

PTP Best Practices for Media and Entertainment and Broadcast Centers

Consistent time synchronization is absolutely critical in a media and broadcast facility. Without equipment locked to the same timing source, lip sync issues occur as well as top of frame alignment and switching are impossible. For example, video could be shifted on the viewing device, or video and audio could be out of sync - each of which result in a poor experience for the viewers. In the baseband world, devices would lock to reference black, or genlock (among others). With the adoption of IP, and more specifically SMPTE-2110 (or SMPTE-2022-6 with AES67), a different mechanism of providing timing is required - Precision Time Protocol, or PTP, also referred to as IEEE 1588 (PTP V2). There are various profiles for PTP, such as SMPTE 2059-2 and AES67, among others. Unlike traditional methods of synchronizing devices throughout a facility, since PTP is fully network based, it can be carried over the same data network connections already being used for transmission or reception of essence flows - or over parallel networks.

There are many different ways to distribute PTP across a facility (or multiple facilities) for the purpose of providing timing to various transmitters/receivers (slave devices). PTP is transported over multicast from a master clock (Grandmaster) and this multicast group propagates throughout the network to all endpoints that are requesting it. There are various types of messages that get exchanged between the master clock and the slaves, each providing necessary timing values to calculate that actual time and offset calculations that the devices will be referenced to.

Some PTP distribution networks are more scalable than others, and others can be more complex to configure. Ultimately, the goal is to ensure that all slave devices can lock to PTP and stay locked. These slaves can include endpoint devices such as SDI-to-IP or IP-to-SDI Gateways, multiviewers, keyers, etc.

More literature is available on the Arista website detailing the intricacies of PTP. This is strictly a document that details configuration possibilities with recommendations for a media network.

Best Master Clock Algorithm (BMCA) Attributes

The Best Master Clock Algorithm (BMCA) is what is used to elect a Grandmaster (GM) for the PTP distribution. The GM is the clock that provides the timing reference for the system.

The attributes configured on the Grandmasters, contributing to the BMCA are evaluated in the following order:

- Priority 1 (lowest number wins)
- Clock Class (GPS, free-run, etc)
- Clock Accuracy (accuracy to UTC)
- Clock Variance (jitter and wander)
- Priority 2 (lowest number wins)
- GMID (similar to mac address)

These messages are advertised in the Announce Message, described below.

PTP Message Types

The messages used to determine the GM, and the timing offset at the slave devices, switch included, are:

Announce Message: Message that advertises the PTP attributes (listed above) to determine what master is the Grandmaster. This is done using the BMCA (Best Master Clock Algorithm). For SMPTE-2110 deployment, this value should be set to 0 (zero) – which translates to 1 message per second ($2^0=1$)

Sync Message: Message that contains the time that a packet was transmitted from the master to the slave. In Two-Step mode, the follow-up message contains the time that its partner sync packet was transmitted from the master. For SMPTE-2110 deployments, the sync interval is set to -3, which translates to 8 messages a second. (2^{-3}) = 0.125

Delay Request (Response): Messages used to calculate the one way delay of PTP. Delay request messages and delay response messages is the same frequency as the sync interval.

Management Messages: These messages are used to maintain, monitor, and configure the PTP network.

Announce Timeout: Number of announce messages the device can miss before triggering a new BMCA to find a new GM. Generally, this value is set to 3 (Arista default)

The table below represents the different range of values per message type or domain based on AES67 and 2059-2 profiles, and their overlap. However, the most common values for these message types are:

Announce Interval = 0

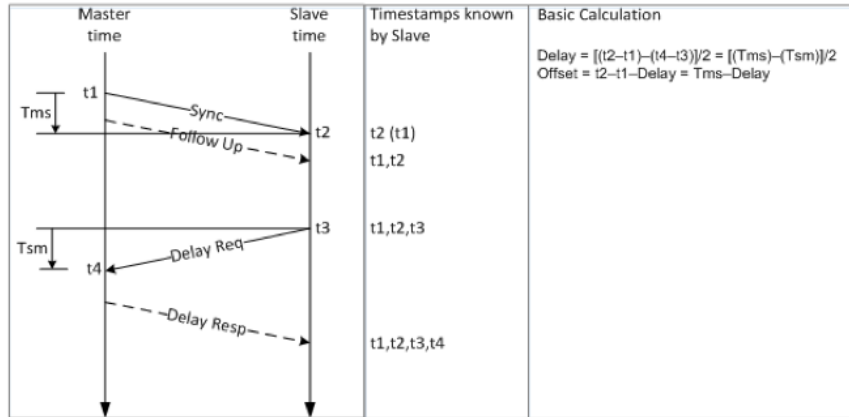
Sync Interval = -3

Delay Request Interval = -3

Attribute	AES67	SMPTE2059-2	Common settings	
domain	0 to 255	0 to 127	Limited	0->127
Announce interval	0 to 4	-3 to 1	Limited	0 to 1
Sync interval	-4 to 1	-7 to -1	Limited	-4 to -1
Min delay interval	The configurable range shall be -3 to 5 or Sync Interval to Sync interval + 5, whichever is the most restrictive	Sync Interval to Sync interval + 5	Limited	Sync Interval to Sync interval + 5
Min Pdelay req interval	0 to 5	Sync Interval to Sync interval + 5	Limited	0 to 4
Announce receipt timeout	2 to 10	2 to 10	Unlimited	2 to 10

All of the PTP messages above have the destination address of 224.0.1.129, though Event Messages utilize UDP Port 319 and General Messages utilize UDP Port 320. Of the messages above, Sync and Delay Request Messages are Event messages, while Announce, Delay Response, and Management Messages are General Messages.

The message exchange between the master and slave devices can be seen below:



Once the slave has all four time values (t1, t2, t3, t4) it can accurately calculate the offset between the slave and the master clock, as well as the propagation delay between the two. Therefore, it is clear that the slave uses the Sync (and Follow-up), Delay Request, and Delay response messages to synchronize to its master.

What is a Boundary Clock?

Within PTP, a Boundary Clock (BC) is a device that allows for the synchronization of multiple devices, spanned across multiple subnets, while at the same time preventing the back and forth messaging between each and every slave. Switches supporting PTP boundary mode synchronize with grand masters and provide timing and drift information to upstream and downstream devices.

A Grandmaster is connected to a switch, and that switch interface will become a slave to its grandmaster. All other interfaces on the switch that are PTP enabled, will be a master for its connected endpoint device. This endpoint device can be a transmitter, receiver, or even another switch. A downstream switch will follow the same pattern mentioned above, whereby the interface joining the two switches together (on the GM switch) will be the master to the downstream switch's slave.

See drawing below for assignment:



They key points of this topology are:

- All switches should be in Boundary Clock mode, and that Grandmaster A has a lower value for P1 than Grandmaster B.
- All interfaces connected to only slave devices should have a config on them to set the interfaces to master mode only. The command is 'ptp role master'.
- Notice that PTP A and PTP B only have one connection to each side of the network. Even though some PTP Grandmasters are equipped with dual interfaces, architectures using dual interfaces have not been thoroughly tested and should be avoided.
- The connections from switch to GM, and from switch-to-switch, are in dynamic mode (ie, not master only – the interface can be elected as master, slave, or passive, depending on the BMCA election).
- The reasoning behind connecting the two switches (SMPTE-2022-7) together is such that in the even that PTP A is the GM, but a primary link on an endpoint is lost, then the device's secondary connection will acquire PTP A's timing information. All devices, regardless of which connection receives PTP will be locked to the same clock. It is important to note that PTP messages (in either Boundary Clock or Transparent mode, described below) will not traverse the peer-link in an MLAG domain. This is currently not supported. In such a case, it is advised to add an extra connection (or two) between the switches, in addition to the peer-link, that will be dedicated for PTP transport. Another important point is that, by default, all PTP messages that are generated by the CPU are untagged, therefore, if the traffic must traverse a trunk, a native vlan must be configured on the trunk interface.

What is a Transparent Clock?

Transparent Clocking devices remove PTP message jitter (the result of their own queues) differently. The multicast or unicast PTP messages are passed through using entirely normal unicast or multicast forwarding, but the precise time that each packet takes to transit the switch is recorded, and this time (correction factor) is included in the PTP messaging as it passes to the intended recipient. The end-point can now use the correction factor to compensate for the delay through the switch (or multiple Transparent Clock switch hops as well).

There are different use cases for creating a PTP Distribution network in TC mode, such as modular switches with dual supervisors, or the front-end switch (such as a 7020TR) connected directly to both the GM and the data network.

Recommendation: Boundary Clock

Wherever possible, configuring a PTP distribution network utilizing Boundary Clock mode is recommended for the following reasons:

Security: For each interface on an Arista switch, there is a configuration option called "ptp role master". This ensures that the devices connected to that interface can never become the GM for the overall network in the event that someone were to accidentally set a higher priority 1 for its BMCA (or that the interface cannot be set to slave mode. It could be disastrous for the system if a rogue device were accidentally connected into the network with a lower P1 value than the GM. PTP Role Master protects the system from this scenario.

Message Suppression: Since each interface is a master only to the device connected to it, the delay request messages (sent from slave to master) do not leak to other switch interfaces on the local or connected switches. All PTP messages terminate at the boundary clock. This is important because all PTP messages have the same multicast address – 224.0.1.129. Therefore, when every device transmits 224.0.1.129 around 8 times second, all of those messages would be received at every other device that is also PTP aware. The number of messages grows exponentially as more and more PTP aware devices get added to the network. This can cause some devices to get overloaded. Boundary clock restricts this.

Scalability: PTP aware devices have a limit to the number of messages that can be received per second, including the grandmasters. Thousands of messages a second could cause devices to lock up. Since the Boundary Clock absorbs many of these messages, the overall PTP distribution network can grow and expand without the concern of increasing the number of messages traversing the network. Boundary mode helps to expand the capacity of the broadcast network by distributing the time synchronization workload across the switches supporting BC.

Ease of Configuration: As Boundary Clock is operated by a PTP agent, it operates independently of routing protocols, therefore expansion of the PTP distribution network does not require manipulation of any routing statements.

Some customers will opt for a Transparent Clock, which is acceptable and will be able to work provide accurate timing to endpoint devices. Transparent Clock can be useful when the number of slave devices connected to a switch, such as a 7500R, grows to an amount that the supervisor is unable to handle in Boundary Clock. This number can be around 400. Ensure that the mode of operation for PTP is in Hybrid mode so that the GM does not receive every PTP message from every slave device.

Clock Recovery: Boundary switches continue providing timing data even if a master clock becomes incapacitated. Once the master clock has recovered or a backup has taken over, the network recalibrates without losing a step, thus ensuring consistent, reliable broadcast operations. Since transparent mode devices don't interact with end stations, they can't provide timing support in case of a master clock failure. Also, by manually adjusting the local Priority 1 and Priority 2 values on the switches, a PTP hierarchy can be created in the event that if all GMs fail, a specific switch will assume the role the reference clock for the entire system. This will remain in place until a GM with a better Priority 1 or Clock Class participates in the BMCA.

How to configure PTP (in Boundary Clock) and Status Commands

Global config:

```
localhost(config)#ptp mode boundary
localhost(config)#ptp domain 127 (or other value, which must match the domain on the GM)
localhost(config)#ptp source ip <ip address> (create a Loopback address and define this as the source IP - thereby a unique IP address of the system matches the PTP Source IP)
localhost(config)#ptp ttl 32 (or other value, default is 1)
```

NOTE: The PTP Source IP is what the switch uses to restamp the PTP messages from the switch to the endpoints. On a per switch device, there can only be one PTP domain, and in addition to that, all devices in the facility should be locked to the same PTP clock.

The below interface values are specific for SMPTE-ST-2059-2

Per Interface Connected to a Non-Slave Device

```
On interface level, L2 or L3, or port-channel,
localhost(config-if-Et1)#ptp enable
localhost(config-if-Et1)#ptp announce interval 0
localhost(config-if-Et1)#ptp sync interval -3
localhost(config-if-Et1)#ptp delay-req interval -3
```

Per Interface Connected to a Slave Device

```
localhost(config-if-Et1)#ptp enable
localhost(config-if-Et1)#ptp announce interval 0
localhost(config-if-Et1)#ptp sync interval -3
localhost(config-if-Et1)#ptp delay-req interval -3
localhost(config-if-Et1)#ptp role master
```

(the above command 'ptp role master' ensures that the device connected to the interface will never become the GM for the PTP domain. This is considered a best practice for PTP security.)

Useful Status Commands and Outputs

show ptp (displays which interfaces are in the following modes: Master, Slave, Passive, Disabled ; Grandmaster ID, # steps away from GM, etc)

```
PTP Mode: Boundary Clock
Clock Identity: 0x28:99:3a:ff:ff:61:dc:ed
Grandmaster Clock Identity: 0x08:00:11:ff:fe:22:f8:1f
Number of slave ports: 1
Number of master ports: 4
Mean Path Delay (nanoseconds): 17683
Steps Removed: 1
Skew (estimated local-to-master clock frequency ratio): 0.99999967997
Last Sync Time: 16:04:00 UTC Feb 13 2019
Interface   State   Transport Delay
              Mechanism
-----
Et3/1/1     Master  ipv4     e2e
Et4/1/1     Master  ipv4     e2e
Et4/1/4     Slave   ipv4     e2e
Et4/3/1     Master  ipv4     e2e
Et4/3/2     Master  ipv4     e2e
```

Key points to note of:

1. Every device in the network should have a read only status of the Grandmaster Clock Identity. This ID should be the same across all switches and endpoints
2. Clock Identity is the ID of the specific switch that this command is run on.
3. Slave port is identified in two locations. This is the interface that is providing this switch with PTP
4. Steps Removed indicates how many hops the switch is from the Grandmaster. In the case above, the GM is connected directly to the switch, therefore the value is 1.

show ptp parent (displays ID and parameters of PTP parent)

```
Parent Clock:
Parent Clock Identity: 0x08:00:11:ff:fe:22:f8:1f
Parent Port Number: 3
Parent IP Address: 192.168.100.100
Observed Parent Offset (log variance): N/A
Observed Parent Clock Phase Change Rate: N/A

Grandmaster Clock:
Grandmaster Clock Identity: 0x08:00:11:ff:fe:22:f8:1f
Grandmaster Clock Quality:
  Class: 248
  Accuracy: 0x31
  OffsetScaledLogVariance: 0x3d24
  Priority1: 120
  Priority2: 121
```

Key points to note:

1. Parent Clock Identity is the ID of the clock sourcing this switch. In this case, it is the Grandmaster. If this command was run on a switch connected to this one, the ID would be the Clock Identity shown in the 'show ptp' output
2. Parent IP Address, in this case, is the IP of the GM itself, since it is directly connected to the switch. In the case of a switch connected to this one, the Parent IP Address will be the Source IP entered in the Global Configuration of the switch.
3. Grandmaster Clock Quality is based on attributes in the BMCA. This GM is in freerun since the class is 248.

show ptp clock (similar to 'show ptp')

```
PTP Mode: Boundary Clock
Clock Identity: 0x28:99:3a:ff:ff:61:dc:ed
Clock Domain: 101
Number of PTP ports: 258
Priority1: 128
Priority2: 128
Clock Quality:
  Class: 248
  Accuracy: 0x30
  OffsetScaledLogVariance: 0xffff
Offset From Master: -92
Mean Path Delay: 17859 nanoseconds
Steps Removed: 1
Skew: 0.999999935993
Last Sync Time: 16:04:32 UTC Feb 13 2019
Current PTP System Time: 16:04:32 UTC Feb 13 2019
```

show ptp interface ethernet # counters (displays counters of required PTP messages, such as sync messages sent/received, delay-req sent/received)

```

Interface Ethernet3/1/1
PTP: Enabled
Port state: Master
Sync interval: 0.125 seconds
Announce interval: 1.0 seconds
Announce interval timeout multiplier: 3
Delay mechanism: end to end
Delay request message interval: 0.125 seconds
Transport mode: ipv4
Announce messages sent: 1271273
Announce messages received: 0
Sync messages sent: 10170176
Sync messages received: 0
Follow up messages sent: 10170176
Follow up messages received: 0
Delay request messages sent: 6874020
Delay request messages received: 0
Delay response messages sent: 0
Delay response messages received: 0
Peer delay request messages sent: 0
Peer delay request messages received: 0
Peer delay response messages sent: 0
Peer delay response messages received: 0
Peer delay response follow up messages sent: 0
Peer delay response follow up messages received: 0
Management messages sent: 1271254
Management messages received: 0
Signaling messages sent: 0
Signaling messages received: 0

```

show ptp monitor (displays up to 100 entries of offset from master, mean path delay, and skew)

Note that this command is only available on EOS 4.21.1F and later. For more specifics on settable thresholds, please view the EOS manual.

```

PTP Mode: Boundary Clock
Ptp monitoring: enabled
Number of entries: 100
Offset from master threshold: not configured
Mean path delay threshold: not configured
Skew threshold: not configured

```

Interface	Time	Offset from Master (ns)	Mean Path Delay (ns)	Skew	Seq Id
Et5	21:24:51.534 UTC Mar 09 2019	19313	0	1.000013127	293
Et5	21:24:50.534 UTC Mar 09 2019	-11335	0	1.000013096	292
Et5	21:24:49.534 UTC Mar 09 2019	18859	0	1.000013098	291

Analysis of PTP Packets

Arista switches, being linux based, allow for on-board tcpdump captures that can be exported and read in packet analysis tools such as Wireshark. This is a fantastic tool for verifying that the PTP packets have the appropriate data to ensure a steady system timing lock.

To perform the analysis of PTP for an interface, issue the following commands:

```
localhost(config)#monitor session test source eth1
localhost(config)#monitor session test destination cpu
localhost(config)#show mon session test
```

Take note of the mirror# next to the session name. This will be used for the next step.

```
localhost(config)#bash
[admin@localhost] tcpdump -i mirror1 host 224.0.1.129 | more
```

(the above will show the output on the terminal. Other linux notations can be used, as -v, -n, -XX etc.

or

```
[admin@localhost] tcpdump -i mirror1 host 224.0.1.129 -w <filename>
```

This will provide a file that can be opened in wireshark to examine PTP packets. This file, by default, is in /home/admin and can be obtained through USB or SCP

NOTE: After the capture is complete, exit out of bash back into config terminal and type 'no mon session test'


Tips for Troubleshooting PTP Distribution

1. Confirm that all devices (slaves and possible GMs) are sharing the same announce rate and timeout - otherwise the BMCA process can lead to multiple active (or bouncing) GM's.
2. Note that Arista BC's do not perform an IGMP join on 224.0.1.129. If the BC is not directly connected to another BC, or a GM, a static IGMP join will be required to receive the PTP traffic
3. Use the PTP counters to troubleshoot whether sync and delay request messages are being received / sent. In Hybrid mode (multicast syncs / unicast delay requests/responses), confirm the global config contains a "ptp source ip <ip>" command
4. Port mirroring can be quite useful - but be mindful that different platforms mirror at different places in the pipeline, so beware that this may give you a misleading view.
5. If a trunk is being used to transport either BC or TC PTP messages, ensure that there is a native vlan assigned on the trunk interface

Configuring PTP with Non-PTP Aware Platforms

Without Boundary Clock, configuring the switches to transport PTP traffic is similar to configuring the switches for regular multicast data transmission, with one key difference. If the GM is in a different subnet than any endpoints, and on different switches than the endpoints, then PIM with a Rendez Vous Point (RP) is required. The reason for this is that devices will never be able to send out a Source Specific IGMP join for PTP (since no one knows what the GM is at any given point), therefore an RP, by definition, is required for routing multicast traffic – and in non-ptp switches, PTP is just standard multicast with a destination address of 224.0.1.129.

Case 1: Spine (PTP aware) & Leaf (non-PTP) L3LS w/ ECMP and Hashing



This deployment is tricky due to the spines (Arista 7504R) supporting Boundary Clock while the leaf switches (7160) do not. The original configuration called for every 7504-7160 interface on the 7504 side to have ptp enabled, with ptp role master set as well. Since the 7160s do not support boundary clock mode, and the leaf switches are all configured for different subnets from each other and from the 7504, an RP was required to route the PTP traffic from the 7504 to the 7160.

The above topology and configuration would produce the following problems:

- Delay request messages from endpoints connected to Leaf switches would leak into other Leaf switches, and therefore other devices
- Resolve by creating a local RP on each Leaf switch, only for 224.0.1.129
- On the L3 connections between the leaf and spine, it is possible that the PTP Master interface would not match the mroute for 224.0.1.129 on the leaf. For example, link on the spine would be the master, but link 3 on the leaf would be listed as the OIF for 224.0.1.129. This would prevent locking of all devices on that leaf. The resolution for this includes the following steps:
 - Enable PTP on only link 1 for each spine to leaf grouping
 - Create a script on the spine (one dedicated for each leaf) that is triggered by an event handler on loss of connectivity of link 1
 - When link 1 is no longer connected, enable PTP on link 2
 - On the leaf switch, create two static mroutes back, with the PTP source IP as the destination - with a higher priority on link 1. Also create a static igmp join for 224.0.1.129

All of these steps are required to ensure proper PTP lock on an L3LS network topology, where the spine supports Boundary Clock and the leaf switches do not. Choosing the correct platforms, ones with Boundary Clock support, will reduce the complexity of the network configuration and also simplify the troubleshooting process.

Case 2: 1x7500R With Dual Supervisor Modules

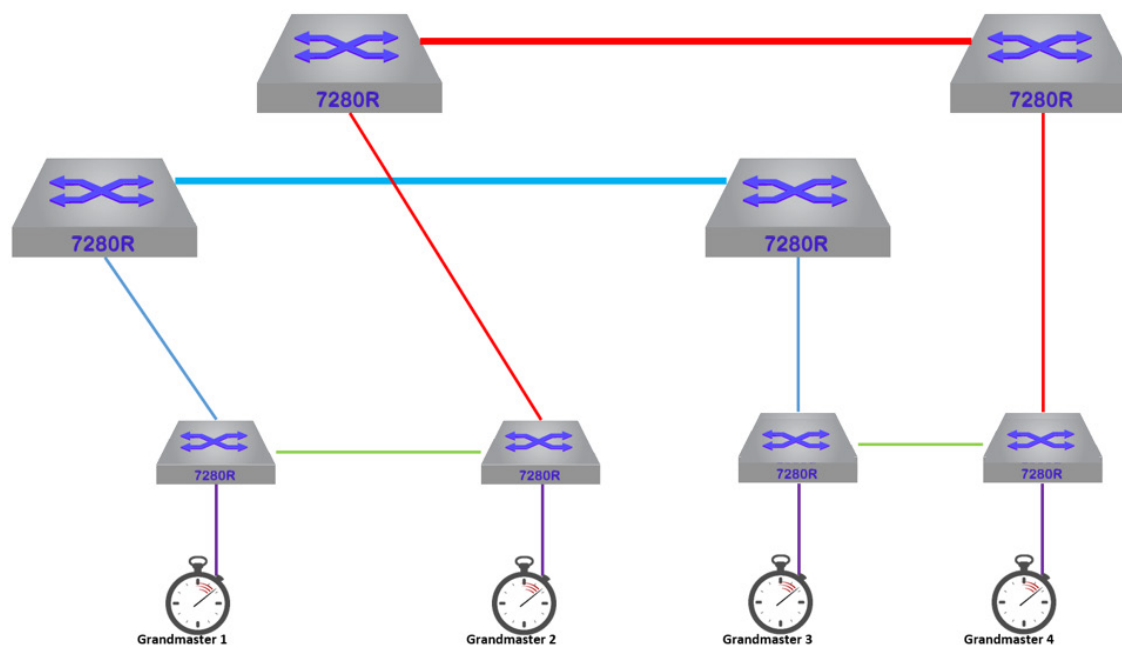
Although Boundary Clock is a recommendation for most SMPTE-2110 deployments, there are use cases for selecting Transparent Clock (TC), while using Hybrid messaging. When the switch is in Transparent mode for PTP, it calculates and adjusts for the packet delay in the switch. In addition to this mode, the endpoints can be set to Hybrid mode to ensure that the Announce and Sync messages remain as multicast from the GM, but the Delay-Request messages are transmitted as unicast packets. This will prevent each slave from seeing every other slave devices messages, thus reducing the amount of PTP messages on the network.

The reasoning behind configuring the 7500R switch in TC mode is that the PTP agent used for implementing BC is hosted in the supervisor. Therefore, if one supervisor were to failover, or be removed, the PTP agent would be forced to restart and negatively impact all of the end devices locking to it. TC mode does not use this PTP agent, therefore the failure of a supervisor would not be damaging to the timing of the overall system. Message suppression is still active due to the Hybrid mode described above, and a correction factor is also present for more accurate offset calculations.

For increased redundancy, be sure to connect two GMs to the switch in TC mode.

Important Note: Not all platforms support Transparent Clock. This must be confirmed

Devising a PTP Deployment



The image above reflects a proposal for a project with Arista and a partner. The criteria were that GMs 1 and 3 would be the “main PTP GMs” unless there was both a break between the two buildings and a failure of either 1 or 2.

Is the above criteria possible to implement? Yes – it is.

Assuming the following in **BOLD** are equal:

Priority 1

Clock Class

Clock Accuracy

Clock Variance

Priority 2

GMID

Use Priority 2 as the critical attribute to break the tie as to which GM will be elected Grandmaster based on the requirements noted above.

Therefore:

GM 1: Priority 1 = 1 ; Priority 2 = 2

GM 2: Priority 1 = 1 ; Priority 2 = 3

GM 3: Priority 1 = 1 ; Priority 2 = 4

GM 4: Priority 1 = 1 ; Priority 2 = 5

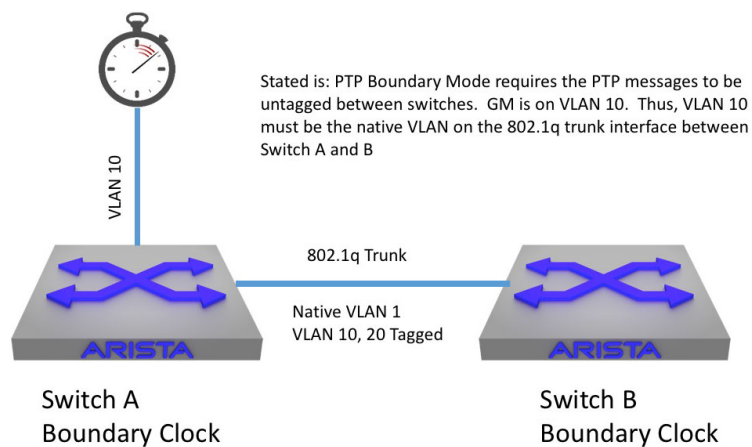
The above attribute settings will ensure GM 1 is the highest priority for being elected GM, followed by GM 2, 3, and 4. Note that it is vital that all possible GMs be configured with the same Announce Interval, otherwise one GM may inadvertently be elected GM for the system if its creation of Announce Messages are quicker than other GMs.

Supported and Not-Supported PTP Distribution Networks

The scenarios below document various network designs that are common in the Media and Entertainment space. The figures are also labelled as either Supported or Not-Supported in terms from a PTP distribution (Boundary Clock) perspective.

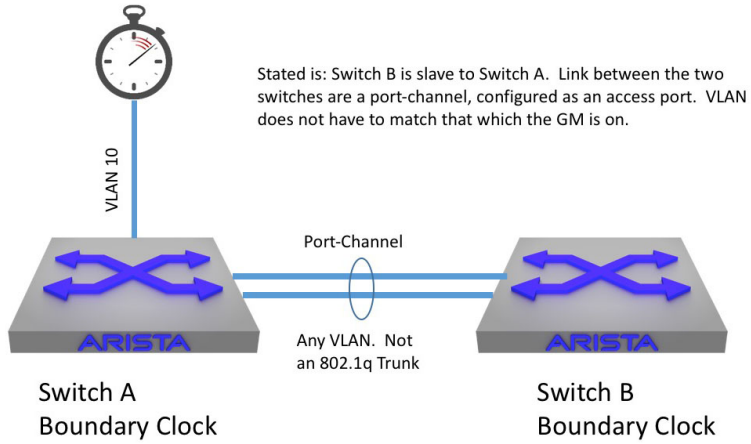
Scenario #1 – Trunk Link Between Switches Supported

Grandmaster



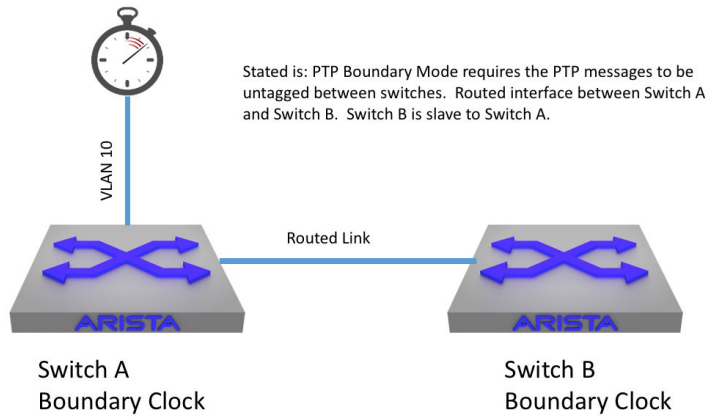
Scenario #2 – Port Channel Between Switches (Non Trunk) Supported

Grandmaster



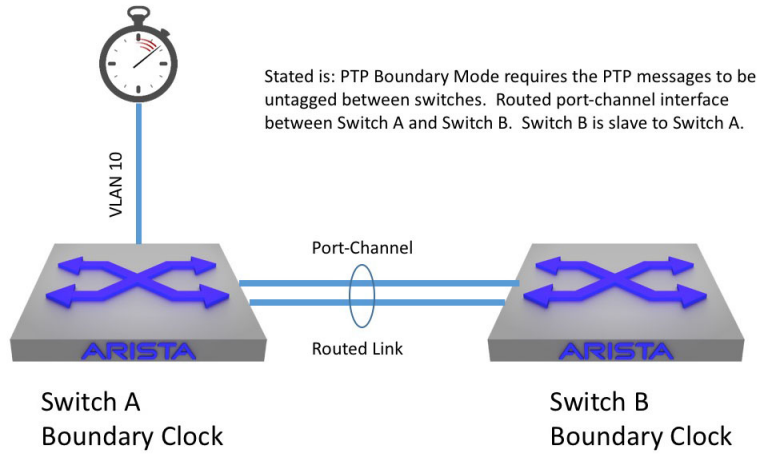
Scenario #3 – Routed Link Between Switches Supported

Grandmaster



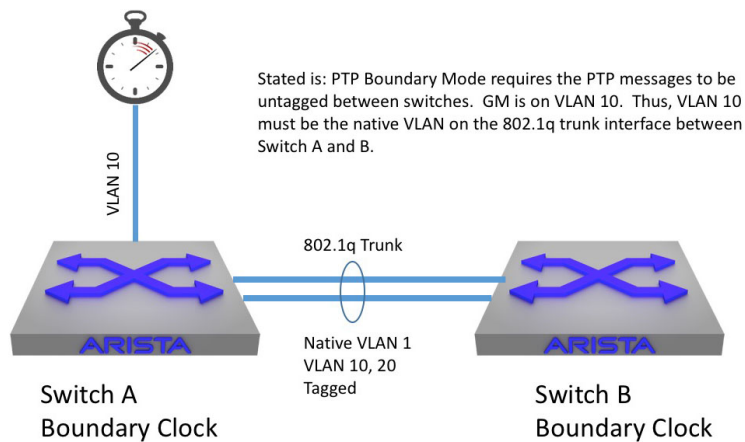
Scenario #4 – Routed Port Channel Between Switches (Non Trunk) Supported

Grandmaster



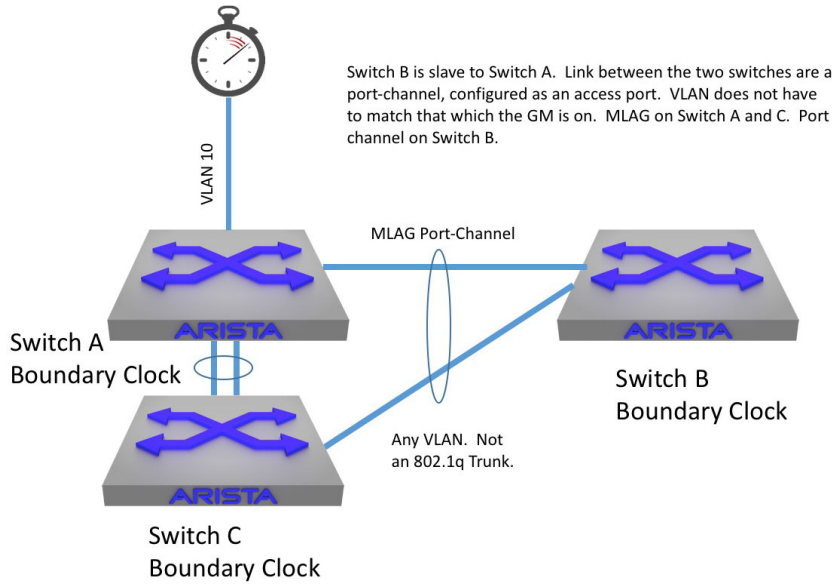
Scenario #5 – Port Channel Trunk Between Switches Supported

Grandmaster



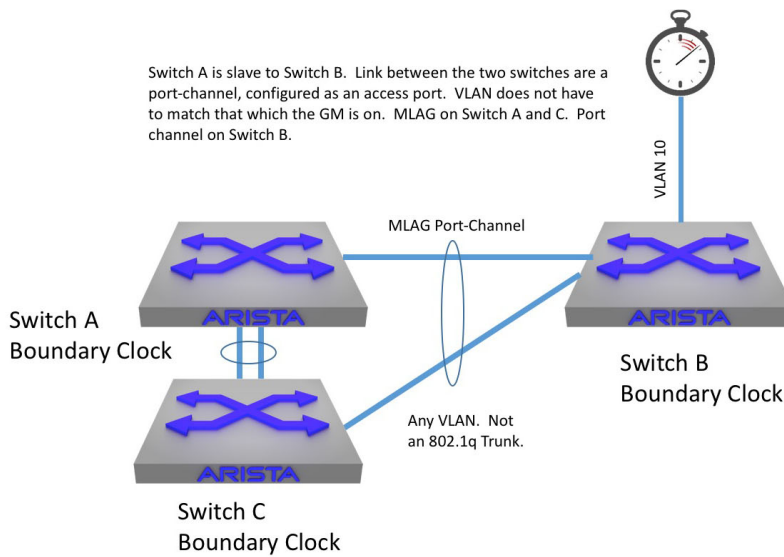
**Scenario #6 – MLAG Port Channel Between Switches (Non Trunk)
(Not Supported)**

Grandmaster



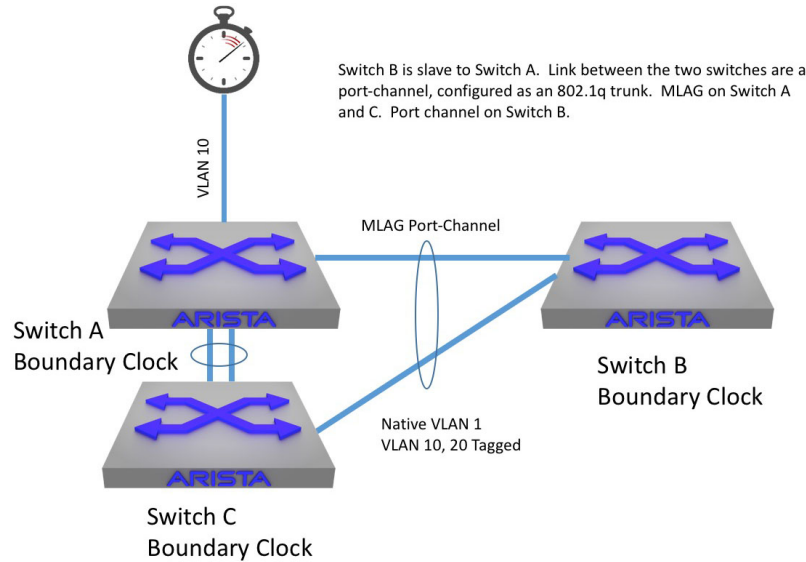
**Scenario #7 – MLAG Port Channel Between Switches (Non Trunk)
Not Supported**

Grandmaster



Scenario #8 – MLAG Port Channel Between Switches (Trunk) Not Supported

Grandmaster



Conclusion

There are many possible variations to a functional PTP distribution network. However, when designing these types of networks, especially in the M&E space where there are possibly hundreds of slave devices, scattered throughout a facility, scalability and CPU utilization of the endpoints must be taken into account. It is recommended that the platforms selected for an M&E deployment, and more specifically SMPTE-2110, support Boundary Clock and that this feature be enabled. In addition to those reasons, there is a simplistic approach to relying on Boundary Clock, both in configuration and troubleshooting, as shown in the first Case Study. Without Boundary Clock, in a routed environment, an RP will be required, which inevitably means that Anycast or a Bootstrap Router will be required - adding to the complexity of the design. Also, ensuring that the network design supports boundary clock is vital and should be the first topic of discussion when designing a PTP distribution network.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. Apr. 02, 2019