Best Practices for Implementing PTP in the Power Industry

Larry Thoma

Introduction

Power system applications such as Sampled Values (SV), synchrophasors, and traveling-wave fault location (TWFL) require time synchronization accurate to 1 microsecond or better for proper operation. The Global Navigation Satellite System (GNSS) uses satellites and clocks to provide high-accuracy time-source synchronization to any location capable of locking to a GNSS constellation. However, what happens when the GNSS signals are unavailable or compromised by a spoofing or jamming attack? Precision Time Protocol (PTP) timing solutions can use connected networks to maintain application timing and synchronization during events when GNSS signals are unavailable.

When building new generation plants and substations or expanding existing sites, the cost of IED cabling and the distance limitations imposed by cable systems often hinder site development. One method to reduce costs is to use the same communications medium for IED control and timing. Initially, this method used serial communications, but newer IEDs use Ethernet-based networks and communications protocols. The IEC 61850 standards are based on using a shared Ethernet communications backbone for multiple applications. By conforming to the standards, IEDs from different manufacturers can interact without custom or proprietary links. PTP leverages the shared Ethernet network to distribute submicrosecond timing from timing sources to end devices and applications that require accurate time.

There is a strong link between communications network synchronization and time distribution. When two devices communicate over a communications network, a method of synchronization is essential to successfully transfer data.

The power industry's PTP implementation is still in its infancy. This paper explores the implementation of PTP across different Ethernet network architectures and looks at methods to provide PTP redundancy. In addition, the paper discusses the different types of PTP clocks that are available and makes recommendations on which PTP clocks to select for specific applications.

Background

Over the past few decades, timing requirements for power industry protection systems have changed drastically. Protection applications have transitioned from not requiring a time reference to needing submicrosecond source timing accuracy and precision.

The need for timing in the substation began with the introduction of digital relays and their ability to provide Sequence of Events (SOE) reporting. Timing allows for the correlation of event reports across the power network and provides data to reconstruct and determine the cause of the power system event. Major outages, such as the 2003 Northeast Blackout, led to standards that require IEDs to synchronize to a common time reference. One-millisecond timing accuracy is adequate for SOE reporting.

Newer substation applications require a 1-microsecond or better timing accuracy from the time source. Dedicated IRIG-B timing and cabling solutions can provide this accuracy to IEDs, with control and data traffic on separate cabling. However, with the advent of PTP, dedicated (separate cabling) timing solutions are no longer the only methods to distribute submicrosecond timing to end devices. PTP is an Ethernet-based protocol that allows timing information to share the same communications infrastructure as control and data traffic.

There are several benefits to using PTP. PTP removes the need for dedicated cabling for timing solutions, thereby reducing the cabling between clocks and IEDs. PTP also removes the distance limitations imposed by IRIG-B cable-based timing solutions. Fiber-optic Ethernet networks can reach over 100 kilometers. Prior to PTP, a clock had to be placed close to its antenna and the antenna needed close to a 360-degree view of the sky. In many cases, this meant the clock had to be placed far away from the equipment requiring the time reference. The maximum distance that an IRIG-B signal can transfer over a coaxial cable is 500 feet (150 meter), which creates challenges for environments with large distances between the clock and the end equipment. A hydroelectric dam is a good example of this type of situation. The protection devices are typically at the bottom of the dam close to the turbine generation equipment. The timing equipment, which includes the clock and antenna, has to be at the top of the dam to receive and lock on to a GNSS signal. With PTP, it is possible to distribute high-accuracy timing over a much larger distance using a communications network, thereby solving applications where distance or antenna obstruction were an issue. The PTP protocol supports redundancy and allows for the configuration of multiple grandmaster clocks in a network. PTP selects a primary clock from the available potential grandmaster clocks on a network, and the primary clock is the ultimate time source for the PTP network. Additional potential grandmaster clocks can automatically assume the role of grandmaster if the current grandmaster becomes unavailable. PTP supports the redundant paths from the grandmaster to the IEDs that are engineered in the Ethernet network topology. Using redundant network topology designs helps mitigate the loss of timing due to a network path failure.

PTP and Sampled Values

Consider an SV measurement solution with merging units (MUs). Figure 1 shows SV IEDs and MUs connected to the substation network, with the MU process and station bus traffic sharing the network backbone. All traffic (measurement, management, timing, and so on) uses the same network. The SV IEDs and MUs synchronize to the PTP timing information provided by the shared network.

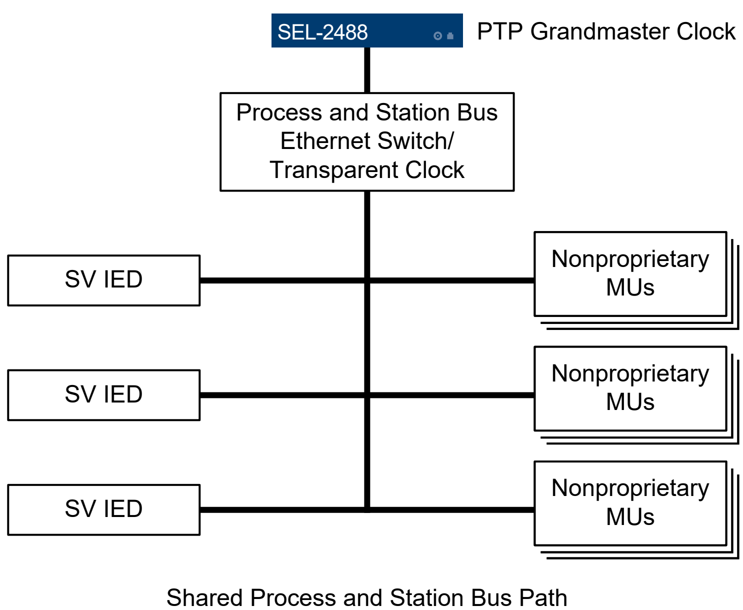

Figure 1 SV solution with network-connected MUs and a shared bus

Typically, network specifications define the host network topology, and the SV solution superimposes on that network topology. PTP, as shown in the examples in this paper, provides timing to the appropriate devices over these networks without special cabling or configuration.

Applying PTP

PTP, as an Ethernet-based protocol, inherits the topology and characteristics of the host network. In many applications, the network topology already exists or is determined by specification requirements. This section discusses how to integrate PTP into an existing network architecture. It identifies several common network configurations and demonstrates how to add PTP to these configurations. It also provides recommendations for how to increase the robustness of the timing solution. In addition, this section includes configuration examples to provide synchronized time from the same source to multiple buildings or locations.

Basics

IEEE 1588v2-2008 is the core PTP standard [1]. It is a broad standard that defines different profiles for specific applications. The IEEE 1588 Power Profile defines PTP for use in the power system and provides requirements for device accuracy and precision. Other profiles defined in the standard include the International Telecommunication Union (ITU) Telecom Profile and the IEEE Power & Energy Society (PES) Power System Relaying and Control Committee (PSRC) C37.238 Power Profile. For the best performance, IEEE 1588v2 recommends that all active devices in the PTP communications path be PTP-compliant. PTP-compliant devices recognize the PTP messages and add delay information to the packets as they egress the device. IEEE 1588v2 does not define performance requirements for the devices and applications [2].

PTP Power Profile

The IEC and IEEE profile committees eliminated many of the IEEE 1588v2 options when creating their power profiles to mitigate problems caused by mismatched settings. The chosen values optimize PTP operation in the power system. Reducing options also simplifies compliance testing. PTP clocks used in the power system should be certified to the IEEE C37.238 standard.

IEC and IEEE each released standards for PTP power profile requirements. IEEE PSRC initially released IEEE C37.238-2011 [3], and it is the original power profile for PTP. It defined clock accuracy requirements; additional type, length, and value structures; and settings limits. In 2016, IEC released their version of the power profile (IEC/IEEE 61850-9-3 Edition 1) [4]. This version adjusted the profile to include specific requirements deemed critical for use in power system timing by IEC. IEEE updated the C37.238 standard (C37.238-2017) to align with the IEC standard and provide additional requirements [5]. The revised IEEE standard allows slave IEDs designed to meet the IEC standard to acquire time from clocks set to either standard. The IEEE C37.238-2017 profile is backward-compatible with the 2011 version, however, some items that were previously required are now optional.

Network devices (i.e., switches and routers) used to consume and forward PTP timing information should be certified to IEEE C37.238 for master clock, boundary clock, and transparent clock (TC) operation. This ensures that the network can provide the PTP timing accuracy necessary to meet the demanding requirements of power system monitoring.

With the extensions defined in IEEE C37.238-2017, PTP contains the information required to generate other timing signals. PTP clocks, through converters, can generate high-accuracy IRIG-B signaling to support legacy equipment based on the information in the PTP timing signals. Some boundary clocks provide PTP-to-IRIG-B conversion without separate hardware. However, for the configuration diagrams in this paper, assume that the PTP-to-IRIG-B converter is a separate component.

Grandmaster Selection

PTP delivers timing information through multicast messages. These messages describe characteristics of the master (announce messages), send timing information (sync and follow up), determine path delay (delay request – response, peer delay request – response), and manage the system. PTP uses the concept of timing domains to define the group of clocks synchronized to each other. Clocks set to one domain generally ignore messages from other domains. TCs are one exception; they process messages for all domains.

The master and slave clocks in the PTP timing solution use these messages to determine the hierarchy of the timing network. The Best Master Clock Algorithm (BMCA) continuously executes on master, boundary, and slave clocks to determine the grandmaster clock. This algorithm also determines the hierarchy of the remaining PTP clocks and redefines it when it detects changes.

The master clocks, including the grandmaster and potential grandmasters, listen to the announce messages from their PTP domain. When a potential grandmaster determines, through the BMCA application, that it has better time credentials, it challenges the current grandmaster by broadcasting announce and sync packets. The current grandmaster, having lesser credentials, stops broadcasting announce and sync messages. The remaining clocks hear the new messages and switch to that master for synchronizing. This process is fully automatic.

PTP clocks continuously listen for the announce messages from their current master or grandmaster. If the messages are not received for a specific period of time, other potential masters start broadcasting announce and sync messages. The BMCA selects a new master or grandmaster for the timing domain. Masters that are not the active master are in passive mode.

Most IEDs operate in slave-only mode and do not attempt to become a master on the domain. If the master clock stops operating or the slaves cannot hear it, the slaves start listening for another master to take its place. With the power profile, the slaves start looking for a new master after 3 seconds. Assuming a new master is available, the slaves lock on to it after a couple of seconds. The IED slaves operate in a free-run timing mode for about 5 seconds before they synchronize to the new master.

General Best Practices

There is a real benefit in having a second potential grandmaster clock available to the PTP domain. It automatically becomes the grandmaster during an outage or if the primary grandmaster is out of service for maintenance. This maintains the IEDs' timing and synchronization rather than each IED free-running until the grandmaster is back online. When the offline clock is restored to service, the BMCA application determines whether it should be the grandmaster or a potential grandmaster. Normally, if it was the grandmaster before, it resumes the grandmaster position after it is restored to service.

The IEEE and IEC power profiles include clock performance requirements not defined in the IEEE 1588v2 standard. A rogue grandmaster clock that does not meet the performance requirements could be very detrimental to the timing solution.

The standards committees are starting to discuss the concept of timing islands. This concept is a methodology for maintaining device timing over a wide area when primary time sources (e.g., GNSS) are not available. IEDs that compare signals with each other over distances identify when the timing solution is in holdover and further evaluate their time source. If the time sources are different, these IEDs recognize that the distance measurement functions may not be valid and remove them from service. The Multiple Buildings and Locations subsection of this paper discusses a PTP network that allows continued operation because the locations remain synchronized to one master.

Using VLANs

IEEE C37.238 introduces the concept of using VLANs to segregate the PTP timing traffic for the power profile from other traffic on the network. Using VLANs does not affect PTP timing accuracy. It may provide some protection from a rogue PTP master trying to become the grandmaster for the solution.

Slave Connections

IEDs have different accuracy and precision requirements for different applications. SV and TWFL applications require submicrosecond timing accuracy, while SOE recording and other applications require millisecond timing accuracy.

The PTP backbone components (when certified to IEEE C37.238) are very accurate and individually introduce minimal unaccounted error into the PTP timing solution. The PTP message initiates at the grandmaster, traverses each TC in the path, and finally arrives at the slave. Each TC the message traverses is a hop, and each hop introduces a small amount of timing error. Over a number of hops, this error can add up. The power profile requires that there be no more than 15 hops in a PTP solution to maintain 1-microsecond accuracy.

The best practice is to connect IEDs that require higher accuracy to TCs that are within three hops or fewer of the grandmaster clock. IEDs with less-stringent timing requirements can connect to downstream TCs where the hop count is higher.

Minimal Network

The minimal network, as shown in Figure 2, is the basic starting point for most installations and typically includes a few switches. It is not configured for redundancy or protection against a single point of failure on the network backbone.

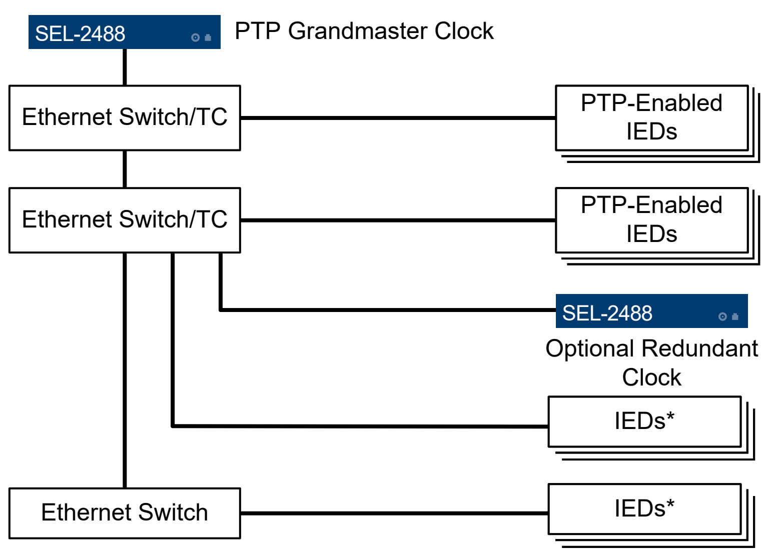

*Timing Signals Provided to IED Through IRIG-B, Pulse Per Second (PPS), or Network Time Protocol (NTP)

Figure 2 Minimal network configuration

In Figure 2, the SEL-2488 Satellite-Synchronized Network Clocks and PTP-Enabled IEDs connect to Ethernet switches configured as TCs. IEDs that are not PTP time-synchronized can connect to the TCs or non-PTP switches. This allows the solution to continue using legacy or installed devices. PTP components can be added as the solution expands or is retrofitted.

As shown in Figure 2, the network does not have redundancy built in. However, SEL recommends installing a second clock with a separate antenna system. The second clock, by connecting to a different TC when available, continues to provide time if the first TC experiences an outage. Installing a second PTP master clock (shown as the Optional Redundant Clock in Figure 2) enhances the robustness of this implementation in multiple ways. First, it provides redundancy if the primary grandmaster or its antenna system fails. Second, the primary grandmaster and its antenna system can be out of service without reconfiguring the slave devices.

Ring Network

The ring network, shown in Figure 3, is popular for limited-redundancy schemes. In the event of a single switch failure, it allows the rest of the network and devices (minus the failed switch and the devices directly connected to it) to continue working.

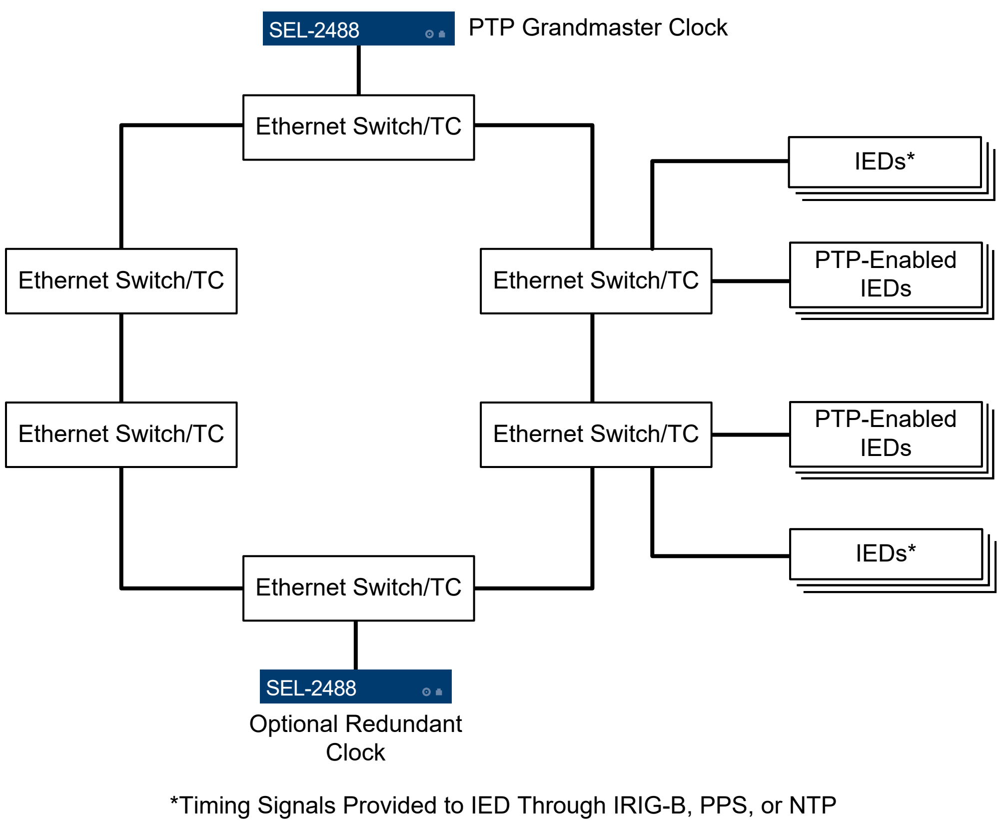

Figure 3 Ring network configuration

To provide accurate PTP timing on a ring network, all switches must be IEEE C37.238-certified and operate as TCs. Any IED can connect to any switch because all switches are TCs. SEL recommends adding a second PTP clock to the network, as shown in Figure 3, preferably to one not connected to the same switch as the primary clock. This maintains the timing solution if the primary clock or the connected switch experiences an outage.

The Rapid Spanning Tree Protocol (RSTP) reconfiguration time may affect the end device timing and selection of the grandmaster clock. A network that is slow to reconfigure has a period of instability while it determines the new configuration. During this period, the IEDs may change their choice of master/grandmaster clock several times (if the reconfiguration takes several seconds).

This switching between multiple grandmasters can cause temporary IED time jitter until the network settles.

Ladder Network

SEL recommends the ladder network for redundant network configuration with the fastest RSTP reconfiguration time and the least impact for any switch outage [6].

Figure 4 demonstrates a typical ladder network with a PTP timing solution. Each "ladder rung" that includes PTP IEDs or clocks requires that both switches be IEEE C37.238-certified TCs. This maintains the PTP timing accuracy, regardless of the packet path from the clock to the PTP IED.

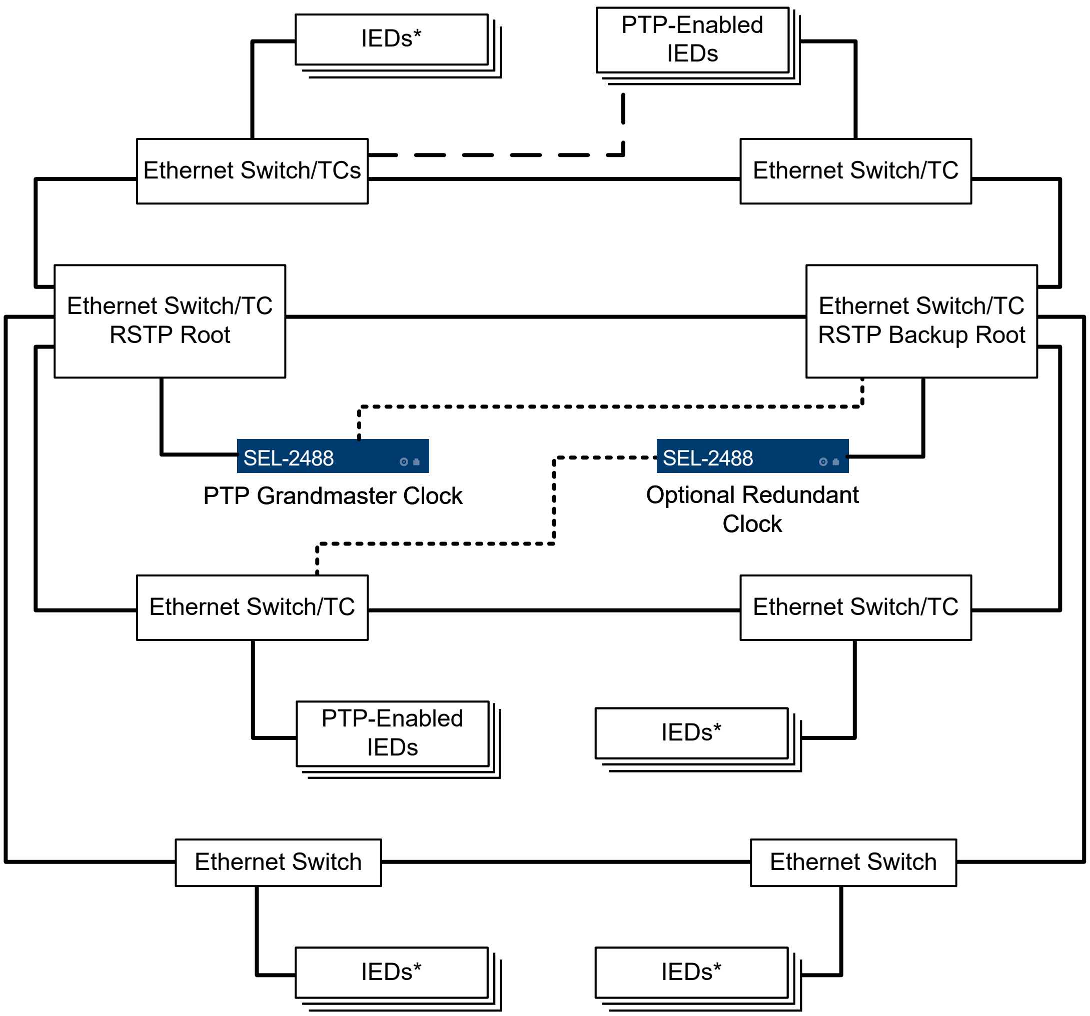

*Timing Signals Provided to IED Through IRIG-B, PPS, or NTP

Figure 4 Ladder network configuration

As shown on the bottom rung of Figure 4, rungs supporting non-PTP equipment do not require TC switches. This allows continued use of existing equipment in the ladder network while maintaining the fastest RSTP reconfiguration and, therefore, the least impact to the PTP timing solution. RSTP reconfiguration time, as discussed in the Ring Network section of this paper, may cause temporary IED time jitter if the reconfiguration takes several seconds.

The SEL-2488 has four Ethernet ports that can function as master/grandmaster PTP clocks. Connecting one port to a TC can provide PTP timing to the network. In addition, this port can provide management access to the clock and NTP time to the network if configured to do so. One PTP power profile requirement is to perform PTP timing messaging at Layer 2 (IEEE 802.3 Ethernet). To leverage the Layer 2 requirement, enable a second PTP Ethernet port in Layer 2 mode in the SEL-2488. This port is connected to an alternate TC, as shown by the dotted lines in Figure 4. This provides additional redundancy for the PTP timing signals.

Many IEDs have the option to failover to an alternate Ethernet port. The alternate port and connection (shown by the dashed line in Figure 4) activate if the primary port experiences a connection failure.

Redundancy becomes a question of criticality of operation versus cost. Primary substations or remote locations may require additional redundancy. The network design determines the redundancy needed to maintain control, automation, and monitoring. As stated, the PTP solution superimposes the network and leverages that design.

PRP Network

Parallel Redundancy Protocol (PRP) networks, such as the one shown in Figure 5, provide two local-area network backbones that are completely independent. IEDs specifically designed for PRP networks use an Ethernet port to connect to each network backbone. A PRP Redundancy Box (RedBox) allows a non-PRP clock or IED to connect to a PRP network. Duplicate messages are simultaneously sent on each backbone and collected by the target device. The first message to arrive at the target is processed, and the duplicate is discarded.

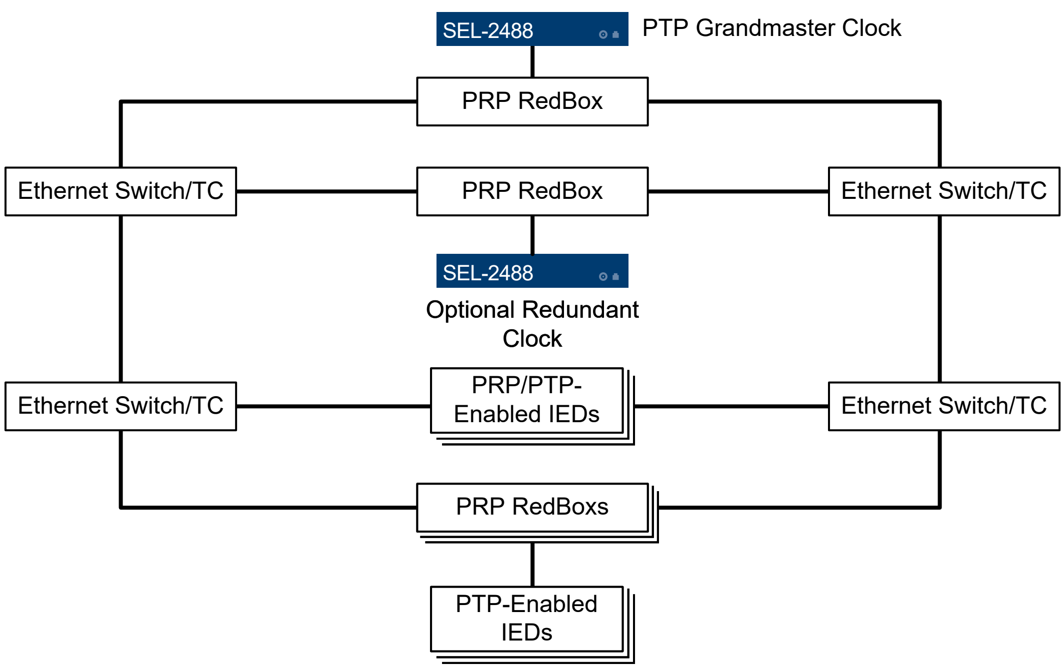

Figure 5 PRP network configuration

Messages are not lost if one backbone experiences an outage. The duplicate message sent on the other backbone is processed as the first arriving message. Non-PRP IEDs connected through the RedBox gain the benefit of the redundant networks for most of the backbone. As shown in Figure 5, the non-PRP clocks also connect through a PRP RedBox.

HSR Network

High-Availability Seamless Redundancy (HSR) is a network strategy in which the IEDs interconnect without switches, as shown in Figure 6. The IEDs forward packets received on one port to the other port until the packets reach the originating IED. HSR networks send duplicate packets in opposite directions around a ring. The first packet received at the destination is used, and the other packet is discarded. Any device directly connected to an HSR network must be compatible with Doubly Attached Node with HSR (DANH). Non-DANH devices must connect through an HSR RedBox.

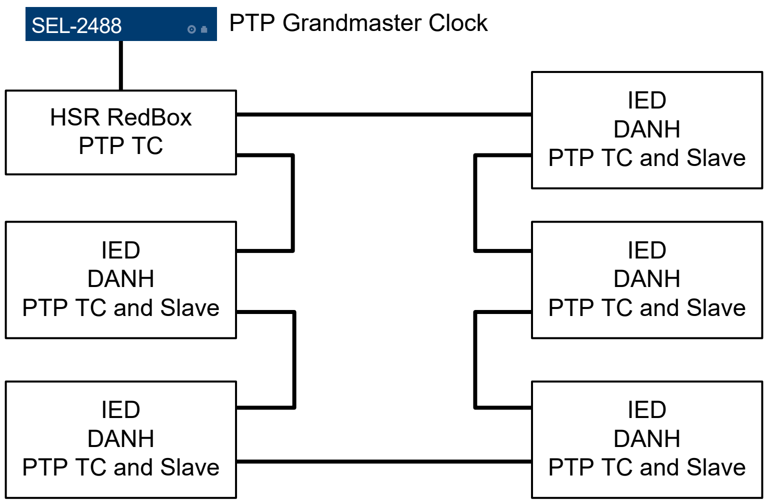

Figure 6 HSR network configuration

For PTP to work on an HSR network, all DANH nodes must be PTP TCs. IEDs synchronizing to PTP also need to be slave clocks. The clock must connect through an HSR RedBox if it is not DANH-compatible. Integrating PTP with an HSR network scheme may be cost-prohibitive because each of the IEDs and HSR RedBox units must support PTP TC functionality.

Multiple Buildings and Locations

Larger installations may have multiple buildings and locations or buildings segmented by purpose, such as transmission control versus distribution control. The challenge with multiple buildings or locations is maintaining timing accuracy over longer distances.

Using a dedicated cabling timing solution is a popular and solid solution when used within one building. However, problems arise when attempting to extend the solution to other buildings or locations, which may exceed cable length restrictions. Additionally, copper cabling solutions are not immune to ground potential energy differences between buildings or electrical interference and can lead to an expensive cabling installation for a timing solution.

Another option for multiple buildings and locations is to install multiple clocks so that each building or location has its own time source. This partially addresses the dedicated cabling issue but adds extra clocks and antenna systems to maintain. Each additional clock becomes a single point of failure for that segment of the system. Separate clocks may also have slight timing variations that can affect substation applications.

Synchrophasors and TWFL are two applications where timing differences between sites can skew measurement results. Multiple locations connected by Ethernet networks present opportunities to use PTP timing while maintaining accuracy and synchronization between sites.

The option presented in Figure 7 implements a PTP timing solution over Ethernet connections (possibly redundant) between buildings and locations. As discussed earlier, PTP automatically compensates for topology changes without user interaction.

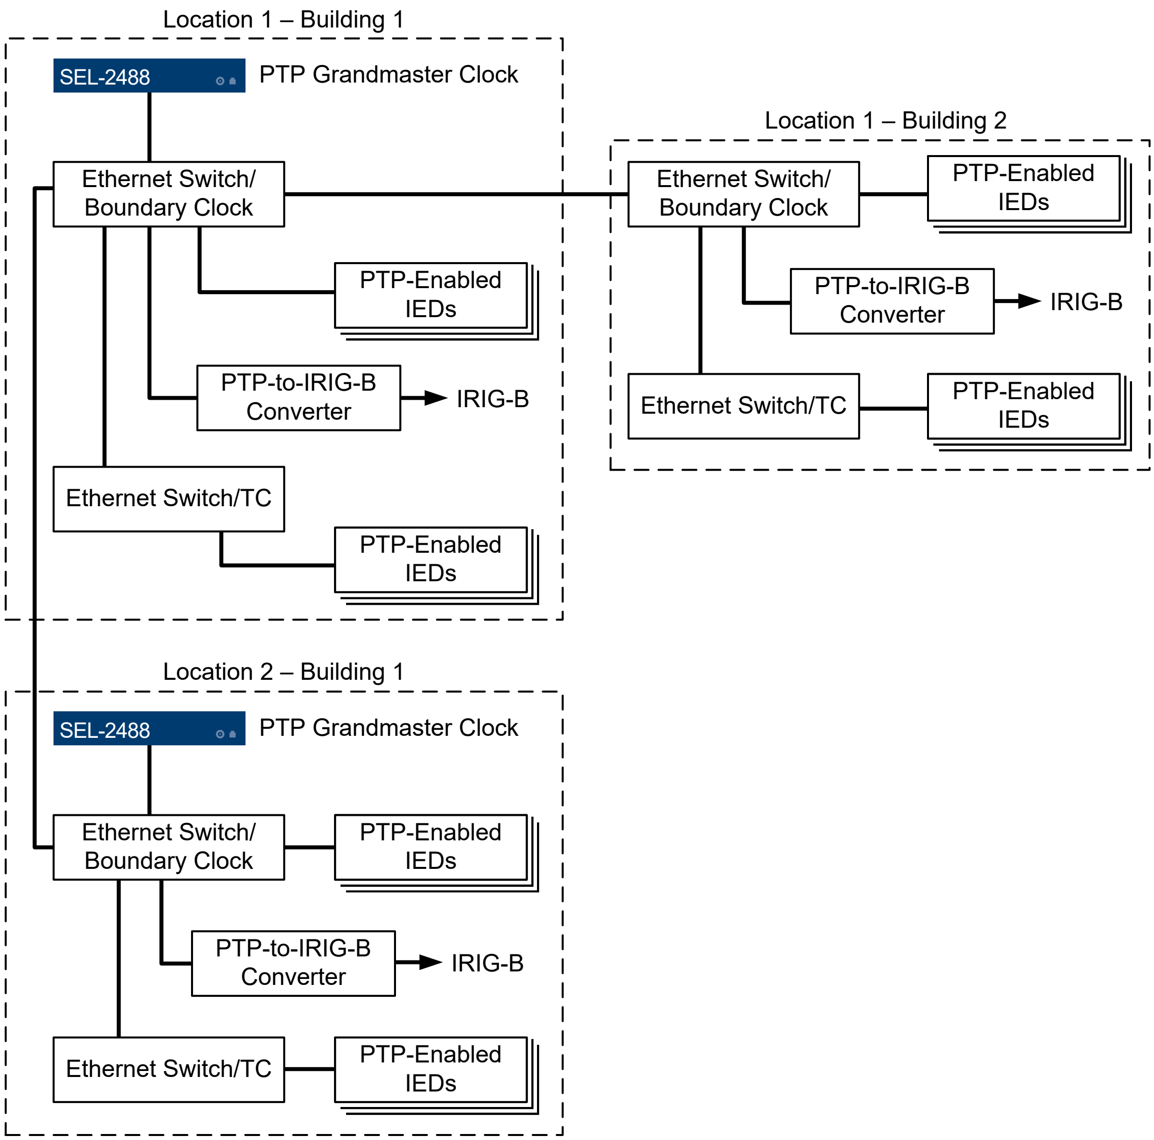

Figure 7 Multiple buildings and locations

In Figure 7, Location 1 – Building 2 does not have a clock, possibly due to physical conditions. It receives time over the network from Location 1 – Building 1. The SEL-2488 PTP grandmaster clocks at Location 1 and Location 2 negotiate and select a grandmaster for the timing domain. If the selected grandmaster becomes unavailable or its time quality changes, the other clock becomes the grandmaster. PTP path delay calculations automatically compensate for the timing delays on the near and remote networks. Any change to the grandmaster or the network is addressed without user interaction to maintain the timing solution accuracy.

Installing boundary clocks (with holdover), as depicted in Figure 7, greatly increases the timing solution robustness. By installing boundary clocks in each building, the solution maintains timing through multiple outage scenarios. If a building or location loses connectivity with the other buildings, the boundary clock in that building maintains the timing. The boundary clock holdover option keeps the IEDs synchronized instead of free-running. This scenario allows for maintenance on the grandmaster with minimal timing effect to the buildings.

Multiple interconnections between the sites for redundancy may significantly enhance the physical network configurations between sites. PTP switches with connections to other sites should be boundary clocks to maintain time in the event of link failure. Otherwise, the automatic network reconfiguration may provide a path that does not maintain the PTP timing solution.

PTP-to-IRIG-B converters provide IRIG-B timing signals to legacy equipment based on the information in the PTP timing signals. There are several benefits to using PTP-to-IRIG-B converters to supply the IRIG-B signals to legacy devices. First, it eliminates the need for dedicated cabling between buildings to support legacy equipment. Second, further reduce in-building cabling by placing the converter close to the equipment it is supporting. Third, legacy devices leverage the benefits of multiple PTP grandmaster-capable clocks the same way PTP devices do. Lastly, because the IRIG-B cables do not directly connect to the grandmaster clock, the grandmaster clock can be out of service and alternate grandmasters will continue to maintain the timing synchronization without cabling changes.

Conclusion

PTP works effectively over most Ethernet architectures and provides mission-critical timing accuracy and redundancy. PTP can reduce or eliminate dedicated cabling for timing solutions. By applying known best practices, PTP provides scalable timing solutions that are suitable for small to large substations and multibuilding or multisite applications.

References

[1] IEEE Standard 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.

[2] J-L. Ferrant, M. Gilson, S. Jobert, M. Mayer, L. Montini, M. Ouellette, S. Rodrigues, and S. Ruffini, *Synchronous Ethernet and IEEE 1588 in Telecoms: Next Generation Synchronization Networks*. Wiley-ISTE Ltd, London, 2013, p. 56.

[3] IEEE Standard C37.238-2011, IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications.

[4] IEC/IEEE International Standard 61850-9-3-2016-05, Communications Networks and Systems for Power Utility Automation Part 9-3: Precision Time Protocol Profile for Power Utility Automation.

[5] IEEE Standard C37.238-2017, IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications.

[6] J. Dearien, “Understanding RSTP and Choosing the Best Network Topology,” SEL Application Guide (AG2017-21), 2017. Available: selinc.com/literature/application-guides/.

Biography

Larry Thoma has over 20 years of experience in emergency service and commercial industries, designing, implementing, and maintaining computing systems and network infrastructures. Larry joined Schweitzer Engineering Laboratories, Inc. (SEL) in 2012 as a test engineer for time and communications products, managing the testing effort for several new products. Larry is currently a product engineer at SEL in the communications group, researching and defining new precise time products.

Making Electric Power Safer,
More Reliable, and More Economical

Schweitzer Engineering Laboratories, Inc.
Tel: +1.509.332.1890 | Email: info@selinc.com | Web: www.selinc.com

LWP0025-01