

IEEE 802 LAN/MAN STANDARDS COMMITTEE (LMSC)

CRITERIA FOR STANDARDS DEVELOPMENT (CSD)

Based on IEEE 802 LMSC Operations Manuals approved 17 March 2022
Last edited 20 May 2025

1. IEEE 802 criteria for standards development (CSD)

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

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.

Item a) is applicable to this project. This project will develop necessary managed objects described by a YANG model.

1.1.2 Coexistence

A WG proposing a wireless project shall prepare a Coexistence Assessment (CA) document unless it is not applicable.

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

NA. This project is not a wireless project; therefore, the CA document is not applicable.

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:

- f) Broad sets of applicability.

The data center market continues to grow rapidly, driven by the scaling of the network for Artificial Intelligence (AI), Machine Learning (ML), and High-Performance Computing (HPC) workloads. As these networks scale in size they often include links with long RTTs. The workloads and many new applications place high transient loads on data center networks which demand fast responding congestion management techniques to prevent packet loss. Beyond congestion, data center networks also benefit from timely backward signaling of other network conditions to enable faster recovery and more efficient resource utilization (e.g., path load balancing). Priority-based Flow Control (PFC, specified in IEEE Std 802.1Q) as well as higher-layer end-to-end congestion control mechanisms are currently used in data center environments. Backward Notification (BN) improves their efficiency by enabling bridges at the point of network events, such as congestion, to generate rapid, structured signals and transmit them toward the source of transmission. These signals convey congestion indications, network state information, and packet-related information to enable timely source-side responses including transmission rate adaptation, load balancing decisions, and fast retransmission needed because of packet loss.

- g) Multiple vendors and numerous users.

Most data center providers are using RoCEv2 and other proprietary Ethernet-based solutions in addition to new high-performance applications. These environments require lossless or very low loss operation. RoCEv2 and other high-performance applications are used by large enterprises and financial institutions, in addition to high-performance computing and AI environments. Further there is strong interest in accessing new high-speed solid-state data storage technologies over Ethernet networks using non-volatile memory express (NVMe) over fabrics (NVMe-oF). The rapid scaling of AI training and inference infrastructure has further amplified these requirements. Multiple networking vendors have developed proprietary implementations of backward notification mechanisms in data center networks, demonstrating broad industry recognition of the problem. Data center operators have expressed strong interest in a standardized solution to ensure interoperability and avoid fragmentation across proprietary implementations. Recent standardization efforts within HPC and AI data centers are specifying a programmable congestion management (PCM) framework that is an applicable consumer of backward notifications.

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.

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

Yes.

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.

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.

There are no other IEEE 802 standards or projects that specify an architecture, bridge forwarding model and service interface for conveying backward notifications of network conditions, exposing the notifications to entities above the link layer.

IEEE Std 802.1Qbb specifies Priority-based Flow Control (PFC), a hop-by-hop, on-off pause mechanism. BN operates multi-hop, carries extensible information content, and specifies service interface to deliver backward notifications to higher layer entities.

IEEE Std 802.1Qau (QCN) specifies a congestion notification mechanism in which the response action at the sender is prescribed as a rate-limiter, and the notification message is limited to Layer-2 bridge domain. BN delivers notifications to higher layer entities that react according to their needs, rather than BN specifying for the higher layers how to react. BN is not limited to a single L2 bridge domain.

This project supersedes the prior IEEE P802.1Qdw effort (Source Flow Control).

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:

- j) Demonstrated system feasibility.

Techniques, similar to the proposed project have been implemented in a proprietary fashion by multiple vendors. Additionally, the mechanisms for generating a specific type of backward notification have been demonstrated by Qau implementations. The proposed project applies to networks of limited bandwidth-delay product and where both bridges and end stations are under the control of a single administration. This reduces

the risk that the benefits of the technique will be eroded by over extended control loops or by end stations acting as adversaries. Backward Notification provides information that is used by higher-layer congestion control mechanisms. PFC and higher-layer congestion control algorithms are broadly implemented in hardware at acceptable cost.

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

The proposed project implements backward congestion signaling which has extensive simulation and analysis over the course of 30 years. In particular, relevant and substantial simulation and analysis of Source Flow Control has been conducted and presented in recent years. In addition, multiple networking vendors have independently implemented proprietary backward signaling mechanisms in production data center networks, demonstrating the practical viability of the approach.

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:

l) Known cost factors.

The proposed standard retains existing cost characteristics of bridges including simplicity of queue structures and will not require maintenance of additional queues beyond the existing per traffic class (priority) queues. In particular, per-flow queuing will not be required. BN frame generation and processing can be implemented within the existing forwarding pipeline with modest additional logic and without additional state.

m) Balanced costs.

The proposed project does not change the cost characteristics of bridges and end stations.

n) Consideration of installation costs.

Installation costs of bridges and end stations within the data center are not expected to be significantly affected; any increase in network costs is expected to be more than offset by a reduction in the number of separate networks required. There are no incremental installation costs relative to existing mechanisms that may be used by the proposed standard.

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

There are no incremental operational costs relative to existing mechanisms that may be used by the proposed standard.

p) Other areas, as appropriate.

No other areas have been identified.