

# Draft Minutes

## T11.3 FC-FS-2 ad hoc work group teleconference

### 27 April 2005 - noon to 1 PM PDT

The FC-FS-2 ad hoc work group of INCITS Technical Committee T11 Task Group T11.3 held a teleconference on 27 April 2005. Attendance was 9 people from 7 companies and is tabulated at the end of this document.

Minutes were taken by Bob Nixon (bob.nixon@emulex.com). Please report any corrections by email to the T11.3 reflector at T11\_3@mail.T11.org.

#### 1 Opening

FC-FS-2 ad hoc work group chairperson Bill Martin opened the teleconference at noon PDT. He thanked Seagate for hosting voice and webex facilities.

#### 2 Agenda

The agenda was declared in the meeting announcement as consideration of document T11/05-229v0, proposing a "source quench" method for controlling class 3 traffic congestion in multiswitch fabrics.

#### 3 Congestion Control in Fibre Channel

T11/05-229v0

DeSanti/Cisco

Claudio DeSanti gave the presentation at the 4 April 2005 meeting, but there was insufficient time to conclude discussion. The presentation describes a likely congestion scenario in fabrics given

- a) Multiple hop paths;
- b) Class 3 traffic; and
- c) Mixed speed end devices.

It then proposes to design a method for feedback from the congestion point within a fabric to the end devices that are determined to be responsible for the congestion, or to their local switches.

Some members of the conference call posed that Virtual Channels already resolve the issue. Others claimed Virtual Channels were only a partial solution, in that congestion could still occur within a single Virtual Channel.

Claudio reported that he had gotten feedback that requiring end device participation in congestion control, other than current buffer to buffer flow control, was undesirable, and that any new congestion control messages should be terminated and processed by the local switch of the source.

It was posed that unless the control messages are forwarded to and processed by end devices, multi-exchange sources might be inefficiently controlled. A response was that sources tended not to be significantly multi-exchange, since the context is prohibitive in many implementations. There was disagreement whether the source of congestion is most likely to be a storage system or a host.

It was posed that in an adequately-provisioned network, congestion issues are transitory, not average; and that the proposal deals with average congestion, not peaks. This was not contested.

Concern was raised that, because of the presumed transitory congestion behavior of fabrics, source backoff mechanisms are likely to cause an overall decrease in fabric capacity. Examples were incorrect source identification and unnecessarily severe or prolonged backoff. It was claimed that transitory congestion and the issues with source quench are typical of networks, not just FC fabrics.

Three questions were posed:

- 1) Is Class 3 congestion a problem?
- 2) If so, is a solution involving end point response acceptable?
- 3) If congestion is a problem but not to be resolved by end points, should it be taken to SW-5?

Without formal response to the questions, the call proceeded as if the answers were "Yes", "No", and "Yes with qualifications". The qualification was that SW-5 would be open to accepting Class 3 congestion control as an investigation, but probably not as a commitment.

It was observed that end points may already participate in congestion control by use of Class 2. It was replied that this would not resolve all congestion scenarios.

It was agreed to recommend to the FC-FS-2 work group that the proposal of T11/05-229 not be included in FC-FS-2. The author was encouraged to bring the issue to FC-SW-5 when that project begins.

## **4 Adjournment**

The teleconference was adjourned at 1 PM PDT on 27 April 2005.

## **5 Attendance**

<b>Representative</b>	<b>Company</b>
Bob Snively	Brocade
Claudio DeSanti	Cisco
Davide Bergamasco	Cisco
Bob Nixon	Emulex
Vinod Bhat	HP
Scott Kipp	McData
Joe Pelissier	McData
Jim Coomes	Seagate
Bill Martin	Sierra Logic