

Congestion Notification Draft PAR (P802.1xx)

5 Criteria

1. Broad Market Potential

A standards project authorized by IEEE 802 shall have a broad market potential. Specifically, it shall have the potential for:

- a) Broad sets of applicability.

Mechanisms to avoid frame loss, of which congestion notification is one, are essential for support of the highly loss sensitive higher layer protocols, which are prevalent in the important applications E.g. data storage, clustering, backplane fabrics.

Back-end data storage networks, clustering networks and backplane fabrics are typically limited in size, making them amenable to a congestion control mechanism that is most effective with a limited network bandwidth-delay product. Each network is typically under the control of a single administrator, so the control technique does not require protection against 'gaming' by separate organizations attempting to acquire an unfair share of the bandwidth.

The data traffic to be controlled by the proposed congestion notification mechanism will be segregated using a VLAN-based technique, thus ensuring that traffic types already supported by VLAN Bridges are not affected and that there is no diminution of applicability to consolidated networks.

- b) Multiple vendors and numerous users

Multiple equipment vendors have expressed interest in the proposed project. There is strong and continued user interest in converting existing networks to Ethernet and in the realization of operational and equipment cost savings through use of a consolidated network. Further there is strong interest in increased use of data storage networks, provided that they can be realized with familiar technology and a consolidated network.

- c) Balanced costs (LAN versus attached stations)

The introduction of congestion notification is not expected to materially alter the balance of costs between end stations and bridges. While the introduction of the congestion notification option may constrain bridge implementation, significant equipment and operational costs savings are expected as compared to the use of separate networks for traditional LAN connectivity and for loss/latency sensitive applications.

2. Compatibility

IEEE 802 defines a family of standards. All standards shall be in conformance with the IEEE 802.1 Architecture, Management and Interworking documents as follows: 802. Overview and

Architecture, 802.1D, 802.1Q and parts of 802.1f. If any variances in conformance emerge, they shall be thoroughly disclosed and reviewed with 802.

Each standard in the IEEE 802 family of standards shall include a definition of managed objects which are compatible with systems management standards.

The proposed standard will be an amendment to 802.1Q, and will interoperate and coexist with all prior revisions and amendments of the 802.1Q standard. The data traffic to be controlled by the proposed congestion notification mechanism will be segregated using a VLAN-based technique, thus ensuring that traffic types already supported by VLAN Bridges are not affected.

Congestion notification frames and frame headers are confined to a domain composed solely of congestion notification capable bridges and end stations, thus preventing interoperability or compatibility problems from arising with either existing end stations and bridges, or with future systems using possible different techniques.

The proposed amendment will not introduce new bridge transmission selection algorithms or rate controls. Proposed end station controls on transmission rate and queuing are intended for use with 802.3 end stations and will be compatible with transmission control mechanisms already developed or under development by 802.3 and subject to liaison with 802.3 using the already established procedures. Such end station controls will be independent of the details of the 802.3 media access control technology and will make use of the existing interface (jointly standardized with 802.3) used by bridges.

The proposed amendment will contain MIBs, or additions to existing MIBs, to provide management operations for any configuration required together with performance monitoring for both end stations and bridges.

3. Distinct Identity

Each IEEE 802 standard shall have a distinct identity. To achieve this, each authorized project shall be:

- a) Substantially different from other IEEE 802 standards.

IEEE Std 802.1Q is the sole and authoritative specification for VLAN-aware Bridges and their participation in LAN protocols. No other IEEE 802 standard addresses congestion notification by bridges.

- b) One unique solution per problem (not two solutions to a problem)

Congestion notification is a reactive (not prescriptive) mechanism, and has not been anticipated by any other IEEE802 specification. It does not require or restrict the use of admission control techniques. It signals congestion through

bridges, unlike mechanisms that are specific to individual media access control methods.

Congestion Notification mechanism (ECN: Explicit Congestion Notification) specified by IETF is applicable to internet-wide topologies and only to TCP/IP applications. This proposal addresses the needs of low bandwidth-delay networks including those carrying non-TCP or non-IP traffic.

- c) Easy for the document reader to select the relevant specification.

IEEE Std 802.1Q is the natural reference for VLAN bridging technology, which will make the capabilities added by this amendment easy to locate.

4. Technical Feasibility

For a project to be authorized, it shall be able to show its technical feasibility. At a minimum, the proposed project shall show:

- a) Demonstrated system feasibility.

Congestion notification techniques have shown to be useful even in networks that are as difficult to control as the Internet. The proposed amendment will be applied only in networks of limited bandwidth-delay product and where both bridges and end stations are typically 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 some of the end stations 'gaming the system'.

The amendment will specify a one way bandwidth-delay product across the congestion controlled domain. The bandwidth-delay product limit is expected to be in the region of 1-5 Mbits (100 – 500 uS control loop delay for 10Gbps network) and simulation and analysis will verify performance characteristics up to the advertised bandwidth-delay product.

It has been shown that end station rate limiting capabilities, suitable for use with congestion notification, can be implemented in hardware at acceptable cost.

- b) Proven technology, reasonable testing.

The proposed amendment is based on extensive simulation and analysis in an area that has been studied for over 20 years.

- c) Confidence in reliability.

In keeping with best practice in this technical area, both end station and bridge behaviour will be specified, and the performance, stability, and fairness of the congestion control algorithm and resulting network throughput simulated and analyzed to the bounds of the specification.

- d) Coexistence of 802 wireless standards specifying devices for unlicensed operation

Not applicable.

5. Economic Feasibility

For a project to be authorized, it shall be able to show economic feasibility (so far as can reasonably be estimated), for its intended applications. At a minimum, the proposed project shall show:

- a) Known cost factors, reliable data.

The proposed amendment will retain existing cost characteristics of bridges including simplicity of queue structures and will not require maintenance of additional queues or queue state beyond the existing per traffic class(priority) queues for conformance to either its mandatory or optional provisions. In particular per flow queuing will not be required.

The proposed amendment may require some functions, specifically the generation of congestion notification frames, at a rate and within a time not practical for some existing and otherwise conformant bridge implementation architectures. However these functions can be performed by some existing bridges with known implementation costs.

The proposed amendment is technically feasible, in the envisaged application environment, with minimal flow state in end stations and will allow for complexity/throughput optimization trade-offs.

- b) Reasonable cost for performance.

The proposed technology will reduce overall costs where separate networks are currently required by enabling the use of consolidated network.

The proposed solution allows network to avoid packet loss without significant throughput reduction.

- c) Consideration of installation costs.

Installation costs of VLAN Bridges or end stations 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.