New outline for the next version (V1.5) of IEEE P802.1DG
Text contribution M. Turner 2023-3-14

1. Overview

2. Normative References

3. Definitions

4. Abbreviations

5. Conformance Modules

5.1. Conformance Module Ingress Selection

5.2. Conformance Module ATS, CBS

Exclude combinations with TAS

5.3. Conformance Module Time Synchronization

Refers to Annex with new Automotive Time Synchronization Profile (Profile is based on Autosar and IEEE 802.1AS, it is NOT a 1588 Profile and will not be called such!)

5.4. Conformance Module Time Aware Shaper

Exclude combinations with CBS & ATS

5.4.1. Bus Mode requirements

5.4.2. Phased Mode requirements

5.4.3. CQF Requirements

6. Automotive in Vehicle Networks

6.1. Device Models

6.2. End Station

6.2.1. Middleware

Applications transfer their structured data to the Middleware for serialization and transmission to the network.

6.2.1.1. Serialization

In order to bridge different memory representations (e.g. endianness), the structured data provided by Applications to the Middleware is serialized before being transmitted over the network. Autosar [AR] specifies two different methodologies. Adding tags and length information (TLV) can significantly increase the number of Bytes which have to be transported over the network in comparison, to what the Application's Structured Data occupies in memory.

Serialized Structured Data is referred to as a Message.

6.2.1.2. Middleware trigger conditions (nPDU feature)

Multiple Messages from one or more Applications can be combined into a single Payload, separated by a Middleware Header (MW-Header). For short Messages the MW-Header can add a significant number of Bytes which need to be transported over the network.

The Middleware provides a buffer where these Messages are stored until the trigger condition to transmit the Frame is fulfilled. While in the buffer Middleware may overwrite a specific buffered message with new data provided by the Application (last-is-best).

The arrival of a certain Message may trigger the transmission of the Frame, independent of the availability of other Messages, potentially leading to differences in Frame size.

The time a message has spent in the Middleware's buffer can also serve as a trigger condition.

The Middleware can serve as a rather granular shaping mechanism, before IP-Packets are put into Frames by the MAC layer.

6.2.1.3. MAC Layer Priority setting

The MAC Layer inserts the PCP into the VLAN-Tag. This can happen at least per VLAN-ID, but usually the upper layers (TCP/IP Stack and or Middleware) can pass per Frame priority information for the MAC Layer to encode into the PCP.

6.2.1.4. Network Shaping

Once the MAC Layer have formed a Frame, mechanisms like Linux's HTB [Linux-HTB] or CBS [Linux-CBS] can perform shaping before the Frames are transmitted onto the Network. The granularity of shaped aggregates is limited by Layer 2, 3, or 4 information, i.e., Middleware information can no longer be considered. If Messages from different Applications have been placed into the same Frame, they are now handled as one.

6.3. Bridge

6.3.1. MAC Relay ([Q] 8.1.1)

The MAC Relay Model as described in IEEE Std 802.1Q-2022 is not intended to be a software implementation specification like Autosar's SWSs. This present standard is intended as a coupling element between the purely abstract model and the actual implementation in hardware and/or software.

Frames shall be processed by the model's elements in the order depicted in the Forwarding Model. The exact process of forwarding between the elements is not specified.

6.3.2. Transceiver Port

Any Bride Port is used for ingress and egress. The Forwarding Process is shown from an ingress port to an egress port, which shall be different. Nonetheless all ports implement ingress and egress functionality.

6.3.3. Bridge Management

The management entity is semi-statically configured.

Dynamic re-configuration can be performed using a well established automotive protocol (SOME/IP(-SD), UDS over DoIP, XCP, ...). The use of non-automotive configuration protocols (LLDP, NetConf, SRP, SNMP, ...) is discouraged but can still be selected by the Implementer.

The LLC Entity or Entities that support Higher Layer Entities from [Q] 8.3.b) are not required by this present standard.

7. Management Port ([Q] Figure 8-8)

8. Frame corruption during physical layer transport ([Q] 6.5.2.a))

The CRC of IEEE Std 802.3 has a Hamming Distance of 4 for frames larger than about 2900 bit.

P. Koopman, "32-bit cyclic redundancy codes for Internet applications," in *Proceedings International Conference on Dependable Systems and Networks*, Washington, DC, USA, 2002, pp. 459–468. doi: 10.1109/DSN.2002.1028931.

9. Support of the EISS ([Q] 6.9)

As part of Frame reception ([Q] 8.5) the EISS shall be supported by tagging functions that in turn use the ISS (IEEE Std 802.1AC). The EISS shall be supported by using the C-VLAN tag (C-TAG), using the 802.1QTagType of [Q] Table 9-1 as the TPID ([Q] 9.4).

The EISS shall implement the Port-and-Protocol-based VLAN classification per [Q] 6.12.

The ingress VID translation table as specified by [Q] 6.9 f) shall be supported.

Upon receipt of an M_UNITDATA.indication primitive from the ISS, the value of the vlan_identifier parameter in the EM_UNITDATA.indication primitive at the EISS shall be set as follows:

a) The value contained in the VID field, translated using the VID translation table, if the frame is C-VLAN-tagged - OR

b) As determined by Port-and-Protocol-based VLAN classification ([Q] 6.12)

if the frame is untagged or Priority-tagged or if the translated VID value (relay VID) is zero.

The default values (8P0D) specify a PCP value equal to the priority value and do not support communication of drop precedence in the PCP field.

10. (Ingress) Stream Identification Functions(s) ([CB] 6.9)

Only passive out-facing stream identification functions shall be configured. Please refer to the informative Annex on Stream Identification for a rationale.

The Stream Identification functions operate on the EM_UNITDATA.indication parameters of the EISS.

All frames passing through the stream identification functions in the UP direction shall be assigned a stream_handle. A stream_handle of value Null (0x00) shall indicate no matching stream identification function was found.

10.1. Number of Stream Identification Functions

Number of ports (*n*)	**Minimum number of Identification functions (*m*)**
less than 5 (2<n<5)	128
more than 4, but less than 10 (4<n<10)	192
more than 9 (n>9)	256

10.2. Identification Mask Length

All filter masks shall be long enough to cover:

– MAC addressing

– one VLAN tag

– the addresses, DSCP, and next header field of an IPv6 and IPv4 header

– the port information of TCP and UDP

11. Active topology enforcement ([Q] 8.6.1)

Learning shall be TRUE for all reception Ports. If ingress filtering ([Q] 8.6.2) would not cause the received frame to be discarded, the source address and VID are submitted to the Learning Process.

Loop-free network configuration shall be ensured through configuration, the Rapid Spanning Tree Algorithm and Protocol (RSTP) is not required and its use is discouraged (see Annex for rationale)

12. Ingress filtering ([Q] 8.6.2)

A frame received on a Port that is not in the member set ([Q] 8.8.10) associated with the frame's VID shall be discarded. This is known as ingress VLAN-Membership check.

Note: The VID may have been altered by the EISS.

13. Frame filtering ([Q] 8.6.3)

The Forwarding Process takes filtering decisions, i.e., reduces the set of potential transmission Ports ([Q] 8.6.1), for each received frame on the basis of:

a) Destination MAC address
b) VID (Egress VLAN-Membership check)
c) The information contained in the FDB for that MAC address and VID
d) The default Group filtering behavior for the potential transmission Port ([Q] 8.8.6)

in accordance with the definition of the FDB entry types ([Q] 8.8.1, [Q] 8.8.3, and [Q] 8.8.4).

13.1. Forwarding Database

The Forwarding Database shall conform to the requirements of OPEN Alliance [Open-TC11]

14. Egress filtering ([Q] 8.6.4)

Any Port that is not in the member set ([Q] 8.8.10) for the frame's VID is removed from the set of potential egress Ports.

This concludes the egress Port selection process. Fan-out is however not yet performed.

15. Stream filtering ([Q] 8.6.5.3)

A stream filter is identified by

a) A stream_handle AND
b) a priority value.

Either one or both can be set to a wildcard value that matches any value.

Frames assigned to the same Stream Filter are processed together for Max. SD Size Filtering, Stream Gating, Flow Metering, and ATS eligibility time assignment.

A Frame with multiple egress Ports will only pass through a single Stream Filter.

16. Maximum SDU Size Filtering ([Q] 8.6.5.3.1)

All Stream Filter shall have Max. SDU Size Filtering configured.

If the Frames are intended to go though a CBS, a configuration is suggested in the “Credit Based Shaper” section.

If the Frames are intended to go though an ATS, a configuration is suggested in the “Asynchronous Traffic Shaper” section.

17. Stream Gating ([Q] 8.6.5.4)

Stream Gating may be disabled for certain stream filters.

All Stream Gates shall always be in state OPEN.

An IPV value of Null (0) shall cause the received frame’s priority parameter to be used as the IPV.

If CQF is not supported a static (non time dependent) IPV assignment shall be supported.

If CQF is supported a time dependent IPV assignment shall be supported.

If CQF is supported a minimum of three (3) IPV assignment slots shall be supported.

18. Flow metering ([Q] 8.6.5.5)

All Stream Filter shall have a Flow Meter configured.

The MEF 10.3 algorithm shall be supported with:

- A configurable Committed information rate (CIR) > 0
- A configurable Committed burst size (CBS) > 0
- Excess Information Rate (EIR) = 0
- Excess burst size (EBS) = 0
- Coupling flag (CF) = False (0)
- Color mode (CM) = color-blind

This represents a Single Rate Two Color Meter. Frames are either

- Permitted to pass (green) OR
- Dropped (red)

If the Frames are intended to go though a CBS, a configuration is suggested in the “Credit Based Shaper” section.

If the Frames are intended to go though an ATS, a configuration is suggested in the “Asynchronous Traffic Shaper” section.

19. ATS Eligibility Time Assignment ([Q] 8.6.5.6)

An certain ATS instance may be configured for one or more certain stream filters.

Frames with different stream_handle may run through the same ATS instance.

A configuration is suggested in the “Asynchronous Traffic Shaper” section.

20. (Egress) Stream Identification Functions(s) ([CB] 6.9)

Egress Stream Identification is performed per egress Port. An ingress Frame with multiple egress Ports must go through fan-out before Egress Stream Identification.

Only out-facing stream identification functions shall be configured.

Since only passive Stream Identification Functions are supported no action is performed here.

Please refer to the informative Annex on Stream Identification for a rationale.

21. Egress VID translation table ([Q] 6.9 g))

Egress VID translation is performed per egress Port. An ingress Frame with multiple egress Ports must go through fan-out before the Egress VID translation.

The egress VID translation table as specified by [Q] 6.9 g) shall be supported.

22. Queuing frames ([Q] 8.6.6)

Queuing frames is performed per egress Port.

The default mapping of the Frame's Priority Code Point to Traffic Class Queue shall be according to [Q] Table 34-1.

Each frame is mapped to a traffic class using the Traffic Class Table for the Port. The priority value used for this mapping is determined as follows:

a) If stream gates are unsupported ([Q] 8.6.5.4), the frame's priority is used.
b) If stream gates are supported and the IPV specification assigned to the frame is the null value, the frame's priority is used.
c) If stream gates are supported and the IPV specification assigned to the frame is an IPV, this IPV is used.

23. Number of Traffic Class Queues

A Relay shall support no less than 8 Traffic Classes per egress port.

Since [Q] does not give any indication on how to handle a Port with more than 8 Traffic Class Queues, it is up to the implementer to follow the concepts laid out in [Q] and this standard in principle.

24. Per TC-Queue Transmission selection ([Q] 8.6.8 a))

The operation of the transmission selection algorithm supported by that queue determines if there is a frame available for transmission or not.

All Traffic Class Queues on all ports shall support the ATS Transmission Selection algorithm of [Q]:8.6.8.5.

At least the 2 numerically highest value Traffic Class Queues (highest priority) on any port shall support the Credit Based Shaper transmission selection of [Q]:8.6.8.2.

25. Port Transmission selection ([Q] 8.6.8 b))

For each Port, frames are selected for transmission on the basis of the traffic classes that the Port supports and the operation of the transmission selection algorithms supported by the corresponding queues on that Port. For a given Port and traffic class, frames are selected from the corresponding queue for transmission if and only if:

a) The operation of the transmission selection algorithm supported by that queue determines that there is a frame available for transmission; and

b) For each queue corresponding to a numerically higher value of traffic class supported by the Port, the operation of the transmission selection algorithm supported by that queue determines that there is no frame available for transmission.

The order in which frames are selected for transmission from the queue shall maintain the ordering requirement specified in [Q] 8.6.6.

26. Limiting Frame lifetime ([Q] 6.5.6)

To enforce the maximum frame lifetime, a Relay may be required to discard frames. Since the information provided by the MAC Sublayer to a Relay does not include the transit delay already experienced by any particular frame, Relays discard frames to enforce a maximum delay in each Relay.

27. Buffer Monitoring

28. Automotive Time Synchronization Profile

28.1. asCapable and asCapableAcrossDomains

asCapable ([AS] 10.2.5.1) and asCapableAcrossDomains ([AS] 11.2.13.12) can be set by the Management Entity

28.2. Announce

Announce Messages ([AS] 10.6.3.1) are not required to be sent.

28.3. Best Master Clock Algorithm (BMCA)

The Best Master Clock Algorithm ([AS] 10.3.1.2) is not used.

29. Credit Based Shaper

29.1. High Credit Monitoring

Indicates that frames have not been transmitted beyond the expected time.

Equivalent to Buffer Monitoring for this Buffer Partition

29.2. Low Credit Monitoring

Is not needed if max SDU size Policing is configured correctly

30. Asynchronous Traffic Shaper

Egress queue level must be monitored

Byte size could be deduced from number of Frames if max SDU size Policing is configured correctly

31. Time Aware Shaper

32. Traffic Classification

Large Frames + High Frequency (5-30µs – depends on line rate and resolution) – Camera (one picture per 30ms) – where to buffer? Last is best

Control: Small frames – per stream low rate (1ms – application driven) – combined results in flows of high frequency last is best

Large Frames + bursty (SotA, Web – Talker and line rate driven) – where to buffer, re-transmit?

Medium Frames low rate (10ms) – combined results in flows of high frequency

Small frames high frequency (20µs=1/48kHz – Audio Frequency!) – audio echo cancellation last is best

Object Lists: Large Frames – Frequency (list size and line rate)?? (one list every 30ms) where to buffer? Last is best

Entertainment Video: (50Mbit/s – Quality driven) – Buffer at source! Re-transmit?

Audio (1.33ms – Audio Frequency aggregated channels – last is best

Very low frequency: NM, SD, - last is best

Diagnostic: (not update) – request response driven – retransmit? When? – Diagnostics timeouts!

Debug: DLT, XCP – last is best? Or buffer at source?

33. Congestion Isolation

The traffic on the in-vehicle network is assumed to be dominated by cyclic communication (see section on In-Vehicle Networks).

[Qdd] 49.2.1 Congesting flow identification An essential step in the process of CI is identifying congesting flows by an Active Queue Management (AQM) scheme that supports Explicit Congestion Notification (ECN) specified in IETF RFC 3168. There are many potential methods of identifying congesting flows and interoperable implementations can exist using different approaches. The CIP Congestion Detection function ([Q] 49.3.1) of the Congestion Isolation Aware Forwarding Process ([Q] 49.3) is responsible for implementing the AQM. This standard defines the CP algorithm ([Q] 30.2.1) for detecting Congestion Controlled Flows (CCFs) in congestion aware bridges. This approach may be used to detect congesting flows in a CI aware system. A number of other possible approaches, including those that support the end-to-end ECN congestion control, are discussed in IETF RFC 7567 [Q][B3].

33.1. Definition of Congestion

The Bride is dropping Frames due to [Q] 6.5.2 Frame loss b)

1) It is unable to transmit the frame within some maximum period of time and, hence, must discard the frame to prevent the maximum frame lifetime ([Q] 6.5.6) from being exceeded. “Limiting Frame Lifetime” – Prevent Buffer Overrun! Also Congestion? - OR

2) It is unable to continue to store the frame due to exhaustion of internal buffering capacity as frames continue to arrive at a rate in excess of that at which they can be transmitted. Buffer overrun! Worst Case Congestion!

A Buffer partition may be congested due to excess ingress or underserved (negative excess) egress.

A Congesting Flow is providing excess ingress – this can be prevented by Policing!

A Congested Flow is underserved on egress – Why?

- Misconfigured Shaper (CBS, ATS, TAS)
- Large lower priority interfering Frame – can be prevented by Policing
- Large higher Priority Burst - can be prevented by Shaping

33.2. Definition of a Congesting Flow

Exceeds the expected ingress rate.

33.3. Definition of a Congested Flow

Can not transmit at the expected rate.

34. Why is Color Mode not supported?

35. Frame discard

1) It is unable to transmit the frame within some maximum period of time and, hence, must discard the frame to prevent the maximum frame lifetime ([Q] 6.5.6) from being exceeded. “Limiting Frame Lifetime” – Prevent Buffer Overrun! Also Congestion?

2) It is unable to continue to store the frame due to exhaustion of internal buffering capacity as frames continue to arrive at a rate in excess of that at which they can be transmitted. Buffer overrun! Worst Case Congestion!

3) The size of the SDU carried by the frame exceeds the maximum supported by the MAC procedures employed on the LAN to which the frame is to be relayed. Ref to IEEE Std 802.3 – size table!

4) Changes in the connected topology of the network necessitate frame discard for a limited period of time to maintain other aspects of QoS ([Q] 13.16).

5) The device attached to the Bridge is not authorized for access to the network. (IEEE Std 802.1X)

6) The configuration of Static Filtering Entries or Static VLAN Registration Entries in the FDB (8.8.1, 8.8.2) disallows the forwarding of frames with particular destination addresses or VLAN classifications on specific Ports.

7) A flow metering algorithm ([Q] 8.6.5) determines that discard is necessary. “Policing”

8) The Bridge supports enhancements for scheduled traffic ([Q] 8.6.8.4) and the size of the service data unit exceeds the value of queueMaxSDU ([Q] 8.6.8.4) for the traffic class queue on which the frame is to be queued.

36. Configuration

Compare AR-XML w/ MIB&YANG

SD snooping

36.1. Why full dynamic reconfiguration is not included?

37. Latency calculations

Comparing IEEE 802.1Q calculations with IEEE 802.1BA and others

38. Informative Annex: Safety

39. Informative Annex: Why is FRER not supported?

40. Informative Annex: Security

41. Informative Annex: Stream Identification

41.1. Why are Active Stream Identification Functions not supported?

An out-facing active Stream Identification function does not check for the header content on the down path, i.e., on egress. In order to ensure modifications are applied to only such frames which carry the intended header, an in-facing passive stream identification function would be required.