

White Paper
Profile for Use of Precision Time Protocol in Power Systems



Abstract

This document lists the requirements for accurate time in Power Systems. This need is driven by Synchrophasors and the supervisory control and data acquisition (SCADA) system. The Precision Time Protocol (PTP), defined by IEEE 1588, is described, including the features that allow it to deliver the required time transfer accuracy. The IEEE Power and Energy Society sponsored the creation of IEEE C37.238-2011, a PTP profile for the Power Industry. The profile describes constraints on the use of PTP to make sure that Power Systems meet their timing requirements. This profile is described in this document.

Introduction

Electric utilities globally have initiatives to modernize their grids. Each utility will create their own path to the smart grid based on their business and regulatory drivers. These drivers include wide area blackouts, renewable energy integration visibility and management, and capacity issues. A common thread in modernization initiatives is the need to provide real time visibility of grid conditions with a goal of automatically driving the grid to optimal performance.

The following large scale outages were all due to cascading events where an outage occurred in one part of the system and was followed by a domino effect on the interconnected system by overloading lines and tripping generation off line as the system became unstable.

- The 2003 blackout in the northeast United States and east-central Canada affected an estimated 45 million people in eight U.S. states and 10 million people in Ontario. This outage was caused by a fallen tree branch over a power line.
- The 2003 blackout in Italy affected a total of 56 million people for 12 hours, and also part of Switzerland for 3 hours.
- The 2005 blackout in Indonesia was a cascading power outage across Java and Bali on 18 August 2005, affecting some 100 million people.
- Over five million people were affected in September 2011 when cascading power outages hit parts of southern California, Arizona, and parts of northwestern Mexico. The outage was started by monitoring equipment that was causing problems at a power substation in southwest Arizona.
- The July 2012 India blackout was the largest power outage in history, occurring as two separate events on 30 and 31 July 2012. The outage affected over 620 million people and was spread across 22 states in northern, eastern, and northeast India. An estimated 32 gigawatts of generating capacity was taken offline in the outage

Providing real-time situational awareness to grid operators will decrease the impact of the outages by isolating the problem areas and avoiding the cascading events. In addition, real-time visibility enables the operators to operate islands reducing the numbers of customers that are impacted.

Similarly, real-time visibility of grid conditions enables management of load and generation fluctuations. This will eventually lead to a self-healing grid that will operate more efficiently and reliably with greater percentages of renewable generation sources connected to the grid. Real-time monitoring of capacity constraints combined with load management tools provides an opportunity to optimize the grid based on actual conditions. Accurate, reliable time synchronization of this data is essential to a more intelligent grid.

Accurate Time in Power Systems

Power systems need equipment upgrades to enable new functionality and reduce the probability of blackouts. Such recommendations include the need for the dissemination of accurate system time.

Time is needed for several power system functions as summarized in the following table, and include synchrophasor measurements, protective line measurements, analog measurements, and time stamping in the SCADA System.

Table 1 • Timing Requirements in Power Systems

Application	Standards	Required Accuracy	Relative Time	Absolute Time
Synchrophasors	IEEE C37.118.1/2 -2011	1 μ s		X
Line differential protection		1 μ s	X	
Sampled analog values	IEC 61850-7-2 (service models)	1 μ s	X	
	IEC 61850-9-2 (mapping to Layer 2 /Ethernet)			
SCADA system		1 ms		X

Synchrophasors, or Phasor Measurement Units (PMU) are instruments that measure the magnitude and phase angle of line voltage and current at multiple locations in the power grid at the same time. Since the measurements are time tagged to a common time reference (Coordinated Universal Time (UTC)) measurements at multiple grid points can be compared. Such measurements provide wide-area situational awareness for operators to detect system instabilities, direct load balancing algorithms, and provide early fault detection and respective correction. The timing accuracy required for synchrophasors is 1 ms—dictated by the need to have the time error be a small fraction of the 60 Hz or 50 Hz period, 16.66 ms or 20 ms [1].

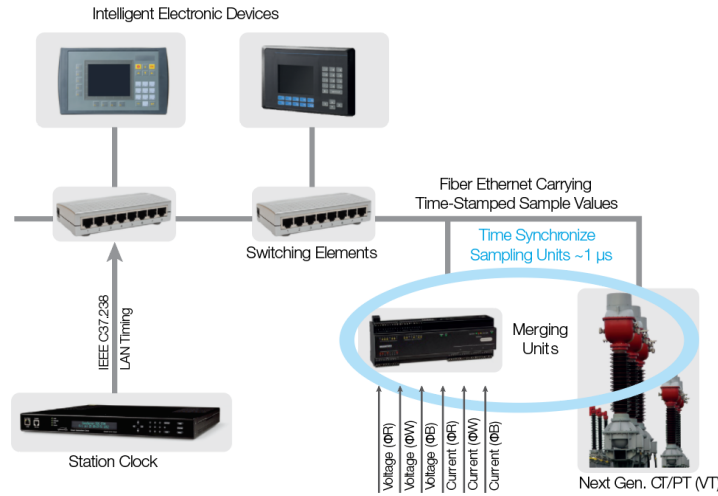
The SCADA system requires that IED events are logged with 1 ms accuracy [2]. Such accuracy is usually not achieved using Network Time Protocol (NTP). However, if IRIG-B or PTP is used for timing distribution, the higher precision of these time transfer techniques are more likely to guarantee that the IED timestamps meet the required accuracy.

The Precision Time Protocol

The Precision Time Protocol (PTP) is an IP/Ethernet based protocol for distributing time in a network from a Master Clock to one or more Slave Clocks. PTP is defined by the IEEE 1588 standard [3]. It is similar to the NTP and Simple Network Time Protocol (SNTP) [4] in that it works by having devices on a network exchange packets to transfer time.

The following illustration shows devices in a power substation.

Figure 1 • Devices in a Power Substation



The following table compares PTP with NTP and IRIG-B.

Table 2 • Time Transfer Technologies

	IRIG-B	(S)NTP	PTP
Accuracy (typical)	1 ms–10 ms	1 ms–10 ms	100 ns–1 ms
Transport media	Dedicated cables	Ethernet cables	Ethernet cables
Protocol style	Master-slave	Client-server	Master-slave
Built-in latency correction	No	Yes	Yes
Set-up	Configured	Configured	Self-organizing, or configured
Update intervals	1 second	Minutes	10 ms–1 s
Specialized hardware	Required	No	Required

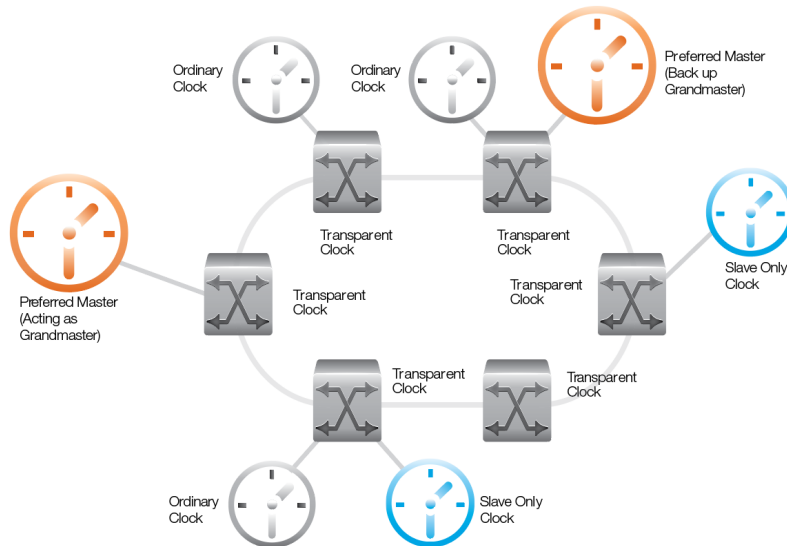
PTP has several important properties that NTP does not have:

- PTP works with a master-slave paradigm, rather than client-server. In other words, time transfer is initiated by the device that has the system clock. This allows some slave devices to have a very simple implementation.
- PTP is a generic protocol that can run over many network types including IP/Ethernet, Ethernet, ProfiNet, and DeviceNet. More network mappings are expected in the future.
- PTP does not specify the slave servoloop behavior, so slave behavior can vary widely from one implementation to another.
- PTP allows a wide range in update rates to accommodate different applications.
- PTP includes the option of hardware time stamping of message arrival and departure rates for greater accuracy.
- PTP defines methods for switches and routers to assist in the delivery of precise time.

Clock Types

The following illustration shows an example PTP network.

Figure 2 • Network of PTP Clocks



The nodes in the network are labeled by the role they play in PTP. The PTP device types are:

- **Grandmaster:** the source of time in the network. The grandmaster is chosen by the Best Master Clock Algorithm. There can be only one grandmaster in a PTP domain.
- **Preferred master:** designed to act as the grandmaster clock of the network. It typically includes a GPS receiver and high quality oscillator such as an OCXO or rubidium atomic oscillator. A high-quality oscillator gives the clock a stable output free from excessive jitter and wander, as well as a long holdover in case the GPS receiver fails or the GPS signal is denied.
- **Slave-only clock:** acts as slave and never as a master.
- **Ordinary clock:** capable of acting as either a master or a slave. In most network implementations, the ordinary clocks remain in the slave state and only become masters when all of the preferred masters in the network are unavailable.
- **Transparent clocks:** network switches that are capable of updating the PTP messages to correct for the time messages spend in the switch. This removes the queuing delay variation that limits the accuracy of NTP.

Note: A real device could combine more than one PTP role in the same device. For example, an Ethernet switch could be both a transparent clock and an ordinary clock.


There is another important type of PTP clock, called a boundary clock, which is not shown in the previous illustration. This is another device that can be an integrated function in an Ethernet switch or a router, or neither. Instead of correcting PTP messages for the time they spend in the switch, a boundary clock terminates the messages and sends new ones. Specifically, a boundary clock has one port that is a slave to an upstream master clock, and all other ports act as master ports for downstream slaves. This type of device is not commonly used with the Power Profile, but is used in other PTP applications.

PTP allows system administrators to define multiple timing domains in the same physical network. These are indicated by a numerical value for the domain in the common message header that all PTP messages carry. Devices in PTP domains use the same multicast addresses, and therefore see the messages from all domains but use only messages from the domain they are configured to operate in. For example, one might test time synchronization in new equipment in a control network that already has PTP running. This could be accomplished by creating a test domain. The test domain would function independently from the control domain.

PTP Messages

PTP message exchanges can be grouped into two categories. The following illustration shows a sequence diagram for the message exchanges among PTP devices.

Figure 3 • PTP Message Sequence



The grandmaster sends three types of messages to the slave, listed as follows. These messages are sent in multicast, so any device on the network that is configured to subscribe to the PTP multicast addresses will see the messages.

- **Sync message** contains the approximate time from the grandmaster.
- **Follow-up message** contains a more precise value of when the sync message left the grandmaster.
- **Announce message** contains the properties of a grandmaster capable clock that is used in selecting the best grandmaster for the network.

To convert the grandmaster time stamp into a slave clock correction, one still needs to account for the time it took for that time stamp to get to the slave. This is accomplished by the exchange of peer delay messages between adjacent devices on the network. More specifically, a device sends a peer delay request message to another device connected directly to it through a cable. In other words, it is a peer on the network. The peer sends a peer delay response and a peer delay follow-up message. At this point, the device that initiated the exchange has the following four time stamps.

- T1 = Time the peer delay request departed from the initiating device
- T2 = Time the peer Delay request arrived at the peer device
- T3 = Time the peer delay response departed the peer device
- T4 = Time the peer delay response arrived at the initiating device

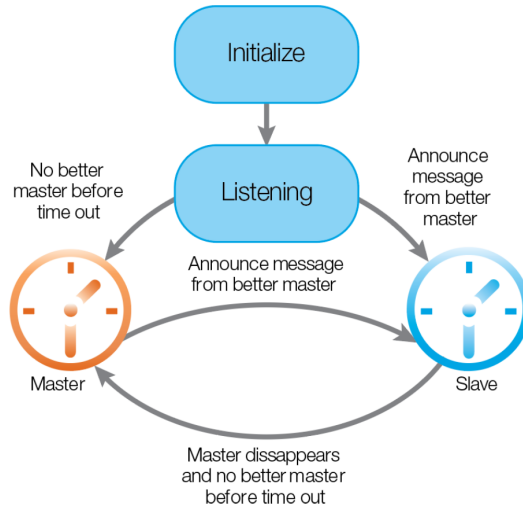
The initiating device can now calculate the cable delay between the two devices as $(T4 - T3 + T2 - T1) / 2$. Four time stamps are needed so that any difference in the clocks between the two devices does not influence the calculation of the cable delay between the devices. This is the approach used in NTP between the client and server, except in this case the delay is only a cable delay, the queuing delays are handled by other means.

When a transparent clock corrects for the dwell time of a sync message, it also subtracts for the upstream cable delay and updates the correction field in the follow-up message. This way, the delays of each cable are also corrected by the transparent clocks. The last cable delay is corrected by the slave. Actually, the last correction by the slave is optional. This is to enable the possibility of keeping the slaves as simple as possible. System designers should be aware of which slaves do not make the cable delay correction and make sure that the last cable delay fits within the time error budget. Alternatively, the slave may have a mechanism to enter the last cable delay manually. For example, RG-58 coaxial cables have a delay of 5.08 ns/m [5].

The Best Master Clock Algorithm

The Best Master Clock Algorithm (BMCA) ensures that there is only one grandmaster in a network. To understand how the BMCA works, consider the behavior of an ordinary clock powering up on the network. A state diagram of an ordinary clock is shown in the following illustration.

Figure 4 • Ordinary Clock States



The first thing an ordinary clock does is listen, which means it waits to see if it is receiving announce messages on the PTP multicast address. If not, then it assumes that the network has no grandmaster, and so it takes on that role. If the ordinary clock does see an announce message from another device, it examines the messages to see if it came from a better clock. If it came from a better clock, then the ordinary clock becomes a slave. If the received announce message is not from a better clock, then the ordinary clock sends its own announce message and begins sending both sync and follow-up messages. The device that earlier sent the announce message will receive the ordinary clock's announce, see that it is better, and cease functioning as a grandmaster.

So how does a PTP clock know if it is better than another clock or not? The answer is that it compared the fields in a clock's announce message with its own master clock properties. These properties are listed in the following table in order of priority.

Table 3 • Master Clock Properties

Master Clock Property	Announce Message Representation	Notes
Priority 1 field	0-255 (decimal)	User-configured. Purpose is to bypass the BMCA.
Clock class	0-255 (decimal)	Lower values means better master clock 6 = Locked primary reference clock 7 = PRC unlocked, but still in spec 13 = Locked to app-specific timescale 14 = Unlocked from app-specific time, but still in spec 52, 187 = PRC, unlocked and out of spec 58, 193 = App-specific, unlocked, out of spec 248 = Default, if nothing else applies 255 = Slave-only clock Other values reserved for future use or for profiles

Master Clock Property	Announce Message Representation	Notes
Clock accuracy	0-FF (Hex)	20 = The time is accurate to within 25 ns 21 = The time is accurate to within 100 ns 22 = The time is accurate to within 250 ns 23 = The time is accurate to within 1 μ s 24 = The time is accurate to within 2.5 μ s 25 = The time is accurate to within 10 μ s 26 = The time is accurate to within 25 μ s 27 = The time is accurate to within 100 μ s 28 = The time is accurate to within 250 μ s 29 = The time is accurate to within 1 ms 2A = The time is accurate to within 2.5 ms 2B = The time is accurate to within 10 ms 2C = The time is accurate to within 25 ms 2D = The time is accurate to within 100 ms 2E = The time is accurate to within 250 ms 2F = The time is accurate to within 1 s 30 = The time is accurate to within 10 s 31 = The time is accurate to within >10 s FE = Unknown accuracy Other values reserved for future use or for profiles
Clock variance (frequency stability)	0-FFFF (Hex)	Lower value indicates a better frequency stability Offset log (base 2) scaled variable Based on Allan Variance over Sync Interval
Priority 2 field	0-255 (decimal)	User-configured. Purpose is to select priority among equivalent clocks.
Clock identity	64 bits	Must be a unique number. Often set to a ports Ethernet MAC address.

System administrators can influence the results of the BMCA by configuring the Priority 1 and Priority 2 fields. Otherwise, the BMCA is determined by the self-reported properties of each clock. Also, the final tie-breaker is the Clock ID number, which is required to be unique for each PTP node. In Ethernet-based networks, the MAC address of a port is often used by manufacturers to select Clock IDs because there is already a standard mechanism for assigning unique MAC addresses to Ethernet ports [6].

One-Step and Two-Step Clocks

Grandmasters and transparent clocks can be either one-step or two-step clocks. The behavior of two-step clocks has been described so far. One-step grandmasters do not send follow-up messages. Instead, the sync message carries a precise time stamp, and a one-step clock flag is set. One-step transparent clocks update the correction field of each sync message, while two-step transparent clocks update the correction field of each follow-up message. Slaves are required to add the correction fields of both the sync message and follow-up message, if there is one. Similarly, one-step clocks do not have a peer delay follow-up message.

In addition, the IEEE 1588 standard defines rules for the case when one-step and two-step clocks are in the same network. Therefore, correct PTP implementations will work together even if one-step and two-step devices are mixed.

PTP Profiles

There are many other optional features allowed in PTP, including an alternative mechanism for determining cable delays. In fact, there are so many options in PTP that there is no guarantee that two PTP devices can work together. The rich list of feature makes PTP useful for a large number of different applications in different types of networks. However, this requires the use of profiles to ensure interoperability. A profile is a set of rules for constraining PTP for a specific application or industry. The IEEE 1588 defines a Standard Profile, often called the Default Profile. Other PTP profiles have been defined by numerous standards bodies. These include the Telecom Profile defined by the ITU [7] and the Test and Measurement Profile defined by the LXI [8]. It also includes the Power Profile defined by the IEEE. Basically, a profile consists of some rules from among the following categories:

- Required PTP features
- Allowed, but not required PTP features
- Forbidden PTP features
- Required non-PTP functionality, which is related to timing
- Performance requirements

The Power Profile

The Power Profile is defined in IEEE C37.238- 2011 [9]. The standards work was sponsored by the Power Systems Relaying Committee and the Substation Committee of the IEEE Power and Energy Society. The profile defined a subset of PTP intended to run over Layer 2 networks (that is, Ethernet but no Internet Protocol). This does not preclude the Power Profile from running over specialized robust Layer 2 network types designed for high availability, such as the High-availability Seamless Redundancy (HSR) protocol or the Parallel Redundancy Protocol (PRP). Indeed, Power Profile implementations have been tested at a multi-vendor plugfest running on a HSR network [10].

The philosophy of the C37.238 standards body was to minimize the number of optional features to ensure interoperability and predictable performance. This is evident in fixing the message rates, as shown in the following table. The only option is that either two-step or one-step clocks are allowed.

Table 4 • Power Profile Message Rates

Message	Message Interval or Trigger
Announce	1 second
Sync	1 second
Follow-up (two-step clocks only)	Triggered by sync message
Peer delay request	1 second
Peer delay response	Triggered by peer delay request
Peer delay response follow-up (two-step clocks only)	Triggered by peer delay response

Message Rates

The announce message time out is also specified. This is the number of missed announce messages before a grandmaster capable clock starts sending its own announce messages. The Power Profile fixes these at two announce Intervals for preferred masters and three announce intervals for ordinary clocks. The shorter time out for the preferred masters is to reduce traffic and master change over time because a preferred master would usually be selected by the BMCA over an ordinary clock.

Transparent Clocks

The Power Profile mandates that all switches in a network are transparent clocks. The transparent clocks must also be capable of peer-to-peer timing, using peer delay request and peer delay response messages. Transparent clocks are a special capability in Ethernet switches specifically designed to support this function. Practically, this limits the Power Profile to specially designed local area networks, or at least enterprise networks designed with PTP in mind. However, a new power substation, or a substation undergoing an equipment upgrade, can take advantage of the PTP Power Profile.

Power Profile messages are also required to carry IEEE 802.1Q VLAN tags [9]. Both the priority field and VLAN ID must be configurable. The default values are a priority of 4 and an ID of 0.

Announce Message Extensions

IEEE 1588 defines a general method for extending PTP messages. The method uses the type length value approach common in TCP/IP network protocols. The Power Profile defines a number of mandatory TLVs for the announce messages. These are listed in the following table. The TLVs include some extra identification fields and also fields for grandmaster and network inaccuracy.

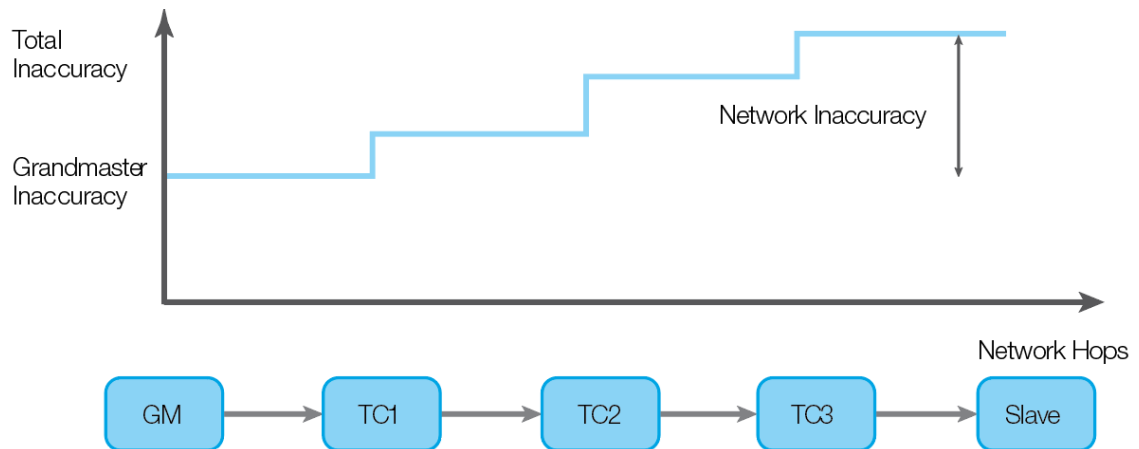
Table 5 • Announce TLV Fields

Field	Value (Hex)	Notes
Organization extension	0003	Denotes a vendor/standards body extension
Length	000C	Number of octets in TLV
Organization ID	1C129D	Denotes Protective Relay Society of the IEEE
Organization sub ID	000001	Denotes IEEE C37.238 standard (Power Profile)
Grandmaster ID	0003-00FE	
Grandmaster time inaccuracy	00000000-FFFFFFFF	Nanoseconds
Network time inaccuracy (ns)	00000000-FFFFFFFF	Nanoseconds
Reserved	2 Octets	

Inaccuracy Fields

Power systems are mission critical systems. Power system failures can lead to serious economic loss and endanger public health and safety. For this reason, much attention in power system engineering is paid to monitoring systems. To support monitoring, the Power Profile Announce Message TLVs include two fields called Grandmaster Inaccuracy and Network Inaccuracy. These values refer to the worst case time inaccuracy for PTP time stamps coming from the grandmaster and contributed by transparent clock corrects, respectively. The concept is shown in the following illustration.

Figure 5 • Inaccuracy Accumulation in a Power Profile Network



Both fields can be configured at the grandmaster or transparent clocks can add their own Network Inaccuracy. Power Profile slaves can examine this field to check if the time they are receiving is sufficient for their applications timing requirements.

Note: Synchrophasors are required to report their estimated accuracy, including timing accuracy.

The Power Profile defines the following performance requirements against which the Inaccuracy Fields can be checked:

- Grandmaster error <200 ns
- Each transparent clock error <50 ns
- Total network inaccuracy from grandmaster to slave <1 μ s

System designers should be aware that the grandmaster by default sets the Network Inaccuracy field to zero, to allow transparent clocks to add their own Network Inaccuracy to this field. However, transparent clocks are not required to have this capability. It is up to the user to make sure that the inaccuracy contributions of transparent clocks without this capability are added at the grandmaster by configuration.

Note: If the slave introduces 200 ns of time uncertainty, then it should be possible to deliver 1 μ s accuracy from the slave over 16 transparent clocks.

Management Information Base

IEEE C37.238 defines an SNMP MIB for Power Profile devices. While the entire MIB will not be reproduced here, it is included with the standard. The MIB is a full read/write MIB and includes several event notifications, sometimes referred to as traps. The event notifications are listed in the following table. Implementation of the MIB is optional for Power Profile devices because not all devices will be sophisticated enough to include an SNMP agent.

Table 6 • Power Profile MIB Event Notification Variables

Notification Event	Description
Change of master	A new grandmaster has been selected.
Master step change	A step change occurred in the grandmaster time.
Faulty state	The clock has entered a faulty state.
Port state change	The port state has changed.
Offset exceeds limits	A clock in the slave state has exceeded its offset limits.
Other profile detected	A non-Power Profile PTP messages has been received.
Leap second announced	A leap second has been announced.
PTP started	PTP services has started.
PTP stopped	PTP service has stopped.

Timescales

PTP time stamps by default use the International Atomic Time scale (TAI). The PTP messages also include an offset field so that TAI can be converted to Coordinated Universal Time (UTC). TAI was chosen for the time stamp values because the timescale does not include leap seconds or other time discontinuities. The Power Profile designers recognize that some network implementers prefer to use a local timescale rather than UTC. For this reason, the Power Profile allows a grandmaster to add an announce message TLV with a local time offset. This is a TLV defined in IEEE 1588. The Power Profile mandates that the PTP default TAI timescale is used for PTP time stamps and any correction to another timescale is accomplished by adding the UTC and local offset field if there is one.

Forbidden PTP Features

A number of features are allowed in PTP, but not in the Power Profile. The following are some of the forbidden features.

- Use of non-TAI time stamps
- End-to-end path delay determination mechanism
- Unicast operation
- Other message rates

Conclusion

Power systems require precise time for use in synchrophasor measurements and the SCADA system. The most stringent requirement is 1 μ s, due to the needs of synchrophasors and sampled values. This accuracy is achievable using the PTP defined by IEEE 1588. The Power Profile was defined to tailor PTP to the needs of the power systems. The following are main features of the Power Profile.

- Message rates of once per second
- All switches must be transparent clocks
- All Power Profile messages must have VLAN tags
- The announce message includes extensions for the grandmaster ID and inaccuracy fields
- The peer-to-peer delay mechanism is used
- An optional Power Profile MIB was defined in the standard

The Power Profile is designed to deliver time better than 1 μ s to the slaves through 16 switches.

References

1. IEEE C37.118.1-2011 Standard for Synchrophasor Measurements for Power Systems, December 2011.
2. IEC 61850-9-2, Ed. 2.0, communication networks and systems in substations – Part 9-2: Specific Communication Service Mapping (SCSM) Sampled values over ISO/IEC 8802-3.
3. IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, August 2008.
4. IETF RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification.
5. <http://ve2azx.net/technical/CoaxialCableDelay.pdf>
6. <http://standards.ieee.org/develop/regauth/tut/index.html>
7. ITU-T G.8265.1/Y.1365.1, Precision time protocol telecom profile for frequency synchronization.
8. LXI IEEE 1588 Profile, Version 1.0, 2008.
9. IEEE C37.238-2011 Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications, July 2011.
10. International Symposium on Precise Clock Synchronization for Measurement, Control and Communication, San Francisco, Sept. 2012.
11. IEEE Std. 802.1Q-2011, Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks.

**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

© 2013–2017 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-01006-1.00-0917