THE BROADCAST BRIDGE
Connecting IT to Broadcast

PTP V2.1 – New Security & Monitoring For IP Broadcast Infrastructures

Essential Guide

EG

ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

Timing accuracy has been a fundamental component of broadcast infrastructures for as long as we've transmitted television pictures and sound. The time invariant nature of frame sampling still requires us to provide timing references with sub microsecond accuracy.

One of the benefits of adopting IP infrastructures is that we can take advantage of the massive innovation we've seen in the IT industry. Finance, manufacturing and power generation need accurate time references and experts in these disciplines have been working to design and perfect timing solutions operating on asynchronous IP networks for over twenty years.

SMPTE's ST 2110 treats timing differently as it removes the sync pulses and inserts video frames and groups of audio samples with a timestamp formed from the globally recognized PTP standard. By replacing the sync pulses we've not only considerably reduced the bandwidth of the overall signal but have made it much easier to process digital video and audio.

As broadcasters continue their IP migration journey, many are starting to think more about security, and it's become a major area of interest as IP networks expand. Engineering is all about balance and compromise, and as we make our networks more accessible, then we must think more about security.

Operability and robustness are critical for any infrastructure especially when considering the stations timing reference. A system that is used by so many industries soon proves its worth as any anomalies in the specification are quickly found and addressed. But industries such as broadcasting have unique and specific operational needs and so we need the ability to fine tweak generic protocols such as PTP.

Monitoring has always been important to broadcasters due to the complexity of the systems we operate. Television would soon degenerate into chaos if we didn't integrate monitoring to optimize signal levels and measure synchronization. IP networks, by their very nature are scalable and dynamic further adding another dimension to the system complexity. For these reasons broadcast facilities are designed with monitoring at the core of the infrastructure and this working practice doesn't look like it's going to change any time soon.

The IEEE 1588 working group have been listening to broadcasters, especially when considering security, operability and robustness, and monitoring. But what is even more important is the need to maintain backwards compatibility.

IP is still an emerging technology for broadcasters and there will be many changes and challenges along the way. However, it's imperative that improvements to core systems such as timing are backwards compatible.

The IEEE 1588 working group understand this and have made backwards compatibility central to the new V2.1 PTP protocol. Equipment using the existing V2 PTP will still work with PTP clocks working to the V2.1 protocol. To benefit from the added security then any equipment complying with the V2 protocol will need to upgrade, but any V2.1 upgrades to PTP clocks or PTP aware switches will still be backwards compatible with the existing V2 equipment connected to them.

Tony Orme.

Timing continues to be central to any ST 2110 broadcast facility and its more important now than ever to understand how PTP works. We no longer have the luxury of pulling out the scope and looking at sine waves and comparing sync pulse edges. The timing signals have changed, and so have the tools to monitor them.

Tony Orme
Editor, The Broadcast Bridge

PTP V2.1 – New Security & Monitoring For IP Broadcast Infrastructures

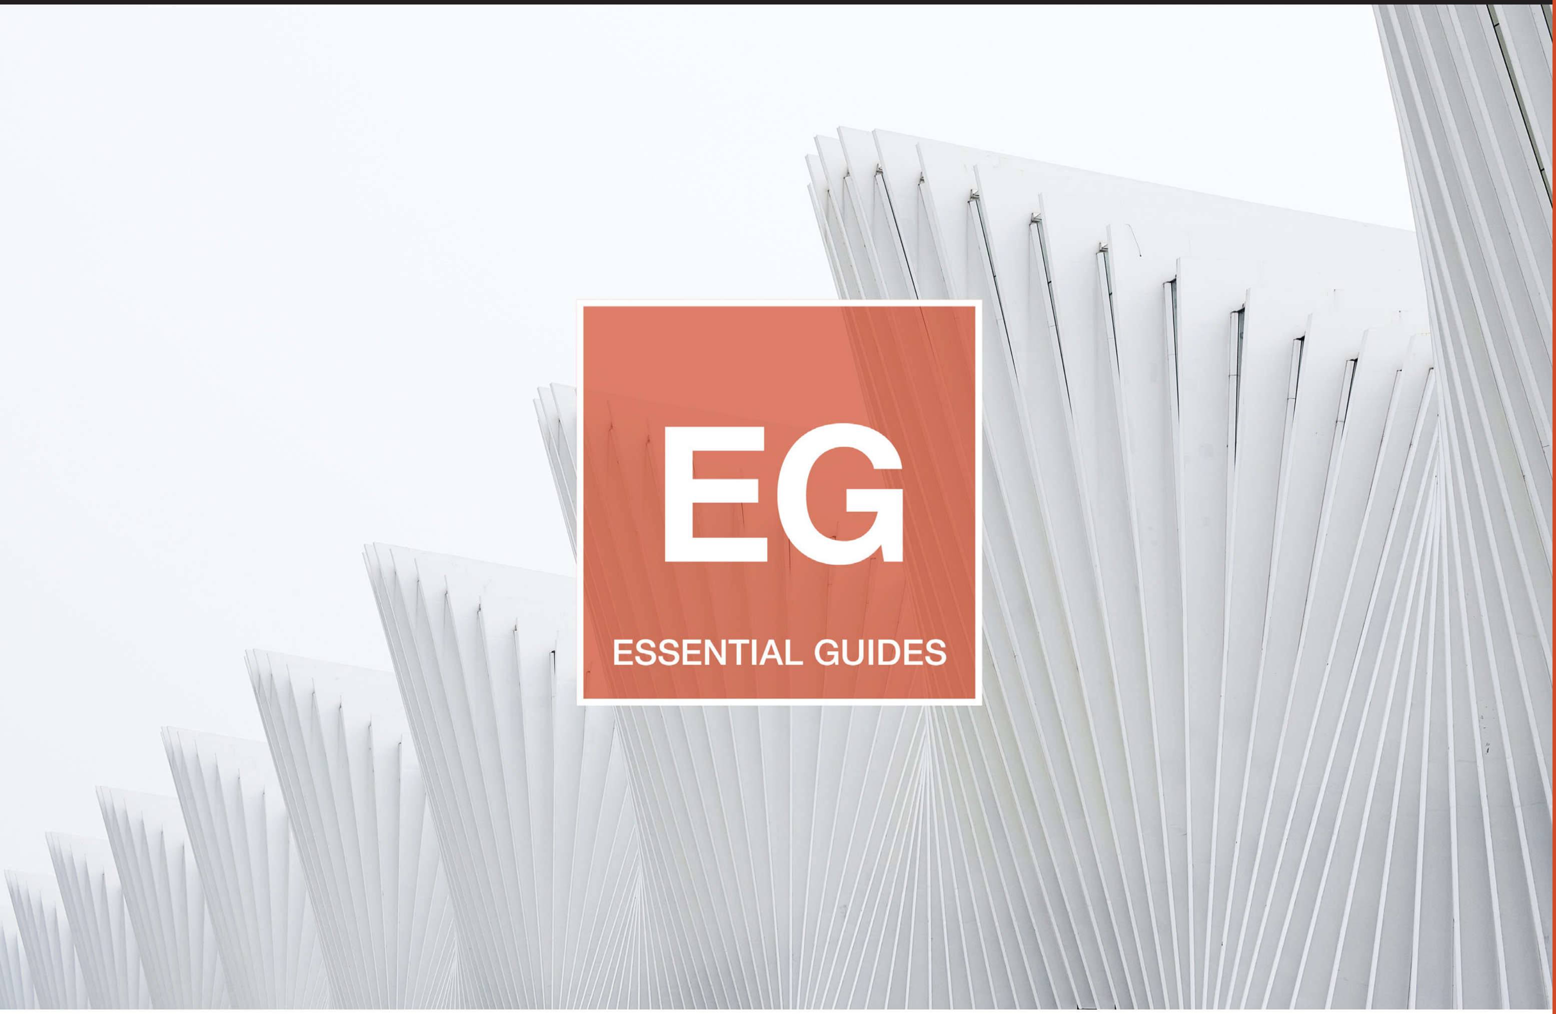

By Tony Orme, Editor at The Broadcast Bridge

As ST 2110 continues to find its way into broadcast facilities, the Precision Time Protocol is gaining greater prominence and with it our understanding of how timing should work in an asynchronous IP environment. The IEEE 1588 working group has been listening to broadcasters during this time and has now released its latest version of the protocol.

SDI and AES have been the mainstay of broadcast infrastructures throughout the world for nearly forty years. They've served us well and provided reliability and consistency, we know how they work, and generations of engineers have grown up with these systems. The downside of these systems is that they are static and difficult to change.

IP seeks to bring new flexibility and scalability into our broadcast infrastructures. We can ride on the back of innovation in the IT industry and learn from them. Although SMPTE's ST 2022-6 was a great stepping-stone for many into the world of IP, the real advance happened when ST 2110 came along.

By abstracting the video, audio and metadata essence from the underlying sync, field, and frame pulse timing information, SMPTE paved the way for moving video and audio from the static SDI and AES systems to the flexible asynchronous IP broadcast infrastructures. Although this system provides us with a great number of opportunities, it also means we have to completely rethink how we do timing.

Fundamentally, television is still a synchronous system based on repetitive evenly gapped video frames and time invariant audio samples. This is a legacy that is yet to be surpassed and is unlikely to change in the near future. Consequently, we need a stable time reference so that the playback speed of the video frames coincides with the record frames of the camera, and the playback samples of the loudspeaker coincide with the record samples of the microphone. In this context, nothing has changed.

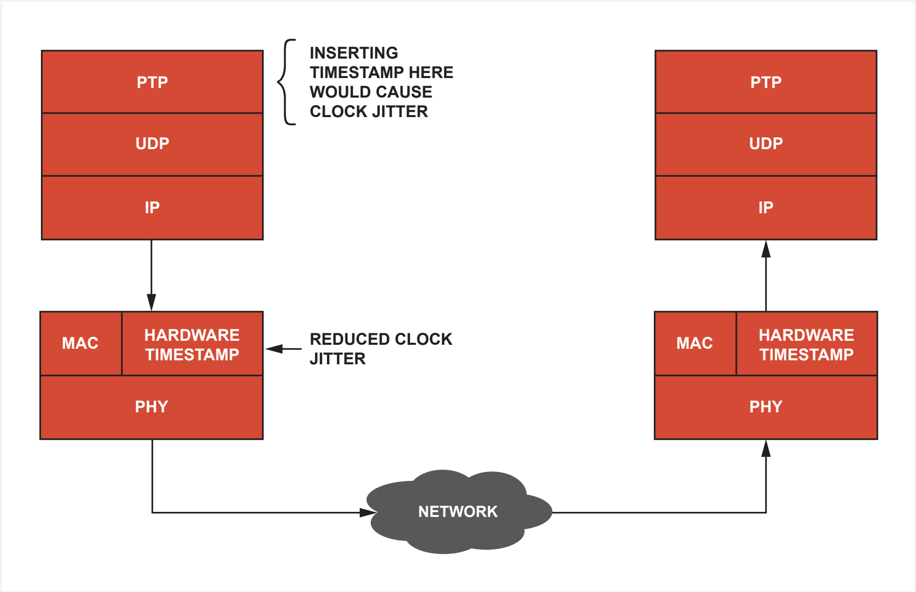

Fig 1 – To achieve maximum accuracy, a PTP timestamp must be derived at the hardware level of the media access layer. In the case above this would be in the ethernet card, although IEEE 1588-2008 or IEEE 1588-2019 does not specify a transport type. If the timestamp was derived and inserted in the software stack, inaccuracies and jitter would result.

PTP has its roots in industry and has proven its ability to deliver a stable time reference to synchronize events. As we progress through our IP integration, instead of thinking of video frames and audio samples in the frequency domain, it might help to consider them as related timed events. This is more intuitive when considering distribution over IP networks as IP is essentially an asynchronous system that facilitates transactional events.

The fundamental reason for using PTP is that devices in a network can share the same time source so that record and playback events can be synchronized. The PTP timestamp is a register that counts the number of nanoseconds from the epoch to the present time, consequently, it provides us with much more than a simple frequency reference.

ST 2110 uses PTP V2, or its formal standard; IEEE 1588-2008. This is version two of the protocol and superseded its predecessor, version one. The latest release, IEEE 1588-2019 is version 2.1. Not only is this backwards compatible with V2 and can be used with ST 2110, but it also provides improved robustness, accuracy and security.

Although V2.1 provides much more functionality than V2.0, the backwards compatibility cannot be over emphasized. Due to the complexity of broadcast systems and the amount of money riding on infrastructures in the way of content, broadcasters like to take small progressive steps. Devices such as cameras, production switchers, sound consoles, etc. using PTP V2, must still work without any anomalies when the network is upgraded to V2.1. The IEEE 1588 working group has achieved this.

Security

Every broadcaster should be concerned with security. For some time, there was a school of thought that suggested security was not really of any concern as broadcaster's networks were protected. This naïve approach is mostly now debunked, especially if we think about the security challenges within any organization and recent well recorded breaches.

Although the IT industry has progressed well with initiatives such as IPsec to address security in networks, their work with network clocks, specifically NTP (network time protocol) and PTP has taken this to new levels. The IEEE 1588 working group has based many of its security motivations on this work to provide well thought out solutions which address and remedy many of the security challenges.

Time protocol attacks can take on many guises, but the most common underlying themes include packet interception, spoofing and data manipulation. It might seem strange that a malicious actor may want to gain access to the timing function but in doing so, and without the correct safeguards, they can cause timing anomalies that are difficult to detect and can cause havoc in a broadcast facility.

The data within the IP packets of V2.0 are sent in the clear. That is, there is no consideration for encrypting or protecting the data within the packet. To maintain backwards compatibility between V2.1 and V2 it is not possible to add encryption to the whole data packet for V2.1. This is because existing V2 compliant equipment would not be able to decode the timestamps. To address this, V2.1 uses the concept of cryptographic integrity check values (ICV).

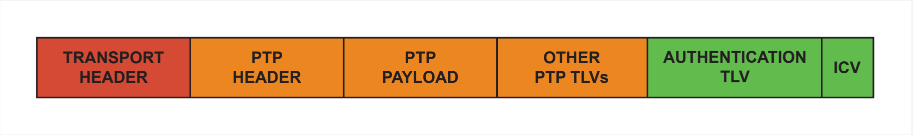

Fig 2 – the PTP payload contains the actual timing information that is sent "in the clear" to maintain backwards compatibility with V2, however, the ICV is calculated and appended to the message using the TLV mechanism so that any V2.1 compliant devices can detect if the timing information in the message has been changed.

If a malicious actor was to access the timing system, then they might want to adjust data such as the timestamp values from the Grand Master in such a way that the internal clock of any downstream equipment, such as a multiviewer, would speed up with respect to the Grand Master. This change might be very subtle, not enough to cause the multiviewer to indicate it has lost sync, but enough for the video and audio to be out of sync with the originator. Resulting problems would be really difficult to detect as there would be no data loss, and the multiviewer would still be showing that it's synchronized. However, as it's synchronized to the maliciously modified values, it might start running faster and would run out of video and audio in its input buffers, thus causing video and audio breakup.

This sort of problem can also happen quite innocently and unintentionally if downstream equipment from the GM was incorrectly configured or a user made a mistake.

ICV will help alleviate this kind of issue. In the ideal world we would encrypt the whole PTP message, but this is not viable as we need to maintain backwards compatibility with the existing ubiquitous V2 PTP protocol. Therefore, instead of encrypting the entire data segment of the message, V2.1 provides a mechanism to provide a form of cyclic redundancy check, or complex parity check on selected parts of the datagram so that we would know if an error occurred or somebody has modified the timestamps. To further improve this, the validation word is encrypted using a secret key. The process of calculating the validity word and encrypting it is the ICV.

This mechanism protects not just against altered PTP messages, but also false messages injected into the network by a malicious agent, since the false messages would not contain an AUTHENTICATION TLV with an ICV created with the correct secret key.

A V2.1 enabled Grand Master would create the ICV and append it through the use of the Type-Length-Value (TLV) to the timestamp message. TLVs are an established method of extending the PTP protocol functionality while maintaining backwards compatibility. Any downstream equipment such as a camera or sound console, would expect to see the TLV-ICV.

When the timestamp packet is received, V2.1 compliant equipment such as sound consoles and vision switchers are able to decrypt the ICV (as they will have a copy of the secret-key) and then compare the validity word to its calculated version of the PTP message. If they are the same, then the timestamp and all the associated data is assumed to be valid and is used. If they are different, then either an error has occurred, or the message has been tampered with in transit causing it to be dropped and an alarm issued.

Addition of the ICV allows any V2.1 compatible downstream equipment to determine if the PTP message is valid and whether it has been maliciously modified or not. Any trusted intermediate equipment such as a PTP aware switch that modifies the data values, for example a transparent clock, will need to know the secret-key so it can decode the incoming message, modify the relevant data parameters, recalculate the ICV using the secret-key, and then re-insert the new TLV-ICV into the PTP message.

Any legacy V2 downstream equipment will still receive the same PTP message with the authentication TLV (indicating an ICV is present), but as it will not recognize the TLV type identifier, then it will simply ignore the TLV and ICV value. This means the V2 equipment will not be able to determine whether the PTP message has been maliciously interfered with or not, but it can still use the PTP timestamp information and all the other associated data to synchronize to the Grand Master, thus maintaining backwards compatibility.

Crucial to this system is maintaining the confidentiality of the secret-key. Although the IEEE 1588 working group does not mandate any particular key management system, the V2.1 specification does provide examples of systems such as Key Distribution Centers and the Group Domain of Interpretation (GDOI) method of maintaining and distributing keys to authorized devices.

As the secret-key is so important to maintaining the integrity of V2.1 PTP messages, the actual key may change daily, or even hourly. This security policy and the maintenance of the keys is outside the scope of the V2.1 protocol but is a system that should be operated in conjunction with the broadcast facility's IT Director.

Robustness

As well as providing better security, the IEEE 1588 working group also wanted to improve PTPs operability and robustness and they achieved this through profile isolation and monitoring.

To maintain maximum flexibility across many industries, PTP uses a system of profiles to quantify many of the parameters that can be configured in the system. For example, although IEEE 1588 specifies the use of the Announce message, it only specifies a generic time interval with which it is sent using the default-profile. This is a base profile common to all PTP devices so they can be tested and measured to the same specification. It's possible to use the default profile but industries such as broadcasting have their own standard; SMPTE's ST 2059-2.

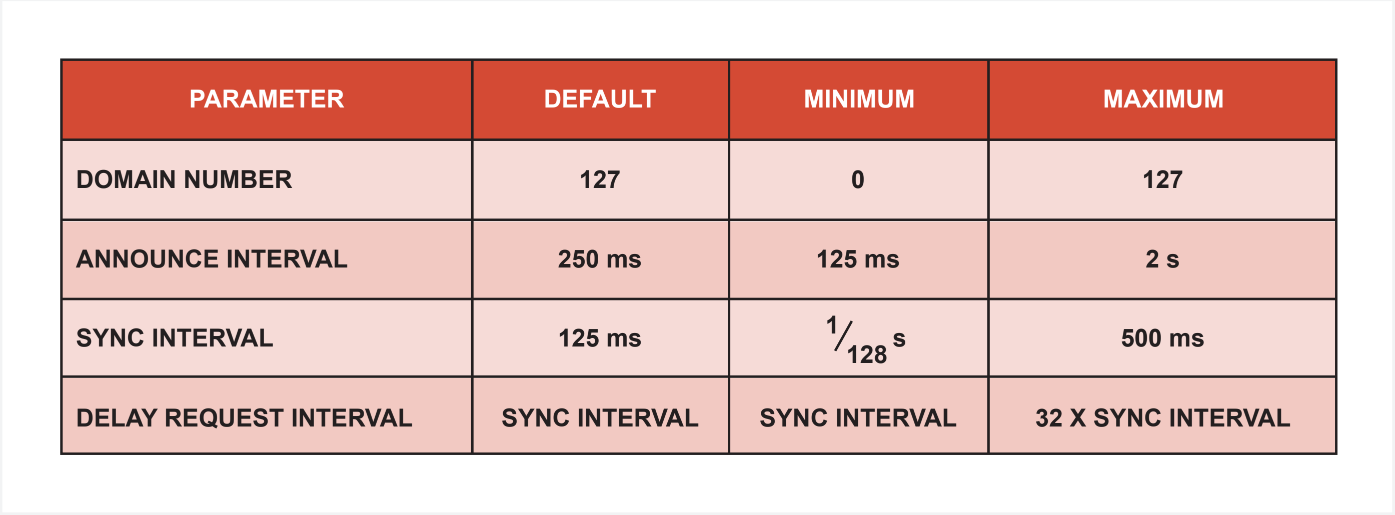

PARAMETER	DEFAULT	MINIMUM	MAXIMUM
DOMAIN NUMBER	127	0	127
ANNOUNCE INTERVAL	250 ms	125 ms	2 s
SYNC INTERVAL	125 ms	$^{1}/_{128}$ s	500 ms
DELAY REQUEST INTERVAL	SYNC INTERVAL	SYNC INTERVAL	32 X SYNC INTERVAL

Fig 3 – the default profile provides a basic configuration to specify the frequency of messages such as the announce interval. Sector specific SDOs fine tweak these values to provide their own profile specifications such as SMPTE's ST 2059-2.

Other industries also have their own profiles such as the IEEE 802.1AS for synchronizing audio and video on bridged networks based on IEEE 802.1. If both these profiles are running on the same network, which is entirely possible, then any receiving device could be confused by the two profiles, especially when resolving master clock status in the Best Master Clock algorithm.

The new profile isolation from V2.1 provides a unique identifier in the PTP header that allows downstream PTP nodes to only process messages with the identifier it recognizes and ignore the other messages. The idea is that a Standards Development Organization (SDO), such as SMPTE, IEC or ITU can request a unique identifier, the SdoID, and use this in the PTP message header.

Backwards compatibility is maintained as the sections of the header that the SdoID uses are either reserved in V2 or a forward extension of the information already in there. SDOs can apply for one of the unique SdoIDs, however, these identifiers are only issued to SDOs and not individual manufacturers or broadcasters. The key advantage is that every SDO can still make use of the full range of domain numbers and other parameters without interfering with other profiles from different SDOs in the same network.

V2.1 also facilitates the use of multiple masters that all send their timing messages to slaves simultaneously. V2 only provided a single master and if it sent out the wrong time message, then all slaves would try and sync accordingly. With the multiple master approach, the slaves can choose a group of masters and dismiss any master that they consider has sent the incorrect time.

Monitoring

Another major addition to V2.1 is that of standardized time accuracy monitoring. Although it was possible to gather the necessary timing data to determine the health of the network time, V2 didn't standardize this. Consequently, any vendor building a monitoring tool had to write specific interfaces for every manufacturers PTP processing equipment.

Time analysis is critical to any network and being able to measure and log the timing data from each PTP enabled device in the network is an absolute necessity. V2.1 provides four new timing statistics integrated over 15 minutes and 24 hours. The basic timing measurements are the average, minimum, maximum and standard deviation.

PTP nodes which contain slave ports such as ordinary (in the slave state) and boundary clocks provide information about their offset from the master. Digging deeper into the network is now achievable to help determine the accuracy of the time within the network.

The data storage format for the timing data is also specified to interleave the 15 minute and 24-hour measurements. This provides a consistent data format making analysis and hence diagnosis, much simpler.

SMPTE's ST 2110 has freed us from the shackles of analogue television. Due to the time invariant sampling nature of video and audio, we still rely heavily on an accurate and reliable time source. PTP V2 provided this for us. But PTP V2.1 has not only improved on V2 to provide security, robustness and improved accuracy, but has made all these new features backwards compatible.

We can easily migrate to the new version to take advantage of the new features and be assured that backwards compatibility to existing PTP V2 equipment will be maintained.

The Sponsors Perspective

PTP In LANs And WANs - An Essential Component In IP Broadcast Infrastructures

By Daniel Boldt, Head of Software Development, Meinberg

PTP - as a precise network timing technology has been available for nearly two decades. It is already widely used in Telecommunication networks, Finance and Trading platforms, substation automation networks and many more industries. Every industry has its own demands such as target accuracy on the end nodes, or whether it should be used locally or via wide area connections. Furthermore, there is often the question of whether existing network components should be re-used or if they will be replaced.

Meinberg has served customers in those industries with PTP solutions for 15 years. Throughout this time, Meinberg gained a huge amount of practical experience in the productive use of PTP.

The migration from an SDI based studio infrastructure to an IP environment according to SMPTE ST 2110 enforces the use of the Precision Time Protocol. Here, PTP clients must keep a maximum time offset of 1 µs between each other to comply with the SMPTE ST 2059-2 profile. Taking into account that there is usually a high traffic load in the network due to the transport of uncompressed video and audio streams, the use of PTP compliant switches is highly recommended, although it is not enforced by the PTP profile. As the migration to IP is usually a greenfield approach in terms of the IP equipment, the availability of PTP compliant switches in those networks is no longer an issue.

As well as local installations, some Meinberg customers have been looking for solutions to provide time synchronization via long distance connections to facilitate remote production scenarios. In those cases, a PTP-aware connection is often not available.

There can be several solutions for this scenario. First of all, you need a nonlinear clock filter algorithm in the PTP Slave to remove the effects of packet jitter via the non PTP-aware network path. Furthermore, in the case where the network path changes, you need a method for detecting any changes in the asymmetry that can occur during the path reconfiguration. Depending on the magnitude of the general asymmetry of the path, a recalibration of the time reference, such as the GPS at the remote location, would be necessary.

In complex network scenarios with a high packet delay variation and a large asymmetry, a GPS-assisted approach at the remote location will become necessary. Meticulous readers may be asking why you would not just restrict yourself to only using GPS as a reference source in this scenario. The answer is because of the rising risk of vulnerabilities from jamming (disturbing) or even spoofing (manipulating) the GPS signal. There is a trend in the industry to avoid the absolute dependence on these kind of satellite-based time sources. This is especially concerning where large-scale remote production events take place, such as the Olympic games, or areas with political instabilities, as it is possible for attacks against the GPS signals to occur. Therefore, PTP can help to keep the timing distribution within a critical infrastructure up and running at all times, even without being available.

Another very important aspect is the monitoring of the PTP infrastructure. Even now there is still no standardized approach to how a PTP node should be monitored and which information should be provided. However, there is currently an ongoing activity at SMPTE to standardize such an approach.

To improve monitoring, Meinberg introduced “SyncMon” a couple of years ago as an extension to the PTPv2 protocol. This approach not only allows PTP node monitoring in terms of status information but additionally allows accuracy measurements of PTP Slaves which support this extension. It was created as the self-reporting sync status of a PTP Slave does not need to be 100% correct all of the time as there may be occasional network asymmetries or failures in boundary or slave clocks. However, with the help of the “SyncMon”, the actual accuracy of a PTP node can be validated as an independent reference.

Speaking of Boundary Clocks, it is essential to monitor them as they have an internal clock that needs to be adjusted. Ideally, all boundary clocks should have their outputs measured using a connected slave unit (a “probe) that reports any offset compared to the time reference. In the scenario where a boundary clock would produce such a drift, all PTP slaves connected to this boundary clock would follow it, although they would report a healthy state. Therefore, actively measuring boundary clock outputs will make troubleshooting PTP issues easier as the root cause of a failure can be identified very quickly.

Meinberg has provided synchronization equipment for a number of ST 2110 projects during the last few years, including BBC Wales and the Swiss SRF. Furthermore, the German public TV broadcaster WDR worked together with Meinberg to build-up a PTP distribution over a wide area network to serve remote studios located in different cities with PTP coming from a central location. After nearly 4 years of building a network that connects several radio studios in the western part of Germany, those radio studios receive the PTP signal without a local GPS installation, and then contribute RAVENNA streams to the radio program broadcast from a centralized studio. For this project Meinberg’s modular synchronization platform “IMS” was chosen because it has the flexibility to provide the timing gateway between telecom and broadcast networks.

For smaller studio environments, the new microSync 7xx/8xx platform for broadcast environments is now available. With the microSync, Meinberg now offers a product line for the broadcasters that provides PTP functionality and legacy base band sync signals at low cost in a compact device. All PTP platforms will soon be upgraded to support the latest PTP standard while keeping support for the previous version.

Daniel Boldt, Head of Software Development, Meinberg.

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

1/2021

Supported by

MEINBERG
The Synchronization Experts.