Evertz IXG

The Internet as a Tool for Broadcast

Version 1.0

The Problem

Professional video users (content providers, service providers and content producers) have a fundamental challenge: the ever increasing bandwidth requirements of video content. Standard definition video uses approximately 270 Mb/s, but now high definition video is ubiquitous (1.5 Gb/s), and 4K is already here, going up to 12 Gb/s. To optimize the transport of this content compression is utilized, but there are limits. An encoder can only compress a professional video signal so much until it no longer looks professional. Video transport generally falls into two categories: infrastructure and remote. Infrastructure is permanent; examples being headquarters-to-affiliate and studio-to-transmitter. Remote includes sports events and news gathering. Then there are hybrids (e.g, a full-time link from the US White House to a news network). Regardless of use, large quantities of video need to be transported. Today, typically video is transported via fiber links, managed Ethernet, microwave or satellite. But each of these has its challenges:

- Fiber is reliable and has massive capacity. It is relatively inexpensive locally in major metropolitan areas (once installed), but over long distances the costs go up exponentially. If fiber is not already in-place, installation often takes months.

- Managed Ethernet is another option. Generally managed Ethernet requires connectivity to the service provider onsite if it exists, but depending on the application can be costly to transmit just a small amount video content. Prices for managed Ethernet connectivity from providers are falling continuously, and the benefits of prioritizing video Quality of Service is compelling, however, for certain applications it can be costly especially if 10GbE or more is required.

- Microwave requires an initial investment in equipment including vehicle costs for transport of the system, plus the ongoing expense of highly trained personnel. The available spectrum for transmission is limited, exacerbating the ongoing frequency coordination problem, along with the challenge of limited range.

- Satellite shares with microwave the need for equipment, remote trucks and trained personnel, but it can span continents and oceans. It also works very well for “one-to-many” applications. But there are limited satellites with finite transponders, and the law of supply and demand means that the cost can be high.

An Alternative: The Internet

Today IP networks are pervasive, and most of them are connected to the Internet. These networks have the capability of moving enormous amounts of information, and thanks to the Internet that data can go almost anywhere. At the same time, professional video infrastructures have already transitioned to IP. This sounds like an opportunity come true for professional video users. Can they just use the Internet and significantly reduce network transmission operating costs?

Historically it has depended on how timely the video delivery requirements were. The Internet does an excellent job of transferring files, hosting web conferences, and email. But these are relatively low bandwidth applications. The Internet is a multidimensional chain and each link in the chain adds latency. For example, trying to transfer an hour of compressed video at 6 Mb over an Internet link with a throughput of 3 Mb would theoretically take two hours. Everyone is familiar with viewing compressed low bandwidth streaming video over the Internet, and the occasional buffering due to one’s limited connection bandwidth, and/or the limitations at the transmitting end. Everyday users will tolerate this inconvenience, but for professional video users this is a non-starter except for exceptional circumstances where there is no other alternative.

Private IP Links vs. the Internet

To avoid the problems of the public Internet, video professionals have moved to managed IP connectivity. The advantage is that users can prioritize Quality of Service over the IP network, and can make sure it has the necessary capacity by directly managing it themselves or from the provider. A consequence of this means hiring highly trained IT professionals in addition to the purchase or lease of IP equipment and access. The advantages of IP video greatly outweigh the costs within a facility, but to expand across a city or between cities means leasing a more costly IP service, and at continental distances the cost becomes unsustainable for certain operators. Conversely, high-bandwidth connections to the Internet are available at a reasonable cost, and with the Internet distance becomes a non-issue.

Managing the Unmanaged with ARQ

TCP (Transmission Control Protocol) is the default Internet layer-3 protocol. It has the advantage that it checks to make sure that every packet is received and is 100% accurate, and if not the packet is resent. This sounds like it solves the problem, but it is not that simple. The basic problem is latency. If a packet fails and needs to be resent, there is an intrinsic time delay. Then there is the fact that TCP is always trying to send data faster. If a packet fails, the sending end assumes that the bitrate was too high, so it cuts the rate in half, and then starts ramping up again. The advantage is that this automatically seeks out the highest usable bitrate without network management. The disadvantage is that a lot of available bandwidth is wasted sending packets that have no chance of being received. Since streaming video cannot tolerate such delays, it traditionally uses the UDP protocol instead of TCP. This uses a "best effort" method, with no checking of a received packet. With traditional streaming video there is no point in checking for packet receipt, since by the time the sending end can determine that a packet was dropped it is too late to resend the packet. In a managed IP network, UDP works fine, since IT professionals are making sure that these video packets have high priority (this is known as Quality of Service, or QoS), and that there is plenty of capacity on the closed network. It is not possible to do this on the inherently unmanaged public Internet using the "standard" protocols.

A family of protocols known as ARQ ("Automatic Repeat reQuest" or "Automatic Repeat Query") solves this problem. ARQ works by a monitor at the receiving end that looks for missing or errored packets, and requesting a resend from the transmit end. (TCP does this as well, but an ARQ system uses more modern, patented code to perform that function much faster.) ARQ uses this data to intelligently adjust the sending speed of the UDP packets, resulting in far fewer lost packets than basic TCP. Thus ARQ provides a degree of management to an unmanaged network.

Adding the “Internet” to the transport conversation

With an ARQ solution, remote production now has a fourth option: the public Internet. This has the advantage of being the lowest-cost connectivity option. In many cases an Internet connection already exists; if not, getting one installed is usually more straightforward than a dark fiber or managed IP solution. Or by replacing a microwave or satellite link it eliminates the need for a truck deployment. While ARQ is not a universal solution, it does give video producers one more tool.

For remote production, there is a simple way to try a new technology: back-up. For a live event it is normal to provide two different paths unless it is simply not practical. These need to be as diverse as possible; typically, two “legs” of the fiber / IP / microwave / satellite triad are used. Using two of just one solution could be a recipe for disaster:

- **Fiber/IP:** A dig-up or inadvertently disconnecting the wrong fiber jumper can take down multiple circuits unless route diversity has been provisioned for.

- **Microwave:** These are often aggregated at a remote location; what happens if that location fails?

- **Satellite:** Solar flare activity can undermine reliability across multiple satellites.

Using ARQ over the Internet provides a fourth possibility, increasing route diversity. A back-up route is the perfect place for a remote video production house to implement a new technology.

Even for video between established locations (such as a network and affiliates), ARQ over the internet offers a disaster recovery mechanism. Typically, a television network distributes its programming to affiliates via satellite; for a one-to-many solution over a continental scale satellite is clearly more affordable than fiber, and due to range limitations microwave is simply a non-starter. Equipment failures happen, but with an ARQ backup system in place, such outage is a nuisance, not a crisis.

Evertz’ ARQ Solution: IXG

The IXG is Evertz’ product platform that offers ARQ management of unmanaged networks. There are two versions: the 7890IXG, which is a plug-in module, taking up one slot of a standard 7800-series frame. Then there is the IXG Cloudbridge, a 1RU platform with much higher capacity. The 7890IXG can handle up to eight video streams (up to four transmit and/or four receive), while the IXG Cloudbridge can handle up to one hundred video streams (up to fifty transmit and/or fifty receive). In both cases, licenses are used to establish the actual available number and directionality of video streams. 7890IXG transmit licenses allow up to four transmit-side video streams, and receive licenses allow up to four receive-side video streams. As used here “transmit” and “receive” refer to transmitting and receiving over the ARQ link.

A 7890IXG has two ASI ports, one transmit and one receive. It also has an SFP cage for IP up to 1 Gb/s. Additional considerations include the number of streams the 7890IXG can process (up to four each direction), the number of transmit and/or receive licenses installed, and the bandwidth capacity of the interfaces (ASI and 1 Gb/s IP).

Broadcast Arrangements using IXG in Cloud

The IXG Cloudbridge has four physical 1 Gb/s IP metallic interfaces, and licensing is the same as the 7890IXG. IXG Cloudbridge can be placed at either end of physical ARQ links, but it can also be used “in the middle,” at a Service provider’s location (In the Cloud). This would be either co-located (owned by the end user) or as part of the Service Provider’s service (owned by the Service Provider, with capacity leased to their customers). This arrangement would typically be used for one-to-many distribution arrangements, such as a television network headquarters distributing program content to affiliates. The advantage is that only the bandwidth required by the initial video stream is used by the originating location minimizing bandwidth costs associated with sending content to multiple destinations.

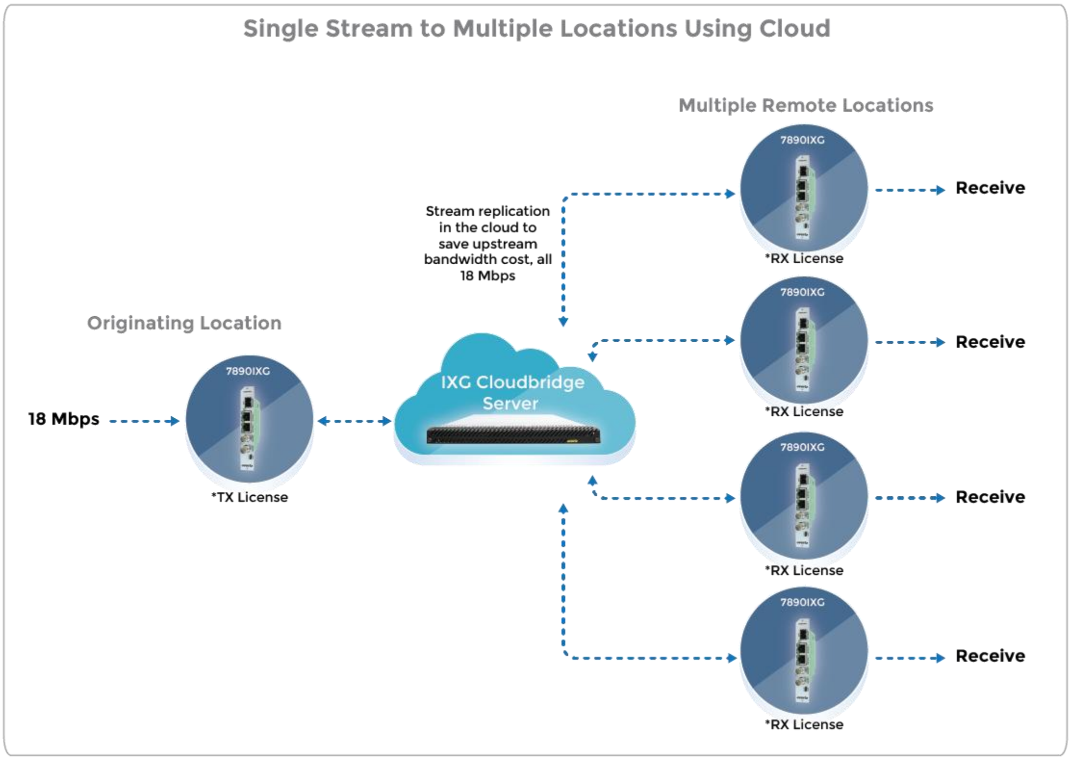

The connection from the origination point to the Service Provider may be either via a leased line (that is, a direct point-to-point connection) or as part regular Internet access. The leased line is better if possible, since that eliminates the possibility of oversubscribing with one’s “regular” Internet traffic, and especially since it avoids having two ARQ connections back-to-back (which would double the latency).

An Internet Service Provider can offer this as a service economically, since a single IXG Cloudbridge could be shared by several of the Service Provider’s customers.

Use Cases:

The simplest case of ARQ using IXG equipment is the basic “point A to point B” transport.

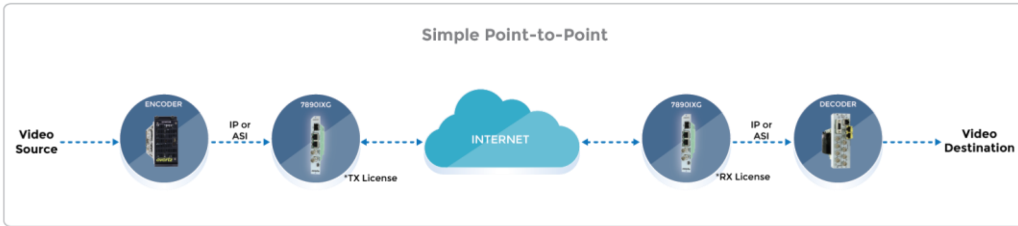

The video encoder is set to use a multicast address, which the transmit-end IXG subscribes to. The transmit and receive IXGs communicate over the Internet (or any unmanaged network) using standard unicast. The receive IXG sends out the packets to the decoder using a multicast address, which the decoder subscribes to. It is important to use the same multicast address at both the transmit and receive ends, for organizational purposes. This same addressing scheme is used regardless of how complex a network of IXG becomes.

Video can also be sent bidirectionally:

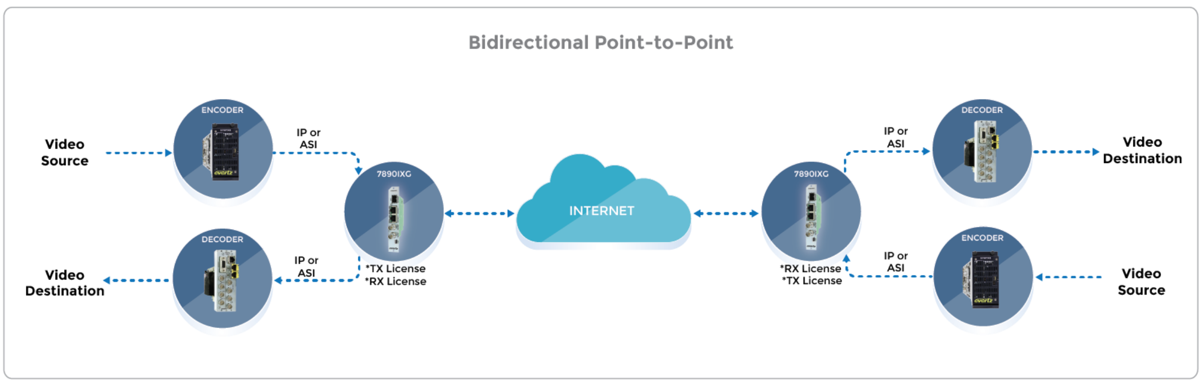

Since a 7890IXG supports up to four video streams, a single pair can handle a bidirectional application.

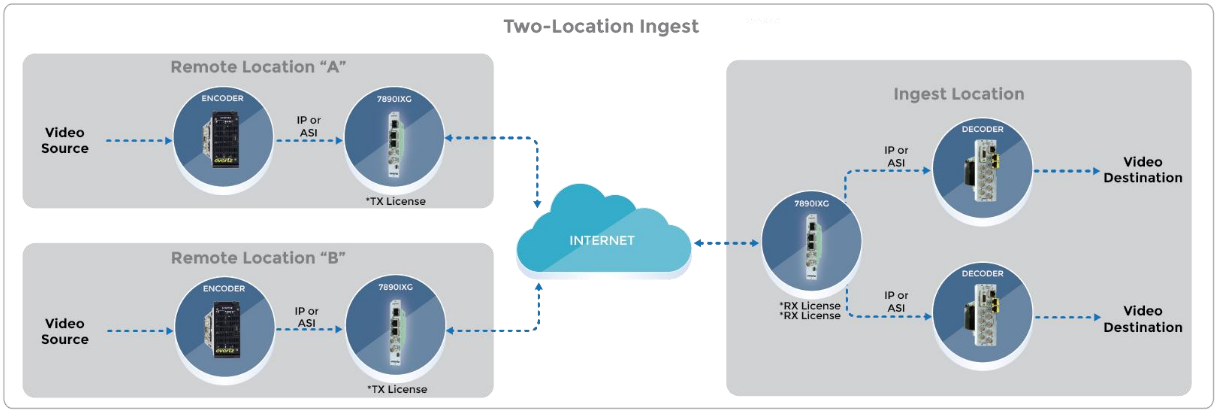

In practice, such ingest is generally used on a much larger scale, such as a cable head-end ingesting local broadcast or PEG channels (Public, Education and Government channels mandated by local governments). For that, the ingest location could use multiple 7890IXGs (depending on the bandwidth required), but in practice often using a single IXG Cloudbridge server would makes more sense:

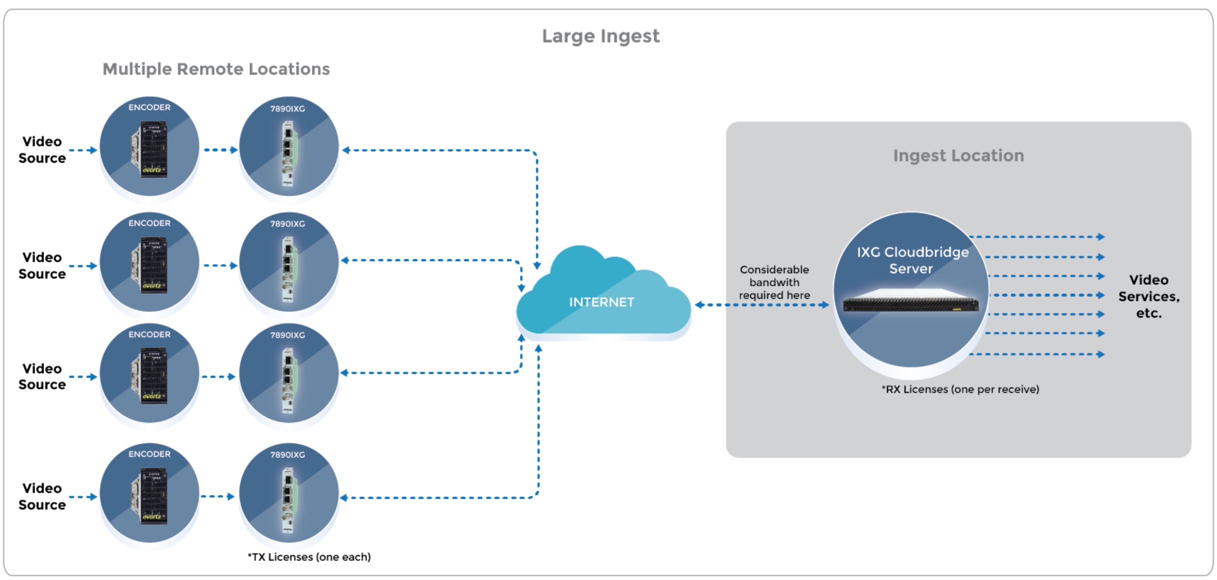

Conversely, a single location can send video to multiple locations.

The above example could also be designed with the Broadcast Location-to-Service Provider link via the Internet using an additional 7890IXG at each end, but as explained above that would double the latency (by using two ARQ links in series to get to each Remote Location).

Or one might be sending <u>multiple</u> video feeds to multiple video locations. A simple example would be one location originating four video feeds, and sending each feed to a different location:

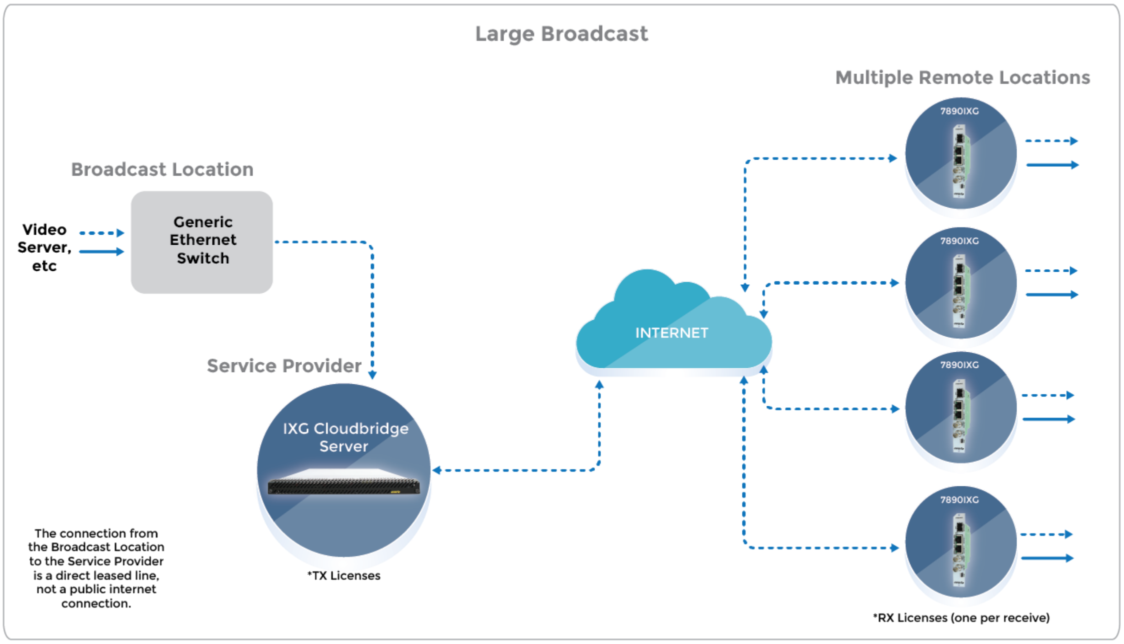

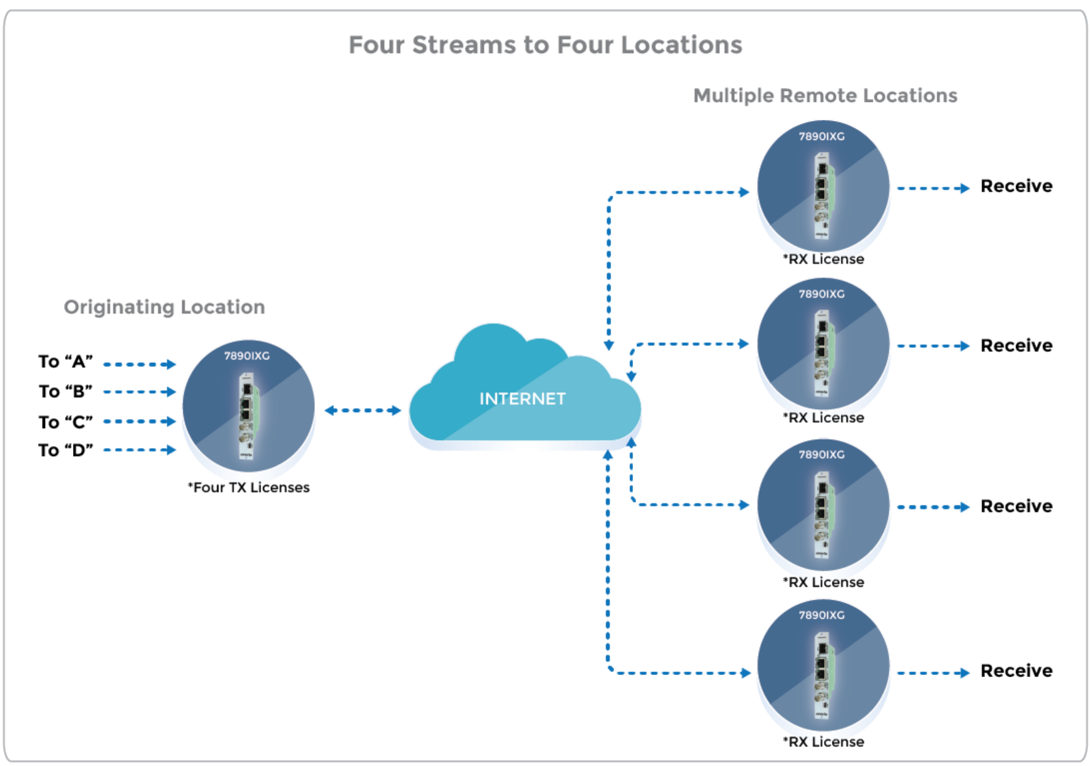

Note that these are simple designs for illustration purposes; there is literally no end of possible video network designs.

There Is Always a Catch – Mitigating the Inherent Issues:

ARQ is not a cure-all ; there are issues that need to be managed. For starters, there is network capacity. Just because an Internet connection rated at a given bitrate does not guarantee throughput at that rate. The actual throughput varies. The further one's ISP is removed from the Internet backbone, the more of a problem this is. Doing speed checks between points should be part of the planning process before settling on a design.

Then there is latency. To give ARQ time to resend packets, there is buffering, which means adding delay in addition to typical compression delay for your codec of choice. How much latency? That depends on the bitrate, quality and reliability of the connection, and the user requirements. The better the Internet connection (both ends and overall throughput), the less buffering one can get away with. Is this sending out broadcast programming from a network hub to affiliates? Then several seconds' latency is probably not an issue. Is it a live remote interview? Then latency should be kept to a minimum.

In the IXG, all of these parameters are easily adjustable using the built-in web-based control system. For users that have an SNMP-based monitoring and control system already set up (such as Evertz' VistaLINK-PRO™), both versions of the IXG also support SNMP-based control. The user has a choice: either dive into the details and set exactly how an IXG link will operate, or just provide generalized guidelines and let the equipment handle the details. Normally these settings are made once when the path is established; after that, the equipment just sits there and works.

Options:

The IXG products have two optional features: Forward Error Correction (FEC) and encryption.

FEC provides the ability for the receiver to recover packets with one or (usually) two bit errors, meaning these packets do not need to be resent. But, there is a serious downside: the bitrate increases significantly. For ARQ, it becomes a trade-off: how often does a packet have a bit error? If there are a lot more bit errors than ARQ can compensate for then FEC is worth the additional overhead if available Actual link testing will determine how the equipment works with FEC on and with FEC off, based on actual performance.

AES128 encryption is available on all IXG products. Since content is being transported across the public internet, security of the video content is a key requirement.

Conclusion:

As technology improves, solutions that were once impractical become practical. Passing professional-quality live video over unmanaged networks such as the Internet is now not only possible, but routine. This drastically reduces transport costs, offering a viable alternative to leased lines via fiber, IP, microwave or satellite. The Evertz IXG products offer an easy way to capitalize on these new capabilities.