THE BROADCAST BRIDGE
Connecting IT to Broadcast

Flexible Contribution Over IP

Essential Guide

EG

ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

Contribution circuits have changed beyond all recognition. Dedicated telco circuits custom designed to transport SDI and AES video and audio across the world are taking a back seat to IP. And even IP circuits have changed in recent years as terms such as managed and unmanaged are gaining a life of their own.

Adoption of IP is probably the biggest change we've seen in broadcasting since the very first television trials with 405-line black and white transmissions. Not only have we fundamentally changed the transport medium from synchronous to asynchronous, but we've also abstracted away the timing plane. Questions of why we still use field blanking, and line timing are constantly being asked, and these are just two examples of the historical legacies that the broadcasting community is still saddled with, and that's before we start talking about interlace.

As interesting as these hypothetical discussions maybe on how we reinvent television to remove the archaic backwards compatibility requirements, we must still work with existing broadcast infrastructures. Few broadcasters have the option to throw away the book and start again. Consequently, we must make IP work in real-world environments where integration and backwards compatibility are not just a nice-to-have but are actually essential.

We have some further challenges for contribution as the SDI and AES circuits of the past have SLAs that are orders of magnitude better than IP circuits. However, the price we pay for this is a massive cost that would make many productions non cost effective. Although we have IP circuits available to us that do have good SLAs through the provision of managed circuits, they still present challenges with security.

Unmanaged circuits, such as the public internet, provide incredible convenience at generally modest costs. But as with all things engineering, there is always a price to pay and a compromise to be sought, and with the public internet, the price we pay is a service that does not guarantee bandwidth, latency, or reliability. But, in most cases, there are work arounds, and when combined with managed services, we have a contribution network that has the potential to reach a level of reliability that would surpass the needs of most broadcasters.

However, security is still an issue that we must take very seriously, especially with the public internet. Even managed services have the potential to be intercepted and viewed. But we must be careful to not fall into a nihilistic depression as our SDI and AES circuits were not as secure as we may have imagined. There was nothing to stop a telco mistakenly routing a signal, or even recording it, I'm not suggesting they ever did this, but these are possible scenarios. The good news about IP is that we have a fundamental mind shift as we assume the IP stream can be intercepted, consequently we have several strategies to reduce the risk of this happening.

Low latency continues to be an area of interest for broadcasters. But again, we must take a pragmatic view on this and although we aim for low latency, we must put this into context. What is low in terms of contribution circuits? The new breed of visually lossless codecs certainly helps keep latency as low as it can be.

Tony Orme.

IP contribution is demonstrating its value not just in terms of cost, but also for reliability and flexibility. Especially when a vendor provides network solutions that keep latency low, maintain scalability and flexibility, and embrace interoperability.

Tony Orme
Editor, The Broadcast Bridge

Flexible Contribution Over IP

By Tony Orme, Editor at The Broadcast Bridge

IP connectivity delivers flexibility and scalability but making the theory work often requires integrated solutions that are adaptable, open, and promote interconnectivity. These challenges are further compounded when we introduce the concepts of managed and unmanaged IP networks, especially as the public internet is becoming increasingly utilized.

Traditional broadcast workflows focused on contribution and distribution often had the advantage of point-to-point connectivity that guaranteed bandwidth, latency, and redundancy.

As we progress to IP, and with it the flexibility and scalability that is provided, the attributes of the static SDI/AES and analog circuits can no longer be taken for granted. Instead, we must look at IP more from the IT perspective to deliver the promised advantages.

IP connectivity is available as managed or unmanaged services. That is, bandwidth, latency, and security are guaranteed to varying levels depending on the service level agreement with the supplier contract. This requires broadcasters to think laterally about how the various attributes and parameters of how the contribution circuits are provisioned.

And the main requirements of many broadcasters are flexibility, low latency, and reliability. Furthermore, security is playing an increasingly important role for unmanaged services.

Security

The high barrier to entry for traditional broadcast systems often resulted in high levels of security being implemented by default. For example, the cost of a VT machines was a barrier for most casual criminals as playing a Digibeta tape often relied on the procurement of a $50K machine and the expertise to go with it. However, in the IP world, the cost of entry for cybercriminals is much lower. Consequently, we need to protect high value media even more.

The media flow as a whole can be encrypted using systems such as AES (Advanced Encryption Standard) and BISS (Basic Interoperable Scrambling System). Both AES and BISS encrypt the video and audio flows directly. This has the advantage of reducing the risk of anybody sniffing and accessing the content but does mean users needing access to the content must be in possession of the relevant keys.

AES is a generic encryption system and uses a symmetric key encryption meaning that both the encoder and decoder use the same key. This has the advantage of being faster and more resource efficient than asymmetric key encryption, but one of the drawbacks is that secure methods of key management must be adopted. Furthermore, new keys must be regularly generated incase one of the key users inadvertently misplaces the key.

BISS2 was developed by the EBU specifically for broadcasters and four modes are specified: Mode 0, Mode 1, Mode E and Mode CA. They vary in their complexity and security depending on their application. For example, Mode 1 was designed specifically for DSNG, fly-away, and emergency type applications and is the fallback mode for BISS2 compliant media exchange.

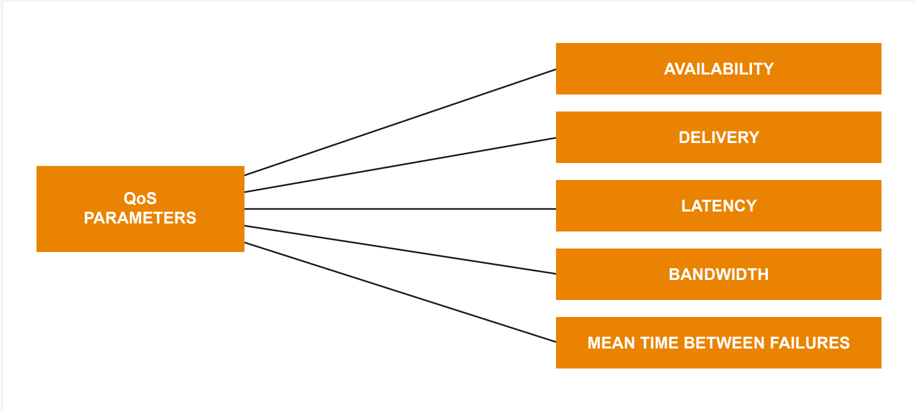

Figure 1 – QoS metrics for IT networks covers many parameters that broadcasters have traditionally taken for granted when working in SDI and AES networks.

A 32-character number is shared by the sender and receiver, known as the Session Word and is manually entered into the encoder and decoder, allowing the media flow to be encrypted. Although this provides a good level of encryption for the media flow, Mode-CA takes this a stage further by encrypting the keys. It uses both symmetrical and asymmetrical keys to combine a complex key encryption and SW exchange system to improve secure exchange of keys and allow an encoder to target specific decoders. Thus, providing a high level of media protection.

IP Opportunities

One of the fundamental strengths of IP is that the data packets are hardware agnostic. That is, they have no "knowledge" of the type of hardware infrastructure they are being transferred over. This could be ethernet or fiber, or RF and WiFi. The ease with which IP packets can be routed between different infrastructures further adds to its strengths.

There are flags within the IP header that indicate the higher level protocols the IP packet is aligned to, such as UDP or TCP, but in the most part, the actual data being caried is independent of the actual IP data section. Again, this further adds to the flexibility that IP has to offer as we can transfer any type of data we like, whether its control, data files, or streamed media.

The power of IP has laid open many opportunities for broadcasters, including the provision of IP services for contribution services used in OBs. Traditionally, OBs have used dedicated SDI and AES circuits provided by Telcos, satellite and RF links. Each of these has their own challenges with lack of flexibility and high costs being at the top of the list. IP services, often provided by Telcos are more flexible and lower in cost than the traditional methods of delivery.

IP and QoS

One of the challenges of the provision of IP services is that they can be either managed or unmanaged. A managed service provides compliance with a much more tightly specified QoS (Quality of Service). This includes packet loss, bit rate, data throughput, latency, and jitter. However, managed services are not always available so a broadcaster may have to work with the unmanaged service instead.

The QoS metrics are important as attributes such as packet loss, latency, and jitter can have a massive effect on the QoE (Quality of Experience) for the viewer. Picture freezing, video dropout, and audio distortion can all have a significant impact on the QoE leading to viewer complaints, resulting in them switching over to another service.

IP Latency

Packet loss not only leads to potential picture break up for video and audio distortion but can have a disproportionate effect on control and monitoring. Although video and audio tend to transfer over contribution networks using UDP/IP, that is fire-and-forget, control and monitoring will use TCP/IP connectivity to guarantee delivery of monitoring and control data. Dropped or delayed packets can initiate the resend-timeout feature leading to large latency occurring between the control and monitoring devices.

One of the unintended consequences of resend-timeout in TCP/IP is that the data rate can appear to be high, as the lost packets are being resent, but the overall data throughput is very low. For any broadcast engineer to effectively utilize IP contribution services they must understand the difference between data rate of packets on the wire, and data throughput provided by protocols such as TCP/IP. Often, they are quite different, which again is a massive difference from how we worked with SDI and AES services.

TCP/IP is an adaption of the ARQ (Automatic Repeat Query) strategy that provides error control in lossy networks such as the internet. ARQ uses UDP/IP packets to exchange data and resend any packets that are lost. Although TCP adds congestion control in addition to ARQs error correction, similar latencies are apparent in ARQ. However, if ARQ is used as part of a custom or proprietary solution, the vendors can tune the ARQ parts of the algorithm to specific applications. This allows them to better stream media and potentially improve on latency.

As a rule of thumb, the tighter the constraints within the QoE metrics, the more expensive the service will be. Although unmanaged services will be less expensive than managed, there is a potential cost associated with this in terms of packet dropout, latency, and jitter, as well as overall reliability. In the whole, it is possible to work with either and possibly both at the same time, but the benefits QoE metrics bring to the contribution network required for an OB must be well understood.

To overcome the QoE limitations of unmanaged services, some form of monitoring is required. This can either be provided manually using network analysis tools, or more productively using automated detection and change over codecs. For example, an OB may be using a managed service for the main contribution feed but an unmanaged service for its back up. An automated system will be able to constantly monitor the networks and switch over appropriately to achieve switchover should one of the services fail without manual intervention.

Unmanaged circuits certainly have their place, it's just that we must be aware of some of the challenges we face when using them. Low level Packet loss and jitter can be overcome using ARQ, FEC (forward error correction) and buffers. However, if a sequence of video is delivered corrupted then the receiver either has to try and fix it, request a re-send, or just flag it as an error. This results in either increased latency or a loss of video quality. Packet jitter has similar challenges, but again this can be fixed with buffers. However, buffers introduce variable latency.

It's important to note that packet dropout, latency, and jitter are a fact of life when working with IP services, even with managed circuits. However, what is important is the predictability of such systems. This is possible with managed networks but less so with unmanaged. It's much easier to work with known and specified latencies and packet jitter within a system.

Determinant Latency

Although latency may be perceived as the enemy of broadcasters, what is more important is determining predictable latency. We can work with 100ms or 200ms of latency, within a few milliseconds of tolerance, what is very difficult to work with is a latency that violently swings between 100ms and 200ms.

It's also worth remembering that we have suffered from dropout, latency, and jitter in television since we broadcast the first transmissions in the 1930s. It's just that the tight timing constraints we've always worked with help keep these metrics so low we barely noticed them. Fast forwarding to the 21st century, there is an argument to suggest that the nanosecond timing developed for SDI is no longer really needed: we no longer use cathode ray tube cameras and televisions, so we don't need to worry about frame accurate timing to within a few microseconds. Modern flat panel televisions and CCD/CMOS cameras are much better at dealing with timing and don't need such tight tolerances.

When providing contribution from OBs we must consider the return path. Video streaming over UDP/IP can theoretically work without a return path, as the data just travels in one direction, the reality is that other applications within the network will be using the return path for ARQ and TCP/IP as well as reverse vision, sound, and IFBs. Again, keeping the latency within tight tolerances helps enormously.

Video Codec Resilience

Video codecs are notorious for introducing latency. It seems that the more efficient the codec, then the more latency that is introduced. This is particularly evident when we use long GOP type compression. For program quality contribution feeds we often use I-Frame only type compression which helps keep the latency more predictable. I-Frame only compression keeps motion artifacts to a minimum and maintains editing and mixing quality in the production gallery as each video frame is compressed in isolation to its neighbors.

The new generation of visually lossless, low latency, and lightweight video compression codecs are helping contribute to the delivery of broadcast quality video over managed and unmanaged networks to studio facilities.

JPEG XS is one such codec that looks to improve upon MPEG and JPEG standards by using wavelet and sub band technologies. Not only does this improve compression performance, but also adds scalability and editability to the feature of tools provided. One of the challenges of JPEG XS compared with J2K or H.264/HEVC is that it requires an increased bandwidth which generally requires JPEG XS to be used on managed circuits.

However, the advantages of lower latency and a lower complexity codec makes software implementation much easier.

Regions of Interest (ROIs) can be defined and encoded to provide a better quality than the rest of the image. The ROI is first decoded before any of the background so that when poor transmission paths are encountered, the decoder can focus on the important areas and fill in the gaps as the data becomes available. Although not ideal, the algorithm works on the principle that it's better to provide data that can create the areas of interest, than providing no image at all, or an image with irrelevant data.

Initially, the image is transformed into the RGB color space using color transforms. Then the images are split into sub bands using block filtering type technology. This creates sub images with varying levels of size and detail to help the codec send the appropriate data for the available network bandwidth. The wavelet transform is then applied to the sub band images to provide image-based coefficients that can be quantized for compression to meet the needs of the HVS (Human Visual System).

Discrete Cosine Transforms (DCTs) are used extensively in JPEG and MPEG compression. Although the DCT doesn't compress the image, it does transform the image from the time domain into the frequency domain. Further processes such as quantization then provides the data reduction to take into consideration the features of the HVS. One of the challenges of DCT is that all the image coefficients must be sent regardless of the available network capacity. This leads to potentially high levels of latency and poor QoE in bandwidth compromised networks.

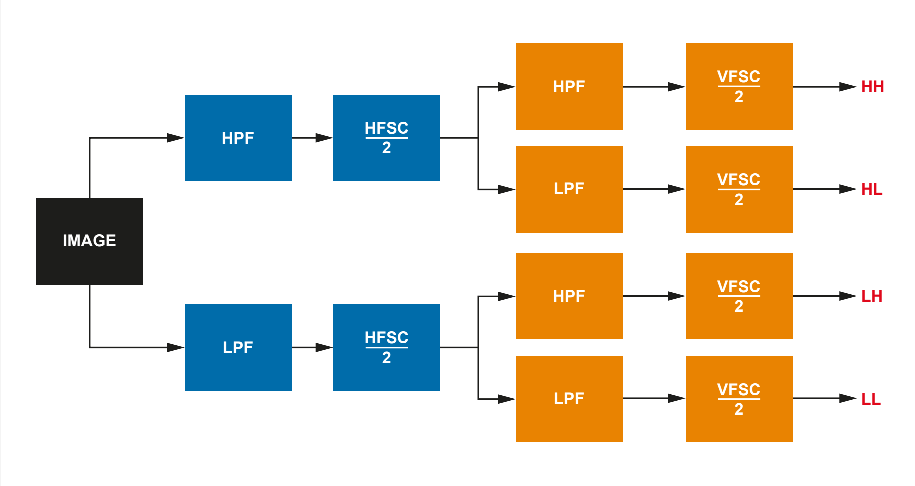

Figure 2a – Four bands are created by filtering the image into high (HPF) and low bands (LPF), then reducing the sampling rate (HFSC and VFSC) by two, to provide the HH (High-high), LL (Low-low) HL (High-low) and LH (Low-high) bands.

DWTs are a form of wavelet analysis that are particularly exciting for broadcasters as they apply their transform to each of the sub band images in isolation. This has the advantage that a complete image, albeit with low resolution, can be sent and decoded with the detail being added as it becomes available. This differs from DCT systems often used with MPEG compression where the whole image is sent as N x N blocks, thus requiring the whole images worth of all coefficients to be received before an image is reconstructed.

The power of DWT for two-dimensional image processing can be fully appreciated when the sub images and hence the sub bands are better understood. The algorithm not only provides multiple sub images with varying degrees of detail, but also reduces the horizontal and vertical resolution throughout the process, thus making better use of the available network bandwidth.

Figure 2a and 2b show how a level-1 sub sampler decomposes the original image into four sub images, all one half the vertical and horizontal size of the original image. The LL (Low-low) band image is the result of a low pass and sub sample in both the horizontal and vertical domains. If this was the only image that was sent to the decoder (because of insufficient bandwidth), then the decoder would have to up-sample the image by a factor of two so that it matches the original size. A viewable image would be provided by the decoder but it would lack much of the detail.

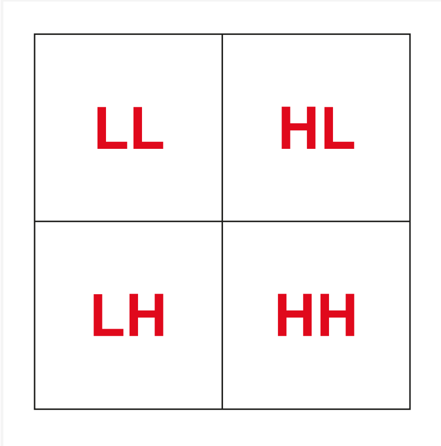

Figure 2b – The original image of N x M pixels is reduced by subsampling the horizontal and vertical samples to provide the four sub band images from figure 2a. LL is the image that is sent first to provide the base image, then HL, LH and HH are sent assuming the network bandwidth is available.

Assuming the three other sub band images could be transmitted, the decoder would double their height and width and then add them to the base LL image, thus providing an image similar to the original.

It's worth reiterating that compression hasn't taken place until the DWT takes each of the sub band images, determines the coefficients of each image, and then applies the quantization. It's the application of the quantization that provides the compression and the resultant coefficients are sent to the decoder. The decoder then reverses this process to create the original image (or a close approximation to it).

Sending each of the sub band images represents part of the true power of JPEG XS and similar wavelet compression systems. The DWT is highly optimized to use the least amount of memory and processing possible, and sending multiple compressed images of the original, each adding a layer of granularity to the image, provides a much lower compression latency. Because of the sub band derivation, wavelet compression systems can adapt to the amount of bandwidth available on the link to optimize the pictures being sent.

Providing flexible contribution circuits over IP managed and unmanaged networks encompasses many different disciplines. From low latency codec choice to secure auto switching managed and unmanaged circuits. With modern developments, this is no longer the onerous task it used to be, and fully automated solutions are available to deliver highly flexible, secure, and low latency contribution circuits.

The Sponsors Perspective

Video And Audio Transportation

By Geoff Bowen, Chief Technology Architect, Appear

The concept of remote production – moving raw content generated at a site event back to the main facility for production and management – has been rapidly gaining popularity in the broadcast world.

Mobile/site content has also been gaining wide popularity with an ability to offer viewers a range of sporting events and other on-site broadcasts via traditional over-the-air, cable, and streaming options.

In parallel to the expansion of consumer portals, the range of content formats in use is expanding with HDR 1080p and UHD contribution becoming common for major sporting events and even 8K production deployments either live or being trialed.

When COVID-19 hit, broadcasters were faced with a dual dilemma – their viewers wanted even more remotely generated content to make them feel more connected while they had to stay at home, but the broadcasters needed to keep their own personnel safe while generating that increase in content that consumers demanded.

Personnel Safety drove two key technology shifts: First, a rapid acceleration of remote production deployments. The 2021 Olympics in Tokyo was a solid case in point. Broadcasters sent 39% fewer people than would be usual for such an event, with much of the production being handled back at their primary facilities or via a remote production company, who in turn had to operate at site with reduced personnel. Sending fewer people not only kept broadcast personnel safe; the broadcasters immediately noted a dramatic reduction in costs as the amount of on-site equipment plummeted and the need for transportation, housing, and other site expenses went down as well. Increasing the volume of on-site broadcasts for viewers also became more viable because of the ability to centralize production, which allows the use of the same production crew to handle multiple events on the same day.

The second key shift was that the facilities themselves were staffed at much lower levels, thanks to the ability to utilize domestic internet connections, compression technologies and secure remote access tools to enable staff to effectively perform operations and engineering functions from home. These 'at home' workflows required a shift away from the compression technologies in use over managed networks towards lower bitrate codecs, protected by ARQ mechanisms such as SRT, Zixi and RIST that enable content to tolerate the packet losses expected over domestic internet connections as well as provide toolsets for encryption and traversal of firewalls within minimal configuration.

While remote production deployments continue to grow, many broadcasters are still puzzling over their options because it isn't as simple as many would have them believe. Everything depends on the telecommunications infrastructure that will be in play for each event. Traditional sports arenas are no problem: there will be dedicated fiber links with all the high-speed bandwidth you need. But what about the non-traditional locations? The problem is that not all sites are created equal: some still have highly reliable infrastructure, such as a pro team sports stadium with dedicated fiber, while others offer different types of connections that may differ wildly in the level of bandwidth and amount of equipment that can be connected at once. No one wants to plug their video and audio feeds into a shaky internet connection and end up with an unusable product (and possible penalties for not providing the contracted coverage requirements). Obviously, a bit of homework is required before planning a remote event to determine how the feeds and backhaul will be handled.

For venues and facilities with access to high bandwidth, high reliability IP connectivity, content contribution can leverage codecs that offer extremely low latency performance, such as JPEG XS, which has seen rapid adoption recently. JPEG XS encode and decode latency, can be so low that a full return path workflow between two sites can introduce less than 1 frame of delay compared to completely uncompressed delivery, making it ideal for productions with talent in multiple geographical locations as conversation flow is more natural, or as a tool for achieving the lowest latency 'glass to glass' workflows.

There are still considerations beyond codec and bitrate selection. JPEG XS compressed video can currently be carried either in a SMPTE ST2110 workflow encapsulated as SMPTE 2110-22, where video, audio and ancillary data are carried as separate essence flows, specified as VSF-TR08, or within an MPEG transport stream, where the essences are multiplexed into a single flow, specified as VSF-TR-07.

The essence-based nature of SMPTE ST2110 has desirable benefits in production workflows. However, it can be complex to handle and monitor and generally requires specialized equipment at both the send and receive site in terms of PTP to provide synchronization of the essence flows, usually from a GNSS locked grandmaster clock and PTP aware switch fabrics. This can present challenges such as antenna positioning with line of sight to satellites in buildings or underground locations. In locations with SDI hand off, provisioning the IP infrastructure may be cost / space prohibitive, especially for small flyaway packs and, as such, may drive a technology decision to utilize the TR-07 transport stream-based encapsulation method. A thoughtfully designed encoder and decoder implementation can be adapted to uphold the low latency characteristics of the JPEG XS codec, enabling TS to be used without latency penalty over ST2110.

SMPTE ST-2110-based deployments popularly utilize SMPTE ST2022-7 diverse path redundancy to protect against packet loss and path failure within the IP fabric. This also needs planning (and testing) to consider the effect of total loss of one path. Can the alternative path deliver the content with zero packet loss, or do we need to apply a degree of FEC to protect against low levels of loss in this scenario?

If diverse paths are not available, a single ended workflow may mandate use of a low latency packet loss protection mechanism such as FEC. While there is a bandwidth overhead to be considered, this doesn't appreciably increase latency in the manner that ARQ-based mechanisms can.

Is the underlying network capable of utilizing multicast for delivery to multiple endpoints from a single encoder? If unicast delivery is required, can the encoder deliver multiple unicast instances of the encoded signal without the need for external NAT or replication services?

Additionally, cloud platforms are growing in popularity for live production workflows, initially driven by necessity during the pandemic and now maturing with software toolsets aligning with capabilities of on prem hardware-based solutions. Contribution of linear video into cloud workflows has traditionally utilized lower bitrate codecs such as AVC, HEVC or NDI augmented by ARQ over public internet. However, increased availability and reduced cost of high bandwidth circuits from venues and broadcast facilities now means utilizing ultra-low latency codecs such as JPEG XS as a method of ground to cloud contribution is also possible.

If reliable bandwidth is available but is constrained, more complex codecs such as AVC or HEVC would enable more content to be carried at lower bitrates while maintaining very high levels of visual quality. The trade off is latency, with widely supported low latency AVC / HEVC workflows adding approximately 700ms of delay to an encode / decode path compared to an uncompressed or JPEG XS workflow.

Even lower latency can be achieved with ultra-low latency implementations of codecs such as HEVC ULL. This approach does not produce a traditional GOP with I/IDR frames, P and B frames. No complete Intra frames are used. Instead, the encoder uses GDR (Gradual Decoder Refresh) as opposed to IDR (Instantaneous Decoder Refresh). This technique is often referred to as stripe refresh. While this approach enables workflows with end-to-end latency of less than 200ms, it also does not leverage the efficiency of GOP-based encoding and, as such, needs to run at higher bitrates compared to traditional AVC or HEVC encoders. Still, it consumes dramatically less bandwidth compared to JPEG XS- or JPEG 2000-based systems.

There is no defined standard yet for HEVC ULL to enable cross vendor interoperability, so for now at least, the same vendor encoder and decoder is required.

Even when reliable and managed bandwidth is assured, contribution functions utilizing public internet can make for a cost-effective continuity strategy and therefore can be utilized to supplement managed deployments.

Many managed contribution networks are dedicated to the purpose of media carriage and, as such, tend not to utilize content protection, favoring the most efficient use of bandwidth and lowest latency achievable. There are however numerous use cases where content traverses a private but mixed-use IP fabric, such as a corporate IT backbone. It may be desirable or even mandatory in this circumstance to apply encryption to the content to prevent possible interception. There are numerous techniques available to achieve this from simple passphrase protected encryption to RSA encrypted session keys for each receiver.

Multiple types of signal hand-off are required at the compression edge. Many facilities and trucks now utilize an uncompressed ST2110 IP routing core and NMOS control layer, while others utilize SDI. Often the contribution solution needs to service both forms of hand-off within the same workflow depending on what is available at a given location, requiring a contribution solution with flexible I/O encompassing both electrical and optical SDI and IP connectivity capable of supporting uncompressed UHD I/O. Other workflows may not decode a compressed signal back to baseband and, as such, require tools to transcode content and perform processing in the compressed domain, police media flows, provide NAT and multicast / unicast conversion capabilities.

Before planning an on-site or cloud production workflow, ask what types of connections are available and what sort of bandwidth can be expected from each. If it's IP-based, find out whether backups exist in case the main feeds fail. Another critical step is to determine the level of security on those connections. If they're visible to hackers, they can be easily disrupted.

One piece of equipment that can help mitigate your bandwidth issues is a good-quality compression platform, which provides low-latency compression and decompression functions to make video easier to transport over IP. Some remote video equipment includes cursory encoders, but those encoders may not be up to the task when presented with varying IP speeds or other bandwidth issues, and they may not have a comfortable level of physical IP security or content encryption capability. Stand-alone encoders, while admittedly increasing the amount of equipment going to the on-site location, are usually the best choice for equipment connection, IP connection, and high levels of security and content protection tools.

When selecting an encoding platform, make sure it has the versatility needed to handle the available bitrate, video resolution capability, capacity and interfacing needed for the available connection. A good platform will offer standards interoperable format support, control APIs, encryption solutions, physical and content redundancy models, and should also provide robust firewall and traffic policing capabilities. Appear has made all of these needs a "must" in our solution portfolio.

THE BROADCAST BRIDGE
Connecting IT to Broadcast

CI
CORE INSIGHTS

EG
ESSENTIAL GUIDES

MEDIA

ARTICLES

For hundreds more high quality original articles and Essential Guides like this please visit:

thebroadcastbridge.com

8/2022

Supported by

appear