Discussion of Comments on the RAP White Paper

IEEE 802.1 TSN TG Call
Oct. 2 2017

Overview of the Comments

A total of 15 comments from Rodney Cummings

- Architecture: [#1 ~ #5]

- Data model: [#6 ~ #7]

- Domain boundary detection: [#8 ~ #11]

- Seamless redundancy: [#12 ~ #14]

- Switchover redundancy: [#15]

Topics for the next version

- L2/L3 collaborated reservation

RAP Architecture – Comment #1

[RC#1] - Page 1 line 11: It is true that LRP will exchange large databases. Therefore, "I want to save a few bytes" is not a valid rationale to revert to the single-PDU design of MSRPv0. It is true that the TLV-based design of MSRPv1 does take a few more bytes for type/length, but that is a small compromise for the flexibility that it delivers. In my opinion, "RAP" is nothing more than SRPv2, using new TLVs for an MSRPv2, where MSRPv2 is designed to take advantage of LRP, and MSRPv2 is an enhancement of the TLV-based design of MSRPv1.

[FC#1]: please refer to my previous email regarding the naming. In short, from perspective of data model, there is nothing wrong with using the term "MSRPv2" to emphasize the close relation in the data model between RAP and MSRP.

[Note]: The MSRP/MRP architecture and its similarity/difference to RAP have been discussed a lot previously in the email (http://www.ieee802.org/1/private/email2/msg26517.html). The current consensus is that **"the creation of a well-designed architecture for the future (i.e. RAP)"** should be the focus.

Oct. 2 2017

IEEE 802.1 TSN Call

RAP Architecture – Comment #2

[RC#2] - Page 2 line 32: This statement is incorrect. For fully-distributed (i.e. no CNC), MSRPv1 supports stream transformation for use of IP headers in talker/listener (see P802.1Qcc/D1.6, 35.2.2.10.5). Fully-distributed MSRPv1 also supports the MaxLatency requirement, which dates back to the "Acceptable Latency" use cases presented in:

http://www.ieee802.org/1/files/public/docs2014/cc-cgunther-acceptable-latency-0314-v01.pdf

MSRPv1 is the evolution of SRP. It did add support for CNC, but that is only one feature.

[FC#2]: It is true that MSRPv1 contains enhancements for both fully-distributed and centralized configuration models, but many of them relating to the TSN features are only provided to the centralized case. That is why I used the term "primarily" in the white paper. Even for the MaxLatency feature, it occurs to me that the MSRP service primitives originally specified in Table 35-8 and 35-9 of 802.1Q-2014 have not yet been extended by Qcc to allow user applications to use this feature.

Oct. 2 2017

IEEE 802.1 TSN Call

RAP Architecture – Comment #3

[RC#3] - Page 6 line 23: In my opinion, this figure shows an incorrect design. Please refer to tsn-cummings-link-local-reg-concepts for additional concepts. MAD is not a block of code, but merely an abstract interface between the MSRP block and the MRP block. This distinction is important, because it means that the compatibility goals of this figure are delivered simply by treating RAP as MSRPv2.

[FC#3]: Based on the current draft of LRP and the results of the ballot comment resolution, I would make the following assumptions:

- Record is used as abstract container and exchanged between LRP and its applications, instead of attribute.

- LRP Apps can pack one or multiple attributes into one record.

- Rules of packing attributes into records are defined by each LRP app.

- LRP synchronizes records by using checksum mechanisms comparable to ISIS with CSNP/PSNP-messages

Oct. 2 2017 IEEE 802.1 TSN Call

RAP Architecture – Comment #3 (cont.)

[FC#3]: The advantages of using record are:

for LRP: (in comparison with MRP)

- completely application-independent

- improved efficiency in support of big database, e.g. with a huge number of attributes

  - reduced data on the wire by using one checksum per record

  - reduced computing power by using checksum for record comparison

for LRP applications:

- Each LRP App can define its own rules of packing its attributes into records to meet its specific optimization goals.

In general, through the use of record, LRP can provide much improved efficiency especially for synchronization of a huge number of small attributes, while MRP always periodically transmits attributes.

Oct. 2 2017 IEEE 802.1 TSN Call

RAP Architecture – Comment #4

[RC#4] - Page 7 line 8: This figure demonstrates that RAP's changes to MSRPv1 are minor, and can therefore be handled easily as MSRPv2. Specifically:

[FC#4]: Yes, from perspective of data model, the changes are minor. But from perspective of encoding, further investigation is needed.

- SRclassID: I'm not sure why you need this, but it can easily be added as a new TLV, or to an enhanced StreamID TLV.

- MinAccumulatedLatency: "Max" AccumulatedLatency is already in MSRPv1. The "Min" was rejected during P802.1Qcc comment resolution, but you are welcome to request that again. Min can be added in a TLV.

- Domain enhancements: These were also proposed in past 802.1 meetings, but with insufficient rationale. Once we reach consensus, they can be added as an MSRPv2 TLV.

[FC#4]: I'll provide more details for the above items (as proposals) in next WP version and suggest to discuss such technical details after the project is started.

Oct. 2 2017

IEEE 802.1 TSN Call

RAP Architecture – Comment #5

[RC#5] - Page 7 line 8: I don't see the equivalent of MVRP and MMRP in this RAP proposal. Do you plan to use MVRP and MMRP on top of LRP? If so, we need the abstract interface of tsn-cummings-link-local-reg-concepts.

[FC#5]: The current focus is only on stream reservation protocol. Does anyone see performance and scalability problems in MVRP and MMRP?

Oct. 2 2017

IEEE 802.1 TSN Call

RAP Data Model – Comment #6

[RC#6] - Page 7 line 8: Your previous proposals insisted on distinct attributes for "dynamic" (fast-changing) and "static" (slow-changing). What happened to that? We may need quantitative rationale for that dynamic/static proposal in order to reach consensus, but it did seem like a reasonable proposal.

[FC#6]: This relates to packing attributes into records, while the discussion of record in LRP was pending at the time of writing the WP. This is a really important issue and I expect joint discussion with LRP.

Oct. 2 2017   IEEE 802.1 TSN Call

RAP Data Model – Comment #7

[RC#7] - Page 7 line 11: This figure is misleading, because it implies a large difference between the requirements of MSRPv1 and the desired data for RAP. The difference is small.

- EndStationInterfaces: It is true that this is specific to CNC, but it is optional (see P802.1Qcc/D1.6, 35.2.2.10.4).

- TSpec_Scheduled (now TSpec_TimeAware): It is true that this is specific to CNC, but it is optional (see P802.1Qcc/D1.6, 35.2.2.10.1).

- FailedBridgeInterfaces: This is FailureInformation of MSRPv0, so it is not new to MSRPv1.

- InterfaceCapabilities and InterfaceConfiguration: These are optional (see P802.1Qcc/D1.6, 35.2.2.10.1). They are used if you want stream transformation for fully-distributed (not specific to CNC).

[FC#7]: Yes, from perspective of data model, the difference is not big (excluding the optional ones intended only for centralized). But from perspective of encoding (organization of RAP attributes), further investigation is needed.

Oct. 2 2017

RAP Domain Boundary Detection – Comment #8&#9

[RC#8] - Page 9 line 11: What is the rationale <u>for negotiating the transmission selection algorithm</u>? This parameter has been proposed, then removed from proposals, then added again. We need to come to a clear consensus as to why it is needed.

[FC#8]: Negotiation of transmission selection algorithm exists already in MSRPv0 (for CBSA) implicitly. I just want to keep it and make it explicit for use of other shapers.

[RC#9] - Page 9 line 22: This is the first proposal to add this. What is the rationale? Why do we need to negotiate less than the max MTU?

[FC#9]: Negotiation of maxMTU (less than the MAX_MTU) can provide further optimization of latency calculation.

Oct. 2 2017

RAP Domain Boundary Detection – Comment #10&#11

[RC#10] - Page 9 line 31: When you say "topology-independent per-hop latency", I assume you mean that the transmission selection can provide a bound to its next hop, but not for the entire topology (at least for CBS and strict). If so, this phrase can be misinterpreted to imply the opposite. I think "bounded latency to the next hop" would be more precise.

[FC#10]: yes, you are correct. I'll use the term you suggested.

[RC#11] - Page 9 line 33: The phrase "network administrator" implies that there is a human who will calculate the bounded latency for the entire topology on a piece of paper. I think we all know that is not the reality. This computation is complex, and it will obviously be performed by a centralized software component. What we are talking about is a CNC without a UNI, and the software investment for that CNC computation will be re-used across fully-centralized and fully-distributed. I think it would be helpful to point this out in informative text, because it helps to clarify that TSN/DetNet's configuration models aren't as radically different as they might seem.

[FC#11]: for SRclass-based stream reservation, if the used transmission mechanisms provide topology-independent and bounded per-hop latency, e.g. Qch, the upper value of the bounded latency for an assumed network diameter (i.e. targetLatency) can be easily calculated.

Oct. 2 2017   IEEE 802.1 TSN Call

RAP for Seamless Redundancy – Comment #12

[RC#12] - Page 11 line 1: ISIS-PCR is great, but as we've discussed in past 802.1 meetings, it will be challenging to mandate implementation of IS-IS in industrial 3-port bridges (i.e. constrained end station with a bridge). Since you are already using a CNC-without-UNI for calculation of "SRclass target maximum latency", and that calculation already requires the physical topology, why not use the same CNC-without-UNI to configure the redundant VLANs with management (i.e. YANG)?

[FC#12]: How trees are established is not the task of RAP, which is the same as MSRP. Either management or control protocol for tree configuration can be used.

RAP for Seamless Redundancy – Comment #13&#14

[RC#13] - Page 12 line 3: How does the Talker know to transmit RAP on redundant VIDs? The same question applies to the bridge in other examples. In other words, let's say that the MST Configuration Table lists ten VIDs... which pair do you use for RAP? I can think of two possible answers:

a) Add the pair of VIDs to RAP's domain, to be negotiated on each link (even if that link uses just one VID). The disadvantage of this answer is that AVB end stations (using MSRPv0) cannot take advantage of redundancy.

b) Add managed objects for the pair of VIDs. This implies CNC-without-UNI.

Maybe we specify both answers, and let profiles decide.

[FC#13]: we might need new managed objects to support this feature.

[RC#14] - Page 12 line 14: For the VID that "this" RAP path is using, where is it in the RAP data? Is it the VID in DataFrameSpecification? If so, that is potentially problematic, because the listener won't see the talker's VID.

[FC#14]: could you provide more details about your concerns.

RAP for Switchover Redundancy – Comment #15

[RC#15] - Page 16 line 5: This "switchover redundancy" is not seamless, but it is "better" than relying on re-configuration of the active topology (i.e. MSTP/IS-IS followed by RAP). The issue is that "better" is always subjective. How much better? Is the better worth the standardization effort? I recommend extensive simulations answer this sort of question, so that we can try to make a quantified decision.

[FC#15]: currently this feature is described in the WP only as a note; more feedback regarding interest in this feature is expected .

L2/L3 Collaborated Reservation (for next white paper version)

q parallel mode or serialized model or both?

P802.1XD
Stream Reservation Protocol v2

WHAT

● A new stream reservation protocol that runs over P802.1CS Link-local Registration Protocol

● Perhaps suitable for processing by routers, label switches, and L2 and L3 end stations

● P802.1XD or P802.1Qxd?? (If suitable for routers, .1XD)

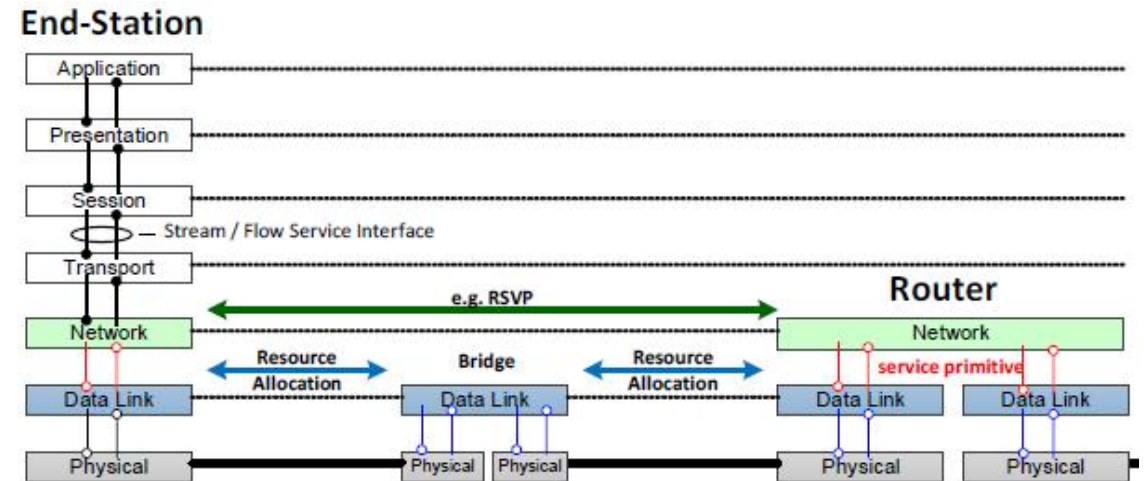● We must decide whether to keep separate protocols for L2 reservations and L3 reservations.

http://ieee802.org/1/files/public/docs2017/new-finn-detnet-needs-0917-v01.pdf

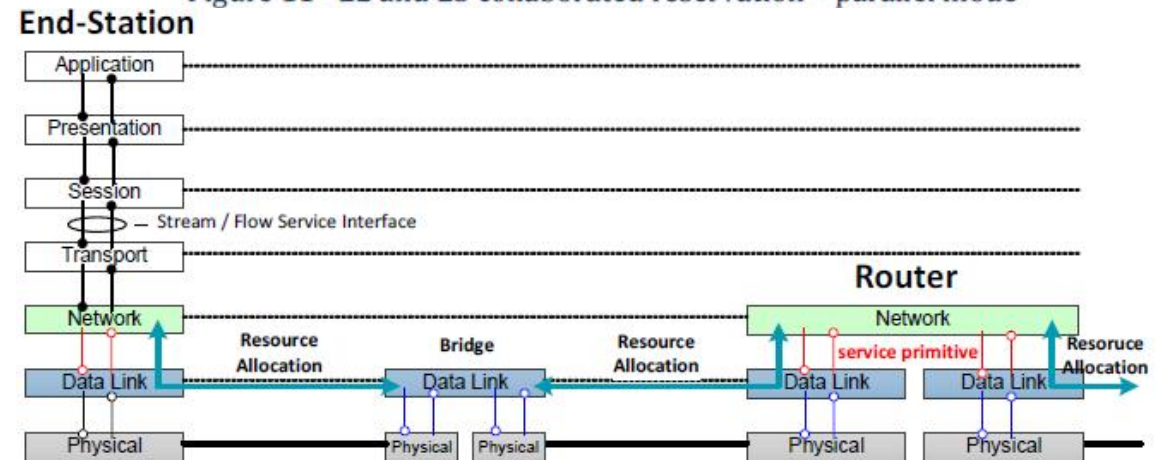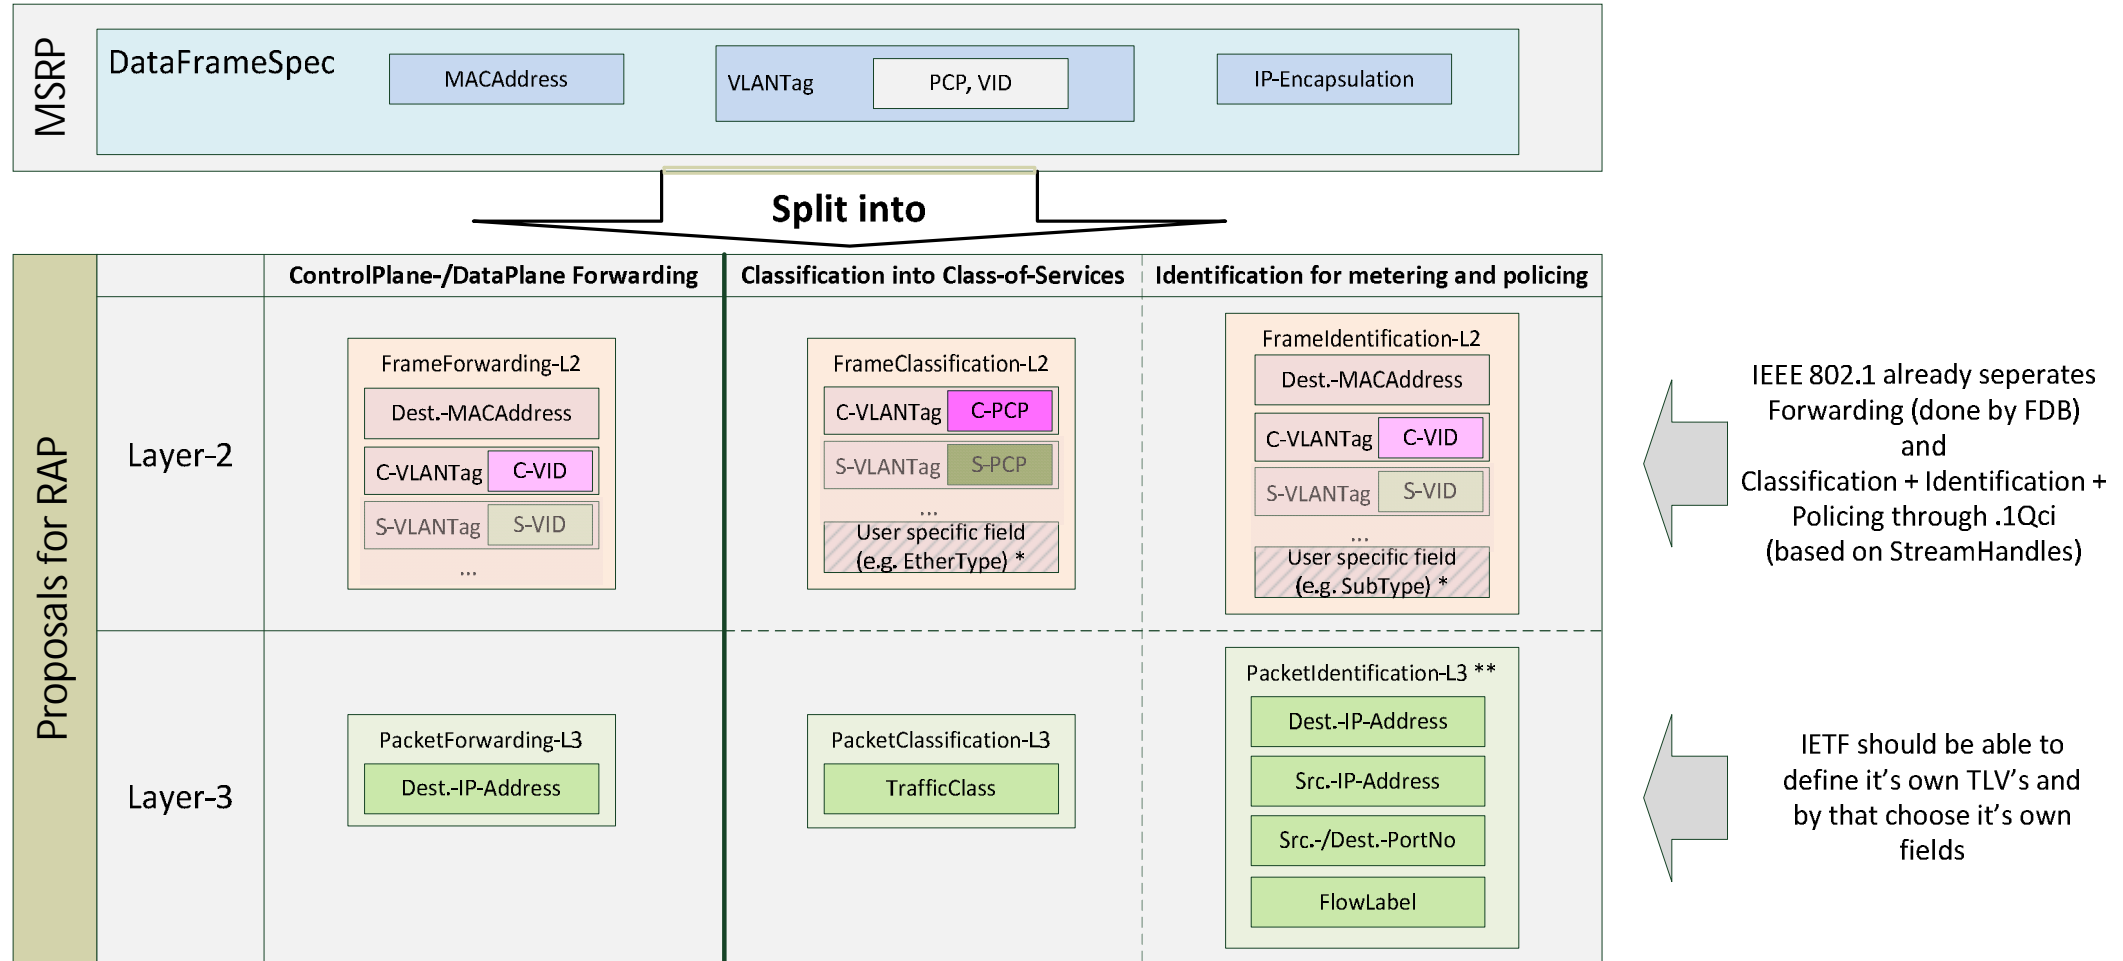Figure 11 - L2 and L3 collaborated reservation – parallel mode

Figure 12 - L2 and L3 collaborated reservation – serialized mode

L2/L3 Collaborated Reservation (for next white paper version)

q RAP data model for support of L2/L3 collaborated reservation

*http://www.ieee802.org/1/files/public/docs2017/new-mangin-extension-stream-id-0717-v01.pdf   ** comparable to FILTER_SPEC of RSVP

Thank You!

Questions?

Oct. 2 2017