THE BROADCAST BRIDGE
Connecting IT to Broadcast

Keeping Studio IP Media Streams Reliable

Essential Guide

EG
ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

IP broadcast infrastructures are delivering on their promise of flexibility and scalability. But as we progress on our IP journey, we are now discovering that an integrated monitoring systems is essential.

SDI and AES systems relied on the transport stream providing the timing components for the video and audio essence. Luckily, this made monitoring relatively straightforward as we could usually assume the transport layer was fine as we were more interested in the video and audio quality. Transport issues did arise, especially if somebody stood on the coax or a dry joint dropped one of the balanced pairs, but these would often manifest themselves immediately as errors on the television monitor or loudspeakers, making them relatively easy to fix.

One of the reasons that we may well have taken the transport layer for granted was that the SDI and AES networks used point-to-point connectivity, and had synchronous clocks embedded in them so they were highly predictable but not particularly efficient. Packet congestion certainly wasn't an issue, and we were rarely concerned with latency. However, the price we paid for this was a lack of flexibility. We could build scalability, but this often relied on duplicating key parts of the infrastructure resulting in a very inefficient design that required regular exhaustive testing.

IP packets are data agnostic and independent of the transport stream. Although this provides unprecedented flexibility, it does so at the cost of increased complexity. IP networks soon increase in complexity to the point where we can no longer take the relationship between the media essence, IP packet, and underlying transport stream for granted. In fact, they are so separated that we have had to adopt the PTP timing system. Consequently, monitoring IP broadcast facilities has taken on a whole new challenge.

Just looking at the video essence on a monitor or listening to the audio on a pair of loudspeakers only skims the surface of the monitoring needed for an IP broadcast facility. QoS still plays an important role as packet jitter and latency have a major impact on media essence, but the network as a whole greatly influences how packets are distributed and delivered. This leads to the potential for errors that may not be immediately obvious and may be even more difficult to diagnose and fix.

Combined with the disjointed transport stream and media essence, we have the added complexity of the number of media essence streams traversing a network at any one time. This could easily be hundreds of streams and may even lead to thousands. When combined with the number of SDP files needed to identify each media stream, the potential for problems amplifies greatly. Human error, software bugs, and misconfigurations are difficult to detect and engineers working in modern IP broadcast facilities could easily see themselves overwhelmed.

Tony Orme.

To deal with these challenges we must look at the new generation of integrated monitoring systems designed specifically for IP broadcast infrastructures. The massive number of media essence streams makes the traditional method of manual monitoring almost impossible. And the flexibility that IP delivers does so at the expense of added complexity resulting in broadcasters needing to maintain the relationships between the essence descriptor SDP files, and the media essence itself.

Keeping modern IP media facilities reliable relies on the new integrated monitoring systems that helps us understand and maintain the strict relationships between the IP network, timing, essence, and descriptor data.

Tony Orme
Editor, The Broadcast Bridge

Keeping Studio IP Media Streams Reliable

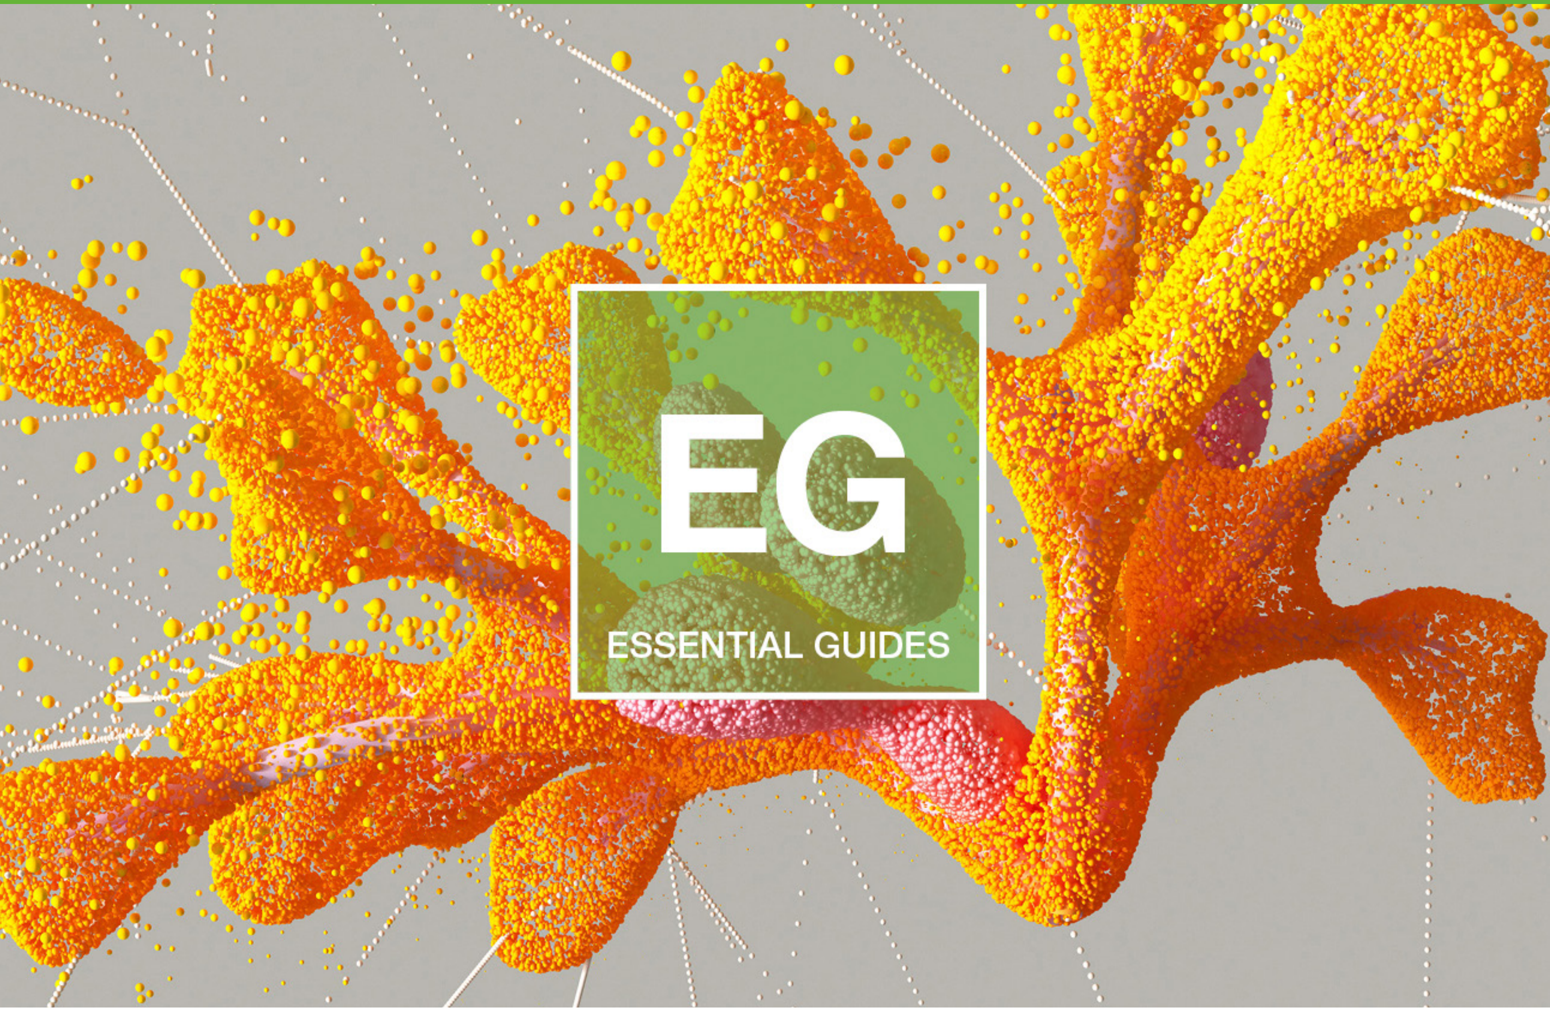

By Tony Orme, Editor at The Broadcast Bridge

Modern broadcast facilities adopting video and audio over IP have found themselves working with thousands, or tens of thousands, of IP streams. Expecting humans to keep track of all these flows, monitor their quality and efficiently fault find is a virtual impossibility.

IP allows broadcasters to abstract away the media essence from the transport stream to deliver flexibility, scalability, and resilience. But in doing so, forces engineers to pay much more attention to the underlying network and transport stream.

This is no trivial job, and the interconnectedness of modern networks demands an automated solution that allows engineers to focus on keeping the video and audio quality high while at the same time having access to tools that allow engineers to observe the system both at high and granular levels.

One of the most powerful aspects of moving to IP is that it allows broadcasters to separate the hardware from the software. Consequently, broadcasters have much more flexibility in how and where they operate their facilities.

COTS and public cloud infrastructures provide an industry standard hardware base thus allowing broadcasters to concentrate on building flexible workflows that meet the needs of the business. The software-based components provide the ultimate in flexibility as they can be enabled or disabled as required. With modern public cloud and COTS infrastructures, all the broadcast media essence and support data can be processed including video, audio, ancillary data, control, and monitoring.

Although software processing is providing incredible opportunities for broadcasters, two challenges manifest themselves, these are human error and software bugs. With flexibility comes responsibility and just because an operator can change a parameter doesn't mean they should, however, this is not always the case.

Building Flexibility

Modern broadcast facilities consist of many interconnected workflows that are often dynamic. Software processing goes a long way in simplifying the operation but there are still opportunities for operators to make mistakes. Often, these go unnoticed and don't become apparent until downstream equipment has difficult decoding the video, audio, and metadata. And these types of errors can be difficult to detect and remedy.

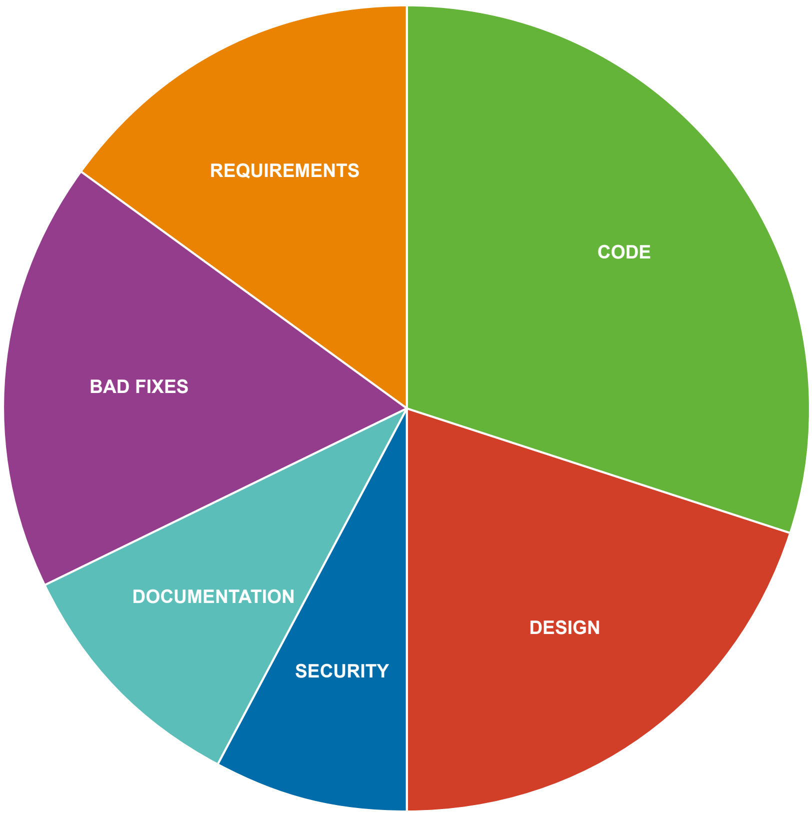

Fig 1 – Providing bug free software relies on much more than just the developers writing the code. Multiple agencies collaborate and even with the best intention, there is scope for human error at many levels.

Software infrastructures are providing broadcasters with massive opportunity for change, thus meeting the ever-demanding needs of the business. However, this places even greater pressure on the development teams which occasionally result in bugs entering the designs. This may not be due to coding issues but due to the potential for mistakes in the whole development process such as an incorrectly specified functionality. Although automated testing systems used by vendors as part of their deployment are progressing quickly, it still takes time to develop test vectors that can test the whole system exhaustively. Microservices have helped enormously with testing as small standalone programs can be tested in isolation leading to improved reliability. But the harsh truth is that software does occasionally contain bugs.

Agile development methodologies have driven software to a rapid deployment model where new features are released within a few weeks. This improves flexibility and delivers greater opportunities for broadcasters but does mean there is a constant risk of new bugs entering the software. Bugs should be thought of as a fact-of-life, not something we should get angry about. It's fair to say that in the ideal world there would be no software bugs, but such an aspiration would mean development times would be shifted back into years, and not the weeks we're now accustomed to. Essentially, we're trading time with convenience and flexibility. And as in all things engineering, there is always a compromise.

One solution to this conundrum is to adapt our mindsets so that we expect bugs to exist and have strategies to deal with them when they manifest themselves. Continuous monitoring of both the audio and video essence, as well as the networks will help with this and in a way, this forms part of the agile mindset.

IP Network Complexities

SDI and AES networks had the advantage that they consisted of point-to-point connections that were relatively easy to understand, debug and verify. Just patch in a scope in one of the SDI segments and whether the segment works or not can be immediately seen. The bandwidths and latencies were well-defined leading to networks that were highly reliable, but static. It is possible to increase routing flexibility with these networks, but they are always limited to a few transport streams, such as HD-SDI and UHD-SDI.

IP networks deliver untold flexibility, not only with their topology choices, but also because IP packets are data agnostic. That is, the payload does not form an integral part of the underlying transport stream, unlike SDI and AES. Consequently, the IP packets can carry any data including media, monitoring and control information. The combined elements of multiple topology options and data agnostic packets form the true power, flexibility, and resilience of IP infrastructures.

However, these opportunities present new challenges in terms of interoperability. For an SDI infrastructure, due to the well specified standards from SMPTE, it's often straight forward to work out the format of signal we are working with. Unfortunately, this is not the case with IP networks. Although ST2110 may have specified the media contents, we cannot be sure the IP stream we're looking at is even carrying ST2110 data. It could equally likely be carrying monitoring information or control data. Even before we start looking at the essence, other issues arise at the network layer including packet dropping due to over subscription of a link, packet jitter, and even packets being sent in the wrong direction.

SDP And Essence

Session Description Protocol (SDP) files help alleviate this as they describe the format of data in a media stream. They describe such parameters as the video frame rate, image size, and source and destination IP addresses, to name but a few. SDP files help the broadcaster's management system understand where the streams are in a network and the format of video and audio they are providing.

The management system or broadcast controller requests the SDP/essence information from a centralized registry (such as the NMOS registry) where all the devices such as the cameras, microphones, production switchers, and sound consoles etc. store and register their information.

Although SDP files prove incredibly helpful for system management, they can inadvertently contain errors due to either human error or software bugs. The main challenge occurs when the SDP file does not match the media essence due to an error in the video and audio workflow configuration. Modern broadcast facilities have potentially thousands of media essences with corresponding SDP files all moving from one stage to another in the workflow. For example, a transcoded media fill will generate a new SDP file, this file and the new transcoded media essence file(s) must be associated otherwise a mismatch will occur potentially leading to decoding errors in downstream equipment.

Essentially, when using these software components, the broadcaster has created a software defined workflow. Any software bugs or errors caused by misconfiguration within this workflow will cause the media essence that is either generated or manipulated by one of the processing stages to be out of specification. For example, a narrow-gapped ST2110 source device may develop excessive jitter due to a software bug which in effect will create a wide-gapped sender style, this will cause any downstream equipment expecting a narrow-gapped ST2110 essence to fail as the video buffers will be sized according to the SDP narrow-gapped information but will probably overflow due to the actual essence wide-gapped packets. The failure is therefore caused by the SDP vs essence mismatch.

To add to this challenge there are potentially thousands of video and audio signals traversing a network at any one time. If we think of a studio camera, this could easily have six input and output video feeds, possibly double that if we include SD and HD outputs. A production switcher may easily have a hundred input and outputs, especially if we consider the aux busses, international and clean feeds. And then there's the multiviewers, as playout channels snowball, the number of multiviewer tiles must keep up. Audio has similar challenges and often surpasses the number of video streams.

Monitoring

The flexibility that IP infrastructures deliver is unprecedented in broadcast facilities. Not only can we distribute video and audio over COTS infrastructures, but we do not need to make special provisions for each of the media, control, and monitoring information being distributed throughout the network. For example, not only would a traditional broadcast facility have to provide separate video and audio routing matrices, but there would also need to be provision for RS422 control signals, GPIOs, and timecode. All these systems comprised their own dedicated sub-infrastructure requiring custom hardware that adds to the cost and complexity of the broadcast infrastructure. Changing these systems is often difficult and time consuming.

Modern broadcast and IT infrastructures predominately use ethernet and fiber transport streams to form the basis of their IP networks. WiFi will be used for bring-your-own-devices, but the streamed media will be limited to highly compressed proxy streams as it's unlikely anybody would try and stream an ST2110 HD stream across a WiFi link due to the latencies and packet jitter involved. Transferring media to the outside world would often require MPLS or dedicated layer-2 networks to provide the greatest flexibility, resilience, data throughput and lowest possible latency. Furthermore, all these physical interfaces are standard connections within the IT industry and so will not require custom connectivity such as that found with SDI and AES.

Monitoring IP streamed media requires more than just monitoring a view of QoS elements such as packet jitter or packet loss. Although these are important and still the primary concern for most IP networks, there are other systems within the infrastructure that need equal attention. Essence layer monitoring, prior to video and audio decoding, includes checking for malformed RTP packets where out of sequence numbering could occur, incorrectly specified ST2110 specifications where the video frame rate in the essence and the SDP file do not match, and ST2110-21 issues where the wide and narrow gap definitions contradict each other causing issues with the transmission traffic shaping. Other conflicts such as poorly defined primary and backup packet timing deltas within ST2022-7 are difficult to detect.

```
s=st2110 0-0-0
m=video 20000 RTP/AVP 96
c=IN IP4 239.10.3.46/64
a=source-filter: incl IN IP4 239.10.3.46 192.168.9.1
a=fmtp:96 sampling=YCbCr-4:2:2; width=1920; height=1080, exactframerate=60000/1001;
depth=10; TCS=SDR; colorimetry=BT708; PM=2110GPM; SSN=ST2110-20:2017;
TP=2110TPN; interlace;
a=mediaclk:direct=0
a=ts-refclk:ptp=IEEE1588-2008:08-00-11-ff-f0-21-91-aa:0
```

Fig 2 – Example SDP file showing the various video parameters such as colorimetry and image size. The parameters in this file must match exactly those in the associated media stream.

With all this in mind, it soon becomes clear that automated and exception monitoring are critical tools for the broadcast engineer. Having the ability to constantly monitor the high-level parameters and compare them to SDP files provides incredible insight into how reliably a network is operating as well as helping to find faults should they occur.

Deep Monitoring

When broadcasters are monitoring their IP networks, they not only have to look at the transport and streaming levels, but also must think about the intrinsic video and audio essence. It's perfectly possible to have a valid IP data stream in terms of the transport and media but have unviewable or distorted pictures and sound.

Integrated monitoring systems that combine network and transport stream monitoring with video and audio essence monitoring allows broadcasters to hone in on a suspect media stream quickly and effectively. This is similar to the using traditional waveform and vector monitors, but they are combined with the media streams in an IP network.

The power of integrated monitoring becomes clear when we realize that the system can monitor tens, if not hundreds of media essence streams at the same time. The alternative would be to configure round-robin monitoring for waveform/vectorscopes across all of the streams, just one at a time. But with integrated monitoring, most, if not all the essences can be monitored and validated simultaneously.

Critical to all media streams is timing. ST2110 employs PTP and any discontinuities in the time base could easily cause picture and sound breakup and distortion. Continually monitoring the PTP reference and its association with the ST2110 packets will quickly flag any problems to the engineer.

PTP must be stable and any adverse fluctuations in the timing reference will cause buffer under and overruns, resulting in dropped packets and distorted video or audio. Having a system that is constantly monitoring the PTP parameters and determining their stability and accuracy is critical for broadcast IP infrastructures. Exception monitoring of timing and PTP will remove the need for the engineer to continuously look for anomalies within the system. Quite often the monitoring system will find problems that the engineer has no chance of detecting.

Systems such as ST2022-7 provide seamless protection for RTP based feeds including ST2022-6, ST2110-20/30 or MPEG compressed streams. Two identical streams of packets are provided by the sender using RTP sequencing. If one stream loses packets, then the receiver can switch to the secondary stream and use the RTP sequence number to seamlessly switch between the two. Although a small buffer is required at the receiver, ST2022-7 proves an incredibly useful solution where high resilience is needed. Also, when considering ST2022-7, a limit for how far apart the same two packets can be on the primary essence and the backup essence must be established. Again, this is difficult to achieve manually but forms part of the core of an integrated monitoring system.

An engineer would find it almost impossible to detect if the receiver was switching between the streams. When this is known, it may point to a more serious issue that must be addressed. Also, as the system is designed to work with lost packets, it's not unreasonable to expect packets to be lost. But if too many packets are lost then an alarm should be raised and the best method of achieving this is not by having an engineer constantly looking at screens trying to determine if a fault has occurred, but instead, using an integrated monitoring system that can be set up to interrogate the receiver to determine if a certain number of losses have occurred so an alarm can be triggered, thus providing another level to exception monitoring.

Also, human errors and software bugs can be detected and dealt with by continually monitoring the contents of the SDP file and media essence parameters. Any differences between the two can be quickly flagged to the integrated monitoring system thus facilitating a speedy resolution.

Conclusion

Broadcast IP infrastructures are flexible, scalable, and resilient, but the price we pay for this is increased complexity. It's unrealistic to expect an engineer to deeply understand all aspects of the IP network, IP flows, and media signals as the systems are just too complex. However, integrated exception monitoring allows errors to be automatically detected and alarms triggered so that the engineer can get on with diagnosing the problem quickly and efficiently.

The Sponsors Perspective

Proactively Monitor IP Video Networks & Essences With Inspect 2110 & PRISM

By Agostino Canepa, Director, Product Management at Telestream.

For over two decades Telestream has streamlined the ingest, production, and distribution of digital video and audio. Today, compared to its SDI/AES-based predecessors, IP video adds exciting new challenges to these workflows.

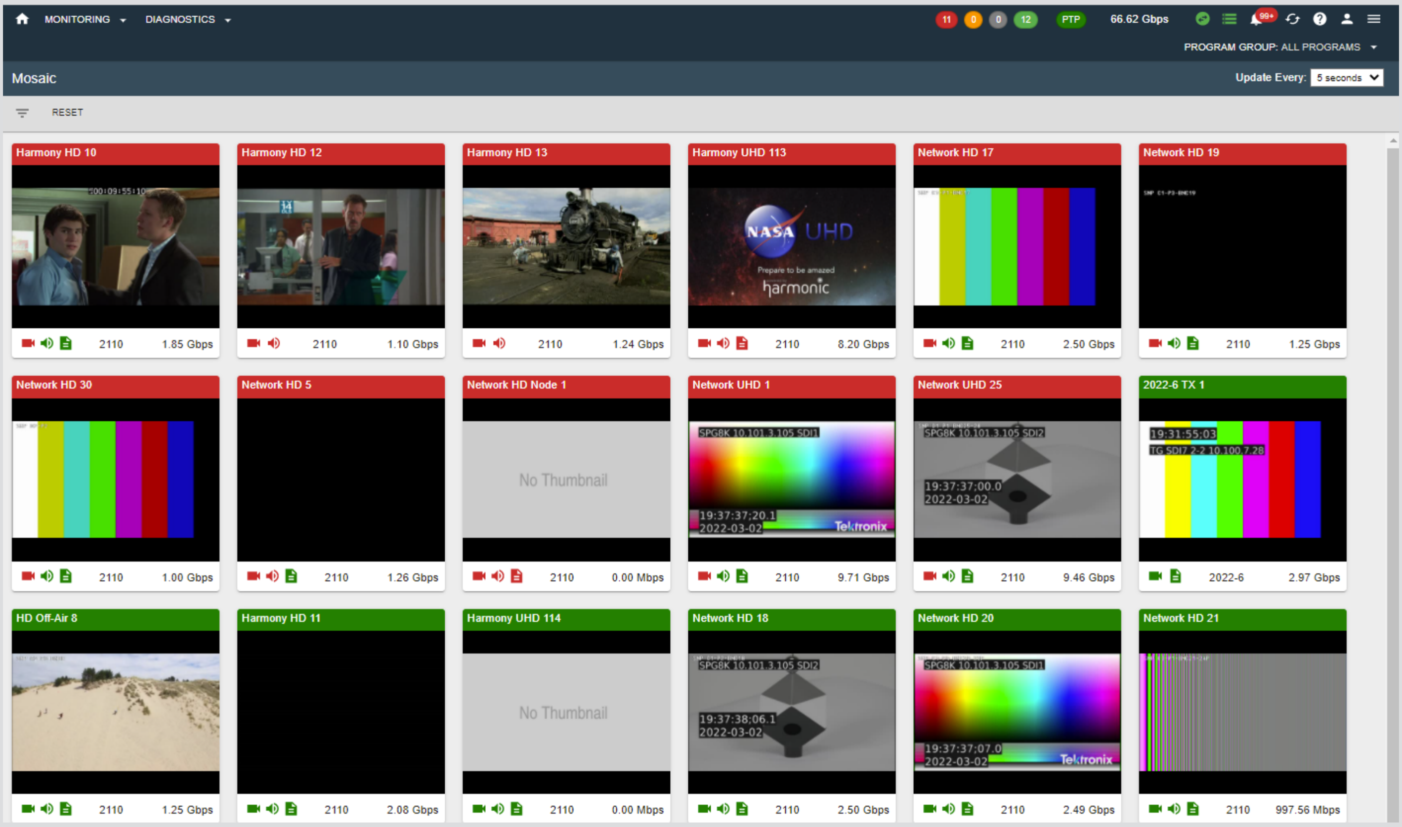

As a protocol alone IP is data agnostic. When used for video, though, there are specific requirements to minimize latency and jitter. A common standard for IP media transport is SMPTE ST 2110. Diagnosing and resolving faults in an ST 2110 IP media network requires network-centric, versus point-to-point, root-cause analysis. When using ST 2110, the video, audio, and ancillary essences are captured and transmitted independently.

These essences must, thus, be tightly and very accurately synchronized.

Stream synchronization is accomplished using the RTP / ST 2110 timestamps present in each ST 2110 packet. This requires higher resolution timing accuracy than a protocol like NTP (network time protocol) which is accurate to about 0.01 sec.

All devices generating, processing, and playing out ST 2110 essences must be tightly synchronized with a precision greater than what NTP can provide. NTP has thus given way to the more accurate PTP (precision time protocol) protocol which is accurate to 1µs.

It is important to properly oversee the quality of content and infrastructure even as these higher speeds and more precise timing requirements challenge the limits of manual human perception, so automating this oversight is crucial. The Telestream Inspect 2110 has been designed to do just that via a “set it and forget it” methodology. Inspect 2110 will tell you when there is a problem. Until it does, an engineer is not necessarily required.

The most streamlined approach to ST 2110 monitoring, diagnostics, and root-cause analysis is to allow engineers to focus on resolving problems quickly rather than manually searching for them. Automating oversight of an ST 2110 Audio/Video network enables engineers and operations staff to perform their day-to-day tasks while software monitors the facility. The monitoring software only alerts staff members when there is a problem and then it directs them straight to the area of concern. Once notified, engineers must be able to seamlessly invoke an IP diagnostics tool to identify and repair root causes. This is the key to reducing fault isolation and remediation time. A properly designed workflow thus allows highly skilled engineers to spend their time fixing problems rather than looking for them. Inspect 2110 enables this by providing the tools needed to maintain an IP video network’s quality and performance.

- **Network health** – Inspect 2110 monitors network traffic to identify bottlenecks and potential bottlenecks, so infrastructure issues can be addressed before they affect stream quality.
- **Timing and synchronization** - Inspect 2110 detects discontinuities within the PTP time base that can lead to breakup or distortion, and it also flags fluctuations in the time base that lead to buffer overruns or underruns.
- **Essence validation** – Once the connection between two video devices is established the sending device transmits an SDP (session description protocol) message, alerting the receiving device of the video and audio parameters of the RTP stream to be sent. Inspect 2110 validates every RTP packet against the SDP and alerts the user to any identified discrepancy.
- **Redundant stream monitoring** – It’s not enough to know you have redundancy built into your system. Inspect 2110 verifies that the redundant stream perfectly matches the primary stream, it is perfectly in sync with it, and it is available. All this is done in real-time.

Many big IP transport issues begin as minor anomalies that are invisible to the human eye. These eventually escalate and result in noticeable quality loss – or worse. Deploying a robust monitoring and alert solution like Inspect 2110 allows video and network engineers to address issues before they are visible to the audience.

While identifying video problems is critical, it’s even more important to get them fixed. Once Inspect 2110 identifies an issue, a single button click allows an engineer to deeply dive into the situation by calling up the Telestream PRISM waveform monitor interface. PRISM includes all the familiar video signal measurement displays and has IP-specific apps that allow for detailed analysis of ANC session data, Dolby status, detailed PTP statistics, and packet interval time (PIT). Only a few seconds elapse between issue detection and the engineer’s ability to perform a detailed analysis of the situation.

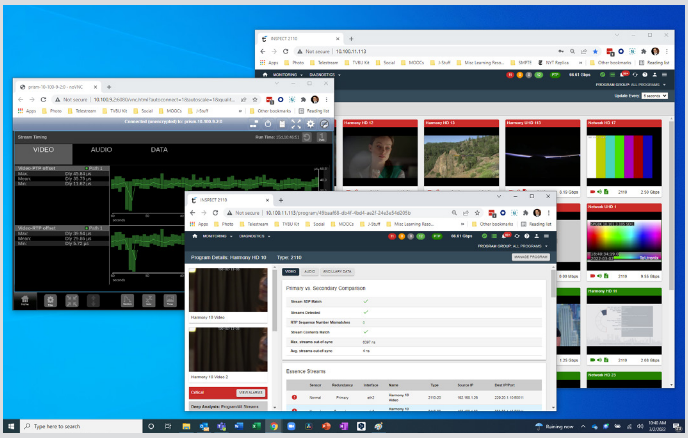

Ease Of Adoption

It used to be straightforward to diagnose SDI issues because you could trace a point-to-point signal path to isolate a problem. Overseeing everything was generally impractical because the cost of routing signals just for monitoring was prohibitive. With IP the reverse is true. You can have visibility into everything that is happening, but without point-to-point connections, and when signal flows are non-deterministic, diagnostics are more challenging. To do the same job in an IP world as in an SDI one requires engineers and their managers to learn and adopt many new concepts and methods. The Telestream Inspect 2110 and PRISM products minimize training time.

Inspect 2110 monitors multiple IP video streams across the production chain or contribution network, and a notification is triggered if any of the video, audio, or ancillary essences does not match the expectations or does not comply with all the applicable standards. Thumbnails provide visual feedback of the video quality, with status icons that include mouse-over alert messages and a direct link to diagnostics – whether in Inspect 2110 or PRISM. Details are but a single click away.

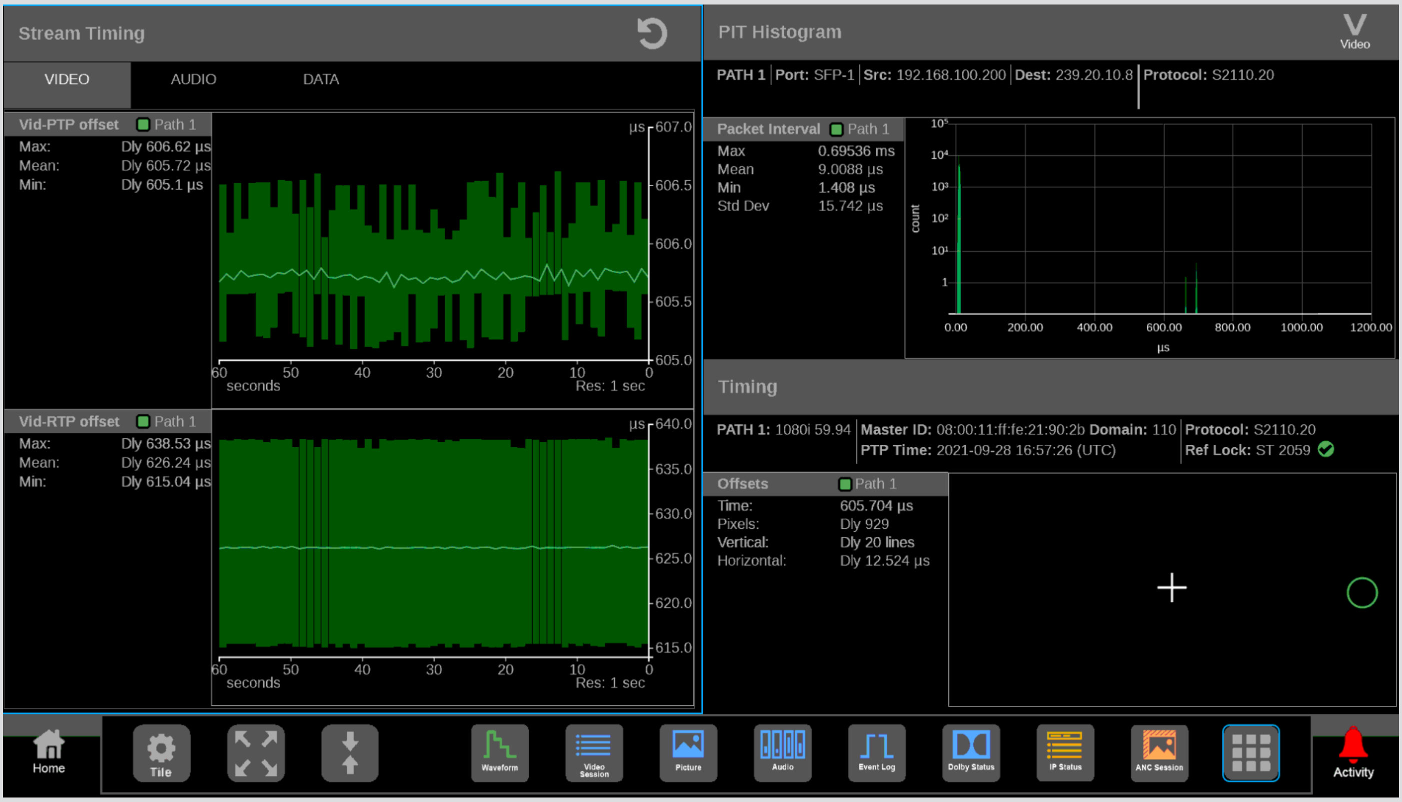

When developing the PRISM applications, Telestream drew upon many of the most recognized interfaces in video engineering circles, further lessening the training time for operators.

Supported Standards For Now And The Future

Inspect 2110 handles all common professional video formats from SD through 8K. Beyond SMPTE ST 2110, the product supports SMPTE ST 2022 (-6 and -7), AMWA NMOS IS-04 and IS-05, and AMWA BCP 002-01.

In February, when the world's first 8K programs were aired for the winter games, broadcasters relied on Inspect 2110 and PRISM for real-time network and signal monitoring. With this proven track record and the ability to update products in the field via simple software updates, the Telestream suite will continue to meet broadcaster's growing needs as formats and standards continue to evolve.

Key Integration Features

The migration to IP-based workflows allows facility designers the freedom to choose best-of-breed solutions. The Inspect 2110 REST, API-first, design allows programmatic access to all the product functionality and all the data collected by the product including essence metrics and video thumbnails. Integration with third-party tools is straightforward for in-house development teams and system integrators. Alerts can be transmitted by most of the common methods including SNMP, Web Push API, Web Sockets, and Amazon SNS.

With such broad support for key video formats and protocols, Inspect 2110 can be used to validate the output of any component in the facility, or broadcast truck, that streams audio and/or video. It eliminates much of the trial and error of diagnosing problems when everything looks like it should be working.

Having a suite of tools from a trusted, standards-compliant vendor to manage and troubleshoot your IP video network is crucial to reap the cost savings and benefits of COTS on-prem solutions and cloud-based delivery systems. The alarm indications provided by Inspect 2110 are a significant driver of efficiency because exception notifications allow your technically savvy engineers to focus on specific issues instead of watching monitor walls. The result is, quite simply, time-to-resolution efficiency.

Telestream has over twenty-five years of experience supplying the world's largest media companies with video and audio transcoding and streaming tools, and we've been providing the world's most trusted video signal measurement tools for even longer. Inspect 2110 represents the industry's most comprehensive set of checks performed on media essences and the most bandwidth-capable platform available today. When working in concert with the recognized leadership of the PRISM family of IP and SDI waveform monitors, the combined solution provides a solid, future-proof investment for any facility making the transition to IP video production.

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

telestream