Proposal to generate a PAR for Enhanced Transmission Selection
19 July 2007

Background

- Priorities/classes are being used to separate traffic with different QOS characteristics
- It is desirable to enable sharing network bandwidth between classes
- For example best effort traffic and congestion controlled traffic
  - Neither class should be able to lock the other out of network access so strict priority shouldn't apply
  - Each class should be able to have a share of network bandwidth allocated to it.
  - Without the ability to share bandwidth between "equally deserving" classes of service, it may be difficult to deploy such classes.
- Some classes of service such as management and voice/video streams may continue to need strict priority over others
  - Provides for lowest latency
  - Most appropriate for traffic that is bandwidth limited by design or reservation hence these classes can also follow bandwidth allocations to avoid starvation of lower priority traffic classes.

Objectives

- Applicable to both end node and bridge ports
- Provide for some strict priority classes
  - Consider including configuration of bandwidth allocation and monitoring for these classes
- Provide for a set of classes with each class allocated a share of remaining bandwidth (i.e. the bandwidth not used by the strict priority classes)
  - E.g. Weighted Fair Queuing (WFQ) or Deficit Weighted Round Robin (DWRR)
- Provide managed objects to configure shares/weights without dictating implementation architecture.

PAR Fields

Amendment Title: Enhanced Transmission Selection

Scope: This standard specifies enhancement of transmission selection to support allocation of bandwidth amongst traffic classes. It will include managed objects to support bandwidth allocation.

Purpose: Networks prioritize traffic to provide different service characteristics to traffic classes. It is desirable to be able to share bandwidth between these priorities rather than servicing them with strict priority. For example, IEEE P802.1Qau will specify congestion management. Congestion managed traffic classes can share a network with traditional best effort LAN classes. Enhanced transmission selection will provide uniform management for the sharing of bandwidth between congestion managed classes and traditional classes on a single bridged network.

Need for Project: There is significant customer interest and market opportunity for Ethernet as a consolidated Layer 2 solution in high-speed networks such as data centers, backplane fabrics, single and multi-chassis interconnects computing clusters and storage networks. The differing service needs of applications supported on a consolidated Ethernet are supported by separate traffic classes. Support of these classes on the same links requires the ability to allocate a guaranteed share of bandwidth to each class. Use of a consolidated network will realize operation and equipment cost benefits. This project allows a uniform management of bandwidth allocation between classes.

Stakeholders: Developers and users of networking for data center environments including networking IC developers, switch and NIC vendors, and users.

Five 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.

Bandwidth sharing amongst classes is important to allow support for data storage, clustering, and backplane fabrics. This will allow deployment of traffic classes to segregate traffic needing differing service characteristics while sharing network bandwidth amongst the classes. This will allow the traffic classes to co-exist, thus enabling network consolidation.

b) Multiple vendors and numerous users

Many switches and end nodes for data centers already support proprietary implementations of bandwidth allocation amongst classes. It is expected that this standard will provide a framework for management of that bandwidth allocation that will be compatible with a range of current switch and NIC architectures.

c) Balanced costs (LAN versus attached stations)

The introduction of enhanced transmission selection is not expected to materially alter the balance of costs between end stations and bridges. Significant equipment and operational cost 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 enhanced transmission selection will continue to allow for traffic classes with strict priority to support traffic that requires minimum delay.
The proposed amendment will contain MIB modules, or additions to existing MIB modules, 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 priority aware Bridges and their participation in LAN protocols. No other IEEE 802 standard addresses transmission selection in bridges and end nodes.

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

Currently strict priority is the only transmission selection mechanism in the IEEE802 specification; consequently, this proposal is the only solution to the problem of allowing bandwidth allocation amongst traffic types.

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

IEEE Std 802.1Q is the natural reference for transmission selection, 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.

Similar techniques are widely deployed as proprietary enhancements in Ethernet bridge and end node products today as well as in other networking technologies, such as InfiniBand. The proposal is a natural extension of the transmission selection capability defined in IEEE Std. 802.1Q and widely deployed in bridge products.

b) Proven technology, reasonable testing.

The technique of this proposal has been proven in real world deployments of Ethernet, InfiniBand, and other networking technologies. These techniques have been shown to be reasonably testable.

c) Confidence in reliability.

The techniques of the proposal have been proven reliable in real-world deployments of Ethernet, InfiniBand, and other networking technologies.

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.

The proposed amendment may require some functions, specifically the distribution of bandwidth to queues not practical for some existing and otherwise conformant bridge and end node implementation architectures. However these functions are performed by many existing bridges and end nodes with known implementation costs.

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.

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. The proposed amendment is expected to reduce installation costs by providing a uniform management for transmission selection.