Remark: For convenient review the poster is formatted as a document: The final poster will use these text paragraphs, bullet points and graphs

PTP Deployment Strategies for Large Broadcasting Networks a Case Study

Nikolaus Kerö, Oregano Systems, Thomas Kernen, Mellanox, Heiko Gerstung Meinberg

As the broadcasting industry is moving from traditional SDI infrastructure towards the All-IP studio, a common frequency or notion of time is no longer a given. Instead network wide synchronization has to be accomplished using a packet-based time transfer technology namely the Precision Time Protocol (PTP). Different network topologies which are typical for broadcasting applications are presented and compared with each other with respect to effective deployment and operation of PTP according to ST 2059-2 – a PTP profile for the broadcasting industry published by the SMPTE. Special emphasis will be put on how to make best use of fully redundant networks which are mandatory for large broadcasting studios. After presenting deployment strategies for time transfer in large media network highlighting crucial operational aspects are highlighted such as continuous monitoring of all PTP devices (end nodes and network elements alike) with respect to synchronization accuracy.

Rationale for the All IP Studio

Until the very recent past, the broadcasting industry has been relying exclusively on legacy infrastructure for media transport and processing, which cannot handle the ever-growing bandwidth band driven by new video technologies. Rather than specifying, developing and deploying yet another generation of legacy devices, the broadcasting industry decided somewhat reluctantly to move to Ethernet as the only transport medium using IP as the transport protocol.

Existing Infrastructure in a typical broadcasting Studio

- All media information is transported via Serial Digital Interface (SDI)
- Dedicated SDI infrastructure (lots of coaxial cables)
- Common system wide frequency is a given: all devices are by design syntonized
- Common phase relationship is accomplished via dedicated analogue sync signals
    - Black Burst
    - Tri-Level Sync
    - More often than not distributed via separate physical medium/network
- Accurate transfer of absolute time is difficult
    - Transmission delay is not accounted for inherently
    - Out-of-Band measurements / estimation
- Pros:
    - Limited monitoring and maintenance efforts especially of timing
    - Quality of frequency transfer independent of operating conditions
    - Highly experience engineers and technicians
- Cons:
    - Static, single purpose equipment
    - Not suited for ever growing bandwidth demands of new video technologies

Moving to IP

It was a huge leap of faith to move to IP because it is nothing less than transitioning from a well-known technology with highly predicable and bounded latencies to an inherently asynchronous medium which is notorious for its unpredictable temporal behaviour under heavy network load. One undeniable advantage outweighs, however, any hesitancy: IP technology serves a total market several orders of magnitude larger than the broadcasting industry, allowing it to benefit from technological advances. Furthermore, other industries haven proven that IP can satisfy even stringent requirements with respect to bandwidth and latency for mission critical applications.

- IP-based Media transfer has been tightly standardized: SMPTE ST 2110-xx series
- All Media is transferred over IP in separate flows
    - Video
        - (U)HD, 8K, compressed and uncompressed
        - High frame rates of up to 120 frames / sec and beyond
    - Multichannel high-resolution audio
    - Various ancillary data channels
- Capturing and processing is done separately (independently)
- Different data/media streams have to be re-combined correctly for final broadcasting
    - Temporal relationship has to be preserved
- Accurate timestamps have to be added to the media essence
    - 1 per video frame
    - 1 per group of Audio sample
    - 1 per ancillary data packet
- ➔ All devices have to synchronize tightly and reliably

PTP profiles for broadcasting

PTP v2.0 is a highly generic time transfer protocol, however it can be tailored to the needs and constraints of a specific application domain be defining a PTP profile according to the rules laid out in the IEEE1588-2008 standard. A PTP profile specifies ranges for all PTP message rates as well as permitted transport mechanisms and protocols. For broadcasting applications two profiles have been defined by SMPTE and AES respectively. While ST 2059-2 focuses on video, AES67 defines time transfer requirements for audio applications. Both profiles overlap sufficiently well in all relevant parameters that a common set of PTP parameters can be selected.

ST 2059-2

- < 1µs accuracy between any two devices
- Purely Layer 3 (IPv4 or IPv6)
- High Message rates
    - Event messages at 8/s
    - Announce messages at 4/s
- Multicast
    - All messages are sent in multicast
    - Network with 500 devices, assuming standard ST 2059-2 rates
        - Every device receives more than 8,000 messages per second
- Mixed Mode
    - Sync and Announce sent as multicast
    - Delay_Request and Delay_Response sent as unicast
    - ➔ dramatic decrease of PTP message load for all devices
        - Every device receives less than 50 messages per second
- Short lock times (< 5-10 seconds)
- Full, partial or no PTP support for network devices permitted
    - Boundary Clocks and Transparent Clocks allowed
- Sync Metadata TLV (via PTP Management messages)
    - Convey broadcasting specific data
    - Daily Jam time
    - Default video frame rate

    - Local time
    - …
- Updated/modified message exchange for Management messages
    - No Acknowledge for Sync Metadata TLV
    - Otherwise Acknowledge in unicast
- ➔ Keep overall message load low for every PTP device

PTP Synchronisation for Audio broadcasting

In the professional audio broadcasting realm, a number of standards and technologies for audio transport over networks have evolved quite independent from each other. Unfortunately, in most cases, they are not interoperable requiring special bridging devices to interconnect such networks. Dante, Ravenna, Livewire+, Q-LAN and Wheatnet IP are among the most common ones. With AES67, an open interoperability was published with the goal to remedy this situation. The Media Network Alliance, Ravenna and others help promoting the adoption of AES67.

With respect to time transfer within an All-IP studio, it has to be considered that commonly audio and video broadcasting divisions are separate entities within a larger organization, however, they need to interoperate to a certain extent, at least with respect to using a common reference time.

- PTP V2.0 required by
    - AES67
    - Ravenna
- PTP V1.0
    - Mandatory for Dante
- PTP V1.0 and PTP V2.0 are NOT interoperable.
    - Different message format
    - PTP V1.0 has no notion of Announce message
    - Conveys clock information with sync messages
- Dante sub-systems have to be attached via dedicated Boundary clock
    - Port connected to Dante has to operate in PTP V1.0 mode

Broadcasting network infrastructure

Even for small and medium sized broadcasting facilities, the requirements for the underlying network infra structure are quite complex calling for data centre grade network equipment to be considered. The favoured network topology in most broadcasting deployments is a Spine-Leaf architecture, where end devices are attached exclusively to Leaf switches which in turn are attached central Spine devices via very high bandwidth links. This topology can be easily extended without increasing the number of hops between any two devices to more than four.

A broadcasting studio network comprises a diverse set of devices ranging from high resolution cameras various display and storage modules all linked to video and audio mixing and processing stations.

- High bandwidth
    - Multiple parallel (U)HD video streams
- Large number of Nodes
    - > 500 for medium sized studios
    - Cameras
    - Audio recording equipment
    - Multi-viewer
    - Video and audio processing devices
    - Intermediate and long-term storage
- Frequent topology changes in the network

    - Studios are re-configured for every production
    - Outside Broadcasting-vans are attached to the studio

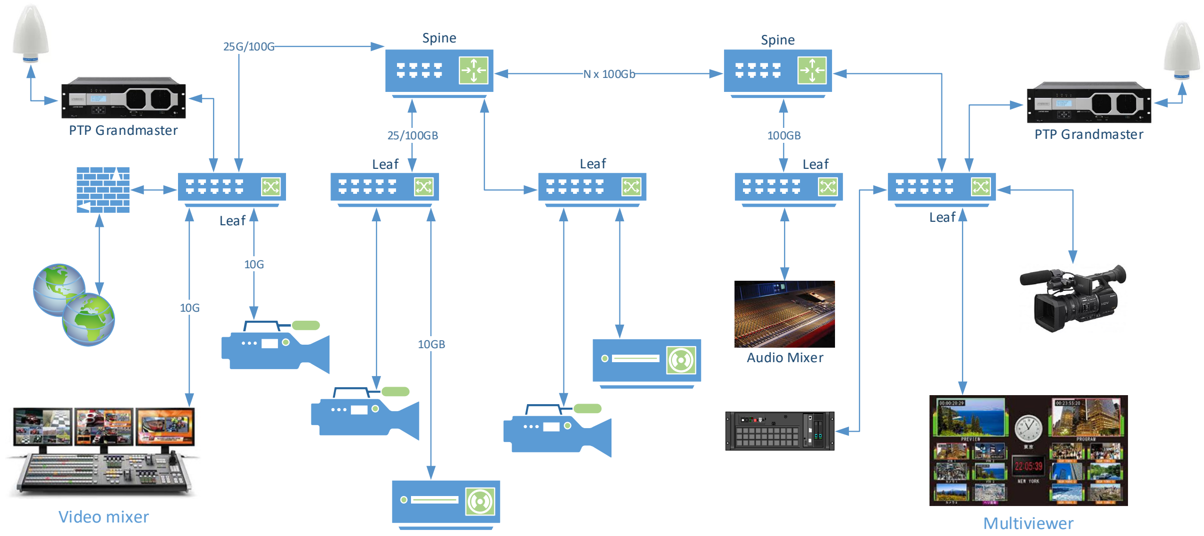

Video sync signals are still alive

The most common approach when transitioning to IP is to deploy new infrastructure gradually. As a consequence, existing SDI device and networks have to interoperate with the All-IP studio equipment. Aside from converting media data back and forth, PTP time information has to be cross linked to traditional analogue synchronization mechanisms.

- Re-generation of black burst from absolute time
- ST 2059-1 defines alignment points for all common video standards
    - 50Hz and 59.97 Hz
- Straightforward Algorithm
    - Calculate closest alignment point in the future ⇔ trigger time
    - Compare current time with the trigger time
    - If trigger time is reached generate sync signal
    - Advance trigger time to next alignment point
- Can be implemented in FPGAs which are usually present for
    - Hardware assisted media packet processing
    - Hardware assisted PTP packet processing

Fault Tolerance in Broadcasting (SMPTE ST 2022-7)

For every mission critical application exchanging data via Ethernet, the network architecture has to be designed to cope with network related failures such are broken connections and even malfunctioning network devices. This is accomplished by providing redundant network paths between critical devices. With the help of lower layer network protocols like RSTP (Raid Spanning Tree Protocol) a path through the network is continuously maintained while avoiding network loops at the same time. Selecting a new path in case the current connection between two devices is broken, however, requires a certain amount of time during which the data flow is transiently interrupted. For a number of use cases this limitation is acceptable. If, for example, two devices exchange information via a TCP (Transmission Control Protocol) connection, where every packet is accounted for by both partners, a transient network fault would not cause data to be lost it would rather be delayed. If, on the other hand, data is transmitted using UDP (Universal Datagram Protocol) packets may very well get lost. Many broadcasting applications have to revert to UDP rather than TCP, especially, if Multicast is used, where a single data packet is transmitted to multiple receivers at the same time.

- Extensive fault tolerance is mandatory
    - Typically, the loss of a single UDP Packet is not tolerated
    - May disrupt (compressed) video streams
- Mission critical equipment is deployed at least twice
    - Time reference / PTP Grandmaster
    - Cameras
    - Production studio module
    - **Interconnecting Network**
- ST 2022-7 redundancy model
    - Network infrastructure is doubled
    - Ideally two air-gap networks
    - Every end device has two network ports
    - Packet are sent over both networks
    - The first packet to arrive is used by the receiver, the late comer is discarded
    - ➔ zero packet loss in case of single failures

By default, PTP offers a considerable level of fault tolerance. First of all, it can cope with transient packet losses i.e. several missing packets will reduce the accuracy of an end device only marginally, let alone cause it to lose synchronicity. Furthermore, a PTP network can cope with the loss of the reference node autonomously without any user interaction (assuming that stand-by reference nodes are present). The ST 2022-7 redundancy model, however, cannot be applied to PTP out of the box:

- PTP for redundant networks
    - PTP has to be able to cope with transient packet loss
    - Switching on a per packet basis is **NOT** possible and not necessary
        - Varying absolute transmission delays introduces noise
        - PDV ⬄ Packet Delay Variations ➔ Accuracy is deteriorated
    - Alternate schemes for taking advantage of redundant networks
        - Process PTP message on both networks
        - Use time information of either one network
    - How to switch between the two networks?
        - Manually or automatically in case of permanent network/link failure
        - Select the better network by consciously evaluating the quality of the time information
    - Implementing Boundary Clocks at the end devices is not recommended.
        - Complex verification caused by large number of BCs in the network
    - Multiple PTP GMs have to be commissioned
    - If every PTP GM is connected to both networks, the GM selection (BMCA Best Master Clock Algorithm) has to be verified under all fault conditions.

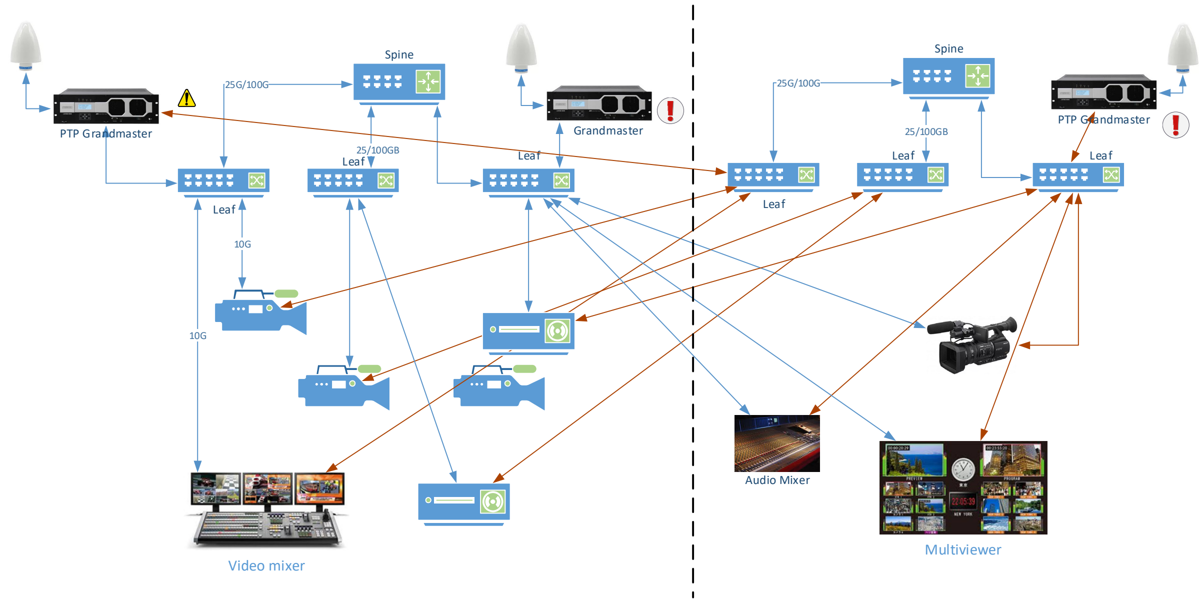

PTP deployment constraints

- Green field vs. brown field
- Full PTP support desirable but not always feasible
- Existing expensive IP infra-structure should be used at least transiently
- Existing IP islands have to considered
    - Audio over IP (Dante vs. Ravenna)

PTP deployment strategies and caveats

- EVERY PTP network is a multi-vendor system!
- ALL devices contribute to time transfer and thus can influence/deteriorate its quality
- Test under real world network load conditions
    - Regardless if BCs/TCs are used or not
    - These devices can malfunction causing the network to revert to non-PTP mode
- Restrict PTP Master role only to GM capable devices
- Restrict role of ports of Boundary Clocks to Master role, except upstream to GM(s)
- Deploy Acceptable Master Tables
- Introduce faults during commissioning / field acceptance tests
    - Disrupt GNSS reception
    - Disconnected active GM
    - Shutdown network segments
    - Add adverse PTP device sending PTP message at high rates
    - …
- In case of non-PTP aware network device even transient network overloading will have serious impact on the accuracy
- Continuous monitoring of network load

Monitoring – a mandatory precaution against PTP disasters

Although PTP offers a significant level of fault tolerance, there are still a series of error conditions which may remain undetected because they would not trigger any state changes within PTP devices, which are triggered only via Announce messages or to be more precise their content and the lack thereof. If messages carrying time information are corrupted or not forwarded at all any longer, no action is taken as per the standard. Such a situation can only be remedied if all PTP data is carefully monitored

Extensive monitoring is also crucial to protect a network against permanent loss of synchronicity caused by malfunctioning devices sending wrong messages or messages at incorrect (far too high) rates. All malicious attacks aimed at tampering with time transfer or hindering it all together can be detected reliably. If respective counter measures have been set in place beforehand their effects on the network can be limited to merely transient failure, which modern PTP stack implementation should be able to handle.

- Gather as much data as possible (not feasible)
- During commissioning AND normal operation
- Use both in-band and out-of-band methods
  - In-band: record offset reported by the PTP Slave
  - Out-of-band: compare clocks by external means (1PPS)
- Log state changes for every PTP port
  - Post-mortem analysis of transient PTP faults
- SMPTE ST32NF study group is currently defining a unified approach for a set of PTP parameters to be monitored by every PTP port.
  - PTP state (changes)
  - Message counters
  - Accuracy information
  - …

Conclusions

After the respective standards for time and media transfer over Ethernet networks have been adopted by the broadcasting equipment vendors, the transition towards the All-IP studio is gaining more momentum and with it the uptake of PTP. it has proven its effectiveness and resilience in several deployments around the world. Lessons learnt from these early adopters clearly show that PTP can deliver the required performance even under adverse and highly dynamic conditions typical for broadcasting studios. Careful planning is crucial for all steps of the deployment and commissioning process. To provide the required service level for time transfer, all PTP ports have to be configured accordingly limiting their roles to the absolute minimum. Furthermore, optional PTP features like the Acceptable Master Table are extremely useful additions to improve resilience and limit the effects of misbehaving devices as well as attack. Finally, extensive monitoring it a cornerstone to successful deployment and operation of any large PTP network.