Coexistence of IEEE 1588, C37.238 and 802.1AS,

Issues and Recommendations

Prof. Dr. Hubert Kirrmann
ABB Research Center,
Baden, Switzerland
hubert.kirrmann@ch.abb.com

Dr. Jean-Charles Tournier
CERN,
Geneva, Switzerland
jean-charles.tournier@cern.ch

***Abstract*— Profiles and extensions of IEEE 1588 have recently been standardized as IEEE C37.238 (for electrical power systems) and IEEE 802.1AS (for audio-video bridging). The working groups focused more on innovation in their application domain than on keeping compatibility with the original default profile of IEEE 1588 J.4. As a result, the profiles cannot share the same network infrastructure, hampering the engineering of mixed system, e.g. integrated power and automation networks. Nodes that support several profiles become complex and have to support three MIBs for SNMP, while one would be sufficient.**

This position paper presents the peculiarities of IEEE C37.238 and IEEE 802.1AS and makes propositions to ensure at least coexistence on the same network, and possibly allow a reuse of a common trunk with compatible extensions. These propositions imply that all standards should be amended simultaneously at their next revision, with 1588v3 taking over the bulk of the additional features and the other becoming true profiles. This contribution is limited to a mapping on IEEE 802.3, which is the most frequently used technology in industrial automation.

Keywords PTP; IEEE 1588; IEC 61588; IEC 61580; C37.238; IEEE 802.1AS; AVB; IEC 62439-3, protocol coexistence.

I. INTRODUCTION

IEEE 1588v2 [1] ("1588" in the sequel) defines a profile as: "*§19.3.1.1 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*". This definition implies that, while nodes of two different profiles are not necessarily interoperable with each other, they both should remain conformant to a subset of 1588. Features not foreseen in 1588 are considered as extensions. Compatible extensions allow coexistence with the original protocol, while incompatible extensions modify the original protocol.

Shortly after 1588 was published, PTP "profiles" specific to application domains were drafted. In this paper, we consider only three profiles, all based on the 1588 Annex F transport (Ethernet) and on the default profile in 1588 Annex J.4 (peer delay) since this is where incompatibilities are expected to first emerge.

Based on the 1588 profile definition, some working groups interpreted it in the sense that interoperability is limited to devices of their own profile but did not consider coexistence with another profile using the same transport protocol. In particular, IEEE C37.238 [2] defines a "profile" for power system applications, and IEEE 802.1AS [3] defines a profile <u>and</u> an extension for Audio-Video Bridging (AVB). Both working groups made the strong assumption that all nodes on the LAN obey only to their protocol. However, profiles can be mixed or connected, for instance, in power plants that integrate automation and power segments, e.g. Profinet (synchronized with 802.1AS) for the process part and IEC 61850 (synchronized with C37.238) for the substation part. Such configurations lead to the following technical issues:

- Due to mandatory extensions to 1588, a device fully conformant to 1588 cannot be configured to operate according to the two domain specific "profiles".
- Transparent clocks in bridging devices would need to support multiple protocols simultaneously, but cannot distinguish them properly.
- A (costly) reference clock per profile is needed, no benefit exists from redundancy.
- Several grandMasters could be sending Syncs at the same time, which would be acceptable if the clocks could distinguish to which master they listen.
- Networks must be engineered entirely out of devices supporting the same profile but homogeneity of the configuration cannot be checked. Wrongly configured multi-protocol devices are difficult to detect.

To overcome these limitations, changes are needed in the different "profiles", with the objectives of:

1) *Compatibility: to distinguish protocols in engineering and operation so that the different protocols do not disturb each other*
2) *Coexistence: to allow nodes to support multiple protocols simultaneously.*
3) *Reuse: to allow reuse of a common core, with the aim of having only one protocol machine that can be configured for a specific profile.*

The rest of this paper is organized as follows: Section II gives an overview of C37.238 and Section III an overview of 802.1AS. For each profile, we distinguish between the parts that define a proper profile, i.e. subset of 1588, compatible extensions, i.e. backward compatible on lower functionality, and incompatible extensions. Section IV focuses on the separation and coexistence aspects of the two profiles, while Section V presents recommendations for both profiles and 1588 in order to improve reuse, applicable to IEEE 1588v3.

II. IEEE C37.238

A. History and objectives

C37.238 was drafted by the IEEE "Power System Relaying Committee" for power system applications. It aims at providing an accuracy better than 1 µs after a maximum of 16 transparent clocks for any slave clock. C37.238 specifies mapping to IRIG-B and time-stamped values such as IEEE C37.118, IEC 61850-7-2 and IEC 61850-9-2.

B. C37.238 profile features (see Table 1)

C37.238 specifies a profile conform to 1588 Annex J with the following characteristics:

- Profile identifier: = 1C-12-9D-00-00-00;
- Device types: slave-only and master-capable;
- Grandmaster clockClass = 6 or 7 (degraded)
(directly connected to reference, may not be a boundary clock except after holdover time)
- Grandmaster accuracy = better than 0,2 µs
(clockAccuracy 20 or 21), degradable.
- Holdover of master-capable clocks = maintain inaccuracy within 2 µs during at least 5s.
- Inaccuracy of transparent clocks = smaller than 50 ns.
- Domain number (default value) = 0;
- Time scale = PTP only;
- 802.1Q-tagging = mandatory for all sent messages, with a default priority of 4 and a default VID of 0;
- Transport mapping = Ethernet layer 2 only (no UDP);
- Messages sent to multicast address (no unicast)
- for Pdelay = 01-80-C2-00-00-0E,
- for all other messages = 01-1B-19-00-00-00;
- Peer-to-peer delay are calculated (no end-to-end);
- Both 1-step (recommended) and 2-step are allowed;
- Best Master Clock for all, except slave-only clocks;
- Intervals specified:
Announce message period = 1 s,
Pdelay_Req message period= 1 s,
Sync message period = 1 s,
announceReceiptTimeout for preferred master = 2 s,
announceReceiptTimeout for all other devices = 3 s,
(allowing a faster recovery for the preferred master);
- Default settings:
priority 1 (slave-only = 255, master-capable = 128)
priority 2 (slave-only = 128, master-capable = 128)
- LocalTime TLV:
A master appends to its Announce messages the TLV ALTERNATE_TIME_OFFSET_INDICATOR defined in 1588 that carries the time offset with respect to UTC. This eases setting the HMI of the IEDs at the beginning and end of daylight saving time periods;

C. Profile specific features: organisation-extension TLV

1588 allows each profile to define organization-specific TLVs. Coexistence is maintained as long as nodes that cannot interpret these TLVs ignores them, as 1588 §14.1 states.

C38.237 specifies that every Announce message shall carry a first TLV that contains two fields: grandMasterId and timeInaccuracy, as illustrated in Figure 1.

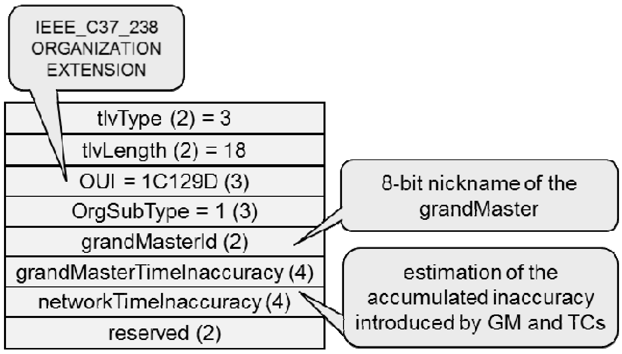

Figure 1. C37.238-specific TLV (tlvLength corrected).

1) grandmasterID

The 8-bit grandMasterID is a nickname for the 64-bit ClockIdentity, intended for slave-only devices with domain specific protocols unable to deal with a 64-bit clock identity. The use case is that the IEC 61850-9-2 frames have only an 8-bit field (SmpSynch) to identify the clock source.

C37.238 specifies that a master shall not send Announce messages if its grandMasterID has not been configured. Therefore, if a non-C37.238 master is present simultaneously to a C37.238 master without its grandMasterID configured, all slaves will be synchronized to the non-C37.238 master regardless of its priority and clock quality. Since a slave is not required to check if the grandMasterId is present in the Announce message, a time jump can occur when the master gets configured.

2) timeInaccuracy

The timeInaccuracy carries an estimate of the clock deviation from the reference time. It consists of two fields: grandMasterTimeInaccuracy and networkTimeInaccuracy.

a) grandMasterTimeInaccuracy

The C37.238-specific TLV carries the clock inaccuracy of the grandmaster as a nanoseconds value, which is a more precise value than the clockQuality in the Announce message. This value is a property of the device depending on environmental conditions and must be less than 0,2 µs.

b) networkTimeInaccuracy

networkTimeInaccuracy should express the inaccuracy introduced by all clocks along the synchronization chain.

In 1588, each transparent clock corrects the received time by its own residence delay and an estimation of the path delay to the peer node from which it received the Sync message. Due to clock domain effects in the hardware, PHY delays, asymmetry in the path to the peer and internal clock drift especially before syntonization, both residence and peer delay have an inaccuracy, expressed as "localTimeInaccuracy" in ns.

The original concept of "dynamic timeInaccuracy" was to let an ordinary clock estimate its inaccuracy at run-time. Each transparent clock would receive the timeInaccuracy from its upstream peer, adds its current (dynamic) localTimeInaccuracy and send the accumulated timeInaccuracy downstream in the Announce message. However, including dynamic timeInaccuracy requires the Announce messages to be modified in transit, although they are no event messages. Transporting this information in the Sync message would have been easier. For lack of consensus on how to compute and use "timeInaccuracy", C37.238 declared the TLV as experimental.

Rather, a static approach has been taken. We argue that network time inaccuracy should be evaluated while designing a network rather than discovering that it is insufficient at run-time. Indeed, an automation network should be engineered so that the clock inaccuracy over a path to any ordinary clock cannot exceed the maximum value permitted by the application on that clock. Knowledge of the topography (considering all alternate paths after reconfiguration) and of the maximum clock inaccuracy of each transparent clock is sufficient to evaluate time inaccuracy off-line. The engineer locates the nodes and links so that time-critical applications are close enough to the grandmaster clock and its back-up(s).

C37.238 requires that a transparent clock introduces an inaccuracy of less than 50 ns under all conditions.

The use case is a phasor measurement device located at the end of a chain of 16 transparent clocks. It expects a maximum accumulated inaccuracy of 16×50 ns = 0,8 µs, that added to the maximum inaccuracy of the grandmaster clock (= 0,2 µs) yields 1 µs.

A transparent clock's worst case inaccuracy can be smaller 50 ns and if so, this value must be part of the device type descriptor and entered into the engineering tools.

The manufacturer guaranteed clock inaccuracy is also part of the MIB of each clock, which is only readable when the devices are in place. Discovering at installation time that a substation network presents intolerable inaccuracy accumulation can cause high costs in re-engineering.

D. MIB

C37.238 specifies a comprehensive MIB with about 100 objects, more precisely $73 + 14 \times p$ (p being the number of ports) objects that list all 1588 datasets plus the additional elements of C37.238. The objects are separated for ordinary and transparent clocks.

The MIB is a compatible extension of 1588 (which has no official MIB), it is not aligned with any other 1588 MIB, and therefore a C37.238 device requires its own SNMP object repository. This MIB is mandatory for grandmaster-capable devices, which includes all nodes that can act as back-ups.

The MIB contains a ProfileId for each port, which is not identical to the profileIdentifier of 1588 and whose use is unspecified. It should indicate if the port supports 1588, C37.238, 802.1AS, lxi or the telecom profile. How this should be used to distinguish protocols is not specified, since other protocols do not support this object.

E. Incompatibilites C37.238 – default profile

We consider as an incompatibility that a C37.238 device cannot work in the presence of non-C37.238 devices and vice-versa.

C37.238 specifies that a master shall appends the two mandatory TLVs to its Announce messages, the TLV ALTERNATE_TIME_OFFSET_INDICATOR and the organization-specific TLV. This was done on purpose to oblige all nodes in the network to be conformant to C37.238.

Indeed, these TLVs are the only way for other devices to distinguish C37.238 masters from default profile masters, since both use the same subtype (=0) and the time domain is not fixed (only the default is).

This prevents synchronization of a network in which any 1588 master is present. Indeed, C37.238 masters will not accept a higher priority 1588 master since it does not carry the two TLVs, while the 1588 master will not accept the lower-priority C37.238 master. In this case, the network is flooded with Announces and no Syncs are sent.

The use of VLAN to avoid this situation requires that the Transparent Clocks support two different protocol machines, i.e. they must be aware that two protocols share the network. However, a bridge does not map the protocol to the VLAN ID.

Checking the configuration to ensure that all nodes obey to C37.238 is difficult. The PTP messages carry no C37.238 identifier. A Transparent Clock cannot know that the Pdelay_Req or Pdelay_Resp it receives was sent by a 1588 J.4 or a by C37.238 node. Moreover, C37.238 dispenses a slave-only device from sending Pdelay_Req (this is a deviation from 1588 since it changes the state machines definition). Since C37.238 also does not require a slave-only node to check that its grandmaster is C37.238, a slave can be synchronized by any non-C37.238 grandmaster and there is no way to report this situation.

III. IEEE 802.1AS

A. History

IEEE 802.1AS names itself gPTP ("generalized Precision Time Protocol"). It was developed by the Audio/Video Bridging Task Group of IEEE 802.1. It allows use of cost-effective, lower-quality local oscillators with a precision better than 100 ppm and a drift of less than 1 ppm/s. The innovation of 802.1AS is the transmission of the fractional clock deviation in Follow_Up messages.

B. 802.1AS Profile (see Table 1)

802.1AS defines a profile according to 1588 Annex J with the following parameters:

- Profile identifier: 00-80-C2-00-01-00
- Device types: "end stations" and "bridges" (ordinary clocks and boundary-transparent clocks)
- Domain number: 0 (single time domain only)
- Multicast messages (no unicast messages) all using MAC address 01-80-C2-00-00-0E
- VLAN tagging prohibited

- Peer-to-peer delay calculation (no end-to-end)
- 2-step only
- Uses "boundary clocks", which 802.1AS claim to be "mathematically equivalent" to "transparent clocks"
- Intervals specified:
  Announce message period = 1 s (default)
  Sync message period = 0.125 s (default)
  Pdelay_Req message period = 1 s (default)
  announceReceiptTimeout = 3 s
  offsetScaledOffsetLogVariance interval = 0.125 s
- Default settings:
  priority 1: slave-only = 255, master-capable = 128
  priority 2 = 248 (default)

C. Non-Profile features

- Best Master Clock "simplified" (incompatible with the default 1588 BMCA);
- All messages use the "link-specific" address 01-80-C2-00-00-0E. Other profiles do not specify whether to accept or ignore such messages. 802.1AS does not specify what an 802.1AS node should do when receiving messages over another address;
- Participation of the bridges in the selection of the masters is a compatible extension.
- Syntonization method based on transmission of frequency ratio in a TLV appended to the Pdelay_Resp_Follow_Up and Follow_Up messages as Figure 2 shows. This can be considered a compatible improvement of 1588's experimental frequency scale factor. 802.1AS does not specify how a clock should behave if it does not receive the TLVs, but 802.1AS does not exclude that it synchronize as in 1588, with degraded accuracy. A 1588 clock from another profile ignores these TLVs.

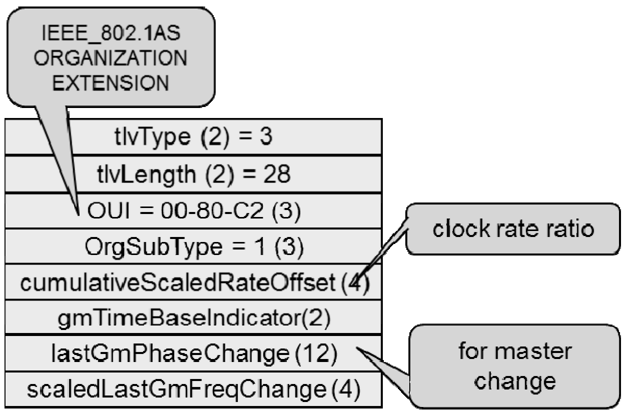

Figure 2. 802.1AS-specific TLV.

- Signaling messages. These messages are optional in 1588, not foreseen in C37.238, but used in 802.1AS to negotiate message rate and to track the path in low-power applications.

D. 802.1AS MIB

802.1AS defines a MIB in the 802.1 OID branch that describe only the Default Data set. It consists of $40 + 40 \times p$ objects (p being the number of ports). Even if IEEE requires each standard to have a MIB, this definition is not aligned with any other 1588 MIB, and therefore 802.1AS requires its own SNMP object repository.

E. Incompatibilites 802.1AS – 1588 J.4

802.1AS messages carry a protocol identifier in the subtype. However, neither 1588 nor 802.1AS specify what a device has to do when the subtype is not its own, so in principle devices can be mixed.

802.1AS does not specify how devices should behave when the upstream device is not 802.1AS. Although 1588-style syntonization could still work, 802.1AS does not foresee a degraded mode and could stop working.

The behaviour of a master in a mixed system is not specified and could lead to multiple masters.

IV. Towards Coexistence and Separation

A. Objectives

This section addresses the objectives stated in the introduction section in order of increasing difficulty, namely co-existence and separation. The third objective, reuse, is the results of the recommendations presented in Section V.

B. Separation

Separation consists in avoiding that devices of different protocols are directly connected and checking that it is so.

For instance, Profinet (802.1AS) and IEC 61850 (C37.238) networks can be kept separated with a boundary clock in between, but some unintentional connection may exist.

So, even with physical separation, the presence of non-conformant devices that could ruin synchronization must be checked. A network observer can detect the presence of 802.1AS devices by their subtype = 1, but a C37.238 master cannot detect that its slave clocks all obey to C37.238 since they do not send Announce frames and need no MIB.

802.1AS does not support separation by boundary clocks.

Protocols could be distinguished using:

- VLANs, but 802.1AS prohibits the use of 802.1Q tags while C37.238 mandates them. VLANs should not be misused to distinguish protocols since they can be removed by intermediate bridges.
- an own multicast domain, but bridges cannot assign a multicast domain to a protocol.
- an own time domain, but it is cumbersome to assign time domains to protocols.
- the transportSpecific field, i.e. the subtype, and an extended field in which profiles can insert their profileID (based on their OUI).

So an explicit protocol identifier is needed in each message, also to ensure that the configuration is homogeneous.

C. Coexistence

Coexistence considers a situation in which devices obeying to the different protocols send their messages over the same transport and understand a subset of each other. For instance, a 802.1AS bridge could receive messages from a 1588 J.4 or a C37.238 upstream or downstream clock.

The condition is the ability to distinguish the protocols or to ignore the distinction when irrelevant. Currently, only 802.1AS can be identified through its subtype (transportSpecific) in each message. C37.238 can only be identified by the TLV in the Announce message.

This is not sufficient, and coexistence requires that the behavior of a clock is defined when it receives a message that is not of its native type, and this in every protocol.

V. Recommendations

A. C37.238

To keep the qualifier of "profile", C37.238 should be reduced to the list of settings and options of 1588v3.

The organization-specific TLV should be optional – its only purpose is to identify a master as obeying to C37.238. All other information can be gathered at engineering time.

C37.238 should suppress the mandatory 802-1Q tag. The use of VLAN is a network engineering, not a clock decision.

If C37.238 insists on keeping any incompatibility, it should receive an own subtype, for instance transportSpecific = 2. This would make the identification of the master through the two TLVs in the Announce frame redundant.

B. 802.1AS

802.1AS should include a backward-compatible frequency ratio correction in case the TLVs are missing.

C37.238 and 802.1AS should use BMCAs that lead to the same result.

One-step clocks should be accepted.

C. IEEE 1588v3

IEEE 1588v3 could allow coexistence of 1588 J.4, 802.1AS and C37.238 clocks on the same Ethernet with relatively small changes to the specifications.

As a principle, degraded synchronization should be better than no synchronization at all.

The nature of the connected devices should be checked to achieve the best possible synchronization. To this effect, IEEE should assign a unique identifier to each profile. A mechanism to report misconfiguration should exist.

Messages received from other protocols should be either rejected or processed up to the possibilities of the receiver.

All bridges and transparent clocks should accept all protocols, and send traffic with their native protocol.

It is a waste to calculate the peer delay with several different messages. A pinged node should ignore the subtype and respond with its own peer delay. The peer delay mechanism should be compatible for all three protocols. 802.1AS devices would benefit from a better syntonization.

The approach taken by C37.238, that each node accepts both one-step and two-step and can generate either one-step (preferably) or two-step should be adopted by the other protocols. Indeed, it is easy to convert a pair of 2-step Sync + Follow_Up into a one-step Sync (respectively for Pdelay).

802.1AS should be considered the most precise clock since it includes a frequency correction not present in the other protocols, but should revert to 1588 if the TLVs are missing. This should be reflected in the clock quality only.

To achieve a common clock domain, a grandmaster clock should have both an 802.1AS and a C37.238 interface, but a safe method should ensure that the two domains always chose the same GM, even after reconfiguration of the network.

1588v3 should not limit the BMC algorithm to the master capable devices, but extend it to slaves, like in IEC 62439-3 [4], which can choose among different GrandMasters. This would allow an ordinary clock to listen specifically to one type of grandmaster or have several active grand masters.

1588v3 should envision several active masters and several profiles active on the same network.

It would be helpful to close the TLV list with a zero-TLV to avoid interpreting padding or trailers as TLVs.

D. MIB (SNMP) issue

1588 has no official MIB.

The MIB of C37.238 being the most comprehensive, it could serve as a common base for 1588v3,

The C37.238 and 802.1AS specific fields could be added as separate branches in the same OID tree.

The MIB should be optional for all clocks. Management messages could provide a reduced service at lower cost.

VI. Conclusion

This position paper studied the possibility of different IEEE 1588 "profiles" being active simultaneously on the same network, in particular 1588 J.4, IEEE C37.238 for power systems and IEEE 802.1AS for audio-video bridging.

It was shown that the latter profiles include specificities that make them unnecessary incompatible with the 1588 J.4 default profile and prevent profiles to run simultaneously on the same network. As future systems may need to support several profiles, this paper makes several propositions towards this end without removing the benefits brought by each profile.

References

[1] IEEE 1588-2008 "Precision Time Protocol"

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

[3] IEEE 802.1AS-2011 "Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks".

[4] IEC 62439-3-2012 "High-availability Automation Networks, PRP and HSR, Annex A.

TABLE I. SYNOPTICAL CHART

	1588 default profile Annex F & J.4	C37.238	802.1AS
Profile identifier	00-1B-19-00-02-00	1C-12-9D-00-00-00	00-80-C2-00-01-00
Media	Ethernet	Ethernet (IEC 61850)	Ethernet, (Radio, EPON,CSN)
Topology	full-duplex 802.3	full-duplex 802.3	full-duplex links
Clocks	boundary clock and transparent clocks	ordinary and transparent clocks (boundary to other protocols)	Boundary-transparent clocks
Path delay	peer-to-peer and end-to-end	peer-to-peer (mandatory) optional for edge clock	peer-to-peer
Steps	one step and two step	one step and two step	two steps mandatory
Layers	layer 2	layer 2	layer 2
Clock domains	multiple, default = 0	one, default = 0	one, default 0
VLAN	not specified	mandatory, default VID = 0	prohibited
Multicast address	01-1B-19-00-00-00 (all but pDelay) 01-80-C2-00-00-0E (peer delay)	01-1B-19-00-00-00 (all but Pdelay) 01–80–C2–00–00–0E (pDelay)	For all messages 01-80-C2-00-00-0E
Ethertype	0x88F7	0x88F7	0x88F7
Subtype	0	0	1
TLVs	Optional	local time offset mandatory, organization-specific mandatory	Follow_Up TLV mandatory
Announce period	(1 s ..16 s, default 2 s	1 s (mandatory)	1 s (default)
Sync period	(0.5 s ..2 s, default 1 s	1 s (mandatory)	125 ms (default)
Pdelay period	(1 s ..32 s, default 1 s	1 s (mandatory)	125 ms (default)
Announce TO	2 s..10 s, default 3 s	2 s for preferred back-up, else 3 s	2 × Announce period
τ	default 1 s	1 s	125 ms
BMCA	default and unspecified alternate	BMCA also for slaves in PRP/HSR	“Simplified” BMCA
Backup master	alternate master	BMC algorithm with shorter time-out for preferred master	unspecified
Priority1	128	128 (GM), 255 (slaves)	255 (slave)
Priority2	128	128 (GM), 255 (slaves)	248 (default)
Precision	unspecified	1 µs over 16 transparent clocks	unspecified
Clock precision	100 ppm (grandmaster)	50 ppm (each transparent clock)	100 ppm
Syntonization	foreseen, unspecified	mandatory	mandatory by frequency ratio-
Holdover time	not specified	5s (master-capable clocks, 2µs)	unspecified
Signaling	Mandatory for node manager	unused	mandatory (trace / negotiation)
Management	Optional	mandatory for timeInaccuracy and traceability	mandatory
MIB	not specified	mandatory for MC, 73 + 14 × p objects (p = number of ports)	specified, 40+ 40 × p objects (p = number of ports)
Security	Annex K	not specified	not specified