

CSD for P802.1Qcz

Bridges and Bridged Networks Amendment: Congestion Isolation

IEEE 802 LAN/MAN STANDARDS COMMITTEE (LMSC)

CRITERIA FOR STANDARDS DEVELOPMENT (CSD)

Based on IEEE 802 LMSC Operations Manuals approved 15 November 2013
Last edited 20 January 2014

1. IEEE 802 criteria for standards development (CSD)

The CSD documents an agreement between the WG and the Sponsor that provides a description of the project and the Sponsor's requirements more detailed than required in the PAR. The CSD consists of the project process requirements, 1.1, and the 5C requirements, 1.2.

1.1 Project process requirements

1.1.1 Managed objects

Describe the plan for developing a definition of managed objects. The plan shall specify one of the following:

- a) The definitions will be part of this project.
- b) The definitions will be part of a different project and provide the plan for that project or anticipated future project.
- c) The definitions will not be developed and explain why such definitions are not needed.

This project will develop necessary managed objects described by a YANG model as part of the project.

1.1.2 Coexistence

A WG proposing a wireless project shall demonstrate coexistence through the preparation of a Coexistence Assurance (CA) document unless it is not applicable.

- a) Will the WG create a CA document as part of the WG balloting process as described in Clause 13? (yes/no)
- b) If not, explain why the CA document is not applicable.

A CA document is not applicable because the standard will have no effect on wireless coexistence.

1.2 5C requirements

1.2.1 Broad market potential

Each proposed IEEE 802 LMSC standard shall have broad market potential. At a minimum, address the following areas:

- 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.

- 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.

1.2.2 Compatibility

Each proposed IEEE 802 LMSC standard should be in conformance with IEEE Std 802, IEEE 802.1AC, and IEEE 802.1Q. If any variances in conformance emerge, they shall be thoroughly disclosed and reviewed with IEEE 802.1 WG prior to submitting a PAR to the Sponsor.

- a) Will the proposed standard comply with IEEE Std 802, IEEE Std 802.1AC and IEEE Std 802.1Q?
- b) If the answer to a) is no, supply the response from the IEEE 802.1 WG.

The review and response is not required if the proposed standard is an amendment or revision to an existing standard for which it has been previously determined that compliance with the above IEEE 802 standards is not possible. In this case, the CSD statement shall state that this is the case.

The amendment will be in conformance with IEEE Std 802, IEEE Std 802.1AC, and the existing provisions of IEEE 802.1Q.

1.2.3 Distinct Identity

Each proposed IEEE 802 LMSC standard shall provide evidence of a distinct identity. Identify standards and standards projects with similar scopes and for each one describe why the proposed project is substantially different.

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.

1.2.4 Technical Feasibility

Each proposed IEEE 802 LMSC standard shall provide evidence that the project is technically feasible within the time frame of the project. At a minimum, address the following items to demonstrate 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).

1.2.5 Economic Feasibility

Each proposed IEEE 802 LMSC standard shall provide evidence of economic feasibility. Demonstrate, as far as can reasonably be estimated, the economic feasibility of the proposed project for its intended applications. Among the areas that may be addressed in the cost for performance analysis are the following:

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 congested flows 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.

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.

e) Other areas, as appropriate.