

[June 2023]

ETHERNOVIA

CMLDS vs. “per Domain” pDelay

IEEE Contribution

Problem Statement

There is currently no obvious indication to the reader on how to invoke the pDelayRequest/Response state machines, if the CMLDS is not provided. Section 11.2.17.1 only talks about CMLDS and does not make any mention of another method, although it is described implicitly in later subsections.

Both 11.2.19.1, 11.2.20.1 start out with “*This state machine is invoked as part of the Common Mean Link Delay Service (CMLDS).*” - luring the reader into interpreting this as an “only”, particularly given what is said in 11.2.17.1 and although it is described implicitly in later subsections.

Problem Statement - Text from AS-2020

- 🔖 11.2.17 Common Mean Link Delay Service (CMLDS)
- 🔖 11.2.18 Common Mean Link Delay Service (CMLDS) global variables
- 🔖 11.2.19 MDPdelayReq state machine
- 🔖 11.2.20 MDPdelayResp state machine
- 🔖 11.2.21 LinkDelayIntervalSetting state machine

11.2.19 MDPdelayReq state machine

11.2.19.1 General

This state machine is invoked as part of the Common Mean Link Delay Service (CMLDS). There is one instance of this state machine for all the domains (per port). As a result, there also is one instance of each of the state machine variables of 11.2.19.2, state machine functions of 11.2.19.3, and relevant global variables of 10.2.5, 11.2.13, and 11.2.18 for all the domains (per port). None of the variables used or functions invoked in this state machine are specific to a single domain. However, the single instances of all of these objects or entities are accessible to all the domains.

11.2.17 Common Mean Link Delay Service (CMLDS)

11.2.17.1 General

Each port of a time-aware system invokes a single instance of the MDPdelayReq state machine (see 11.2.19) and the MDPdelayResp state machine (see 11.2.20). If the time-aware system implements more than one domain, these two state machines shall provide a Common Mean Link Delay Service (CMLDS), as described in this subclause, that measures mean propagation delay on the PTP Link attached to the port and the neighbor rate ratio for the port (i.e., the ratio of the frequency of the LocalClock entity of the time-aware system at the other end of the PTP Link attached to this port, to the frequency of the LocalClock entity of this time-aware system). The CMLDS makes the mean propagation delay and neighbor rate ratio available to all active domains. If the time-aware system implements one domain (the domainNumber of this domain is 0; see 8.1), these two state machines may provide the CMLDS; however, if they do not provide the CMLDS (i.e., if only the PTP Instance-specific peer delay mechanism is provided), they shall be invoked on domain 0. In other words, if the domain number is not 0, portDS.delayMechanism (see Table 14-8 in 14.8.5) must not be P2P.

11.2.20 MDPdelayResp state machine

11.2.20.1 General

This state machine is invoked as part of the Common Mean Link Delay Service (CMLDS). There is one instance of this state machine for all the domains (per port). As a result, there also is one instance of each of the state machine variables of 11.2.20.2, state machine functions of 11.2.20.3, and relevant global variables of 10.2.5, 11.2.13, and 11.2.18 for all the domains (per port). None of the variables used or functions invoked in this state machine are specific to a single domain. However, the single instances of all of these objects or entities are accessible to all the domains.

Reasoning in [AS]:11.1.2

Since the propagation delay measurement is made using timestamps relative to the LocalClock entities at each port at the ends of the PTP Link and the resulting mean delay is expressed in the responder timebase (see 11.2.19.3.4), there is no need to measure the mean delay for the PTP Link in each domain because the mean delay is the same in each domain. In addition, the quantity neighborRateRatio (see 10.2.5.7) is the ratio of the responder to requester LocalClock frequency and is also the same in all domains. Therefore, the propagation delay and neighborRateRatio measurements are domain-independent. Single instances of the

10.2.5.7 neighborRateRatio: The measured ratio of the frequency of the LocalClock entity of the time-aware system at the other end of the link attached to this port, to the frequency of the LocalClock entity of this time-aware system. The data type for neighborRateRatio is Float64. **There is one instance of this variable for all the domains**, i.e., all the PTP Instances (per port). The variable is accessible by all the domains.

10.2.5.8 meanLinkDelay: The measured mean propagation delay (see 8.3) on the link attached to this port, relative to the LocalClock entity of the time-aware system at the other end of the link (i.e., expressed in the time base of the time-aware system at the other end of the link). The data type for meanLinkDelay is UScaledNs. **There is one instance of this variable for all the domains**, i.e., all the PTP Instances (per port). The variable is accessible by all the domains.

NOTE—The variable meanLinkDelay was named neighborPropDelay in the 2011 edition of this standard.

Conclusion:

The present standard assumes there to **always** be only **one single instance of the peer-to-peer delay mechanism** per port independent of how many Domains/Instances are active in the time-aware system.

One single instance of the peer-to-peer delay mechanism - Always in AS-2020!

- either the **CMLDS is provided**, which is mandatory if multiple PTP instances (Domains) are present ([AS]:11.2.17.1 "... If the time-aware system **implements more than one domain**, these two state machines **shall provide a Common Mean Link Delay Service (CMLDS)**, ...")
- or there is only one single PTP instance, which then must be Domain 0 ([AS]:8.1 "A time-aware system **shall support the domain whose domain number is 0**, and that domain number shall not be changed to a nonzero value."): [AS]:14.8.5 "... **If the domain number is not 0, portDS.delay mechanism must not be P2P** ..." i.e. the pDelay mechanism must not be per domain on any domain other than Domain 0.

P802.1ASdm/D1.1 changes on Domain 0

P802.1ASdm/D1.1 on page 33 changes (section 8.1):

19 A time-aware system shall support one or more domains, each with a distinct **domain**
20 **number**domainNumber in the range 0 through 127. ~~A time-aware system shall support the domain whose~~
21 ~~domain number is 0, and that domain number shall not be changed to a nonzero value.~~ Unless otherwise
22 specified in this standard, the operation of gPTP and the timescale in any given domain is independent of
23 operation in any other domain.

But no change on page 105 of P802.1ASdm/D1.1:

42
43 **14.8.5 delayMechanism**
44
45 | The value of delayMechanism indicates the mechanism for measuring mean propagation delay and neighbor
46 rate ratio on the link attached to this PTP Port and is taken from the enumeration in ~~Table 14-8~~Table 14-8. If
47 the ~~domain number~~domainNumber is not 0, portDS.delay mechanism must not be P2P (see 11.2.17).

Delay Asymmetry is per Domain?

Why is there “... one instance of this variable for each domain that uses the instance-specific peer-to-peer delay mechanism.”, if only one single instance (who’s domainNumber=0?) can run the peer-to-peer delay mechanism?

This seems inconsistent with everything stated elsewhere.

10.2.5.9 delayAsymmetry: The asymmetry in the propagation delay on the link attached to this port. If propagation delay asymmetry is not modeled, then delayAsymmetry is zero. The data type for delayAsymmetry is ScaledNs. There is one instance of this variable for CMLDS (see 11.2.17), and there is also one instance of this variable for each domain that uses the instance-specific peer-to-peer delay mechanism. The instance of this variable for CMLDS is relative to the local clock. The instance of this variable for a domain that uses the instance-specific peer-to-peer delay mechanism is relative to the grandmaser time base.

from AS-2020

(the change in units is not considered relevant here!)

Only one peer-to-peer delay mechanism is needed.

The statement from 11.1.2:

“Since the propagation delay measurement is made using timestamps relative to the LocalClock entities at each port at the ends of the PTP Link and the resulting mean delay is expressed in the responder timebase (see 11.2.19.3.4), there is no need to measure the mean delay for the PTP Link in each domain because the mean delay is the same in each domain.

In addition, the quantity neighborRateRatio (see 10.2.5.7) is the ratio of the responder to requester LocalClock frequency and is also the same in all domains.

Therefore, the propagation delay and neighborRateRatio measurements are domain-independent. *Single instances of the respective state machines that cause these measurements to be made are invoked, rather than one instance per domain, and the results are available to all domains.”*

Is fundamentally true! And also applies to delayAsymmetry (11.2.5.9).

The standard needs **only ONE SINGLE way** to do perform the **peer-to-peer delay mechanism!**

Why are there currently two (2)?

The Rationale for CMLDS

[1588]:7.1.4 Assigning values for the domainNumber and sdoId for a domain

The sdoId attribute is the primary mechanism for providing isolation of PTP Instances of Profiles specified by a QSDO from other PTP Instances operating under a PTP Profile specified by a different QSDO (see 16.5). The value of the sdoId attribute is represented in the common header (see 13.3).

[AS-2020]:11.2.17.1 CMLDS-General

NOTE 1—In the above sentence, the condition that the time-aware system implements only one domain implicitly assumes that IEEE 802.1AS is the only PTP profile present on the respective port of the time-aware system, i.e., no other PTP profiles are implemented on that port. If other PTP profiles that use the CMLDS are present on the port, the CMLDS must be provided.

The value of majorSdoId for the CMLDS shall be 0x2. The value of minorSdoId for the Common Mean Link Delay Service shall be 0x00. As a result, the value of sdoId for the Common Mean Link Delay Service is 0x200.

NOTE 2—The above requirements for majorSdoId and minorSdoId are for the CMLDS. The requirements for gPTP domains, including instance-specific peer delay messages, are given in 8.1.

Support other QSDO's (PTP) Profiles on the time-aware system and still run only one single instance of the peer-to-peer delay mechanism.

CMLDS compatibility mode

[AS-2020]:11.2.17.1 CMLDS-General

If a PTP Port that invokes the CMLDS receives a Pdelay_Req message with majorSdoId value of 0x1, minorSdoId value of 0x00, and domainNumber value of 0, the PTP Port shall respond with PTP Instance-specific peer delay messages (i.e., the Pdelay_Resp and Pdelay_Resp_Follow_Up corresponding to this Pdelay_Req) using the instance-specific peer-to-peer delay mechanism. These instance-specific messages have majorSdoId value of 0x1, minorSdoId value of 0x00, and domainNumber value of 0.

So a “CMLDS Port” can run in an “*instance-specific*” mode and still make the resulting values (*neighborRateRatio* and *meanLinkDelay*) available to multiple Domains or multiple other Profiles on the time-aware system!?

How to bring them back together?

Given

- that with AS-2020 all pDelay messages are transmitted with domainNumber=0 (instance-specific and CMLDS),
- CMLDS uses majorSdold=0x2, instead of 0x1 for instance-specific,
- the differences in calculations described in [AS]:11.2.17.2 between CMLDS and instance-specific,
- that ASdm will remove the requirement for Domain 0 to be present, and
- the desire to keep things backwards compatible ...

Proposed Resolution

We can introduce a single peer-to-peer delay mechanism, which can operate in

- "transportSpecific" (Sdold=0x1) mode, or in
- "CMLDS" (Sdold=0x2) mode.

Separated through:

Table 14-8—delayMechanism enumeration

Delay mechanism	Value	Specification
P2P	02	The PTP Port uses the transportSpecific peer-to-peer delay mode
COMMON_P2P	03	The PTP Port uses the CMLDS peer-to-peer delay mode

Proposed Resolution - Header Fields

All **pDelay Messages shall use domainNumber = 0** in their header!
(Make sure this gets expressed properly in section 8.1.)

Even if Domain 0 is not present in the time-aware system!

- The value of majorSdold of pDelay messages in transportSpecific mode shall be set to 0x1.
- The value of minorSdold of pDelay messages in transportSpecific mode shall be set to 0x00.
- As a result, **the value of sdold of pDelay messages in transportSpecific mode is 0x100.**

- The value of majorSdold of pDelay messages in CMLDS mode shall be set to 0x2.
- The value of minorSdold of pDelay messages in CMLDS mode shall be set to 0x00.
- As a result, **the value of sdold of pDelay messages in CMLDS mode is 0x200.**

pDelay messages with a domainNumber NOT equal to 0 shall be discarded silently!

Backward Compatibility

change [AS-2020]:11.2.17.1

A PTP Port receiving pDealy messages with sdold = 0x100 (majorSdold value of 0x1, minorSdold value of 0x00, and domainNumber value of 0x0) shall always, independent of its set mode, respond in transportSpecific mode, with pDelay messages (i.e., the Pdelay_Resp and Pdelay_Resp_Follow_Up corresponding to a received Pdelay_Req) using the transportSpecific header fields of: majorSdold value of 0x1, minorSdold value of 0x00, and domainNumber value of 0 and perform any further calculations ([AS]:11.2.17.2) in transportSpecific mode.

Further Actions

It the WG agrees with this proposal a vehicle and methodology to make further changes would be needed!

e.g.:

- Verify the delayAsymmetry change of units and ensure there is only one instance per Port.



ETHERNOVIA

Thank you for your attention!