Credit-based Flow Control
PAR, 5 Criteria & Objectives

Results of 1/23/14 DCB TG discussion

PROJECT AUTHORIZATION REQUEST

Scope

This standard specifies the protocols, procedures, and management objects for Credit-based Flow Control in IEEE 802.1Q networks.

Update 802.1Q Data Center Bridging to support a per-priority Credit-based Flow Control including:

- Credit advertising mechanism
- Credit processing and handling mechanism
- Credit loss / error handling
- Enhancements to Data Center Bridging eXchange protocol for capability sharing
- New or updated MIB modules

Completion dependent on another standard?

No

Purpose

Converged networks are being deployed to support LAN/Storage/HPC over single LAN physical infrastructure

Several of these protocols benefit from a lossless network behavior: FCoE, RoCE & iSCSI

802.1Qaz/bb defined PFC so networks could offer lossless operation for these traffic types but PFC has inherent inefficiencies in buffer utilization. Especially as link speeds increase beyond 100G

To address these inefficiencies, this project enhances IEEE 802.1 to support Credit-based Flow Control.

Need for project

There is significant customer interest and market opportunity for data center bridged networks that operate with link speeds of 100G and higher.

DCB defined PFC so networks could offer lossless operation for certain traffic types (such as FCoE, RoCE, iSCSI) but PFC has inherent inefficiencies in buffer optimization and potential bandwidth utilization, especially above 100G

- For a given buffer size and link speed, beyond a certain cable length, it becomes impossible to provide lossless operation

Ethernet needs a more predictable, dynamically controllable, and granular flow control mechanism

5
CRITERIA

Broad market potential

a) Broad sets of applicability: Data center bridges are used in many data centers transporting protocols such as FCoE, iSCSI, RoCE, etc. The deployment of these continues to grow.

b) Multiple vendors and numerous users: There are many vendors that build silicon and systems for bridges and end stations that would benefit from a standard definition of Credit-based Flow Control

c) Balanced costs (LAN versus attached stations): The introduction of this standard will optimize the cost dynamics of bridges and end stations for buffer/memory and bandwidth utilization.

Compatibility

The proposed standard will be an amendment to IEEE 802.1Q, and will interoperate and coexist with all prior revisions and amendments of the IEEE 802.1Q standard.

Any protocol changes and frame formats will be backwards compatible.

Additional MIB objects will be required and these would be backwards compatible with the existing MIB module.

Distinct identity

a) Substantially different from other IEEE 802 standards: There is no other standard that defines Credit-based Flow Control for 802.1 networks.

b) One unique solution per problem (not two solutions to a problem) : The proposed standard is the only standard that provides robust flow control between 802.1 stations (end stations or bridges) that operates correctly regardless of round trip delay.

c) Easy for the document reader to select the relevant specification: IEEE 802.1Q is the natural reference for vendors and users of data center bridges.

Technical feasibility

a) Demonstrated system feasibility: Other networking technologies (IB, FC) use various forms of credit-based flow control. These other networks are widely deployed. Contributions to 802.1 have shown mechanisms suitable for 802 networks and data center bridges and end stations

b) Proven technology, reasonable testing: Mechanisms similar to what is being proposed exist in other networking technologies and have been shown to be reasonably testable.

c) Confidence in reliability: Other network technologies supporting credit-based flow control are widely deployed.

d) Coexistence of 802 wireless standards specifying devices for unlicensed operation: Not applicable.

Economic feasibility

a) Known cost factors, reliable data: The proposed amendment provides robust flow control that operates correctly regardless of round trip delay. The cost of implementing and testing for similar mechanisms in other networking technologies is known.

b) Reasonable cost for performance: The proposed amendment provides robust flow control that operates correctly regardless of round trip delay. The cost of implementing and testing for similar mechanisms in other networking technologies has proven to be reasonable.

c) Consideration of installation costs: The proposed standard will have no effect on the cost of installation.

OBJECTIVES

Draft Objectives (1)

a) Will be backwards compatible with 802.1Q
b) PFC, CFC or no flow control operation can be negotiated
c) Flow control mechanism can be individually enabled or disabled on a per-priority basis
d) Only a single type of flow control mechanism will be operational at a given time on any physical link
e) Flow Control mechanism will be hop-by-hop and operate between peer ports
f) No New tags
g) IF flow control is supported:
    - Below 400G, PFC is default, credit-based flow control optional
    - 400G and above, Credit-based flow control is default and PFC is optional

THANK YOU