PTP System Commissioning: Grandmaster Changeover

Whitepaper

A series of grandmaster changeover tests for PTP systems in IP media networks

Contents

Introduction

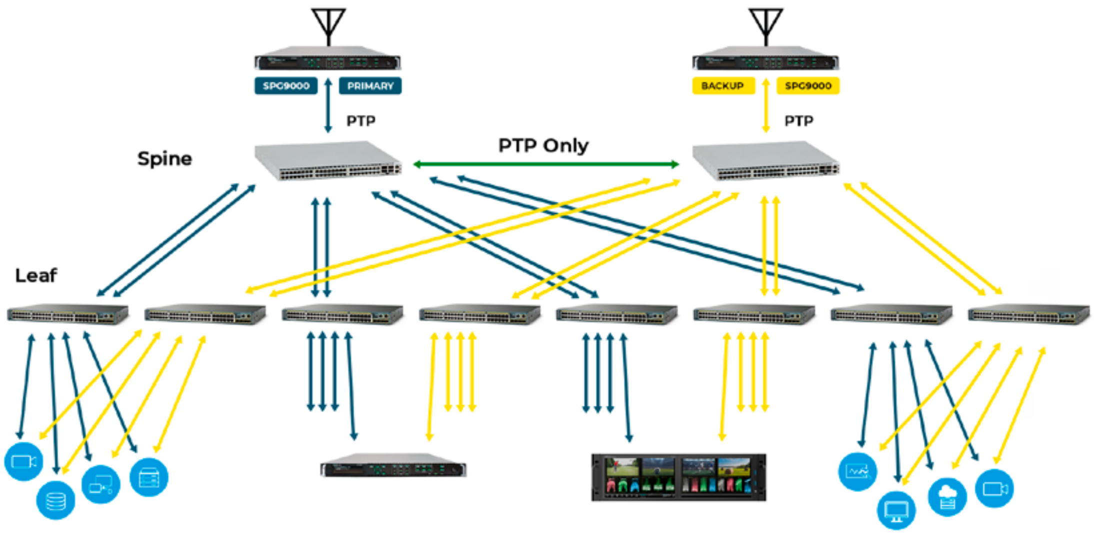

Figure 1. Hybrid SDI/IP Video Facility with Primary and Backups SPG9000

Precision Time Protocol (PTP) provides IP media networks with timing accuracy within the tolerance required for video timing within a facility. PTP has specific syntax and configurations required for correct operation and therefore it is critical to understand these requirements and carry out adequate testing of your facilities system timing during installation to understand the nuances of the protocol and to ensure fault-free operation of the system. This should be done during commissioning of the IP media network rather than waiting until the system is operational to find out the potential pitfalls of the IP media system architecture.

Most professional A/V PTP systems have multiple PTP Grandmasters (GMs) which are expected to provide seamless redundant operation. Therefore, it is important to test the behavior of the entire system when conditions cause the network to switch from one GM to another. Figure 1 shows a block diagram of a possible hybrid SDI/IP video facility, but a variety of configurations are possible to meet requirements of your network, effectively every system is unique. Therefore, it is important to test your PTP architecture during design validation or commissioning. This will help ensure the system works well if faults occur in actual operation. Additionally, this type of testing provides training and insight for the personnel who will have to support the equipment once the system is operational.

GM changeovers are implemented using the PTP Best Master Clock Algorithm (BMCA). This algorithm uses the clock quality parameters in the PTP Announce messages to choose the best lead clock. Figure 2 shows the SPG9000 web browser status page displaying the PTP syntax and the state of the PTP source along with the clock identity and the Grandmaster BMCA values (Priority 1, Clock Class, Clock Accuracy and Priority 2) that are used to determine the lead clock. There are several ways to trigger a GM changeover and they have different effects on the system. A PTP system may have a Primary SPG9000 and Backup SPG9000. It is recommended to make a note of the Clock Identity of each PTP source to aid in identification of the current grandmaster when a changeover occurs.

Several tests can be defined as shown in the list of recommended GM changeover tests. Although some of these may not make sense in all systems, users are encouraged to try as many as possible. Sometimes this will reveal unexpected behavior of the network and devices. Most of these tests have an initial failure which causes a changeover to occur, and then a second changeover will happen when the fault is fixed and the system recovers to the original state.

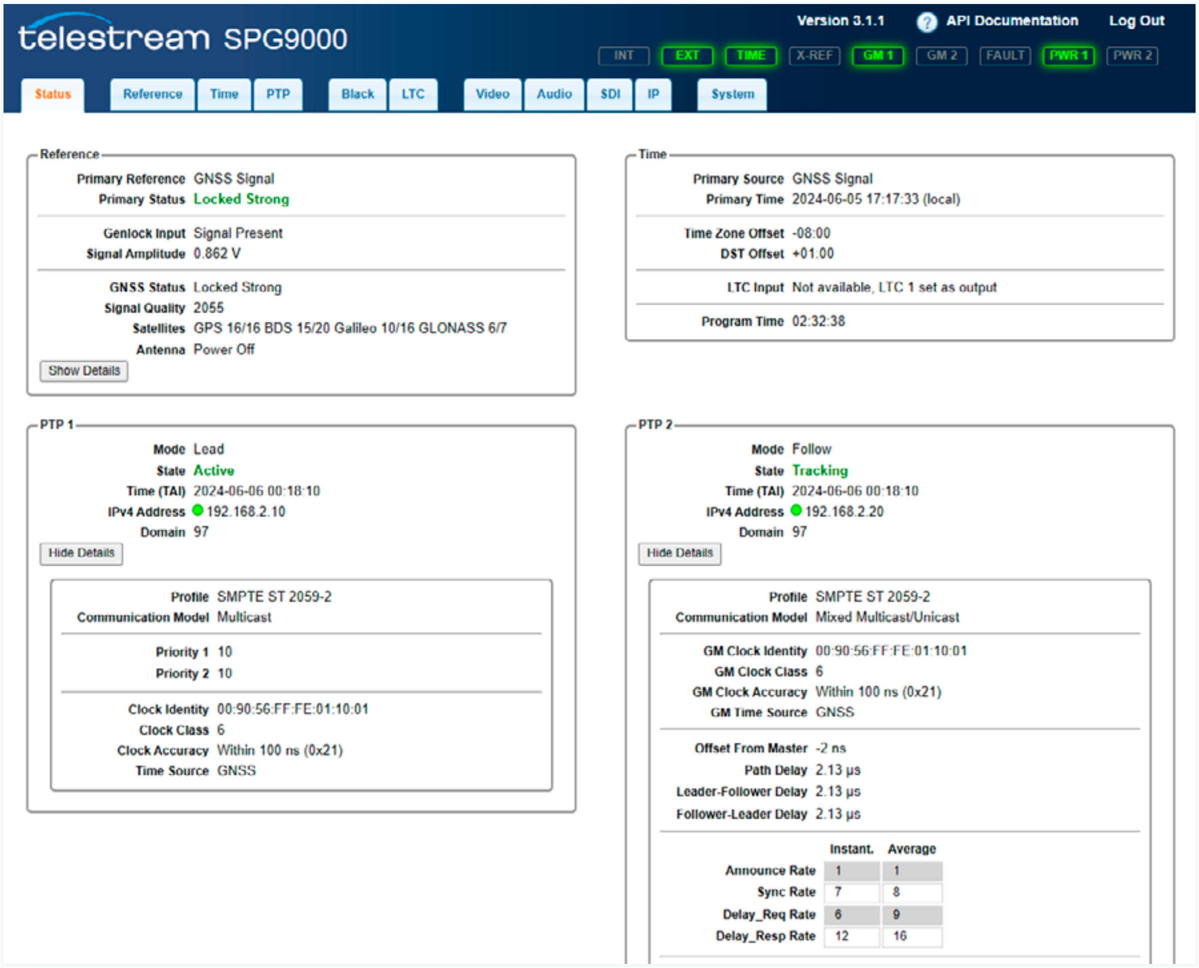

Figure 2. SPG9000 web browser status page.

During these tests it is important to monitor the operation of the system. Are all the video, audio and ancillary data (ANC) streams working correctly? Are there any clicks, pops, drop-outs, flashes, or freeze frames that occur in the program? Are all the end-point devices reporting a good PTP lock? Is NMOS running and providing connectivity? Is the scheduling system still on track? It can be challenging to verify these changes manually so often it helps to run stream captures during the test to analyze what is happening during the changeover.

A PTP monitor such as the Telestream PRISM can display graphs of how well the system is tracking the PTP, this will help in seeing if the PTP is shifting and which device is the active GM. Analyzing the switch and device logs can be useful as well.

Recommended GM changeover tests:

1. Graceful changeover with continuous PTP messages
2. Abrupt changeover on loss of PTP from GM
3. Abrupt changeover on failure of link carrying PTP
4. Loss of power to the active GM
5. Loss of power to parts of the network
6. Changeover with controlled timing offset between primary and backup GMs

PTP System Test Details

Test 1 - Graceful changeover with continuous PTP messages

During operation, one way this type of changeover will happen is if the active lead clock loses its GNSS or other reference input and goes into holdover. In this mode, the PTP messages from the primary GM are still present but the clock quality is degraded. The backup GM applies the BMCA and determines that it is better and so it starts transmitting Announce and Sync messages. Briefly, both GMs are sending messages, then the primary GM detects the “better” messages from the backup and goes to passive state. Because the messages from the two GMs overlap, there is minimal disruption to the system. This test can be initiated by changing priority of one of the GMs, or by disabling the antenna input to the primary. The recovery from this condition also has continuous messages and minimal disruption. One exception would be if the GM is recovering from loss of the GNSS antenna signal or other reference and starts sending PTP before it is truly stable.

Preferred test method:

Remove the antenna signal from primary GM. Observe the system behavior during the changeover. Wait five minutes then reconnect the antenna. Observe the system behavior during the recovery changeover.

Alternate test method:

Modify the Priority 2 on either GM to cause a changeover, observe the system behavior during the change-over. Wait five minutes then restore the original settings to recover and observe the system behavior during the recovery changeover.

Figure 3 shows the PRISM PTP Graphs for Leader-Follower and Follower-Leader delay. The Leader-Follower delay shows the graph of the forward path messages and indicates the difference in the time stamps for the PTP Sync messages that go from the lead clock to the follower. After the loss of GNSS there is a slight change to the delay as the backup master clock takes over as the Grandmaster. Notice how after the changeover that the graphs show the difference in the two clocks, but then the PTP servo loop adjusts the follower clock to match the leader clock. Resulting in the offset from the leader centering around zero as the two delay plots converge to a common value. This variation is due to the slight difference in time between the two GNSS receivers and does not cause a loss of PTP lock as this disturbance occurs. The user should note the difference in Grandmaster ID from the primary to the backup as the changeover occurs. The Follower-Leader graph shows the difference in the time stamps for the PTP Delay Request messages that go from the Follower to the Leader. A similar disturbance is shown around the same time within this graph.

Figure 3. PRISM PTP Graphs for Test 1.

PTP System Test Details

Test 2 - Abrupt changeover on loss of PTP from GM

During operation this type of changeover might happen if the GM stopped working or was disabled or reconfigured. A change of domain or changing from lead mode to follow mode are also examples which would cause this.

In this case, the messages from the leader abruptly stop. By the BMCA, all the other devices wait for the number of missing Announce messages configured in the "Announce timeout count" parameter. This is configurable from 2 to 10 but is usually set to 3.

Many systems use an announce period of one second and a timeout count of 3. Therefore, after three seconds, the other devices will assume the original GM is lost and start sending their own Announce and Sync messages. For some systems the backup GM will start sending messages after the three second delay. However, for systems using switches in boundary clock mode, only the boundary clock port connected to the GM will see the loss of PTP messages. The boundary clocks will continue to provide all the messages on all the other ports. Thus after the three second delay, only the boundary clock will react and take over as the GM. Once that happens, then the announce messages will reflect the lower clock quality of the boundary clock and then the backup SPG will detect that and take over as the GM. Depending on the topology this can extend the BMCA process and increase the opportunity for disruptions. The recovery is more graceful. When the original GM becomes available it will send out messages and take over the network without interruption to the PTP messages.

Preferred test method:

Disable the PTP from the primary GM, observe the system behavior during the changeover. Wait five minutes and then re-enable the GM and observe the system behavior during the recovery.

For this test the PTP port 1 was disabled on the primary SPG9000 that caused an abrupt loss of the Grandmaster and the PRISM media node showed a loss of lock in the system as the changeover to the backup occurred via the IP media switches. Figure 4 shows the Offset from Leader graph that indicates the error in the follower clock phase as measured by the PTP time stamps. The Offset from Leader shows the phase error of the phase-locked loop, that controls the follower clock. The Offset from Leader is a reasonable estimate of how well the PTP follower is locked to the master clock. Figure 5 shows more of an significant change of several milliseconds that might occur with older boundary clock implementations that would need to be addressed.

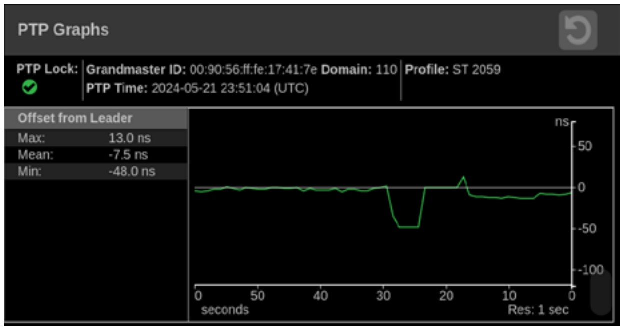

Figure 4. PRISM Offset from Leader PTP Graph for Test 2.

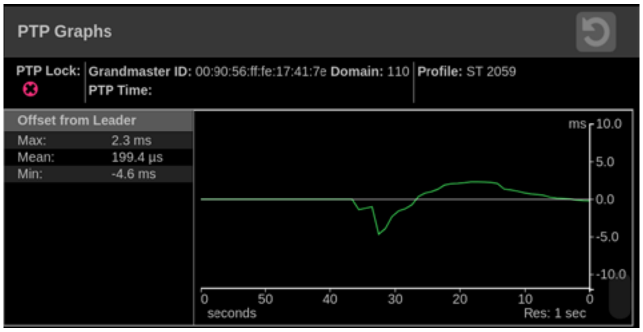

Figure 5. PRISM Offset from Leader PTP Graph for Test 2 with significant shift.

PTP System Test Details

Test 3 - Abrupt changeover on Loss of link carrying PTP

Examples of how this could happen during operation include a fiber or cable link getting disconnected, or an SFP (Small Form-factor Pluggable) failure. In some cases, the resulting changeover can be identical to Test 2. However, in other cases the switch can detect the link down and start the BMCA process sooner. Generally, this test is equivalent or less stressful than Test 2, but it can appear quite different if one is examining the PTP traffic or log messages.

Preferred test method:

Disconnect the PTP link at the GM. Observe the system behavior during the changeover. Wait five minutes then reconnect the link to initiate the recovery and observe the system behavior during the recovery.

Test 4 - Loss of power to the active GM

Examples of how this could happen during operation include GM failure, overloaded circuit breakers, or other power disruptions. This changeover from loss of power is equivalent to Tests 2 or 3 depending on the GM and network. However, the recovery when the GM reboots has some unique challenges and that is what makes this test important. When the GM restarts, it is important that the device gets fully stabilized and correctly timed before it takes over as the active GM. If this is not true, then there can be a significant timing shift when the GM restarts.

Preferred test method:

Remove power to the primary GM. Observe the system behavior during the changeover. Wait five minutes then re-apply power to the GM. Observe the system behavior as the GM reboots and resumes the GM role.

Test 5 - Loss of power to parts of the network

Examples of how this could happen during operation include a switch failure or loss of power. Depending on the design of the network, PTP GMs may hook to a leaf, spine, or other type of switch. For a redundant network, the loss of one switch should not block the PTP, but it may cause it to re-route or switch to a backup GM. This test entails de-powering the switches in the PTP path and ensuring that PTP is still available to all the devices, and that operation is not disrupted.

Preferred test method:

Examine the block diagram of the system and identify the elements critical to the PTP operation. One by one, remove and restore power to these elements while observing the system behavior. In some network topologies it may be appropriate to de-power multiple elements simultaneously to force the system to reroute or use the available backups. For example in Figure 1 above, removing power to the blue network spine switch will force the system to use the backup PTP path.

PTP System Test Details

Test 6 - Changeover with controlled timing offset between primary and backup GMs

In an ideal system all the GMs would have the exact same time so that there would be no timing shift during a changeover. However, in a reality, each GM may have slightly different timing or may be connected via a different network path with a different delay. The intent of this test is to add known offsets between the primary and backup GM and see how the system responds during GM changeovers. By testing with a few offsets, it is usually possible to determine if some devices in the system are sensitive to this type of timing shift.

Figure 6. PRISM Offset from Leader PTP Graph for Test 6.

Preferred test method:

Use the capabilities of the GM to offset one GM from the other. If the GMs are GNSS (Global Navigation Satellite System) referenced, then they may have an antenna cable length compensation which can be used to shift one GM. If the GMs use PTP as a reference, then often there is an asymmetry compensation which can implement a shift. If the GMs can reference to black burst, then the timing of the syncs to which they are locked can be adjusted to create a timing offset.

Once the two GMs have a known offset, create a changeover using one of the methods above and observe the behavior of the system. Usually Test method 1 is preferred since it is the least disruptive, so the effects of the timing shift are more apparent. Typical PTP accuracy is in the 0.1 to 1 µs range, so it is reasonable to test with GM offsets of 0.5 µs and 1 µs. Stress testing at higher offsets like 2 µs, 5 µs or 10 µs can also be useful.

Figure 6 shows an offset of 10 µs that was created using an antenna cable delay of -10 µs configured in the Reference tab of the SPG9000 web user interface. The PRISM PTP graphs show this 10 µs offset when the changeover occurs and then converges to a nanosecond range dependent on the delays through the network. Be sure to remove the timing offset when the tests are complete. Otherwise, the system will have to work harder if a real GM changeover happens during live operation.

Conclusion

Precision Time Protocol can provide a reliable timing system for IP Media networks however understanding how a changeover occurs between the primary and backup clocks within system configurations is critical to the reliability of PTP in ensuring an error-free system. The series of tests outlined are one part of the testing that should be considered during commissioning. Engineering staff should also consider test to the switching fabric and IP media node end points along with other GM specific tests. By performing testing during the initial design and commissioning of the installation. Engineering staff can become familiar with the operational performance of PTP and see the behavior of the system in response to changes within the media network. By carrying out this series of tests to the system, they can observe the performance of PTP and the effects of changing over Grandmasters between the primary and backup systems can be validated. A good understanding of PTP and the effects a changeover can have on switches and media nodes in the network can help in quickly resolving issues should they occur.

Learn more

Find out more information on the SPG9000 timing and reference system.

Telestream SPG9000 Timing and Reference System PTP Webinar Series