



Synchronization for Next Generation Networks

Best Practices for IEEE 1588 (PTP) Network Deployment

IEEE 1588 Version 2 means that precise timing and synchronization over Ethernet is now a reality — but the solution is only as good as the sum of its parts

Table of Contents

Not All Clocks Are Created Equal.....	1
Know Your Telecom Profile.....	2
Choosing the Right Clients	3
Managing Your PTP	5
PTP Best Practices.....	6

Abstract

With the recent ratification of the IEEE 1588 Version 2 protocol specification — also known as Precision Time Protocol (PTP) Version 2 — distribution of precise time and frequency over Ethernet network infrastructures enters a new phase. By replacing costly TDM backhaul lines with Ethernet, service providers can reduce recurring backhaul line costs by up to 90% while simultaneously boosting network capacity for bandwidth hungry applications like text messaging, music downloads and video streaming.

Until now, many network service providers have been reluctant to transition from TDM to Ethernet backhaul due to Ethernet's inherent lack of an adequate timing and synchronization mechanism. However, with the advent of PTPv2, providing precise timing and synchronization over Ethernet is no longer an issue.

PTPv2 is a next protocol specifically designed to provide precise timing and synchronization over packet-based Ethernet infrastructures. With the advent of v2, PTP has now been adapted to meet the more sophisticated synchronization requirements of telecommunications applications. PTPv2 captures those requirements by providing a set of added capabilities and protocol extensions that allow service providers to fine tune their packet-based networks for the stringent timing and synchronization requirements of telecom-oriented applications. Some of these capabilities, like frequency synchronization and multicast support, are explicitly supported within the PTPv2 framework while other capabilities, such as support of unicast and telecom profile extensions, provide the higher level of accuracy and performance optimization that telecom requires.

As service providers begin to deploy PTPv2, they need to understand how these different capabilities, or lack of them, may impact their network performance and service level agreements. Optimal results can be achieved if they can anticipate future requirements and choose a carrier-class PTP solution that offers the highest level of resiliency, performance, scalability and management.

Not All Clocks Are Created Equal

PTP employs a hierarchy of clock types to ensure that precise timing and synchronization is maintained between the timing and synchronization source - a grandmaster, boundary or transparent clock - and the numerous PTP clients that are distributed throughout the network. PTP clocks use hardware time stamping to precisely mark each packet which mitigated the delays that are typically incurred as timing and synchronization packets traverse the network.

A grandmaster is the highest-ranking clock within its network domain. There may be a single or multiple grandmaster clocks within the same network. PTP grandmasters can be deployed as either standalone devices or as plug in modules or "blades" that can be integrated into an existing synchronization supply unit (SSU) or building integrated timing supply (BITS) shelf. Grandmasters are just as important as their name implies. They are an essential component of any PTP solution because they are the primary reference source (PRS) for all other PTP elements within their network domain.

Boundary clocks act as “intermediaries” between a PTP grandmaster and its PTP clients. Boundary clocks typically live at the edge of the network and receive their primary reference from an upstream grandmaster while simultaneously serving as a master to the downstream PTP clients in the network. The boundary clock mitigates the number of network hops and resulting delays that occur in the packet network between the grandmaster and its clients. Some PTP clocks can function as either a grandmaster or boundary clock. The primary advantage of boundary clocks is that they eliminate the need to have a GPS timing source at every server location, which in turn significantly reduces the cost of deploying an end-to-end PTP solution.

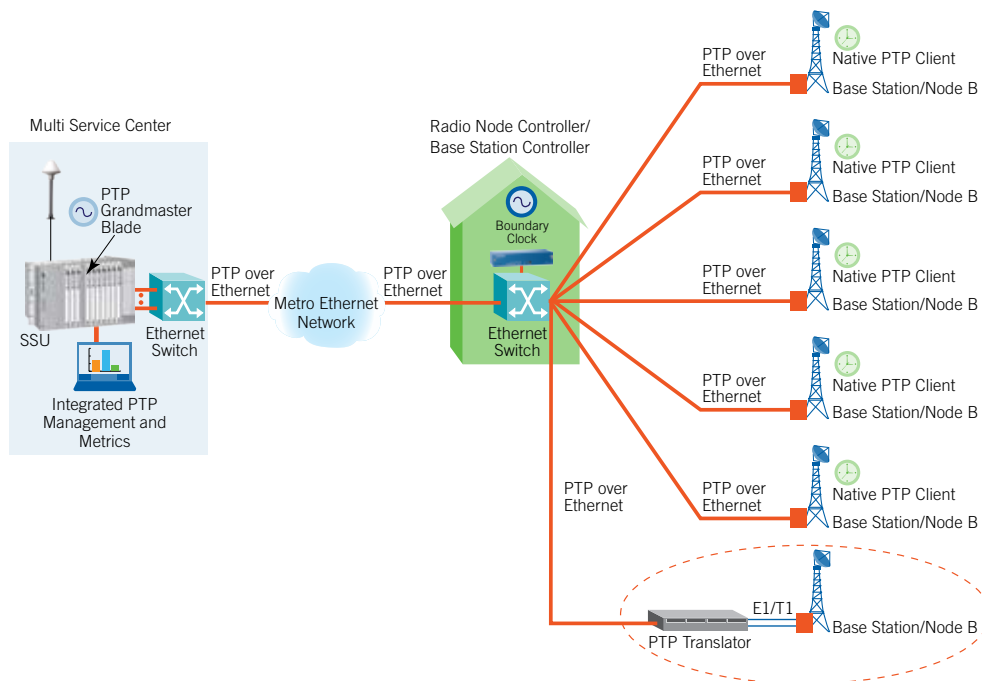


FIG 1 PTP Grandmaster and Boundary Clock Implementation

PTPv2 also specifies transparent clocks. A transparent clock maintains its own precise internal clock that forwards timing packets from one network segment to another. It compensates for packet delays by updating the timestamps of each packet that passed through it.

Know Your Telecom Profile

Telecom profiles are extensions to the PTPv2 protocol specification that define a unique set of capabilities or enhancements that are required to adequately support PTP telecommunications applications such as Ethernet Backhaul, synchronization of IP DSLAMs and Passive Optical Networks (PON).

These enhancements include:

Support for wireless standards: Different wireless standards specify various time and frequency offsets that must be met in order to prevent calls from being dropped as mobile users move from one base station coverage area to another. In particular, ITU-T Recommendation G.8261/Y.1361 (formerly G.pactiming), “Timing and Synchronization Aspects in Packet Networks,” specifies the upper limits of packet delay variation (PDV) that network equipment at the TDM interfaces. The telecom profile also specifies a maximum error level that can be tolerated in order to meet basic quality of service requirements. For example, GSM and CDMA base stations require a threshold of .05 parts per million (ppm) be maintained by their timing source in order to ensure that calls are not dropped.

Unicast in addition to multicast support: When PTP was first introduced it only specified support for multicast message exchanges between PTP clock sources and clients. However, in the telecom environment unicast support is required in order to fine tune and optimize network performance. With multicast, every device in the multicast group must examine every message. This means that each client must listen to every other client’s messages to its master — saturating client processors with messages they throw away. Since clients should be the lowest-cost elements of the synchronization ecosystem, minimizing the amount of processing power required is a priority — and so too is minimizing the amount of bandwidth consumed. Unicast support helps achieve these objectives.

Higher message rates: Based on the first version of the standard, PTP masters and clients send message packets every 30 seconds by default. Telecom’s much higher timing requirements, however, call for commensurately higher message rates (up to 64 times higher). The required frequency is dependent on several factors, for example, the performance of the client device, the stability of the client’s local oscillator, and the amount of noise on the network. A good starting point is 16 messages per second — and to adjust from there based on bandwidth availability and timing performance. A grandmaster clock must therefore be able to both sustain high message rates for potentially thousands of clients, and also allow operators to adjust those rates as needed.

Shorter message formats: The timestamps in Version 1 were 165 octets and included information about the source and quality of the clock. As the source and quality don’t change often, PTP Version 2 separates the message into two components: 1) an announce message with details about the clock, and 2) a smaller 44 octet PDU for synchronization messages. Trimming the message size thus allows the telecom provider to achieve the higher message rates required while minimizing the bandwidth required to do that.

“Best grandmaster” algorithm: It is possible that PTP clients and grandmasters will lose their timing reference due to network outages or other reasons. PTPv2 specifies that clients support a best master clock (BMC) algorithm. This enables the device to scan the network and identify the best possible reference available as a replacement — taking into account the quality of the timestamps received as measured via special message exchange.

Specified transport over more network layer protocols (e.g., UDP/IPv4, UDP/IPv6, Ethernet): Version 2 gives operators the flexibility to specify alternative network layer protocols over which to route synchronization traffic. This allows operators to better manage how they allocate resources according to packet type and network transport availability.

Static and Dynamic Reservations: Clients must be able to support both static and dynamic reservations. In static reservation mode the grandmaster interacts with a specified list of statically allocated clients and controls the allocation of resources to each one. In dynamic mode PTP clients dynamically request resources from the grandmaster or boundary clock and auto-negotiate for allocation of available resources. PTP clients should support both static and dynamic reservations to ensure optimal network performance and efficient allocation of network resources.

PTPv2 telecom profiles provide service providers and network equipment vendors with the ability to “fine-tune” timing and synchronization parameters to meet the unique requirements of their specific applications and are an indispensable component of any comprehensive PTP solution.

Choosing the Right Clients

PTP clients may be hardware and/or software components that reside in an end-point network edge device such as wireless base station, DSLAM, or LTE. Clients continuously exchange synchronization messages with PTP clocks to ensure consistently precise timing and synchronization. These message types include:

- Signaling Messages
- Sync Messages
- Delay Request Messages
- Delay Response Messages
- Management Messages

Figure 2 illustrates the types of PTP messages that are typically exchanged between the grandmaster clock and its PTP clients.

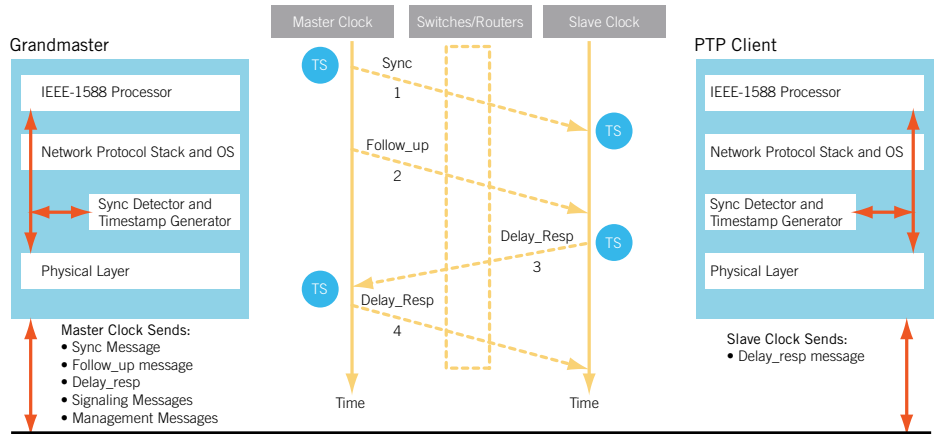


FIG 2 Message exchange between Grandmaster and PTP client.

PTP clients may be implemented in several different ways. The simplest and least intrusive implementation is a PTP translator. A PTP translator is essentially a “client in a box.” It is a standalone device that has a PTP-over-Ethernet interface for communicating with the PTP grandmaster or boundary clock and multiple T1/E1 interfaces for communicating with legacy TDM-based end-point devices such as wireless base stations that do not have a native PTP client on board. The advantage of implementing a PTP translator solution is that it enables network service providers and equipment vendors to rapidly implement Ethernet backhaul at very low cost without having to make any modifications to existing legacy devices. The end point device thinks it is still communicating over a legacy TDM line while the translator transparently performs the TDM-to-Ethernet conversion. Figure 3 illustrates a typical PTP translator solution.




FIG 3 PTP Translators convert TDM signals to Ethernet signals on behalf of network elements that do not have their own native PTP clients.

Alternatively, PTP clients may also be sourced off-the-shelf from third-party client vendors. By virtue of the fact that PTP is a standard, any third-party client should be interoperable with any PTP grandmaster clock. The operative word here is “should.” There may be instances where a vendor updates their PTP client implementation without completing thorough interoperability testing with the v2 grandmaster vendor beforehand which could result in unforeseen interoperability problems and in extreme cases a potential disruption of critical services.

A third means of implementing PTP clients is to purchase a development tool kit and software licensing rights from a client software vendor. In this scenario, client software may be directly ported onto a network endpoint device such as a wireless basestation or DSLAM. As is the case with sourcing any third-party client solution, the service provider or network element manufacturer must assure continued interoperability and compatibility between the grandmaster and client vendors as they introduce new features and new software releases.

In many cases, service providers and network equipment vendors may choose to implement a “hybrid” client solution where a mix of native PTP client devices and PTP translators coexist within the same network infrastructure. Figure 4 illustrates a hybrid PTP solution consisting of both native PTP client and PTP translator. This type of solution is only possible if all PTP network components — from grandmaster to boundary clocks to clients — remain in lock-step with respect to their release levels and supported feature sets.

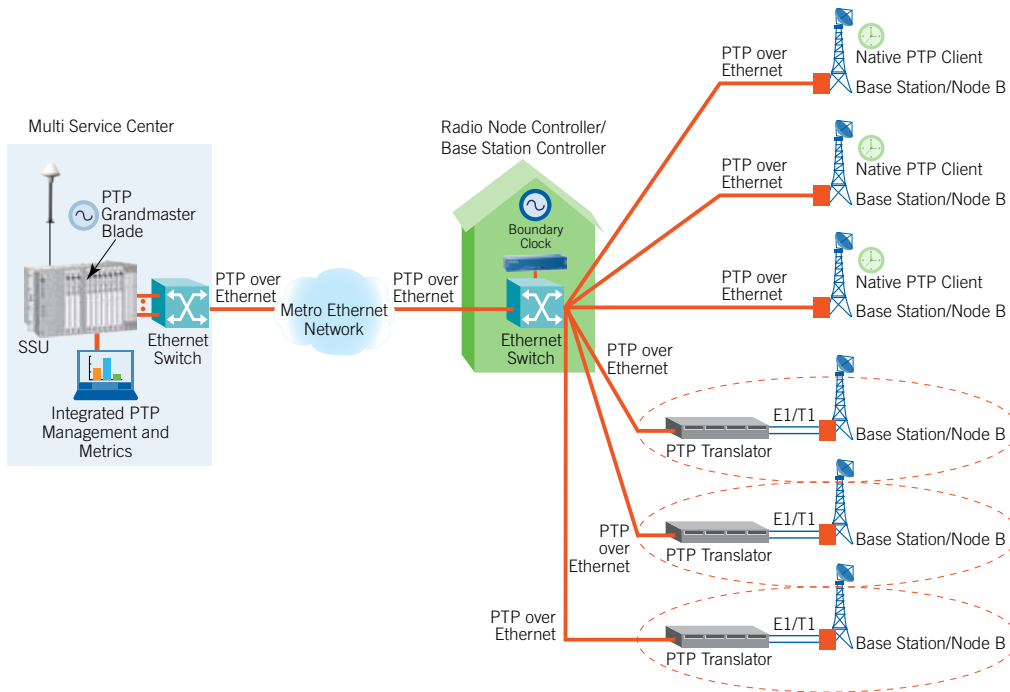


FIG 4 “Hybrid” PTP client/PTP Translator solution.

Managing Your PTP

An obvious consideration when adopting any multi-vendor, PTP-based solution is how to manage it. The risk is that network operators could end up with a hodgepodge of PTP network elements, each with its own unique management interface for configuration, monitoring and control. The optimal approach is to have complete end-to-end visibility across all PTP network elements — from grandmasters to boundary clocks to remote clients — as well as robust set of network performance metrics and monitoring tools.

A unified PTP management solution offers significant operational advantages including the ability to perform end-to-end performance analysis and troubleshooting along with the ability to manage all aspects of the PTP network including:

- Faults and Alarms
- Configuration
- Accounting
- Performance Metrics
- Security
- Inventory Management
- Software Downlpads

Figure 5 illustrates a unified network management solution which provides complete end-to-end visibility of all PTP network elements under a common network management umbrella.

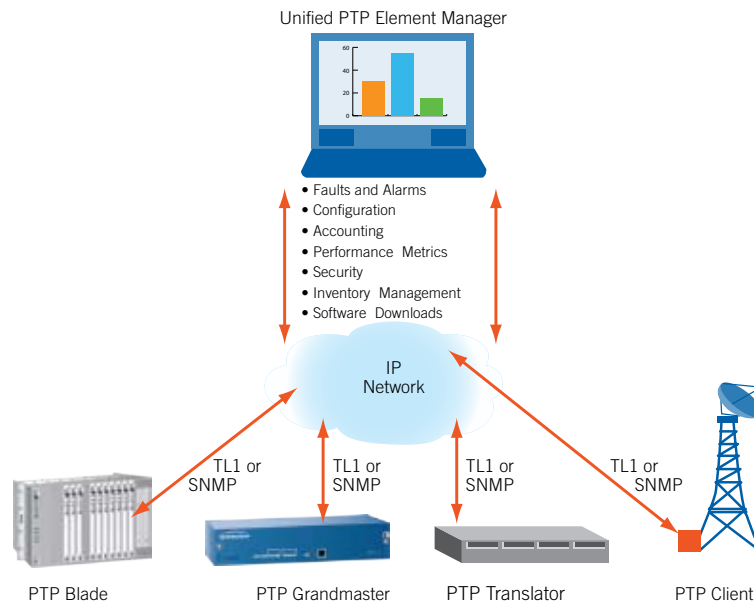


FIG 5 Unified PTP Management Solution

PTP Best Practices

PTPv2 sets a new and improved standard for packet-based synchronization. It also challenges telecom vendors and service providers to employ best practices to meet that standard. To ensure the best network design, service providers and network equipment vendors should consider several other factors when planning out their PTP deployment:

Distributed Grandmasters: As a general rule, PTP network designs should seek to minimize the number of hops between Grandmaster clocks and their clients in order to mitigate packet timing delays. Where excessive hops cannot be avoided, network planners should deploy distributed grandmasters near the edge of the network to guarantee the highest degree of timing and synchronization accuracy.

Redundancy: PTPv2 specifies a best master clock (BMC) algorithm which allows clients to seek out an alternative grandmaster if the original grandmaster fails or degrades. However, dependence on the grandmaster algorithm may be insufficient in some applications where five-nines (99.999%) or better reliability is required. Where carrier-class reliability is required, PTP network components must have redundant power supplies and clock modules in order to guarantee maximum reliability and uptime.

Holdover Options: Holdover is a key component of PTP best practices. Holdover specifies how long an oscillator can remain within a specified accuracy range after losing its external reference such as GPS. Network elements, such as wireless base stations, are often placed in remote locations — so a PTP clock's holdover rate is critical to keeping a network element within operational tolerances. A high-quality quartz oscillator, for example, typically has a holdover period of 24 hours while rubidium oscillators can have a holdover period of seven days or more. PTP best practices demand that PTP network elements provide both quartz and Rubidium oscillator options.

"Purpose-Built" IEEE 1588 v2 Components: The PTP standard specifically calls for a "fit for purpose" method of timing and synchronization in telecom applications. This implies the use of "purpose-built" PTP network components that have been designed from the ground up to provide the highest levels of reliability, scalability and performance.

Clearly, if service providers and network equipment vendors are to meet the challenge of providing timing and synchronization solutions in an Ethernet world they must go above and beyond the minimum requirements called out by the PTPv2 specification. Mere PTPv2 compliance is not enough. Rather they must strive to offer more complete, end-to-end PTP over Ethernet solutions that not only provide timing and frequency but also deliver the robust feature sets, carrier-class reliability and seamless network management capabilities that next generation PTP applications will demand.

More Background Reading on PTP

Please refer to these other White Papers produced by Symmetricom.

1. *"Deployment of Precision Time Protocol for Synchronization of GSM and UMTS Basestations"*, Symmetricom white paper, May 2008
2. *"Advances in Backhaul Synchronization - Maximizing ROI"*, Symmetricom white paper, May 2008



SYMMETRICOM, INC.
2300 Orchard Parkway
San Jose, California
95131-1017
tel: 408.433.0910
fax: 408.428.7896
info@symmetricom.com
www.symmetricom.com