

P802.1Qcz PAR and CSD Review

Paul Congdon

802.1 Interim, Pittsburgh

5/24/2018

Status

- The PAR and CSD were pre-circulated before March and all comments from other 802 WGs were resolved. The latest versions are available here:
 - <http://www.ieee802.org/1/files/public/docs2018/cz-congdon-congestion-isolation-PAR-0318-v1.pdf>
 - <http://www.ieee802.org/1/files/public/docs2018/cz-congdon-congestion-isolation-CSD-0318-v01.pdf>
- March motion to forward PAR and CSD to Nescom narrowly failed – Many abstains
- March motion to authorize May interim to pre-circulate PAR and CSD a second time passed – These slides are to review current text

Basic PAR information

- Current Text
 - **Type of Project:** Amendment to IEEE Standard 802.1Q-2014
 - **PAR Request Date:** 13-Dec-2017
- Questions:
 - Will this now be an amendment to 802.1Q-2018?
 - Do we update the PAR Request date to today (or near future)?

4.2 & 4.3 – Expected Dates

- Current Text
 - 4.2 Expected Date of submission of draft to the IEEE-SA for Initial Sponsor Ballot: 03/2021
 - 4.3 Projected Completion Date for Submittal to RevCom. Note: Usual minimum time between initial sponsor ballot and submission to Revcom is 6 months.: 10/2021
- Proposed Change: Add 3 months to each

5.2.b - PAR Scope

5.2.b. Scope of the project:

This amendment specifies protocols, procedures and managed objects that support the isolation of congested data flows within data center environments. This is achieved by enabling systems to individually identify flows creating congestion, adjust transmission selection for packets of those flows, and signal to neighbors. This mechanism reduces head-of-line blocking for uncongested flows sharing a traffic class in lossless networks. Congestion Isolation is intended to be used with higher layer protocols that utilize end-to-end congestion control in order to reduce packet loss and latency.

5.3 - Dependency

- 5.3 Is the completion of this standard dependent upon the completion of another standard: Yes
- If yes please explain: This amendment will specify a new Link Layer Discovery Protocol (LLDP) Type-Length-Value (TLV) and its associated YANG model. Project IEEE 802.1ABcu is currently specifying the YANG model for IEEE Std 802.1AB which must be completed in order for this amendment to specify its extension.

5.5 – Need for the Project

There is significant customer interest and market opportunity for large scale, low-latency, lossless Ethernet data centers to support high-performance computing and distributed storage applications. Congestion is the primary cause of loss and delay. These environments currently use higher layer end-to-end congestion control coupled with priority-based flow control at Layer 2 to avoid performance degradation from packet loss due to congestion. As the Ethernet data center network scales in size, speed and number of concurrent flows, the current environment creates head-of-line blocking for flows sharing the same traffic class. Isolating flows that cause congestion reduces latency for flows not causing congestion and improves the scale and performance of the Ethernet data center network. This amendment will support the identification and isolation of the higher layer protocol flows that are creating congestion. The amendment will interoperate with existing IEEE 802 and higher-layer congestion management. Use of a consolidated Ethernet data center network will realize operational and equipment cost benefits.

CSD – 1.2.1 Broad Market Potential

a) Broad sets of applicability.

Congestion is the primary reason for loss in data center networks with a limited network bandwidth-delay product. Higher layer congestion control protocols are widely deployed in those networks to reduce performance degradation due to loss. The higher layer protocols are limited in their ability to fully mitigate loss as data center networks scale in size and expand with higher-speed links. To eliminate loss in data center networks, the higher layer congestion control protocols can be combined with priority based flow control; however, negative consequences of head-of-line blocking and congesting spreading have been observed. Congestion Isolation improves the effectiveness of widely used higher layer congestion control protocols by isolating the flows that are causing congestion and providing additional time for the end-to-end protocols to react. Congestion Isolation can be applied to all current data center environments as well as future converged high-performance computing environments.

CSD – 1.2.1 Broad Market Potential

b) Multiple vendors and numerous users.

Multiple equipment and integrated circuit vendors have expressed interest in the proposed project. There is strong and continued interest from data center network operator in converging specialized high-performance networks to Ethernet and in the realization of operational and equipment cost savings through use of a consolidated network. Furthermore, numerous vendors are building new high-speed solid-state data storage access solutions over Ethernet networks, presuming that they can be realized with familiar technology and a consolidated network.

CSD – 1.2.3 Distinct Identity

IEEE Std 802.1Q is the sole and authoritative specification for VLAN-aware Bridges and their participation in LAN protocols. The existing congestion notification in IEEE Std 802.1Q is distinctly different in that it signals across the layer-2 network to a source reaction point and does not perform any flow isolation at the congestion point. Congestion Isolation performs flow isolation locally within the bridge, limits signaling to the next hop neighbor and does not define an equivalent reaction point. In order to support re-use and ease implementation efforts, congestion isolation incorporates and derives applicable procedures and protocols from congestion notification. The proposed amendment is intended to handle the short term effects of congestion while higher layer congestion control protocols, such as those based on IETF's explicit congestion notification (ECN), moderate the sources of congested traffic. No other IEEE 802 standard addresses congestion isolation by bridges.

CSD – 1.2.4 Technical Feasibility

a) Demonstrated system feasibility.

The proposed amendment incorporates techniques for flow identification and traffic scheduling that are currently available in many production data center switches.

b) Proven similar technology via testing, modeling, simulation, etc.

Performance improvements and a reduction in head-of-line blocking have been demonstrated through simulation and analysis of methods consistent with the proposed project. The proposed project can identify flows that are causing congestion by using techniques similar to the existing IEEE Std. 802.1Q congestion notification and IETF's explicit congestion notification (ECN). Flow identifying information can be maintained for congested flows. This information is similar to the information need to support existing features such as Access Control Lists (ACLs).

CSD - 1.2.5 Economic Feasibility

a) Balanced costs (infrastructure versus attached stations).

The proposed amendment does not significantly change the cost characteristics of bridges and does not require additional traffic classes. Implementations of the proposed amendment may need to include a flow table that can recognize flows that have been identified as causing congestion. Simulation data has shown that size of this table can be small and limited because the flows causing congestion are only a small fraction of the overall flows in transit. The proposed amendment operates independently of attached stations and does not impact station implementation

CSD - 1.2.5 Economic Feasibility

b) Known cost factors.

The proposed solution can reduce overall cost of data center networks by improving scaling and consolidating dedicated high-performance computing networks on to a common Ethernet fabric. This allows data center operators to consolidate their storage and computing networks, and to run traffic in a lossless environment that helps minimize flow completion time.

c) Consideration of installation costs.

Installation costs of data center bridges are not expected to be significantly affected.

d) Consideration of operational costs (e.g., energy consumption).

The proposed amendment uniquely incorporates existing technologies and as a consequence is not expected to significantly affect the operational cost of data center networks. A small amount of additional configuration is required, but fewer bridges and links need to be configured because the proposed amendment reduces the need to overprovision the data center network to achieve similar performance.