

White Paper
**Improving Mobile Backhaul Network Reliability with
Carrier-Class IEEE 1588-2008 (PTP)**




Introduction

Wireless backhaul networks have rapidly shifted from circuit switched to packet switched technologies to meet the exploding demand for bandwidth, but whereas circuit switched TDM networks carried precise frequency synchronization throughout the network at the physical layer, packet switched networks based on asynchronous Ethernet cannot provide such deterministic synchronization services. This can cause problems for operators, since access platforms such as cellular base stations, have traditionally relied on synchronization delivered over the network backhaul connection to assure high QoS for end user applications. A key factor in the evolution from circuit based to packet based backhaul in telecom networks is the ability to deliver reliable carrier grade synchronization to wireless base stations and other access platforms.

IEEE 1588-2008 (Precision Time Protocol or PTP) has gained traction as the technology of choice to deliver synchronization to remote telecom elements over Ethernet backhaul connections. The following illustration shows a typical example of PTP synchronization for cellular networks.

Figure 1 • PTP Synchronization for Cellular Networks



Delivery of synchronization to next generation base stations will rely on PTP grandmaster clocks deployed in the network. Sync packets flow from the grandmaster clock to the slave clocks in the base stations.

Mobile base stations including 3G, LTE, and LTE-Advanced all require frequency synchronization of 16 ppb to assure high quality of service and avoid dropped calls. In addition, TDD and LTE-Advanced networks also require tight phase synchronization of 1.1 us. The ITU have issued a complete series of standards and recommendations covering the use of IEEE-1588 PTP to meet these demanding frequency and phase synchronization requirements.

The loss of physical layer sync has generated a requirement for new base station designs incorporating PTP slave clocks that will meet the 16 ppb requirement using packet technology. Such PTP slaves in the base stations rely on access to a carrier-class PTP grandmaster clock deployed in the mobile access network.

Grandmaster Redundancy Considerations

Network deployment and reference network test models have been established to assist carrier engineering staff develop specific deployment rules for IEEE 1588-2008 synchronization solutions in their networks. To provide failover protection, there are two main approaches to consider:

- Network-based redundancy. Two entirely separate grandmasters are deployed, possibly at different geographical locations within the network.
- Built-in grandmaster hardware redundancy. The grandmaster employs an active and a standby card. The standby master shadows the active master with all settings including the IP address, but remains dormant until the active master fails.

This paper discusses the advantages and disadvantages of each technique, and describes features built into Microsemi hardware to provide carrier-class reliability and performance.

PTP Slave Locking Considerations

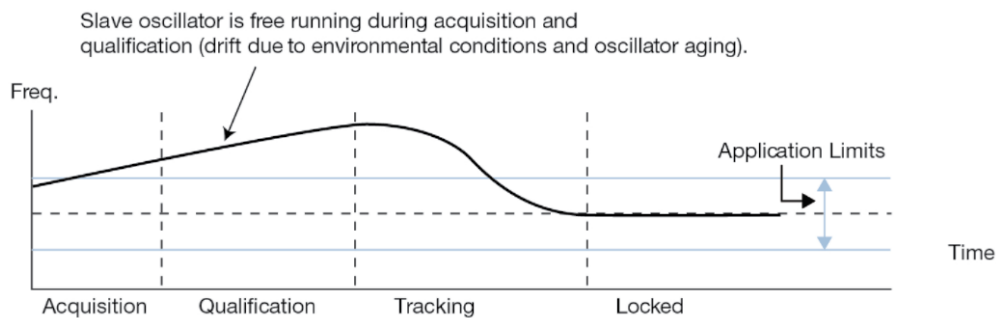
Before discussing redundancy, it is necessary to understand how PTP slave clocks synchronize to a grandmaster. The slave clock establishes communication with the grandmaster by requesting a reservation for a synchronization flow, specifying parameters such as message rate and reservation duration. Once the synchronization flow is established with the grandmaster, the slave clock goes through the stages outlined in the following table to lock to the master.

Table 1 • IEEE 1588 PTP Slave Clock Lock Stages

Stage	Description
Acquisition	Establishes sync flow and determines initial offset.
Qualification	Monitors sync packet stability to qualify the master clock for tracking.
Tracking	Begins reference oscillator alignment process to tune the slave oscillator to the grandmaster.
Lock	Slave oscillator is now locked to the grandmaster and will begin long-term tracking to hold the frequency stability inside the application limits.

The following illustration shows how the process works.

Figure 2 • IEEE 1588 PTP Slave Clock Acquisition and Locking Process

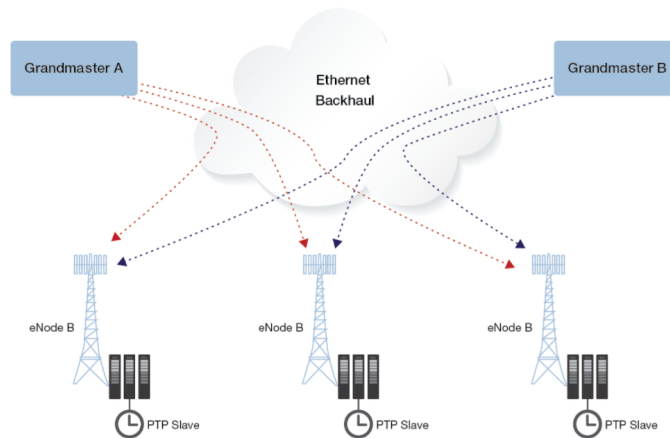


Initially, the slave free-runs at an indeterminate frequency while the acquisition and qualification processes take place. During this period, both the output frequency and time offset may be well outside the specification required by the given application. The process of aligning the frequency and time to the master takes place in the tracking stage. By the time the locked stage is reached, the slave is well within the application limits and able to maintain the target level of performance over the long term.

Network-Based Redundancy

In network-based redundancy, two grandmasters are employed. These may be co-located, but they are most often at different geographic locations to provide maximum protection against failure, as shown in the following illustration.

Figure 3 • Network-Based Redundancy



The PTP slave clock maintains a list of acceptable masters to use and establishes communication with at least one of these masters. If the first master fails, the slave may then attempt to lock to the second master, as shown in the previous illustration.

Network-based redundancy has the property that it protects against network failure in the vicinity of the grandmaster as well as failure of the grandmaster itself, since the second grandmaster is normally connected by a separate access link into a different part of the network. However, there are some significant issues created by switching between two different grandmasters.

First, PTP slave clocks for stringent telecom applications can take up to an hour to achieve fully locked status. This is because the variation in message delay through the packet network is several orders of magnitude larger than the accuracy requirements for telecom applications, and the slave has to build its filter history in order to compensate for the network path behavior. The network path between the slave and the two grandmasters will be different. When switching to the new grandmaster, the slave has to re-build its filter history to match the new path. While the slave goes through the acquiring, qualifying, and tracking process with the new master, the output may drift outside the application limits.

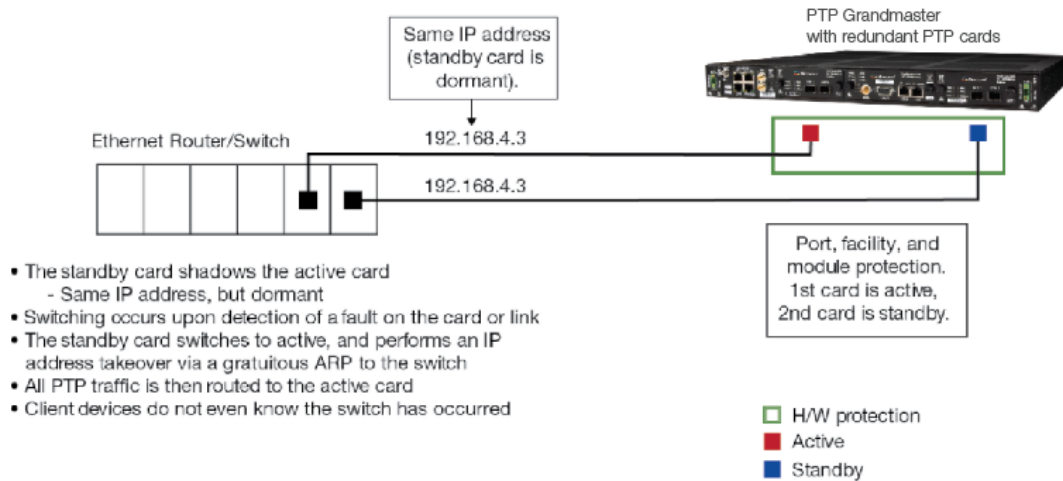
Second, the path delay to the new grandmaster may be so significantly different that it will result in a large phase or time offset between the two. The slave has to identify this offset and compensate for it correctly to avoid building a phase shift into the output clock. It should be noted that a test that compares switching between different network paths has been constructed by ITU-T (G.8261 Appendix VI, Test Case 17), and qualified slaves should pass such a test with an acceptable performance degradation.

Third, there may be a significant phase or time offset between the two master clocks if they are locked to different primary reference sources. This time difference may cause the slave to skew its output frequency outside the application limits during the switchover process in order to align its phase and time to the new master. There are some techniques within the slave that can be used to mitigate these issues. For example, since the frequencies of the two grandmasters can be expected to be almost identical (within 1 part in 10^{-11}), the process of acquiring the new frequency should not take as long as if starting completely from scratch. It may also be possible, given sufficient processing resources in the slave and capacity at the grandmasters, to have the slave track two different masters simultaneously, such that when a switchover is required the slave has already acquired the filter history and offset between the two masters.

Built-in Grandmaster Hardware Redundancy

A more robust method of providing grandmaster redundancy is to implement a carrierclass IEEE 1588-2008 grandmaster clock with built-in hardware redundancy. A fully redundant grandmaster clock employs an active and standby clock, synchronized to redundant primary reference sources (for example, GPS or T1/E1), as shown in the following illustration.

Figure 4 • Built-in Hardware Redundancy



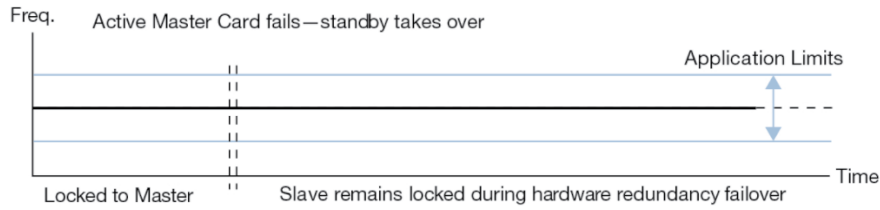
The active clock connects through the network switching fabric to service all the PTP slave synchronization flows in its timing domain.

The standby master shadows the active master with all settings including the IP address, but remains passive until either the active master fails, or the link to the switch goes down. At that point, the standby grandmaster goes active and takes over the IP address with a gratuitous ARP to the switch. All PTP sync flows are then serviced by the standby master with no impact to the PTP slave community.


The big advantage of this configuration is that since the active and standby clocks share a common reference and common network location, the PTP slave devices see no synchronization offset during a failover switching scenario. The slave clocks will all remain locked to the redundantly protected grandmaster, and will not be forced to acquire, qualify and track to a new grandmaster clock with an unknown offset, or connected through a different network path.

The following illustration shows the benefits of hardware-based grandmaster redundancy. IEEE PTP grandmaster clocks that employ hardware redundancy protect slave clocks from potential service impact during failover scenarios.

Figure 5 • Benefits of Hardware-Based Grandmaster Redundancy



Scenario 1: Slave remains locked during hardware redundancy switching failover of the Master. Hardware redundancy protects NGN services.



Scenario 2: Slave switches to a back-up Master. Potential service impact while slave is out of lock while it acquires and tracks to the new Master reference frequency.

The PTP slave clock in scenario 1 remains locked throughout the switchover process, and in practice may not even know that the event has taken place. On the other hand, the PTP slave clock shown in scenario 2 goes out of lock as it switches to its backup master clock, and initializes the acquisition, qualification, tracking, and locking process with the new grandmaster clock. The relocking duration and resulting phase and frequency offsets are functions of the slave clock design, grandmaster performance, and network path and delay differences, much of which can be avoided by deploying a carrier-class IEEE 1588 PTP grandmaster clock with hardware redundancy.

General Redundancy Issues

Some reliability issues affect both redundancy methods. The failure of a primary reference source (PRS) is an important issue in the design of a grandmaster, particularly if accurate time is required as well as frequency. For frequency applications, backup may be provided using a T1 or E1 frequency reference. In the event of a PRS failure, a T1/E1 input reference that is traceable to a Stratum 1 clock source will keep the PTP time stamps from drifting and is an acceptable redundancy engineering practice. The failure of the slave or its access link will only affect the network element(s) dependent on that slave. The impact on the overall network is low. It is not common practice to provide redundant links or clocks to the base station in cellular networks due to the cost of such protection.

Summary

With IEEE 1588 (PTP) slave clocks, consideration must be given to how synchronization is protected during failover events. Deployment of carrier-class grandmaster clocks with built-in hardware redundancy assures that all slave clocks remained locked and fully protected during clock failover events.

Microsemi PTP Grandmaster devices, such as the TimeProvider® 5000 and the PTP server blades for the TimeHub and SSU 2000 synchronization systems, provide built-in hardware redundancy to assure maximum performance from the PTP synchronization network. They can also lock to several different primary reference sources, providing failover protection in the event of a reference failure. Microsemi PTP slaves also enable network-based redundancy by being able to switch between two grandmasters in the event of a failure.



Microsemi Corporate Headquarters

One Enterprise, Aliso Viejo,
CA 92656 USA
Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

© 2017 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at www.microsemi.com.

MSCC-0104-WP-01004-1.00-0917