Aspects of Multiple Domains in 802.1AS-Rev and Backward Compatibility

Geoffrey M. Garner
Consultant

IEEE 802.1 TSN TG
2015.11.09

gmgarner@alum.mit.edu

Introduction and Background

- This presentation provides background information on multiple domains, and how they are modeled and used in 802.1AS-Rev/D2.0
  - This material was requested in the 802.1AS-Rev call of November 2, 2015
- 802.1AS-2011 allows for only one gPTP domain, namely domain 0
  - A time-aware system compliant with 802.1AS-2011 ignores gPTP messages that carry a domain number other than 0
- 802.1AS-Rev has introduced the possibility of more than one domain, for two main purposes
  - Some applications require more than one domain (e.g., industrial networks need a universal time domain for an entire factory and a working clock domain for a group of machines that work together
  - The next edition of IEEE 1588 is expected to allow for redundant paths and redundant GMs, with each redundant synchronization spanning tree in a different domain; IEEE 802.1AS-Rev would like to use this capability as a number of new TSN applications will require redundancy and fault tolerance
- In IEEE 802.1AS-Rev
  - It is required that domain 0 be supported, though it may not always be active
  - Supporting additional domains, with domain numbers other than 0, is optional

Multiple Domains and Pdelay

- In IEEE 1588 – 2011 (and in the latest draft of the next edition of 1588), in a network of BCs and OCs each domain is a separate instance of PTP
  - This means that the peer delay process would run independently on each domain
- In 802.AS, timestamps are taken relative to the free-running, LocalClock entity
  - Mean propagation delay is computed in the timebase of the peer delay responder, and then converted to the GM timebase using the cumulative rate ratio
  - This means that the initial computation in the timebase of the peer delay responder will be the same in every domain, and it is wasteful to run the computation in every domain
- Based on the above, the TSN TG decided that the peer delay process would run outside of any gPTP domain, but that the messages will carry a domain number of 0
- This allows for backward compatibility with 802.1AS-2011, since there only domainNumber 0 is recognized in received messages and inserted in transmitted messages

Enabling of Ports and Domains – 1

- It was decided (as a result of the D1.0 comment resolution), to introduce a variable domainEnabled
  - When set to TRUE, the domain is enabled on all ports for which portOper and ptpPortEnabled are TRUE
  - This variable must still be introduced in the state machines (see item (4) at the bottom of p.(iii)), but its purpose is to allow the entire domain to be brought up and running at once (e.g., to avoid multiple successive reconfigurations that would occur if the ports of the domain were enabled one by one)
- portOper is a per-port global variable
  - Since portOper indicates whether the port is up and able to send and receive any messages (gPTP or otherwise), it should not be per domain (i.e., it has nothing to do with gPTP domains)
    - Note that one comment against D2.0 indicates portOper should be equated to MAC_Operational
  - It is stated in the definition of portOper that the term *port* in this definition is a physical port; it should be further stated that portOper not per gPTP domain and has nothing to do with gPTP domains (it is, in effect, common to all domains)

Enabling of Ports and Domains – 2

- ptpPortEnabled is a per-port, per domain global variable
  - When TRUE, the time synchronization and best master selection functions of the port are enabled in the respective domain; in this case the port will send, receive, and process gPTP messages on the respective domain provided portOper is TRUE and domainEnabled is TRUE for that domain
- With the decision that the peer delay process will be outside of any gPTP domain, none of the per domain ptpPortEnabled variables will apply to the peer delay process; note that the MDPdelayReq and MDPdelayResp state machines are now domain-independent
  - This means that a new mechanism is needed to enable the peer delay process; see the Editor's Note in 11.2.14.2.6 (p.151):
    - **<<Editor's note: This state machine, and the MDPdelayResp state machine, contain the variable ptpPortEnabled as an entry condition. Since both state machines are now domain independent, but ptpPortEnabled is a per domain variable, this must be changed. One possibility is to introduce a domain-independent version of ptpPortEnabled, for use in these 2 state machines. Note that if this is done, the condition that this variable is TRUE will need to be added to most of the state machines in this document. Another possibility is to remove this condition entirely, which means that if a system supports the gPTP protocol and portOper is TRUE, the peer delay protocol automatically runs on full duplex, point-to-point links. The TSN TG must decide how to handle this. Comments are requested. >>**

Signaling gPTP Protocol Capability to a Neighbor - 1

- It was decided, in the D1.0 comment resolution, that a port will signal to its neighbor that it is enabled for gPTP on a particular domain using a Signaling message that will carry a TLV
  - Perhaps 'gPTP Protocol Capability' is the wrong term; possibly we should say 'Signaling that a port is enabled for gPTP on a particular domain'
  - The TLV is new in D2.0; it is the 'gPTP capable TLV'
- Two new state machines are introduced
  - gPtpCapableTransmit state machine (10.4.1)
  - gPtpCapableReceive state machine (10.4.3)
- A new per-port, per-domain global variable is introduced
  - neighborGptpCapable (indicates whether the neighbor is capable of invoking the gPTP protocol on this port in this domain)
    - Actually, it means the gPTP is enabled at the neighbor port in this domain, since portOper, ptpPortEnabled, and domainEnabled must be TRUE to transmit the TLV

Signaling gPTP Protocol Capability to a Neighbor - 2

- The gPTP capable TLV has the following fields:
  - tlvType, lengthField, organizationId – these are usual TLV fields (we need not discuss them here)
  - logGptpCapableMessageInterval – log to base 2 of the mean interval, in s, between transmissions of successive Signaling messages carrying this TLV
    - This is needed so that the recipient knows how often to expect the TLV; if the TLV is not received after gPtpCapableReceiptTimeout number of these intervals, the port sets neighborGptpCapable to FALSE
    - Both logGptpCapableMessageInterval and gPtpCapableReceiptTimeout are currently (in D2.0) TBD; need contributions on what their default values and ranges should be
  - Flags – Currently not used; these were included in case they are needed in the future (defining them now will make backward compatibility easier in the future, as it will not be necessary to define a new TLV field.
- In summary:
  - A port knows that its neighbor is capable of and enabled for gPTP on a particular domain if it receives the gPTP capable TLV on that domain

Signaling gPTP Protocol Capability to a Neighbor - 3

- It was considered that the message interval request TLV might be used to signal gPTP capability (i.e., being enabled for gPTP on the respective domain)
  - However, it was realized that, if this were done, the AnnounceIntervalSetting and LinkDelaySyncIntervalSetting state machines would be invoked at regular intervals; this seemed undesirable
  - While it is realized that the interval between sending successive gPTP capable TLVs will be fairly large, it is not known how large, and therefore seemed best to define a separate TLV
- It also was considered not defining a TLV at all, and simply sending a Signaling message. However, IEEE 1588 – 2008 requires that a Signaling message carry a TLV
  - This makes sense, as if the Signaling message had no TLV its meaning and use would not be obvious
  - We would effectively be defining a default meaning and use for the case where it does not carry a TLV
  - In any case, 1588 does not allow this

asCapable – general

- In 802.1AS-2011, the per-port global variable asCapable indicates whether a time aware system and its neighbor can interoperate with each other via the 802.1AS protocol; this means
  - Both are capable of executing the 802.1AS protocol
  - Both are enabled for gPTP
  - Various media-dependent conditions must be satisfied, e.g.,
    - For full-duplex, IEEE 802.3 links, there must be non-802.1AS systems in between the time-aware systems that introduce sufficient impairments that the performance of Annex B cannot be met
    - For 802.11 links, both time-aware systems must at least support timing measurement (they may also support fine timing measurement)
    - For 802.3 EPON links, the default value of asCapable is TRUE
- In 802.1AS-Rev, since the Pdelay mechanism is domain-independent, and since the Pdelay state machines set asCapable, a media-dependent but domain-independent instance of asCapable is needed
  - This is set by the Pdelay state machines, and it must be TRUE in order for the domain-dependent instance of asCapble to be TRUE

asCapable – full-duplex 802.3 links - 1

- For a full-duplex, IEEE 802.3 link, the per-port, per domain instance of asCapable shall be set to TRUE if and only if the following conditions hold:
  - The value of the per-port instance of asCapable that is common across all domains is TRUE, and
  - One of the following conditions holds:
    - The value of neighborGptpCapable for this port is TRUE and the value of domainNumber is not zero, or
    - The value of domainNumber is zero
- But, see the Note and Editor's Note from p. 140 (reproduced on the next 2 slides)

asCapable - full-duplex 802.3 links - 2

- Note from p.140:
  - NOTE - Condition (b)(2), and the fact that the domainNumber is 0 in the peer delay messages (see 11.4.2.3) ensures backward compatibility with the 2011 edition of this standard. A time-aware system that is compliant with the 2011 edition of this standard will interpret the peer delay messages as being sent on domain zero, and will process these messages since domain 0 is the only valid domain in the 2011 version of this standard. A time-aware system compliant with the current edition of this standard that is attached, via a full-duplex, point-to-point link, to a node compliant with the 2011 edition of this standard will not receive Signaling messages that contain the gPTP capable TLV and will not set neighborGptpCapable to TRUE; however, condition (b)(2) ensures that asCapable for this port and domain (i.e., domain 0) will still be set to TRUE if condition (1) holds.

asCapable – full-duplex 802.3 links - 3

- Editor's Note from p.140 :
  - **<<Editor's note: It is possible that asCapable for a particular port and domain 0 could be set to TRUE if the neighbor port invokes at least one other domain in addition to domain 0, but domain 0 happens to be inactive at the time (e.g., ptpPortEnabled is FALSE on domain 0 at the neighbor). In this case, the neighbor still runs the peer delay mechanism, because in the new edition of this standard the peer delay mechanism uses 0 as the domain number and the results are accessible to all domains. However, if the port attempts to send gPTP messages other than peer delay messages to the neighbor, there will be no response because gPTP is not active at that port on domain 0. The TSN TG should decide if the above is acceptable behavior and, if not, what the alternative should be. Comments are requested on this.>>**

asCapable – 802.11 links - 1

- In the case of 802.11 links there is no per-port, domain-independent instance of asCapable
- Instead, there is a per-port global variable tmFtmSupport that indicates whether (a) both ends of the link support both TM and FTM (tmFtmSupport = 2), (b) both ends support TM but at least one does not support FTM (tmFtmSupport = 1), or (c) at least one end does not support TM (tmFtmSupport = 0)
  - Note that it is required that if FTM is supported, TM must also be supported
- Details are on the next slide

asCapable – 802.11 links – 2

12.3 Determination of timing measurement and fine timing measurement capability

- The per-port global variable tmFtmSupport shall be set as follows:

a) If:

1) the port supports both timing measurement and fine timing measurement, and
2) the timing measurement bit in the Extended Capabilities information element defined in Table 8-103 of IEEE Std 802.11-2012 indicates that the peer 802.11 station is capable of participating in the timing measurement protocol, and
3) the *fine timing measurement Responder and Initiator bits in the Extended Capabilities information element defined in Table 8-132 of Draft P802.11REVmc D4.2 indicate that the peer IEEE 802.11 station is capable of participating in the fine timing measurement protocol,*

tmFtmSupport is 2;

b) Else if:

1) the port supports timing measurement, and
2) the timing measurement bit in the Extended Capabilities information element defined in Table 8-103 of IEEE Std 802.11-2012 indicates that the peer IEEE 802.11 station is capable of participating in the timing measurement protocol

tmFtmSupport is 1;

c) Else tmFtmSupport is 0.

asCapable – 802.11 links – 3

- **<<Editor's note: The intent of the above requirement is that tmFtmSupport = 2 indicates that the ports at both ends of the link support TM and FTM; tmFtmSupport = 1 indicates that the ports at both ends of the link support TM, but at least one of the ports does not support FTM; and tmFtmSupport = 0 indicates the neither port supports TM (and, in this last case, it does not matter whether one or both ports support FTM because it was decided that supporting TM is manditory for asCapable to be TRUE (the value of tmFtmSupport is used in 12.4 as part of the determination of asCapable)).>>**
- **<<Editor's note: Draft P802.11-REVmc_D4.2 is the lastest draft of this document as of D2.0 of the current document. This reference will be updated to an approved, published document prior to sponsor ballot.>>**

asCapable – 802.11 links - 4

- For an 802.11 link, the per-port, per domain instance of asCapable shall be set to TRUE if the following conditions hold:
  - The value of tmFtmSupport is not zero, and
  - neighborGptpCapable is TRUE and the value of domainNumber is not zero
- If the value of domainNumber is zero and the value of tmFtmSupporrt is not zero, asCapable may be set to TRUE
- In all other instances, asCapable shall be set to FALSE

asCapable – 802.11 links - 5

- Note and Editor's note from pp. 177 – 178:
  - NOTE - The above conditions ensure backward compatibility with the 2011 edition of this standard. A time-aware system that is compliant with the 2011 edition of this standard will not process the gPTP capable TLV, and asCapable will be determined as specified in the 2011 edition. A time-aware system compliant with the current edition of this standard that is attached, via an 802.11 link, to a node compliant with the 2011 edition of this standard will not receive Signaling messages that contain the gPTP capable TLV and will not set neighborGptpCapable to TRUE; however, condition (b)(2) ensures that asCapable for this port and domain (i.e., domain 0) will still be set in a manner consistent with that of the 2011 edition of this standard.
  - **<<Editor's note: The TSN TG should review the above to determine whether this is the desired behavior in determining asCapable for an 802.11 link. Comments are requested.>>**

asCapable – 802.3 EPON links - 1

- In the case of EPON links, there is no per-port, domain-independent instance of asCapable
- Instead, the EPON link always invokes MPCP synchronization protocol, and asCapable is TRUE if gPTP is enabled
- For EPON links:
  - For domain 0, the default value of the per-port, per domain instance of the global variable asCapable shall be TRUE
  - For domains other than domain 0, the per-port, per domain instance of asCapable shall be set to TRUE if the value of neighborGptpCapable is TRUE
  - (Note that the wording in 802.1AS-Rev/D2.0 is incorrect; it has been corrected here)

asCapable – 802.3 EPON links - 2

- Note and Editor's Note on pp. 202 – 203:
  - NOTE - The above conditions ensure backward compatibility with the 2011 edition of this standard. A time-aware system that is compliant with the 2011 edition of this standard will not process the gPTP capable TLV, and asCapable will be determined as specified in the 2011 edition. A time-aware system compliant with the current edition of this standard that is attached, via an 802.11 link, to a node compliant with the 2011 edition of this standard will not receive Signaling messages that contain the gPTP capable TLV and will not set neighborGptpCapable to TRUE; however, the ensures that asCapable for this port and domain (i.e., domain 0) will still be set in a manner consistent with that of the 2011 edition of this standard.
  - **<<Editor's note: The TSN TG should review the above to determine whether this is the desired behavior in determining asCapable for an 802.11 link. Comments are requested.>>**

Thank you