



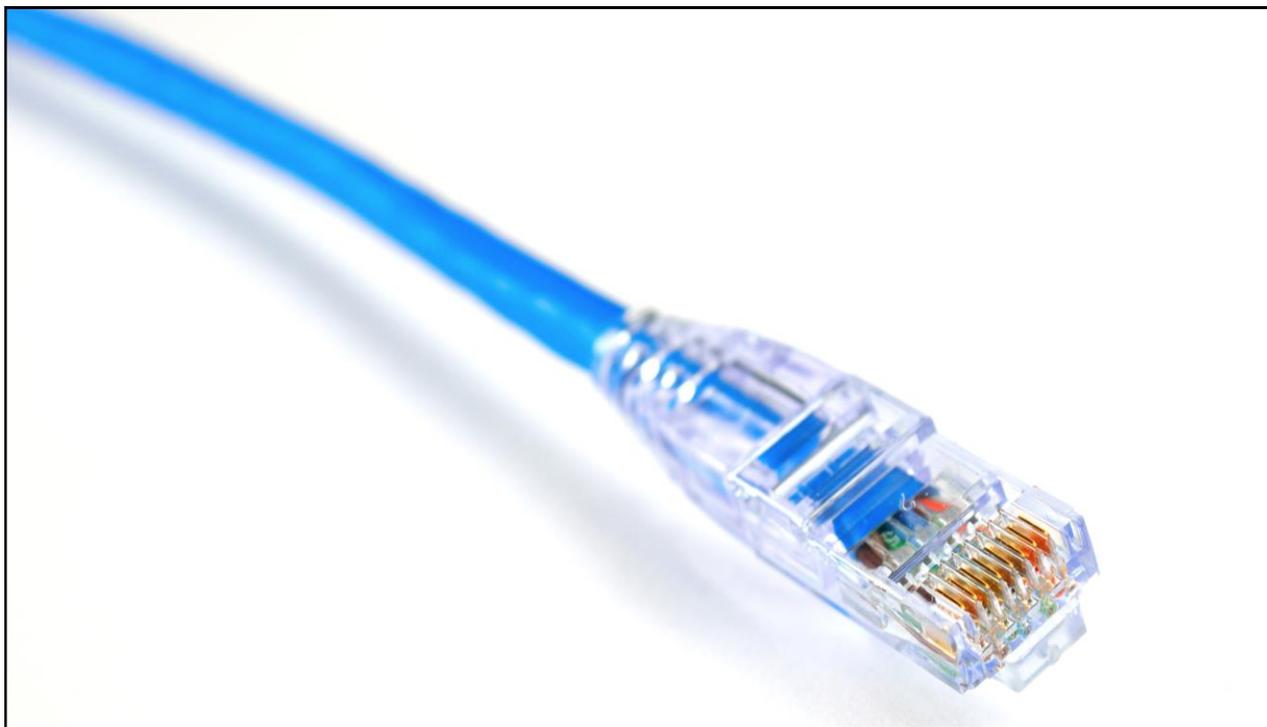
Demystifying Troubleshooting PTP Synchronised Media Fabrics

Thomas Kernen
Software Architect
Mellanox



IP SHOWCASE THEATRE AT IBC - SEPT. 14-18, 2018





Agenda: Things to know

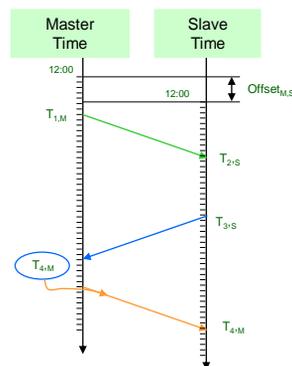
- How are you communicating?
- What are your message rates?
- Did you receive all your packets?
- Why is my PTP port in such a state?
- What are you synced to?
- To protect and serve
- Measure for measure





PTP communication modes

- Multicast
 - All PTP messages are sent to specific multicast group
 - All PTP devices listen to this group and process all the messages
 - Including Del_Req/Del_Resp node specific messages before discarding them
 - Creates additional load for PTP slaves – Beware of impact on implementations
- Unicast
 - Messages are sent individually between Master and Slave
 - Requires specific configuration of each PTP endpoint in the Grandmaster
 - May hinder PTP scalability for large deployments – Not all slaves support this
- Mixed modes
 - Sometimes (mis)labeled as “hybrid”
 - All common messages from Master to Slave are sent as multicast
 - All node specific messages (Del_Req/Resp) are sent as unicast
 - Reduces the PTP message processing load for PTP Slaves – Endpoint friendly



https://tech.ebu.ch/files/live/sites/tech/files/shared/techreview/trev_2018-Q2_PTP_in_Broadcasting_Part_1.pdf (co-authored with Oregano Systems)



PTP Message rates

- Different default message rates and ranges
 - SMPTE 2059-2, AES67, AESr16
- Message rates & respective timeouts
 - How long does a PTP slave wait before it thinks it lost its Master?
- The total message rate per PTP port/device can be handled
 - Messages aren't dropped, delayed, corrupted
- Applies to Slaves, GrandMaster and Boundary Clocks
 - Scaling capabilities of devices in presence of many PTP nodes





PTP Port state

Port State	Definition
Initializing	Port initializes its data sets, hardware, and communication facilities
Faulty	Fault state of the protocol, no PTP messages except management
Disabled	No messages on its communication path
Listening	Waiting for the announceReceiptTimeout to expire or to receive an Announce message from a master
Pre_Master	Behaves as a master except no messages sent, only management
Master	Port is behaving as master
Passive	No messages sent except signaling or management messages
Uncalibrated	Transient state to allow initialization of synchronization servos, updating of data sets when a new master port has been selected
Slave	Synchronized to the selected master port

Some implementations add additional port state



How are you synced and to whom?

- Is the application synced to the real-time clock or the PTP stack?
 - PTP Stack -> Real-time clock -> Application
 - Can cause rounding and offset issues due to multiple time format conversions
- SDP file generated by SMPTE 2110 senders
 - Based on RFC 7273 “RTP Clock Source Signalling”
 - “ts-refclk” field defines the Clock Reference source used by the sender
 - “a=ts-refclk:ptp=IEEE1588-2008:<GrandMaster Clock ID>:<Domain ID>”
 - “a=ts-refclk:ptp=traceable”
 - “a=ts-refclk:localmac=<Ethernet MAC address>”



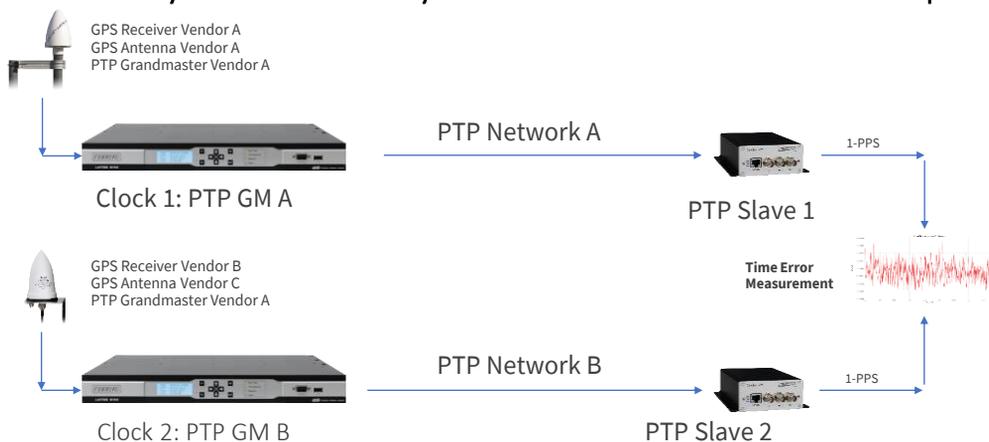


Why is “traceable” important?

- Timestamps from traceable reference clocks can be compared
 - Typically GNSS (GPS/GALILEO/GLONASS) reference locked GrandMasters
 - Reference signals can be directly compared between them for validation
 - Accuracy can be further increased by using multi-GNSS capable GrandMasters
- All senders/receivers can be locked to the traceable sources
 - Source of the Reference clock is common (single or multi-GNSS) and traceable
 - Removes the requirement for locking all devices to a single GM Clock ID
 - Removes perceived PTP network design constraints



PTP Sync Accuracy – Two Clocks Test Setup

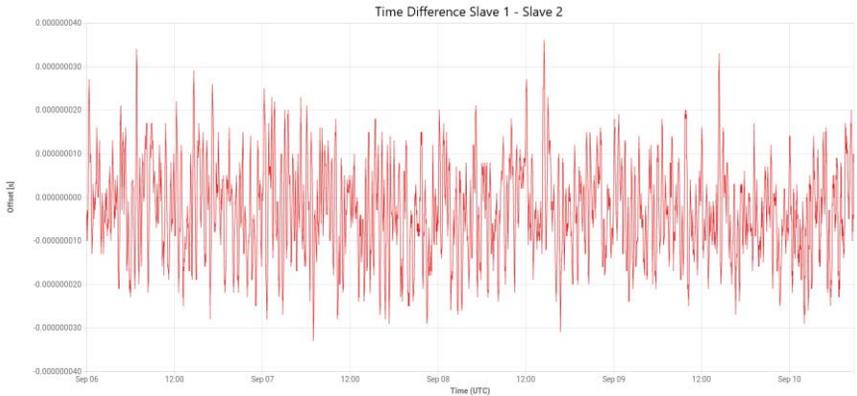




PTP Sync Accuracy – Results 4 days Test



PTP Sync Accuracy – Results 4 days Test





SMPTE 2059 TLV



- SMPTE 2059 TLVs should only be generated by the PTP GrandMaster
- End nodes (PTP Slaves) should either process the messages or ignore them
- If a response to the TLV message is required (ACK, UNKNOWN, ...)
 - MUST only be Unicast to the GM!
 - NEVER to all (via Multicast)
- Boundary Clocks that are SMPTE 2059 compliant MUST forward 2059 TLV messages on all MASTER ports



13



To Protect and to Serve

- Acceptable Master Table (AMT)
 - Check if PTP messages originate from a whitelisted GrandMaster
 - Protect against rogue or misconfigured devices that may take over as GM
 - Apply to Boundary Clocks and Endpoints
- Forced Master role on Boundary Clock ports
 - Prevent endpoints attempting to elect itself as a Master/GrandMaster
 - Apply on a per port, per VLAN basis
- Verify logs to detect misuse of the PTP infrastructure

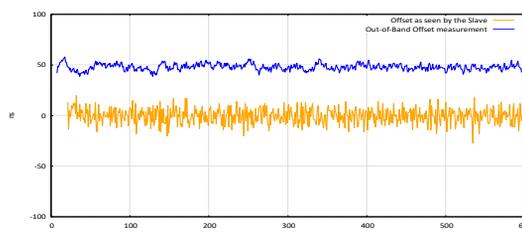


14



Measure for Measure

- Compare 1pps slave output with GrandMaster 1pps
- Provides out of band End to End verification of the full PTP chain
- Sample across the entire PTP infrastructure
 - Acts as a canary in a coal mine
 - Compares in-band vs out of band
 - Sanitize across multiple nodes



15



PTP Slave check list



- Current PTP Port state: Stable or changing state every few seconds?
- Does it support both 1 and 2 step mode? It MUST!
- Is it set the Slave Only mode? It SHOULD (really really SHOULD)
- Are all devices operating with the same PTP Domain ID?
- Do all devices have the same message rates and timeout values?
- Is the communication mode correctly set (Mcast, Ucast, Mixed mode)
- Any dropped/missed PTP messages in the message counters?



16



Final checks

- Perform PTP analysis with Software or Hardware tools
- Verify that all candidate GMs operate with the correct values:
 - Priority1, Clock Class, Clock Accuracy, Priority2
- Check PTP nodes can sustain the aggregate message rate on their interfaces
 - Boundary Clocks, Slaves and Masters!
- If devices don't provide in-depth PTP debug capabilities, request them from the vendor!



Thank You

Thomas Kernen, Mellanox

tkernen@mellanox.com

Come and see us at booth 8.E27!



IP SHOWCASE THEATRE AT IBC – SEPT. 14-18, 2018