White Paper

Best Engineering Practices for Cable Timing Architecture—A Study of DOCSIS 3.1

Power Matters.™

Introduction

As high bandwidth services, such as live video and video on demand, migrate to Ethernet and OTT distribution, cable operators are faced with increasing competition from traditional wireline and mobile wireless service providers. To meet this challenge, they are leveraging their existing penetration into the residential market and the relatively high bandwidth provided by cable systems. However, as they scale up their services, they have to ensure that they maintain backward compatibility with the existing cable plant. One of the fundamental functions they have to consider is the timing requirement.

From the beginning, cable networks required good synchronization: first, because the physical-transmission medium is shared by cable modems and poor synchronization could cause interference and crosstalk; second, services such as T1/E1 circuit emulation and high quality video required good timing.

The DOCSIS 2.0 Integrated Cable Modem Termination System (I-CMTS) architecture was developed to accelerate delivery of data-over-cable and integrate internet services with cable-specific functions into a single unit. To meet the tight synchronization specifications required, the I-CMTS was engineered with an embedded oscillator providing a 10.24 MHz clock synchronization pulse for the edge QAM.

Eventually, the high cost and inherent scaling limits of the I-CMTS led the cable equipment vendors to look for a solution that would maintain legacy functions, including synchronization, but create a modular CMTS architecture to facilitate the development of a more flexible, less costly platform. The result was the DOCSIS 3.0 Modular Cable Modem Termination System (M-CMTS), a more modular platform with enhanced performance capabilities.

In the M-CMTS, the embedded oscillator used in the I-CMTS was replaced with a high-precision, standalone timing server to generate the clock with nanosecond alignment at each connected device. This maintained legacy synchronization between the modem and cable headend while allowing the deployment of a modular platform. DOCSIS 3.0 also added network features such as channel bonding that enabled high definition video, fast real-time gaming, and faster internet access. DOCSIS 3.0 facilitated migration of the cable network to Ethernet transport while maintaining the very high quality services provided by the tightly synchronized access plant. The greater efficiencies that resulted from the new timing architecture were crucial in the considerable success of DOCSIS 3.0. However, DOCSIS 3.0 had some inherent constraints that limited the expansion, scale, and flexibility of the cable plant. In particular, there were restrictions in engineering the two-way time transfer, known as DTP or DOCSIS Timing Protocol (standardized by CableLabs as J211) between the DOCSIS timing server and the CMTS core elements. The need for a dedicated copper twisted pair from the DTI server to the slaves in the RPD QAM devices limited the deployment flexibility and the scalability of the M-CMTS systems.

Therefore, the next step was to find a way to remove the limitations imposed by the use of the J211 protocol over dedicated wire. CMTS vendors began to consider the use of Precision Time Protocol (PTP) to replace J211 where possible. This has led to the development of DOCSIS 3.1 and creation of the second version of the Modular CMTS, known as Modular Headend Architecture, Version 2 (MHA-V2), as shown in the following illustration.

Figure 1 • DOCSIS 3.0 Modular Headend Architecture

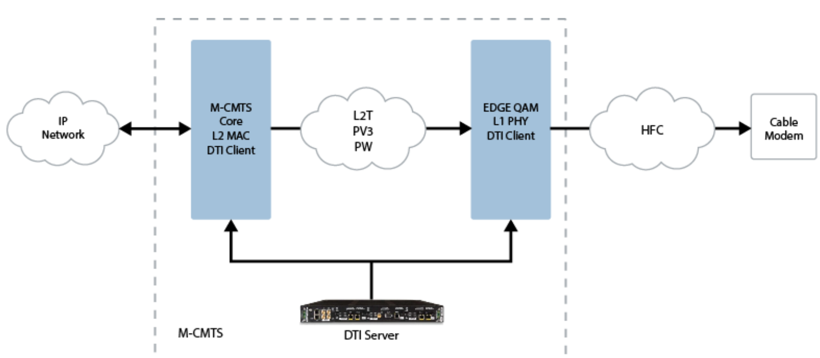

Modular Headend Architecture Version 2 (MHA-V2)

The fundamental changes to the CMTS structure (MHS-V2), also known as DOCSIS 3.1, are intended to extend and enhance flexibility and deployment scale of the Converged Cable Access Platform (CCAP) infrastructure while leveraging the features already built into the previous evolutions of the CMTS. However, to avoid having to change out residential modem and cable access plant, the tight synchronization functions previously deployed at the access layer have to be maintained. In DOCSIS 3.0, the imposition of a dedicated timing infrastructure reduced the cost effectiveness and flexibility of the solution. This is avoided in DOCSIS 3.1 by using PTP to replace J.211 (DTP) wherever possible without affecting the overall timing functions.

The following illustration shows the DOCSIS 3.1 Modular Headend Architecture Version 2.

Figure 2 • DOCSIS 3.1 Modular Headend Architecture V2 (MVA-V2)

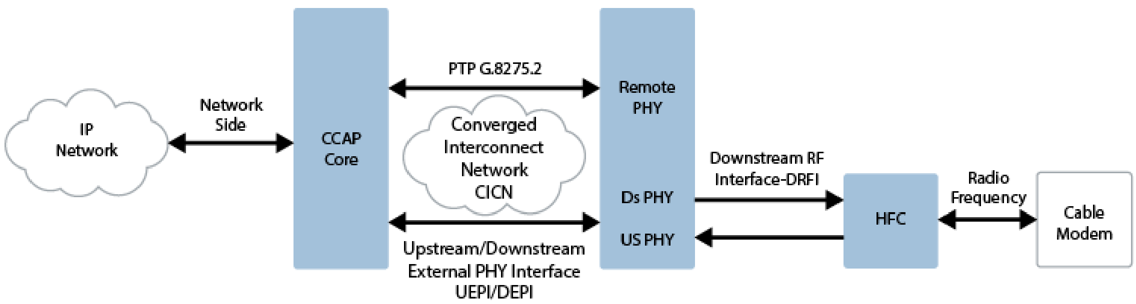

The Advantages of PTP

Modular Headend Architecture Version 2 (MHA-V2) based on DOCSIS 3.1 introduces a key architecture difference compared to DOCSIS 3.0. The timing between the CORE and RPD uses PTP protocols designed to carry phase instead of DTP. PTP runs over Ethernet and allows the cable operator to introduce Ethernet instead of DTI (J.211) over twisted pair between the CCAP core elements and the remote physical devices. This removes the need for a constrained separate physical infrastructure dedicated to timing and allows a more flexible timing architecture based on deployment of PTP.

In this architecture, a PTP grandmaster clock is deployed next to the CMTS core elements or next to the RPD. Each of the network elements is enabled with a PTP slave clock and timing is distributed from the GM to these slaves. Compared to the DOCSIS 3.0 model, this allows further disaggregation of the M-CMTS: the core elements and remote physical devices can be deployed in separate chassis or even in different physical locations. DTP is maintained in the last mile to ensure backward compatibility with the existing plant feeding cable modems. The following illustration shows the timing architecture under DOCSIS 3.1.

Figure 3 • PTP Deployment on the DOCISIS 3.1 Cable Network

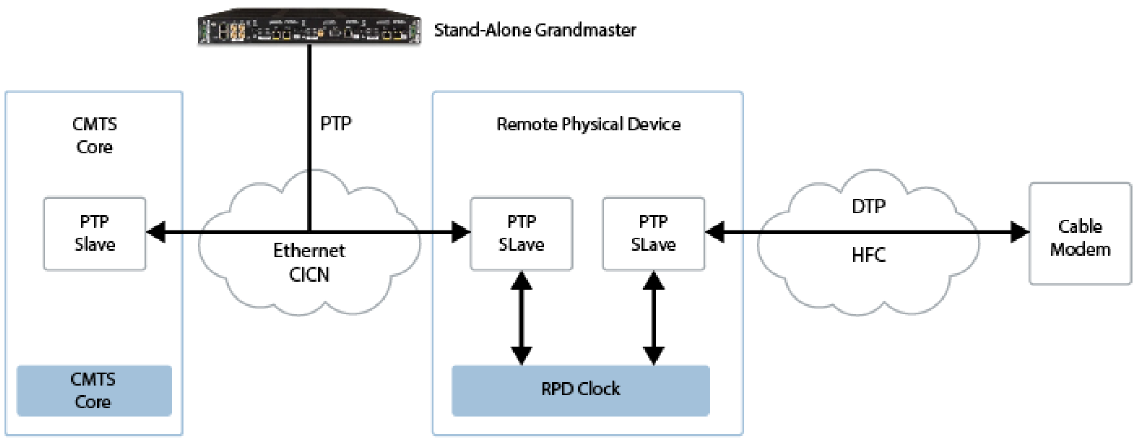

PTP has two phase transport protocol variants, G.8275.1 and G.8275.2. These have different engineering and deployment constraints. G.8275.1 is based on multicast PTP injected directly in the Ethernet frame (between the MAC and the PHY) and is normally deployed on a Layer 2 network. The use of this protocol requires dedicated clocking hardware, known as a boundary clock (BC), on every network element. G.8275.1 then uses a hop-by-hop flow model: each BC terminates an incoming PTP flow and generates another outgoing PTP flow. (A hop is a single network element such as a switch or router.) The maximum number of hops before violation of the time error (TE) limit is gated by BC performance. For example, for the very stringent ±1.5 µs TE requirement used in engineering LTE-based networks, the ITU-T recommends a maximum of 10 hops with a class A rated BC (±50 ns of TE) or 20 hops with a class-B rated BC (±20 ns of TE). However, this level of time error control is not necessary for the DOCSIS 3.1 network where ±1 ms of TE is sufficiently stringent for the network to operate as designed. In this case, we can turn to G.8275.2, which has much less onerous engineering requirements because it does not rely on, nor does it impose, the deployment of BC hardware on every transport node.

For G.8275.1 deployment, the availability of BC hardware in the NE is a major consideration, as it will have to be added to every element in the timing path if it is not already present. Moreover, to ensure the network elements are frequency-aligned, it is recommended to implement synchronous Ethernet (syncE) on the Converged Interconnect Network (CICN), an additional CAPEX and OPEX cost. On the plus side, it can potentially scale to thousands of end clients because it is multicast G.8275.1.

G.8275.2 uses unicast mode over L3 in a client/server model designed to traverse asynchronous networks. As a G.8275.2 unicast PTP flow is not constantly de-jittered by concatenated BC, it is more likely to be subject to random noise accumulated as it traverses the routers. For this reason, it is difficult to determine the maximum number of routers the G.8275.2 flow can cross before the accumulated time error violates the demands of the end application. The hop count will depend on the router implementation. However, for the cable CICN It is important to note that the 1 ms TE limit is relatively easy to achieve. G.8275.2 is also more easily engineered onto a classic (legacy) IP/MPLS network core than G.8275.1, as it does not need additional dedicated BC hardware. PTP servers running unicast mode can usually scale up to a thousand unicast clients.

There are other advantages to the use of G.8275.2. The unicast model enables visibility of the PTP client residing in the PTP receive device; in the case of the CMTS system, that may be the RPD, the core cable network element, or both devices. As the cable architecture expands to thousands of attached RPD, automated end-to-end monitoring and management of the service critical timing system will be necessary—this is not currently possible using a network configured with G.8275.1 and BC transport.

The model also facilitates the deployment of upstream core reference clocks that can provide backup over unicast mode using Assisted Partial Timing Support (APTS) to PTP edge clocks. It is always possible to mix the two protocols, with unicast mode providing the APTS feed from the centralized reference clocks to the clocks at the edge of the network and the edge clocks connecting to the RPD using G. 8275.1 multicast—but again, this prevents visibility of the edge RPD PTP clients and also breaks the continuity, or traceability, of the clock from the core reference to the devices.

Unicast also allows the deployment and engineering of network wide holdover. If a network element or RPD should go into holdover relying on an embedded BC clock, then there will be a very rapid drift, or phase divergence, beyond the maximum allowed 1 ms time error requirement between the CMTS core and the downstream client RPD devices. Because of the inherent relative instability of the internal oscillators used in BC hardware, this phase drift will occur in only a few minutes. For this reason, it is much better to have the RPD and the CMTS continually communicating with the same GM using the unicast mode G.8275.2 profile, which, in contrast to G.8275.1, ensures a direct connection between the remote slave and the source clock.

It should be noted that both syncE and BC are optional with G.8275.2 but the disadvantages in deployment of BC for MHA-V2 outweigh the advantages. For all these reasons, CableLabs has recommended G.8275.2 as the profile of choice for DCOSIS 3.1 PTP deployments. However, this does not exclude the use of G.8275.1 where the necessary hardware is already in place.

The following table lists the main attributes of the two transport protocols.

Table 1 • G.8275.1 and G.8275.2 Attribute Comparision

Attribute	G.8275.1	G.8275.2
Mode	Multicast	Unicast
Transport	L2 (Ethernet)	L2.5/L3 (IP/MPLS)
Boundary clock	Mandatory on all NE	Optional
SyncE	Recommended on all NE	Optional
Assisted partial timing support	No	Yes
End-to-end management visibility	No	Yes
Holdover	Yes on GM No on BC	Yes on GM
Timing traceable to GM	No	Yes

The choice of which PTP phase protocol to deploy, G.8275.1 or G.8275.2, will be driven by the usual network engineering parameters of flexibility and cost depending largely on the current switch router architecture in the Converged Interconnect Cable Network (CICN). In particular, the key parameters of legacy network transport and deployed hardware will gate the decision.

Grandmaster Clocks for DOCSIS 3.1

The move from DTI to PTP for M-CMTS timing requires replacement of the DTI clock servers by PTP grandmaster (GM) clocks. As it is an Ethernet device, the GM may be deployed anywhere within 1 ms of TE of the end clients. Thus, instead of having to be within 200 meters of the slave devices, as in DTI, the PTP GM may be next to the core CMTS, next to the RPD, or across the CICN in a different facility—as long as the phase transport can reach the CMTS and RPD clients within the timing constraint.

Note that a 1 ms time error limit for phase is a relatively easy engineering goal for PTP even over an asynchronous network. This new flexibility in the deployment of the CMTS timing elements enables significant changes in the overall CMTS/RPD system architecture. The following table lists the main differences between DTI server and PTP GMC.

Table 2 • DOCSIS 3.1 PTP Grandmaster Clock and DCOSIS 3.0 DTI Time Server Attribute Comparison

Attribute	DTI/DTP Server	PTP Grandmaster
Standard	J211	G.8272, G.8262
Oscillator	OCXO	Rubidium (high performance), OCXO
Phase transport protocol	DOCISIS Timing Protocol (two-way time transfer)	PTP: G.8275.1, G.8275.2
Transport	Copper twisted pair	Ethernet
Infrastructure	Out of band dedicated	In band
Timing accuracy	±<2 nanoseconds typical	±100 nanoseconds typical
Slave clock	DTI client	PTP client

Microsemi
Power Matters.

Deployment of Grandmaster Clocks for DOCSIS 3.1

Replacing I-CMTS internal oscillators with stand-alone DTI servers was the catalyst for the success of the DOCSIS 3.0 service model. Despite the physical constraints inherent in the use of DTP, there can be no doubt that the shift to DTI contributed to the proliferation of high speed broadband cable services.

For DOCSIS 3.1, the replacement of the DTI server with a GM function can follow two models: embedding the GM into the CMTS network elements; or, a stand-alone GM. There are advantages and disadvantages to each approach.

Embedding the GMC in the cable infrastructure NE such as in a core CMTS element or the RPD appears at first glance to be a logical solution with both CAPEX and OPEX benefits. The operator avoids purchase of stand-alone timing servers and moreover, if the PTP is multicast mode with BC in the CICN switches or routers, there should be no problem meeting capacity requirements. However, there is no significant holdover possible with the small oscillators used in the embedded GM clocks.

Holdover is a critical issue. Delivering timing is in itself not a significant challenge, especially at the 1 ms levels of time error required, but delivering, holding, and maintaining very precise time over a network needs much more than a basic PTP chipset on an IC. CMTS core and RPD network elements are optimized for low-cost service delivery, not for accurate timing, and an embedded solution will necessarily be a compromise between the cost of components and the performance of the clock. On the other hand, standalone timing servers have other attributes not found on embedded clocks, such as the ability to monitor connected PTP clients and assure holdover performance.

The decision to use embedded GM or a stand-alone GM will consider the following:

- Oscillator quality in the embedded clock while keeping the NE cost competitive.
- Holdover performance of the embedded clock oscillator—low-performance oscillators quickly drift from the reference clock, causing a severe degradation or complete loss of service.
- Oscillator options—stand-alone GM can be populated with many different oscillators, including Rubidium oscillators that are able to maintain frequency and phase for significantly longer time periods than even the best Quartz oscillators. Holdover issues are less critical when using a Rubidium oscillator because this implementation allows operations flexibility, especially for a phase-based network. Rubidium is not cost viable on embedded GM.
- Operations and scaling—it is relatively easy to deploy additional stand-alone servers to handle more downstream slave clocks as extra capacity is required compared to the complexities and costs of deploying additional CMTS elements with embedded GM.
- Engineering considerations—synchronization servers are completely different than network elements such as routers, switches, and CMTS elements. Clock system hardware is optimized for the generation, distribution, and maintenance of the timing source, and is a built-for-purpose system. This is a different quality paradigm to the clocking on an embedded system where the priority is to minimize cost by sharing CPU and other resources to better process bearer data, rather than to generate or manage a high quality clock.

Service Critical Functions: Capacity and Redundancy

The goal of MHA-V2 is to ensure that the cable service is cost effective to deploy with maximum availability and resilience. The cost efficiency, scalability, and flexibility of the timing architecture will depend on the capacity of the clock that should be able to support a large number of unicast clients.

High capacity, however, is not enough—the system also has to be high availability to guarantee continued service, with an inherent ability to protect against failure at multiple levels, including power supply, IC and GNSS components, oscillators, and failure of PTP, IP, and Ethernet routing systems. Simultaneously achieving such comprehensive protection at multiple layers with a high precision clock is a difficult engineering problem, and is currently only delivered to the market by the Microsemi TP5000. It is noteworthy that this clock can support NTP and all PTP protocols for both IPv4- and IPv6-based networks.

The following table summarizes and compares embedded and stand-alone GM server attributes.

Table 3 • Embedded and Stand-Alone GMC Attribute Comparison

Attribute	Embedded GMC	Stand-Alone GMC
TCXO/mini OXCO	Yes	Yes
DOCXO	No	Yes
Rubidium	No	Yes
SyncE input/output	Yes	Yes
PTP input/output	Yes	Yes
Extended holdover	No	Yes
High-performance APTS	No	Yes
Flexible deployment	No	Yes
Clock monitoring	No	Yes
Clock management	No	Yes
High redundancy options	In NE	Yes
Power redundancy	In NE	Yes
High-performance NTP services	No	Yes
Gateway clock functions	No	Yes

Whether the operator finally decides on the embedded or the stand-alone model, the clocking infrastructure is critical for cable services. Moving it to PTP, rather than using DTI, in no way eliminates or mitigates the need to carefully design and engineer the timing system, and there can be no doubt that a stand-alone GMC will use significantly better clocking hardware and enable much more flexible deployment and management than is available on an embedded system.

New Applications and Services on Cable Networks

Mobile Backhaul and Timing Considerations

Cable operators are now moving into Mobile Backhaul (MBH) as a way to improve revenues and capture more subscribers. However, MBH services have very different synchronization requirements to those established for a cable network. In cable DOCSIS architectures, timing servers are not tied to an external clock reference such as Coordinated Universal Time (UTC). The CMTS, RPD, and cable modem can be allowed to drift in free-run mode because all elements connected to the same clock will drift with that source and therefore remain synchronized to the server and to each other. There is no requirement for an external reference to act as an anchor for the system or to ensure that the network devices such as RPD or modems stay aligned with other off-network devices.

Mobile services follow a different paradigm. For LTE-A and/or 5G, phase-based services will be a dominant network requirement. Base stations will require stringent phase/time synchronization such that eNB from any operator can be aligned to a common reference clock tied to UTC. If a cable operator wishes to deploy eNB, or provide MBH to a mobile service provider, frequency or phase-based synchronization services tied to that same GNSS UTC reference must be delivered to the eNB. This requires the deployment of telecom grade GNSS systems (antenna, splitters, cabling, and so on) that can plug into the GM, which at that point is now acting as a Primary Reference Time Clock (PRTC). The PRTC has to meet the ITU-T requirements established by G.8272; therefore, it must be connected to GNSS and deliver an output at ±100 ns from UTC. CMTS vendors delivering an embedded GM function will be required to implement a PRTC GNSS chipset, a surge protected L1 input, and a much more sophisticated clock algorithm to synthesize the incoming satellite signals to ensure compliance with G.8272. As with embedded GM, embedded GNSS implementations are built around minimum cost, whereas the stand-alone GM GNSS is built around maximum performance. The result is a very different type of GNSS instance. It should be noted that the TimeProvider 5000 meets all the ITU-T G.8272 requirements.

Virtual CMTS

One of the next steps in the DOCSIS 3.1 evolution is the introduction of virtualization technology into the ecosystem whereby virtual CMTS servers will replace the traditional CMTS component. This will result in further segregation and decomposition of the CMTS technology, where the vCMTS servers will reside in a cloud-based data center location and the physical hardware, such as RPD aggregation switches, will reside in a separate location optimized for the network connection to the RPDs.

In this architecture, the vCMTS servers and the RPD, rather than the RPD aggregation switches, will be populated with PTP slaves. The logical choice for GM placement would be at the vCMTS locations and then provide QoS for the PTP flows within the data center to the RPD locations through the RPD aggregation switches.

An important technical consideration is the configuration of the GM clocks the use of either GNSS/GPS or free run modes of operation. GNSS/GPS is not needed to phase align the CMTS to the RPD. The GM clocks can be configured in the free run mode of operation and this will provide the necessary phase alignment. When engineering the timing architecture, however, careful consideration should be taken when using the free run mode in a GM clock.

When deploying vCMTS servers in a data center type environment, it is important to make sure all the RPDs in this architecture are using the same free running GM clock as the vCMTS servers. If there are RPD using a different free running GM clock than the one serving the vCMTS servers, the phase alignment between the vCMTS servers and the RPDs will be compromised. One way to correct this phase offset beyond the 1-millisecond requirement would be to add GNSS/GPS capability to the GM clock configurations. With GNSS/GPS UTC traceability, it is a very low probability that there will be phase divergence between the GM clocks.

Robust End-to-End Clock Architecture

Server capacity, inherent clocking element redundancy, and economic constraints are not the only considerations when building an end-to-end timing system. If cable operators are to provide mobile backhaul and vCMTS services that require interconnect with other timing systems, there is a need to ensure end-to-end timing consistency relative to the UTC reference clock in line with telecommunications service providers worldwide.

With PTP services now available, the cable operator can ensure that core and edge clocks are tied together and back each other up over traceable PTP timing flows using Assisted Partial Timing Support (APTS) (G.8273.4). APTS is a technique that uses PTP to enable core high-precision, high-availability clocks based on a Cesium reference and high-accuracy GNSS services as backup for clocks deployed at the edge of the network. The following diagram shows a typical APTS architecture.

Figure 4 • Assisted Partial Timing Support

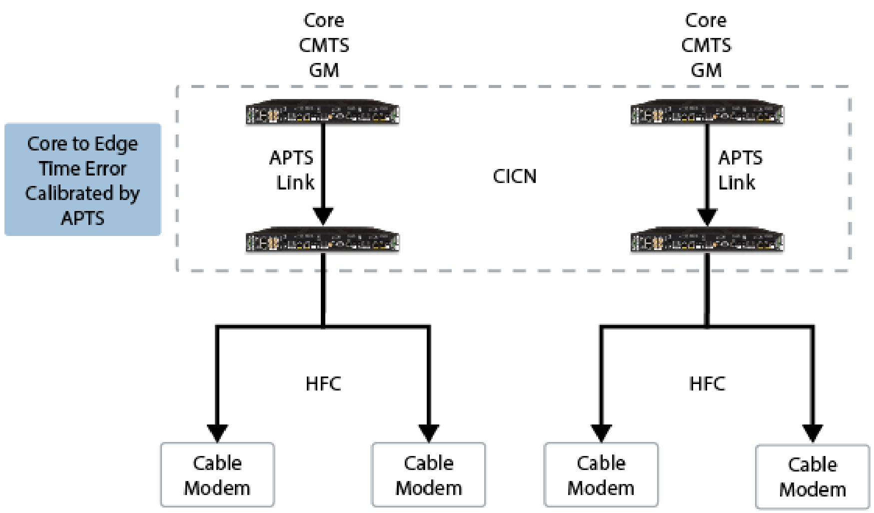

In this model, the UTC GNSS reference clock is connected to the Core CMST GM. This is typically a high-precision clock backed up with Cesium resources. The PTP-based APTS services are configured over the CICN to the edge clocks deployed where RPD connections are aggregated. To ensure that the core and edge clocks are aligned in both phase and frequency, the edge clocks are also enabled with GNSS. We now have a system where the core and edge clocks are phase and frequency aligned and the PTP flow between the two is calibrated by the common GNSS reference. If the GNSS on the RPD clock is impacted for any reason, the edge RPD clock can use the GNSS-calibrated incoming PTP flow as the timing reference. The core GNSS is protected by the Cesium input. This architecture will ensure that all clocks in the network are running off the same reference even if there is a failure of GNSS at any point for any reason.

APTS is a very effective way to use available PTP resources to ensure consistent end-to-end timing services and maintain coherent clocking between core and edge devices. Instead of free running in holdover, the edge RPD will be homed to the core clock that provides a network wide reference clock. The ability to link and mesh the clocks end-to-end across a network is one of the major advantages of PTP, but requires APTS, preferably with robust management and monitoring. The following diagram shows APTS and time error for a cable CICN.

Figure 5 • APTS System with Core Clocks, Edge Clocks, and Allowable Time Error

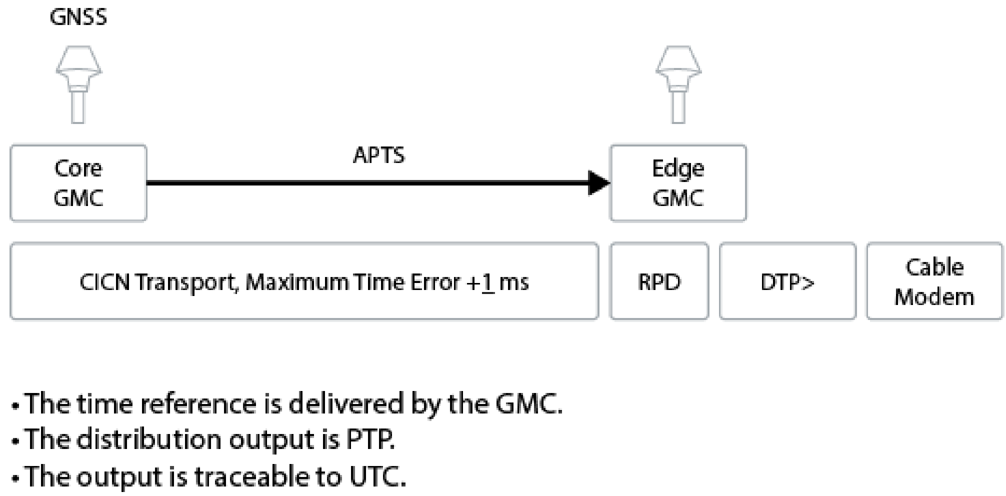

Multiple Timing Protocols in the Cable CO

The introduction of PTP to complement DTI is required for network infrastructure timing. However, there is also a need for high-performance NTP to support the logging and billing functions that are ubiquitous in networking elements of all kinds, from LDAP and provisioning servers to switches, routers, and RPD.

Typically, NTP services are provided by a stand-alone NTP timeserver such as the Microsemi SyncServer S650 series. In addition, Microsemi TimeProvider 5000 stand-alone GM clocks are designed to provide both high performance PTP and NTP services from the same device simultaneously. Both the SyncServer S650 and the TimeProvider 5000 advanced NTP clocks support IPV4 and IPv6 as required. IPv6 will be increasingly required as the networks scale.

Summary

For more than thirty years, Microsemi has been engineering timing services for telecommunications, data centers, enterprise, government, defense, and cable networks worldwide: the company designs the Maser and Cesium clocks that make up 90% of the BIPM standard for global time, built one of the very first GPS receivers (now in the Smithsonian), provides the clocks for the major global financial institutions, stock markets, and mobile networks, and has a well-known pedigree in the development of the global IEEE and ITU-T standards for synchronization and timing, including the specification of the original DOCSIS 3.1 Timing Protocol in collaboration with CableLabs.

Microsemi is the world leader in the development of the PTP systems now being adopted by cable networks. Since 2007, it has supplied PTP GM clocks to more than 400 mobile networks, facilitating the adoption of high bandwidth Ethernet services that are now accepted as the norm. These systems have quietly and efficiently enabled the extraordinary global expansion of internet based mobile services.

Microsemi's innovative engineering and unparalleled experience in timing and synchronization systems will enable the continued evolution of global communications, IoT, connected vehicles, and high speed 5G mobile networking. With a focus on excellence and on customer success, Microsemi PTP solutions will now also drive the next generation cable systems.

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

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-01002-1.00-0917