

White Paper
**Synchronization for Next Generation Networks—The PTP
Telecom Profile**



Abstract

This paper is designed to help network engineers, network planners, and network operations understand how to deploy IEEE 1588 Precision Time Protocol (PTP). PTP is a next-generation, packet-based timing protocol targeted for use in asynchronous network infrastructures based on packet transport technologies.

IEEE 1588-2008 introduces the concept of profiles to specify particular combinations of options and attribute values to support a given application. This document describes the Telecom Profile, recently released by the International Telecommunications Union—Telecommunications Standards Bureau (ITU-T) as recommendation G.8265.1. This profile addresses the application of the PTP to the frequency synchronization of GSM, UMTS, and LTE-FDD base stations.

Synchronization for Next Generation Networks—The PTP Telecom Profile

This paper is one of a series of white papers and application guidelines produced as part of an overall framework for synchronization and timing in the next generation network (NGN). These papers are intended to help service provider network architects, planners, and engineers design and deploy stable, robust synchronization and timing architectures to support applications and services that will be deployed on the NGN.

The PTP defined in IEEE 1588-2008 introduces the concept of profiles to specify particular combinations of options and attribute values to support a given application. The ITU-T recently released the Telecom Profile (recommendation G.8265.1), designed to address the application of PTP to the frequency synchronization of telecommunication systems, primarily cellular base stations. This document describes the Telecom Profile, covering both the main features of the profile and the detailed operation of its more complex elements.

The Telecom Profile

The PTP defined in IEEE 1588-2008 is a complex protocol designed to be used in a number of different applications and environments. Some parts of the protocol are aimed specifically at certain applications and are not applicable to others, making it difficult to understand which sections of the protocol should be used in a particular environment.

Therefore, IEEE 1588-2008 introduces the concept of a PTP profile. The idea of a profile is to specify particular combinations of options and attribute values to support a given application, with the goal of inter-operability between equipment designed for that application. The purpose of the profile is outlined in IEEE 1588-2008, clause 19.3.1.1 [1]: “The purpose of a PTP profile is to allow organizations to specify specific selections of attribute values and optional features of PTP that, when using the same transport protocol, inter-work and achieve a performance that meets the requirements of a particular application. A PTP profile is a set of required options, prohibited options, and the ranges and defaults of configurable attributes.”

ITU-T Recommendation G.8265.1 [2] defines a profile, colloquially known as the Telecom Profile, aimed at distribution of accurate frequency over packet networks. This is primarily intended for use with synchronization of cellular base stations, where the main requirement is to operate the radio interface at a frequency accuracy of within 50 parts per billion.

Aims of the Telecom Profile

The Telecom Profile was created to ensure that PTP can be used in accordance with existing telecom sync practices. The following were key aims of the profile:

- Permit operation over existing managed, wide-area, packet-based telecoms networks.
- Define message rates and parameter values consistent with frequency distribution to the required performance for telecom applications.
- Allow interoperability with existing synchronization networks, such as SyncE and Synchronous Digital Hierarchy (SDH).
- Allow the synchronization network to be designed and configured in a fixed arrangement.
- Enable protection schemes to be constructed in accordance with standard telecom network practices.

These aims led to some important decisions about how the profile works:

- It was decided not to use on-path support, such as boundary clocks and transparent clocks, because this is not available in existing networks.
- IPv4 was adopted as the network layer due to its ubiquity, rather than operation directly over Ethernet or other lower-layer protocols
- Unicast transmission was adopted over multicast because it could be guaranteed to work over wide-area telecoms networks.
- The default best master clock algorithm (BMCA) described in IEEE 1588-2008 was replaced by static provisioning. This allows the synchronization flow to be planned, rather than dynamically adjusting itself.
- The clockClass indication was adapted to carry the quality level (QL) indications defined in G.781 [3], for continuity with SDH and SyncE synchronization status messaging.

Profile Features

The following sections describe features of the profile.

One-Way and Two-Way Operation

Frequency can be delivered by sending messages in just one direction (for example, from master to slave). Time distribution requires messages in both directions, to allow compensation for the propagation delay of the messages through the network. Some frequency recovery algorithms in the slave make use of both directions in order to deliver a more accurate and stable result. Therefore, the Telecom Profile permits both one-way and two-way operation.

A PTP master compliant with the profile must be capable of supporting both one-way and two-way operation (that is, PTP Sync and Delay Request/Response messages). A slave uses either one-way or two-way operation, but is not required to support both methods.

One-Step and Two-Step Clocks

PTP defines two types of clock behaviors: the one-step clock and the two-step clock. In a one-step clock, the precise time stamp is transported directly in the Sync message. In a two-step clock, a Follow_Up message is used to carry the precise time stamp of the corresponding Sync message. Follow_Up messages were invented to facilitate timestamping at the hardware level, improving the accuracy of the time stamp. Their use means that the master does not have to modify the time stamp in the Sync message on the fly as the packet is being transmitted, but can send it in a separate, non-time-critical packet later.

If the master is capable of implementing a one-step clock, this significantly reduces the number of PTP messages it has to send. However, some security mechanisms or architectural features in the master might require the two-clock approach.

Therefore, both one-step and two-step clocks are allowed in the profile. A PTP master compliant with the profile may use either a one-step clock or a two-step clock or both. A slave must be capable of accepting both one-step and two-step clocks.

Unicast Transmission

PTP allows the use of both unicast and multicast modes for the transmission of the PTP messages; however, the Telecom Profile is restricted to unicast only. The main reasons for this are that multicast is more complex to provision and upstream multicast is often not permitted for security reasons. The use of a hybrid model, with downstream multicast and upstream unicast, was considered but left for future definition. The slave must request unicast service from the master for each of the message types required (that is, Announce, Sync, and Delay_Request/Response). The standard unicast request mechanism is used (defined in clause 16 of IEEE 1588-2008), with the address of the master taken from a pre-configured list held by the slave.

Network Layer Protocol

The profile adopts transmission of PTP over user datagram protocol (UDP), running on the IPv4 network layer, as defined in Annex D of IEEE 1588-2008. This is because IPv4 is universally supported, while other possible transport layers such as Ethernet or multi-protocol label switching (MPLS) may not always be present, or exist for the full span of the master-slave transmission path.

Transmission Rates

The transmission rates were chosen to allow stable synchronization over most wide-area networks without the need for on-path support, such as boundary or transparent clocks. The rates are therefore higher than in most PTP applications, where the network reach is smaller and on-path support may be present. The exact rate to be used may vary with the type and size of the network, and on the equipment chosen.

Sync messages (if used, Follow_Up messages will have the same rate)

- Minimum rate: 1 packet every 16 seconds
- Maximum rate: 128 packets-per-second

Delay_Request/Delay_Response messages

- Minimum rate: 1 packet every 16 seconds
- Maximum rate: 128 packets-per-second

Announce messages

- Minimum rate: 1 packet every 16 seconds
- Maximum rate: 8 packets-per-second

Signaling messages

- Rate will be governed largely by the duration of unicast service requests.

Domains

According to PTP, a domain is a group of clocks synchronized together using the protocol, acting independently from any other PTP clocks. In the context of unicast messaging without on-path support, each communication path from a grandmaster to a slave therefore represents a domain, because the protocol operates independently over this path.

Therefore, the domain number can be chosen independently for each grandmaster-slave pair, and could in fact be the same for every path although each forms a unique PTP domain. The recommendation is to choose a number in the range 4 to 23. Any differentiation between domain numbers on separate paths is down to operator preference. One consequence of this is that the standard BMCA cannot be used, as each domain contains only a single grandmaster and single slave. Therefore, each element is statically configured as either a grandmaster or a slave.

Synchronization Traceability

In order to provide traceability in the same way for existing SDH and SyncE-based synchronization, the QL indicators defined in G.781 must be carried over the PTP path. These indicators are carried in the synchronization status messaging (SSM) field of Synchronous Optical NETWORK (SONET)/SDH systems and the Ethernet synchronization message channel (ESMC) messages of SyncE.

Similar information is carried in the clockClass field of the PTP Announce messages, and various ranges of clockClass values have been designated for use by alternate profiles. Therefore, the profile encodes the QL indicators in the clockClass values 80 to 110.

Redundancy and Protection

The profile is intended to work with planned redundancy schemes, similar to how existing synchronization systems work. The slave is pre-programmed with a list of grandmasters, each with an associated priority, and selects one of the grandmasters using a procedure based on the QL-enabled process defined in G.781. The slave selects the grandmaster on the basis of QL value first, then priority, provided that the protocol doesn't indicate any kind of signal failure.

Detailed Operation

The operation of the Telecom Profile is different in some aspects to the normal behavior of PTP systems. This section describes in more detail the operation of some of the key aspects of the profile.

Unicast Service Negotiation

Since the Telecom profile is based on unicast messaging, unicast service negotiation is enabled by default in profile-compliant devices. The slave initiates the communication to the grandmaster by making a request for service using the process described in clause 16 of G.8265.1. This consists of sending signaling messages containing a REQUEST_UNICAST_TRANSMISSION type length value (TLV) to the IP address of the grandmaster. If the master has sufficient capacity to handle the slave's request, it responds with a signaling message containing a GRANT_UNICAST_TRANSMISSION TLV.

The REQUEST_UNICAST_TRANSMISSION TLV contains several parameters:

- **messageType**—The type of service being requested (Announce, Sync, or Delay_Response).
- **durationField**—The duration of the requested service. This has a default initialization value of 300 seconds and a configurable range of 60 seconds to 1000 seconds.
- **logInterMessagePeriod**—The transmission rate of the requested messages.

When initiating unicast negotiation with a grandmaster, the slave should use all 1's as the initial value for the targetPortIdentity of the Signaling message. The correct value should be learnt from the grandmaster's response, and used in subsequent signaling messages.

The GRANT_UNICAST_TRANSMISSION TLV contains the same list of parameters as the request. The values of these parameters should be the same as those in the request. If the master denies the request (because it has no remaining capacity, for example), this is indicated by setting the durationField value to zero. The R flag (Renewal Invited) is not used in the Telecom Profile.

Denial of Unicast Service

While it is permitted in IEEE 1588-2008 for the master to reduce either the duration or transmission rate of the granted service, the Telecom Profile states that the master should deny the request altogether rather than offering reduced service to the slave. This prevents the slave being granted a rate that is too low for it to meet the performance constraints. If the slave is capable of operating with a reduced service (for example, at a lower transmission rate), it may repeat the request with the reduced rate.

In the event of being denied service by a master, or receiving no response to the service request, a slave should wait a minimum of one second before issuing a new unicast service request for that message type to the same master. This prevents the master from being saturated with repeated requests that it cannot service. If a slave has issued three service requests for the same message type with either no response or a grant denied response, it should wait another 60 seconds before retrying the request to the same master.

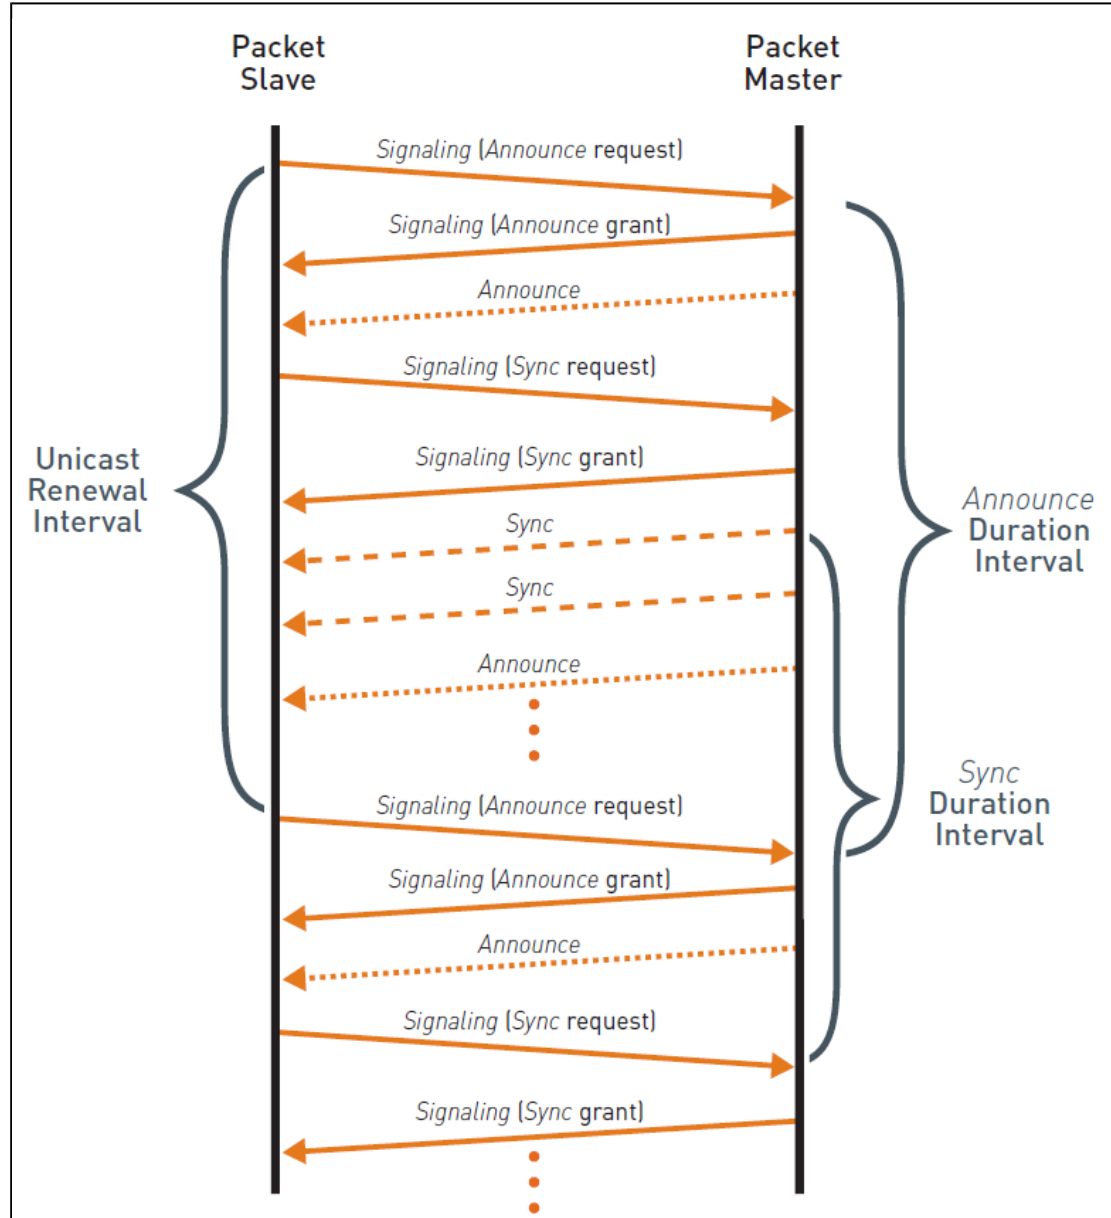
Continuity of Service

As unicast service is only granted for a limited time, the slave must re-request service periodically, before the expiration of the current grant. It is recommended that the slave reissue the request sufficiently far in advance of the expiration to allow the request to be repeated at least twice if no grant is received. For example, the packets containing the request or the grant may be lost in transmission. This shorter interval is known as the unicast renewal interval, and should help to ensure continuity of service.

Order of Service Requests

The following timing diagram shows the order of a typical message exchange between a slave and a master.

Figure 1 • Unicast Negotiation Example (Figure 1 from G.8265.1)



First, the slave requests unicast Announce message service. When this has been granted and the first Announce message received, the slave can check the QL value conveyed in the clockClass field of the message. If this is acceptable, the slave can then request service for Sync and Delay Request messages. As the end of the duration interval approaches, the service requests should be reissued to ensure continuity of service.

The exact order of messages and message requests will depend on type of service the slave requires. It is also possible for a slave to include multiple TLVs in a service request. This may reduce the possibility that a master will grant one request but deny other requests for different types of message.

Traceability

The following sections describe traceability.

SONET/SDH Traceability

In SONET/SDH and SyncE synchronization networks, the ultimate source of the timing signal frequency—or the traceability of the signal—is conveyed by means of the QL value carried in the SSM. These QL values are defined in G.781. For example, if a timing signal is traceable to a primary reference clock, the QL value is set to QL-PRC. This signal will be conveyed all the way down the chain to the end node in the four-bit SSM code.

If the chain is broken immediately before an SDH or SyncE clock (that is, and SEC or EEC), the downstream timing signal is now driven from that clock, and the QL value is set to QL-SEC or QL-EEC1. A subsequent higher quality timing element downstream of the chain such as an SSU will then decide that it can generate a higher quality timing signal than the SEC or EEC, and go into holdover, free-running on its own clock. It will then set its own QL value to QL-SSU_A or QL-SSU_B accordingly.

Therefore, at any point in the synchronization chain, the SSM value reflects the QL value of the synchronization element driving that segment of the synchronization chain.

PTP Traceability

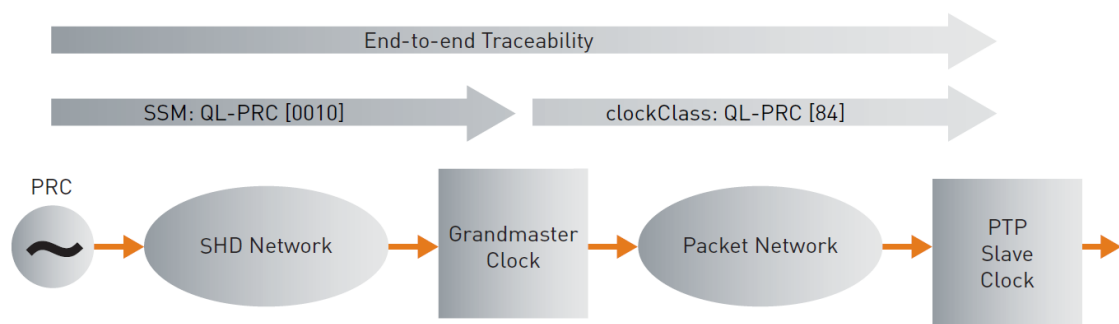
In PTP systems, traceability is carried by the clockClass attribute in the header field of Announce messages sent by the grandmaster. Because PTP was designed as a protocol to carry time rather than frequency, traceability is normally defined in terms of the timescale to which the grandmaster is referenced. For example, clockClass 6 is defined as a PTP clock that is synchronized to a primary reference time source (for example, UTC time), while clockClass 13 is defined as a PTP clock that is synchronized to an arbitrary source of time.

Telecom Profile Traceability

The Telecom Profile is not concerned with time, but frequency. Therefore, the distinction between a primary reference timescale and an arbitrary timescale is irrelevant to the purpose of the profile. Secondly, the clockClass values defined in IEEE 1588-2008 don't indicate the accuracy of the frequency reference, merely the traceability of time.

The clockClass parameter has the range 0 ~ 255, and there are several sections within this range of values that are designated for use by alternative profiles. In order to maintain consistency between various methods of frequency distribution, it was decided to use one of these sections to encode the QL values, rather than attempt to use the existing clockClass definitions. This use of QL values allows interworking between SONET/SDH, SyncE and PTP frequency distribution, and for the end nodes to view consistent traceability information regardless of the means of delivery, as shown in the following illustration.

Figure 2 • Interworking with Existing SDH Synchronization Systems



The mapping between QL and clockClass is shown in the following table (copy of Table 1, G.8265.1). The options are G.781 QL values.

Table 1 • G.8265.1: Mapping of Quality Levels to PTP clockClass Values

SSM QL	Option I	Option II	Option III	PTP clockClass
0001		QL-PRS		80
0000		QL-STU	QL-UNK	82
0010	QL-PRC			84
0111		QL-ST2		86
0011				88
0100	QL-SSU-A	QL-TNC		90
0101				92
0110				94
1000	QL-SSU-B			96
1001				98
1101		QL-ST3E		100
1010		QL-ST3/QL-EEC2		102
1011	QL-SEC/QL-EEC1		QL-SEC	104
1100		QL-SMC		106
1110		QL-PROV		108
1111	QL-DNU	QL-DUS		110

If the grandmaster is slaved to a primary reference frequency, regardless of the traceability of time information, then the clockClass value should be set to QL-PRC [84] (or QL-PRS [80] for Option II). If the primary reference frequency fails, then the clockClass should be degraded to a value consistent with the grandmaster's own internal oscillator. For example, if the internal oscillator is classed as suitable for an SSU type A, then the clockClass should be set to QL-SSU-A [90]. Similarly, if it is Stratum 2 quality (Option II), then it should be set to QL-ST2 [86].

Packet Timing Signal Fail

The Telecom Profile defines the notion of Packet Timing Signal Fail (PTSF), which is not included in IEEE1588-2008. It defines the following three types of PTSF:

- PTSF-lossAnnounce—Loss of reception of PTP Announce messages from a master, carrying the traceability information.
- PTSF-lossSync—Loss of reception of PTP timing messages from a master (that is, Sync or Delay_Response messages).
- PTSF-unusable—Unusable packet timing signal received by the slave, for example, where the packet delay variation is excessive, resulting in the slave being unable to meet the output clock performance requirements.

Loss of the Announce messages means that there is no longer any traceability information available to qualify the accuracy of the signal from the grandmaster. Therefore, the slave should switch to an alternative grandmaster that has valid traceability information. If no alternative master is available, the slave should go into holdover, and set its output QL value consistent with the accuracy of its own local oscillator.

Similarly, in the case of loss of timing messages, traceability of frequency to the master is lost, and the slave should switch to an alternative grandmaster or go into holdover if none is available.

The time taken to recognize these two failures is implementation dependent. If the slave has a very good quality local oscillator, it may be able to implement longer timeouts before declaring a signal fail than if it has a cheaper, more unstable oscillator. The benefit of being able to wait through possible temporary interruptions in service perhaps due to congestion must be balanced against the possible degradation in performance of the slave's output signal due to the interruption.

The durations of the timeouts are defined by the parameters `announceReceiptTimeout`, `syncReceiptTimeout`, and `delayRespReceiptTimeout`. The `announceReceiptTimeout` parameter is already defined in IEEE 1588-2008, while the other two parameters are new in the Telecom Profile.

For PTSF-unusable, the action to be taken is undefined by the profile. If the network is congested, it may mean that signals from other grandmasters are also unusable, but it may not be possible to determine that without attempting to synchronize to them. In that case, it may be preferable to go into holdover for a short period until the congestion has passed.

Grandmaster Selection and Protection

IEEE 1588-2008 defines a method for a slave clock to determine the best master to synchronize to, as part of the Best Master Clock Algorithm (BMCA). This is based first on the priority value, then on the clockClass and the clockAccuracy parameters.

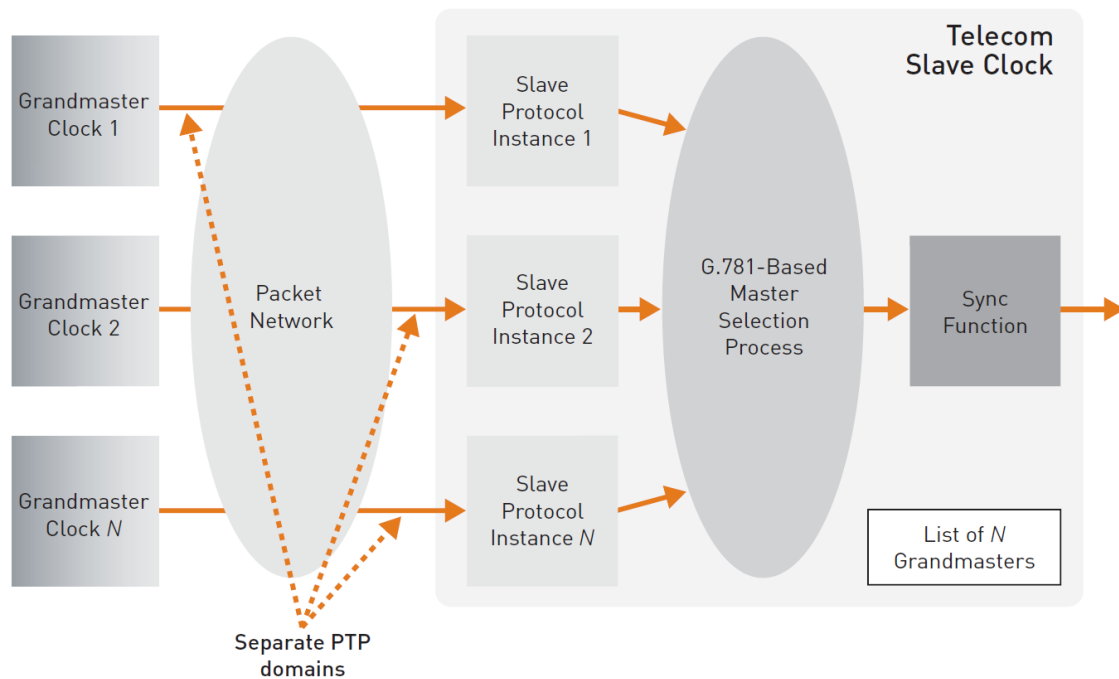
However, G.781 defines a selection procedure for SONET/SDH and SyncE in the opposite order, based first on QL values, then on priority. Secondly, because the unicast PTP communication paths between the slave and each grandmaster are considered as independent domains, the BMCA cannot be used because this would mean one domain is affecting the operation of the other domain.

Telecom Slave Clock

The Telecom Profile defines a new selection procedure based on G.781. This procedure functions outside of the PTP protocol itself. The result of this is that, logically speaking, the slave consists of multiple instances of the protocol operating independently. Each of these instances feeds into the selection process, and the synchronization algorithm operates on the selected instance. The protocol instances obtain the address and priority of each grandmaster from a list of grandmasters. This list is similar to the Acceptable Master Table defined in IEEE 1588-2008 in that it contains both addresses and priority information, but since it operates outside of the protocol it must be considered a separate table.

The collection of protocol instances, selection process, synchronization function, and the list of grandmasters is called the Telecom Slave Clock, as shown in the following illustration.

Figure 3 • Telecom Slave Clock



Initial Master Selection Process

On startup, the process to establish the grandmaster follows a series of steps:

- The list of grandmasters is populated with N entries, each containing the protocol address (that is, IPv4 address) of an acceptable grandmaster and an associated priority. The Telecom Profile does not describe how the information for this list is obtained or entered. The table could be manually populated, or entered through some kind of management system (the means is outside the scope of the profile itself).
- The Telecom Slave establishes a protocol instance for each of the N entries in the grandmaster list.
- Each protocol instance requests unicast Announce service from its respective grandmaster. If service is granted, the protocol instance must obtain the clockClass of the grandmaster and pass this information to the selection process. The means by which the protocol instance communicates with the selection process is implementation-dependent.
- The selection process determines which grandmaster has the highest QL value. If several grandmasters have the same QL value, the priority value is used to select the grandmaster.
- The protocol instance associated with the selected grandmaster requests Sync and Delay_Response service. If granted, the timing information is passed to the synchronization function and an output timing signal generated that is synchronized to the selected grandmaster. If service is not granted for the timing messages, it must either retry the request or the selection process must choose another grandmaster.
- The protocol instances associated with the grandmasters should continue renewing the requests for Announce service, continually monitoring the status of each grandmaster.

If no grandmaster is available, or the QL value of each available grandmaster is lower than the quality of the Telecom Slave Clock's own oscillator, then the Telecom Slave must go into holdover until the conditions causing unavailability have cleared.

Protection and Restoration

The Telecom Slave must trigger the selection process for an alternative grandmaster if any of the following events happen:

- The selected grandmaster declines or terminates a unicast service request from the slave instance.
- The QL value of the grandmaster is degraded for any reason.
- A PTSF-lossAnnounce or PTSF-lossSync condition occurs.
- The QL value of a higher priority grandmaster becomes equal or higher to that of the currently selected grandmaster, providing it is not exhibiting a PTSF condition.

By default, the switch process is revertive. This means that when the condition that caused the change has cleared, the Telecom Slave should re-synchronize to the originally selected grandmaster. The reversion function may be disabled if required by a configuration option on the Telecom Slave.

When one of the trigger events occurs, the following steps occur:

- Each of the remaining protocol instances requests unicast Announce service from its respective grandmaster, passing the clockClass information to the selection process.
- The selection process determines which grandmaster has the highest QL and priority value.
- The protocol instance associated with the selected grandmaster requests Sync and Delay_Response service. If granted, the timing information is passed to the synchronization function, and an output timing signal generated that is synchronized to the selected grandmaster. If service is not granted for the timing messages, it must either retry the request or the selection process must choose another grandmaster.
- The protocol instances associated with the originally selected grandmaster must periodically request Announce service from that grandmaster. When the condition that caused the reselection has cleared, the selection process should operate again, allowing the reversion to occur.

As before, if no grandmaster is available, or the QL value of each available grandmaster is lower than the quality of the Telecom Slave Clock's own oscillator, then the Telecom Slave must go into holdover until the conditions causing unavailability have cleared.

The following several parameters are associated with the protection and restoration.

- Wait to restore time—This is the time from the clearing of the condition that caused the original protection switch before the restoration occurs. If the network is unstable, it may be necessary to increase this time to avoid switching over too early, before the network has stabilized again. There is no value recommended in the Telecom Profile, and the value chosen should represent both operator preference and vendor guidance.
- Non-reversion function—It should be possible to disable the restoration of the original grandmaster. This prevents unnecessary or too frequent switching between grandmasters. The disabling of reversion should be under operator control through a configuration option in the Telecom Slave Clock.
- Output squelch—If the QL value of the selected grandmaster is under a threshold, or the Telecom Slave Clock is in holdover, it must be possible to squelch (that is, disable) the output timing signal from the slave. This prevents the end application equipment from synchronizing to a timing signal that is not synchronized to a reference. The squelch function should be associated with a holdoff time, to allow a re-selection process time to occur in the event of degradation of the QL value, for instance.
- Output QL hold-off—Where the Telecom Slave generates an output timing signal that also carries a QL value, it may be necessary to delay the transition of the QL value in the output timing signal. For example, if the QL value from the primary grandmaster is degraded, time should be allowed for a reselection process to occur and a new grandmaster selected before changing the QL value. This prevents unnecessary switching of frequency references in any downstream equipment.

Summary of Telecom Profile Features

The following tables list the features and parameter ranges specified in the Telecom Profile.

Table 2 • Profile Identification

Feature /Parameter	Note	G.8265.1 Clause
profileName	ITU-T PTP Profile for Frequency Distribution without timing support from the network (unicast mode)	Annex A
profileVersion	1.0	
profileIdentifier	00-19-A7-00-01-00	
Specified by	ITU-T (Study Group 15, Question 13)	
Location	www.itu.int	

Table 3 • General Features

Feature/Parameter	Note	G.8265.1 Clause
Clock identity	EUI-64 (as specified in clause 7.5.2.2.2 of IEEE 1588-2008)	Annex A
Permitted nodes	Ordinary clocks (grandmasters, slave-only clocks)	Annex A
Prohibited nodes	Boundary clocks, transparent clocks	Annex A
Protocol mapping	Both masters and slaves must support IEEE 1588-2008 Annex D IPv4/UDP stack.	6.4, Annex A
Path delay measurement	Uses Delay_Request/Response mechanism, if required (two-way operation) peer delay mechanism must not be used.	6.3.1, Annex A
One-way and two-way operation	Masters must support both one-way and two-way operation. Slaves may support either one-way or two-way, or both.	6.3.1
One-step and two-step clock	Master may support either one-step or two-step clocks, or both. Slaves must support both one-step and two-step clocks, without configuration.	6.3.2
PTP management	Not specified in this version of the profile.	Annex A
Domain number	Range must be supported by both masters and slaves. Range is from 4 to 23.	6.2.1, Annex A
BMCA	Masters must implement static BMCA, with state BMC_MASTER, and state decision code M1 Slaves must implement static BMCA, with state BMC_SLAVE, and state decision code S1	6.7.1, Annex A

Table 4 • Message Rates

Feature/Parameter	Note	G.8265.1 Clause
Sync and follow-up	Range from 1/16 to 128.	6.5, Annex A
Delay_Request/Response	Range from 1/16 to 128.	
Announce	Range from 1/16 to 8. Default value 1/2.	
Peer Delay_Request/Response	Range and default value not used.	

Table 5 • Unicast and Multicast Operation

Feature/Parameter	Note	G.8265.1 Clause
Full unicast operation	Both masters and slaves must support full unicast operation.	6.3.3, Annex A
Hybrid unicast /multicast	For further study; see Appendix I for more details. Not required in first version of the profile.	Appendix I
Full multicast operation	Not required by this profile.	
Unicast negotiation process	Both masters and slaves must support unicast negotiation (IEEE 1588-2008 Clause 16.1). Masters must reject any request it can't completely fulfill, rather than grant less than the requested message rate or duration.	6.6, Annex A
duration field	Duration of unicast service lease. Range is from 60 to 100. Default value is 300.	
logInterMessagePeriod	Masters must accept at least the entire ranges specified. Slaves must operate within these ranges, but do not have to cover the entire ranges. Sync and Follow-up (from 128 per s to 1 per 16 s). Range is from –7 to 4. Delay_request/response (from 128 per s to 1 per 16 s). Range is from –7 to 4. Announce (from 8 per s to 1 per 16 s, default 1 per 2 s). Range is from –3 to 4. Default value is 1.	
logQueryInterval	Slaves must wait a minimum of 1 second before repeating a unicast negotiation request. Default value is 1. Note: If three consecutive requests are denied or not answered, slaves must wait 60 s before repeating the request to the same master.	6.6
number of TLVs	Master must accept multiple TLVs in the same unicast negotiation request. Slaves may include multiple TLVs in the same unicast negotiation request. Range is from 1 to 3.	6.6

Table 6 • Master Selection Process

Feature /Parameter	Note	G.8265.1 Clause
Master selection process	<p>Slaves must select a master from a locally-configured list of acceptable masters according to the following criteria:</p> <p>Master not in PTSF state (Packet Timing Signal Fail), for example, loss of Sync or Announce service, (or optionally, unusable timing signal)</p> <p>Lowest clockClass (SSM QL) value</p> <p>Local priority value</p>	6.7.3.1
clockClass	<p>clockClass value is set according to Table 1 of G.8265.1 (derived from G.781 SSM values). Masters must encode clockClass using this table according to the quality of their input reference.</p> <p>Slaves must use clockClass to select the master to synchronize to.</p>	6.7.3.2
PTSF-lossAnnounce	A slave must set the PTSF-lossAnnounce signal when it has not received any Announce messages for a period longer than the announceReceiptTimeout attribute.	6.7.3.3
PTSF-lossSync	A slave must set the PTSF-lossSync signal when it has not received any sync messages for a period longer than the syncReceiptTimeout attribute, or has not received any Delay_Response Messages for a period longer than the delayRespReceiptTimeout attribute (these are locally-configured attributes, not part of PTP itself).	
PTSF-unusable	<p>A slave must set the PTSF-unusable signal if it is not able to achieve the required performance target.</p> <p>The criteria for establishing this are not specified and for further study.</p>	

Table 7 • Protection Switching

Feature/Parameter	Note	G.8265.1 Clause
Protection switching	<p>A slave must switch to an alternative master under one or more of the following conditions:</p> <ul style="list-style-type: none"> -clockClass (QL value) degraded to a higher value (that is, lower quality) than an alternative master -PTSF-lossSync or PTSF-lossAnnounce conditions on the current master -(optional) PTSF-unusable on the current master (by configuration) <p>Note: If no master is acceptable (for example, all in PTSF condition), the slave must enter holdover or free-run.</p>	6.7.3
Temporary master exclusion	A slave must have the ability to temporarily exclude a master from the list of acceptable masters (by configuration).	6.8.1
Wait to restore time	After a protection switch, when the highest priority master has been restored, a slave must wait before switching back to the primary master. Delay time to be configured (not specified in profile).	6.8.2
Non-reversion	A slave may optionally provide non-revertive switching (by configuration).	6.8.3
Forced traceability	A master must provide the ability to force the QL value to a specified value (by configuration).	6.8.4
QL hold-off	A slave must be capable of delaying a change in QL value to downstream devices. Delay time to be configured (not specified in profile).	6.8.5
Output squelch	A slave must be capable of squelching the output timing signal in case all masters are in PTSF, or QL value of all masters is less than a pre-determined threshold (by configuration).	6.8.6

Conclusion

The Telecom Profile is a key component in delivering reliable synchronization solutions to applications, such as cellular backhaul, that require accurate frequency alignment. It allows operators to design packet-based synchronization systems compatible with existing synchronization solutions, and to maintain existing best practice. It also facilitates the design of inter-operable systems using equipment from multiple vendors.

Microsemi has been active in the development of both IEEE 1588-2008 and the G.8265.1 Telecom Profile. The TP5000 is a carrier-class PTP grandmaster fully compliant with the provisions of the profile, and has been proved interoperable PTP slave devices from many other vendors.

The TP5000 can be deployed with confidence by operators to build synchronization solutions for mobile backhaul applications.

For further information, please contact Microsemi at www.microsemi.com.

References

- [1] IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, IEEE Std.1588TM-2008, 24 July 2008
- [2] Precision time protocol telecom profile for frequency synchronization, ITU-T Recommendation G.8265.1, October 2010
- [3] Synchronization Layer Functions, ITU-T Recommendation G.781, September 2008

**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

© 2011–2017 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-01008-1.00-1117