



FCoE: Flow Control Enforcement

T11/10-522v0

Joe Pelissier

Background

- **The IEEE is completing the standardization of Priority-based Flow Control (P802.1Qbb) and the DCBX protocol that is used to negotiate its use (P802.1Qaz)**
- **P802.1Qaz has one outstanding “no” vote due to the fact that there is no defined mechanism to ensure that flow control is enabled for traffic that requires it**

The comment was rejected and the specification is moving forward with the “no” vote

FCoE is the only known example of such traffic

There was no significant objection to providing such a mechanism in principle; however, no acceptable mechanism was found

- **This presentation proposes a such a mechanism for FC-BB-6**

The IEEE Proposals

- **Proposal 1: Discard all traffic on a given link / priority if the Priority-based flow control negotiates to an inconsistent state**
 - IEEE philosophy is best effort. If an inconsistent state occurs, then allow traffic to flow as best it can
 - This mechanism could prevent the very management traffic required to fix the problem
 - All know traffic types, except FCoE, would work fine in this situation, therefore its an FCoE problem to fix
- **Proposal 2: same as Proposal 1 except the protection mechanism may be enabled / disabled per port**
 - Same as Proposal 1
 - The same functionality can be provided with ACLs
 - Its an FCoE problem
 - T11 has an FCoE specification under development
 - That specification just happens to have an annex that describes the use of ACLs to increase FCoE fabric robustness
- **Proposal 3: In the case of an inconsistent negotiation, transfer all of the traffic to a non-flow controlled priority. This way, and end station can tell that something is wrong on the path since the traffic is arriving on an unexpected priority.**
 - Placing FCoE traffic on a non-flow controlled priority could have drastic effects on the network
 - FCoE does not respond appropriately to lost frames due to congestion
 - Could result in congestion collapse affecting all traffic on the LAN (not just FCoE)
 - First principles: do no harm

FC-BB-6 Proposal

- **In Annex C: Increasing FC-BB E Robustness Using Access Control Lists (Informative), add the following:**

In C.1 Overview, add after the third paragraph:

In addition, it must be insured that the FCoE traffic travels over paths that are not subject to frame loss due to congestion. One mechanism to achieve this is to operate the FCoE traffic over a priority one which Priority-based flow control is enabled (see IEEE 802.1Qbb). Other mechanisms are also possible. If a priority-based mechanism is used, then fabric robustness may be increased by ensuring that the FCoE traffic is isolated to the appropriate priority or priorities.

In C.2.2, add the following to the list:

f) Pri: the priority of the frame

In C.3.5, change the ACE here and everywhere else it appears in the annex to:

SA = assigned VN Port MAC address, DA = FCF-MAC address, Type = FCoE, Pri=FCoE allowed priority, permit;



Thank you!