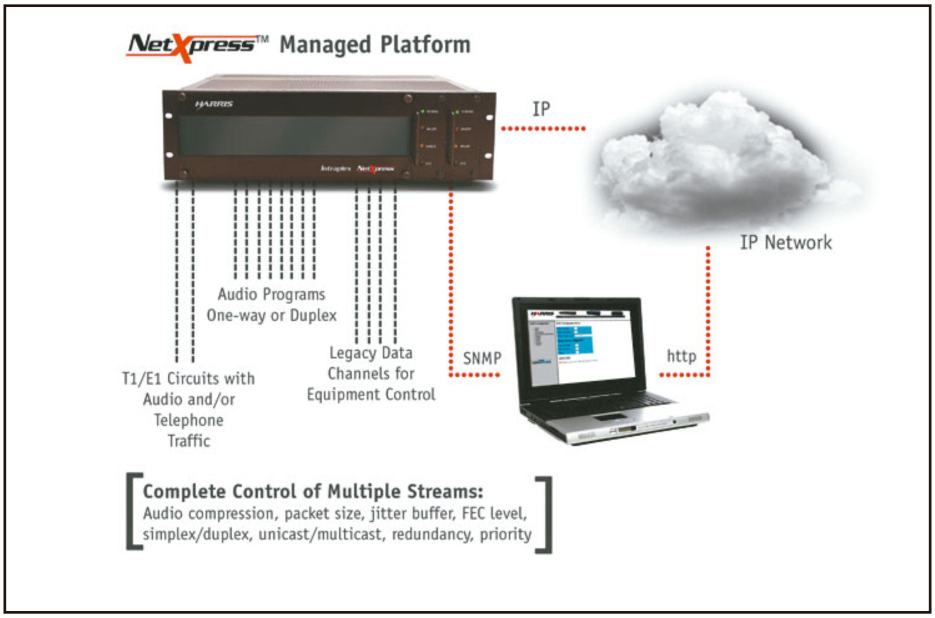

Managed platform approach to transporting audio over IP

A low-level SLA is fine for TCP traffic such as file transfers, where if a packet is missing, the receiver simply asks for it to be sent again. However, in real-time media transport, there is no tolerance for the system to wait while requesting a packet to be re-sent, so a higher-level SLA is required.

Despite everyone's best efforts, there may still be times when for one reason or another some packets do not arrive at their destination. One way to deal with this is through the use of Forward Error Correction.

FEC rearranges the transmitted bits and adds redundant data so that a missing packet can be reconstructed based on nearby packets. The amount of FEC needed depends on the conditions in the network — especially the last mile — and an analysis of traffic over time.

The right amount of FEC can take a reasonable network and make it quite robust. On the other hand, too much FEC is a waste of bandwidth, as increased FEC levels mean increased overhead. Typically, more FEC also means larger buffers, and larger buffers mean larger delays. An ideal system allows the user to adjust the amount of FEC on each stream to best suit the actual conditions.

Finally, there is the issue of unicast vs. multicast. Basically, a unicast IP stream is equivalent to an E1 point-to-point service; packets are sent from one address to another, intended for receipt by that location only.

A multicast stream, on the other hand, is made available to the network on a special IP address that allows an unlimited number of users to gain access to it. This is ideal, for example, for distribution of a program from a studio to multiple transmitter sites. It also is highly bandwidth-efficient; the multicast protocol allows the same program to be made available in multiple locations without the need to reduplicate the bandwidth more than once on any given link.

APPROACHES TO IP AUDIO TRANSPORT

Thus far there have been two basic approaches to IP audio transport: simple standalone audio codecs, and "homegrown" systems solutions, often involving a combination of software and hardware from multiple vendors. Stand-alone codecs do not scale easily as needs and technologies change, while homegrown solutions grow rapidly in complexity and present significant problems in the area of management and control.

Harris has taken a new approach to IP audio transport: the Intraplex NetXpress managed platform.

What, in this context, is a managed platform? It is an integrated combination of hardware and software that allows broadcasters to handle multiple audio programs, each with the appropriate type of audio compression (or none at all); integrate with existing T1/E1 services; carry legacy data circuits; and connect to PABX telephone systems.

It provides T1/E1 circuit emulation, and the ability to carry T1/E1 circuits intact over IP networks, thus allowing users to move in a gradual, controlled manner from legacy fixed-bandwidth services to the new world of IP — on the same managed platform.

GETTING IP RIGHT

As we discussed earlier, IP circuits present special challenges in transporting real-time program audio reliably. The right tools can help overcome these challenges:

Redundancy. Hot-standby power supplies and support for redundant networks, automatically switching traffic to a backup link (where provided) in case of primary network failure.

Packet size. Control of packet size is a key tool in network management. The optimum packet size for any given stream is a tradeoff between latency (delay) and overhead. Where low delay is critical, set the packet size to a minimum, and the packetization delay can be kept below 1 ms — but this will create more overhead. Conversely, where bandwidth is tight, use larger packets. This will incur more delay, but requires less overhead.

Jitter buffering and resequencing. As stated above, packets may not arrive at the receive end in the sequence they were sent. A deep, user-adjustable jitter buffer accommodates the real-world conditions encountered in IP networks, with the ability to retain and restore out-of-sequence packets.

FEC. Forward Error Correction allows the system to rebuild lost or dropped packets, restoring the original audio quality before playout.

Network statistics. Provide cumulative and current statistics on packets sent, received, lost and delayed on every stream, allowing you to see what's going on in your IP network and adapt to changing conditions.

Event logging. Alarms and failures, on both the network and hardware side, are recorded for maintenance and review.

Stream type. Every stream can be set up for one-way or duplex unicast, or for multicast, for maximum network efficiency.

Priority tagging. Support for multiple levels of prioritization, ensuring that real-time audio traffic moves through the network ahead of less-critical data.

Input policing. Provide transport of LAN traffic along with the audio channels, and police the input to limit the amount of bandwidth available to the LAN, ensuring it can never create undue congestion levels that might interfere with the audio packets.

Legacy telephone support. The ability to carry traffic from traditional telephone systems, linking PABXs and establishing Off-Premise Extensions without moving to VoIP telephony. To this end, it incorporates built-in echo cancellation and supports traditional telephone signaling.

MANAGEMENT AND CONTROL

A managed platform cannot be truly managed without an effective means of monitoring and control. Accordingly, the management interface on NetXpress is an element manager based on Simple Network Management Protocol, with access to all aspects of the hardware and network environment.

SNMP is the most widely accepted protocol in IP network management. It provides a means to monitor and control many different types of network devices, allowing the user to change configurations and collect performance statistics.

A big advantage of using SNMP is that it allows the system to communicate with higher-level network management software, such as Harris NetBoss or HP Openview, making it part of a complete operations management system. The basic architecture is that of client/server; the NetXpress acts as a client and maintains a Management Information Base, a database of information about itself and its internal processes, which it shares with the network manager.

NetXpress also contains a built-in Web screen generator, which means it can be viewed and operated via any Web browser, with appropriate security and password control.

SUMMARY

IP networks offer the possibility of highly flexible, low-cost audio transport, and when properly managed can offer the degree of reliability required for use in contribution and distribution networks.

But there are many challenges to be overcome in ensuring this level of reliability. Radio broadcasters armed with a suite of software and hardware tools for redundancy, error mitigation, quality of service and network performance monitoring can more effectively tackle these challenges. ■

For more information, please visit www.netxpress.harris.com