

Draft Minutes

T11.3 FC-BB-5 ad hoc work group regular meeting

6 August 2008 - 9 AM to 5 PM PDT

Bellevue WA

The FC-BB-5 ad hoc work group of the Fibre Channel Protocol (T11.3) Task Group held a regular meeting at Bellevue WA on 6 August 2008, hosted by Microsoft and Bob Griswold. Attendance was 38 people from 23 organizations and is tabulated at the end of this document.

Minutes were taken by Bob Nixon (Emulex) (bob.nixon@emulex.com). Please report any corrections by email to the T11.3 reflector at T11_3@mail.T11.org.

1 Opening remarks

1.1 Introductions

In the absence of work group chair Claudio DeSanti (Cisco), Craig Carlson opened the regular meeting Wednesday, 6 August 2008 at 9:08 AM PDT. He thanked our host, Microsoft and Bob Griswold, and led a round of introductions.

2 Meeting Policy

2.1 Attendance and Membership

Craig Carlson explained that attendance is recorded electronically at www.t11.org/att, and explained the procedure. Attendance at this meeting does not count toward attendance at the plenaries of T11 and its task groups (i.e., being here will not get you out or keep you out of membership jeopardy).

Craig Carlson stated that all persons present are considered members of this meeting and may vote on questions, limited to one vote per company present. He advised that although T11 does not limit participation in the activities of its work groups to representatives of T11 member organizations, T11 encourages organizations that are regularly represented at T11 work group activities to become members of T11.

2.2 Patents

Craig Carlson indicated that among the rules and policies under which this working group operates are the ANSI intellectual property policies as specified in pages 1-3 of http://www.incits.org/pat_slides.pdf. He called for a moment of silence as he displayed these pages without comment or explanation, and directed that questions about the policy should be referred to the questioner's legal counsel or the ANSI General Counsel. The membership observed that the font size assured the slides were unobservable. Those further back than the front row were advised of the web link given above.

2.3 Antitrust

Craig Carlson indicated that among the rules and policies under which this working group operates are the INCITS Antitrust Guidelines. Any member of the meeting is responsible for objecting if he believes discussion in the

meeting violates those guidelines. As examples, there should never be discussion of the following topics at any INCITS or INCITS subgroup meeting:

- Any company's prices or pricing policies;
- Specific R&D, sales and marketing plans;
- Any company's confidential product, product development or production strategies;
- Whether certain suppliers or customers will be served;
- Prices paid to input sources; or
- Complaints about individual firms or other actions that might tend to hinder a competitor in any market.

If such discussion is not immediately terminated, it is the chairperson's responsibility to terminate the meeting. The INCITS Antitrust Guidelines are available at

<http://www.incits.org/inatrust.htm>

3 Administrivia

3.1 Approval of Agenda

An agenda for the FC-BB-5 ad hoc work group regular meeting 6 August 2008 has been posted as T11/08-441v0.

Pat Thaler requested time to present a proposal "Ethernet Link Status".

Dave Peterson requested to defer T11/08-450v0.

Silvano Gai (Cisco) moved and Bob Nixon (Emulex) seconded to accept T11/08-441v0 with the changes noted above as the agenda for this regular meeting. Approved unanimously.

3.2 Review of Minutes

Minutes for the FC-BB-5 ad hoc work group regular meeting 4 June 2008 have been posted as T11/08-325v0.

Bob Nixon (Emulex) moved and Bob Kembel (Solution Technology) seconded to accept T11/08-325v0 as the minutes of the FC-BB-5 ad hoc work group meeting on 4 June 2008. Approved unanimously.

4 Review of Old Action Items

080508-1/A5 Dave Peterson to propose an introductory advisory that data structures in this standard are displayed in Fibre Channel (i.e., "big-endian") format, while specifications originating in the Ethernet community may display data structures in Ethernet (i.e., "little-endian") format.
(Carry)

080604-2 Pat Thaler to define a mapping from Ethernet physical link conditions to FCoE Virtual Link events.
(Closed by T11/08-466v0)

080604-3 Claudio DeSanti to re-publish T11/08-263v1 reflecting the presented T11/08-263v1 and the agreements at the FC-BB-5 ad hoc work group regular meeting 4 June 2008.
(Completed by posting of revised T11/08-263v1)

080604-4 FC-BB-5 editor to incorporate the protocol described in T11/08-263v1 into FC-BB-5.
(Carry)

- 080604-5 Dave Peterson to publish T11/08-324v1 reflecting T11/08-324v0 and the agreements at the FC-BB-5 ad hoc work group regular meeting 4 June 2008.
(Completed by posting of T11/08-324v1)
- 080604-6 FC-BB-5 editor to incorporate T11/08-324v1 into FC-BB-5.
(Completed by posting of T11/08-352v0)
- 080604-7 FC-BB-5 editor to add a sentence qualifying figure 28 to the effect that a lossless Ethernet bridging element, if present, may be connected to multiple lossless Ethernet MACs.
(Carry)
- 080604-8 FC-BB-5 editor to add a requirement that the response to a FIP message use the same VLAN as the request.
(Carry)
- 080604-9 Dave Peterson to publish T11/08-258v2 reflecting T11/08-258v1 and the agreements at the FC-BB-5 ad hoc work group regular meeting 4 June 2008.
(Completed by posting of T11/08-258v2)
- 080604-10 FC-BB-5 editor to incorporate T11/08-258v2 into FC-BB-5.
(Completed by posting of T11/08-352v0)
- 080604-11 FC-BB-5 editor to add text recommending the use of independent VLAN learning if separate VLANs are used to isolate Virtual Fabrics.
(Carry)
- 080604-12 Anoop Ghanwani to publish T11/08-283v1 reflecting T11/08-283v0 with the changes as presented at the FC-BB-5 ad hoc work group regular meeting 4 June 2008.
(Completed by posting of T11/08-283v1)
- 080604-13 David Black to publish the FCoE threat model discussion presented at the FC-BB-5 ad hoc work group regular meeting 4 June 2008.
(Completed by posting of T11/08-338v0)

5 Old Business

No old business was scheduled.

6 Scheduled Business

6.1 Adapting the FC_Port definition to VN_Ports T11/08-250v2 Nixon (Emulex)

A single change was proposed to normalize the definition of FC_Port with the revised definition in FC-FS-3. No inconsistent uses appeared in the rest of FC-BB-5.

Bob Nixon (Emulex) moved and Bob Kembel (Solution Technology) seconded to incorporate the FC-BB-5 changes proposed in T11/08-250v3 into FC-BB-5, reflecting T11/08-250v2 and the agreements at other work groups. The motion passed 18 Favoring, 0 Opposing, 1 Abstaining.

ACTION Bob Nixon to publish T11/08-250v3 reflecting T11/08-250v2 and the agreements at other work groups.

ACTION FC-BB-5 editor to incorporate the FC-BB-5 changes proposed in T11/08-250v3 into FC-BB-5.

6.2 FIP Keep Alive

T11/08-434v0 Ghanwani (Brocade)

The presentation is a revision of an earlier presentation, incorporating recommendations made at its prior review. It reminded the group of the motivations for FIP Keepalive messages, posed some issues raised by the need to balance response versus system load, and recommended some additional FIP functionality allowing better tradeoffs.

One proposed function was to allow an end station to recognize a direct connection to the FCF and turn off FIP Keepalive.

Another was to change Keepalive from one per VN_Port to a single FIP Keepalive per Ethernet port. In the simplest implementation of this, it would not refresh learning tables in intermediate switches. The proposed full solution was to send per-VN_Port FIP keepalives much less frequently (e.g., once per minute).

As an alternative to the low-period FIP Keepalive, ARP was suggested (as IP-consistent). In discussion, it was realized that the preferred ACL configuration for FCoE would block the ARP.

There was further observation that standard Ethernet flooding would obviate the need for the low-period Keepalive. It was countered that flooding is in some ways not desirable, and the low-period Keepalive essentially eliminates flooding for FCoE.

It was agreed that the change to keepalive per Ethernet port was a good thing to pursue. The low-period Keepalive was acceptable as an option, not mandatory.

Given the change in Keepalive sourcing, it was observed there is much less value to the ability of a directly connected end station to turn off FIP.

It was agreed that the no-Keepalive function from the end device was not of interest, and the existing no-Keepalive function from the FCF is a directive, not (as documented), a request (i.e., an end station that sends keepalives to an FCF that set its keepalive period to zero is noncompliant).

6.3 Increasing FCoE Robustness using FIP Snooping T11/08-264v3 Pelissier (Cisco)

This presentation extends its prior FPMA-based proposal to include SPMA as well. It now covers preferred ACL configurations for edge switches in an Ethernet LAN that carries FCoE traffic. It explicitly omits recommendations of ACL configurations for internal switches.

It was agreed that the set of all FCFs may be determined without administrative intervention.

Dave Peterson (Brocade) moved and Bob Kembel (Solution Technology) seconded to incorporate a revision of T11/08-264v3 into FC-BB-5, reflecting the presentation and agreements at the FC-BB-5 ad hoc work group regular meeting 6 August 2008. The motion passed 18 Favoring, 0 Opposing, 0 Abstaining.

ACTION Joe Pelissier to revise T11/08-264v3 to reflect the agreements at the FC-BB-5 ad hoc work group regular meeting 6 August 2008.

ACTION Joe Pelissier to provide source for T11/08-264v3 to FC-BB-5 editor.

ACTION FC-BB-5 editor to incorporate T11/08-264v3 into FC-BB-5.

6.4 FCoE: MIB Module Proposal T11/08-458v0, T11/08-427v0 Pelissier (Cisco)

T11/08-458v0 is an overview of a proposed MIB for management of FCoE aware elements of an Ethernet LAN. T11/08-427v0 is the formal text for the MIB. Joe explained the intent of the MIB and its content. He requested that members' SNMP experts review the text and offer any comments at the next meeting.

6.5 Multiple Fabric support via FIP T11/08-450v0 Peterson (Brocade)

This presentation was deferred.

6.6 Open issues T11/08-451v0 Peterson (Brocade)

It was posed that sending an FLOGI FIP requesting one mode of address assignment and a subsequent FDISC FIP requesting a different address assignment should be prohibited. It was observed that virtualization devices might need this capability.

This recommendation was deferred for further consideration.

It was posed to include a recommendation but not a requirement that the MAC address(es) at an Ethernet port for FCoE should be different from the MAC address(es) used for normal IP traffic. In early discussion, the proposal was extended to cover FIP as well. It was pointed out that in common cases, this would require an additional server-assigned MAC address by which FIP can be conducted.

The recommendation was accepted, but not the extension to cover FIP.

6.7 Ethernet Link Status T11/08-466v0 Thaler (Broadcom)

This presentation examines the possibility of mapping FCoE virtual link states from Ethernet link states. Pat summarized that all media of interest detect and report a common set of link states, but they may not report them uniformly. There is a MIB object that provides a uniform representation.

She also indicated there is a MIB object that provides a common representation for media type.

The discussion expanded into information on link statistics, particularly those used by the RLS and RLIR ELSs.

It was agreed to be sufficient to the goals of this project to provide a mapping of standard Ethernet receiver conditions into a subset of the information in the Fibre Channel Link Status Block. It would be useful future development to extend the FC architecture to give more intelligent information on conditions specific to media other than native Fibre Channel media.

ACTION Roger Hathorn to define a mapping from Ethernet physical link statistics to at minimum some of the counters in the Link Status Block.

7 Unscheduled Business

7.1 Followup on FIP Keepalive T11/08-471v0 Ghanwani (Brocade)

Since his morning presentation, the presenter had prepared a clear statement of what was being proposed earlier.

His new summary was acceptable.

7.2 SPMA and FPMA and FLOGI and FDISC T11/08-470v0 Gai (Cisco)

In followup to an earlier discussion, Silvano presented an example of a case where an FLOGI FIP requesting one mode of address assignment and a subsequent FDISC FIP requesting a different address assignment is valuable. The case was an FCoE N_Port_Virtualizer that concentrates the traffic of several physical FCoE N_Ports via a single N_Port, without doing any address mapping.

The impact of allowing this case on the preferred ACL proposal will be considered.

8 Project Schedule

To reflect the difficulty of several final issues, the date for letter ballot was slipped to October 2008.

Milestone	Expected date
Last Technical Input	June 2008
T11 letter Ballot	October 2008
Forward to INCITS	Month yyyy

9 Review of Action Items

080508-1/A5 Dave Peterson to propose an introductory advisory that data structures in this standard are displayed in Fibre Channel (i.e., "big-endian") format, while specifications originating in the Ethernet

community may display data structures in Ethernet (i.e., "little-endian") format.
(Carry)

080604-4 FC-BB-5 editor to incorporate the protocol described in T11/08-263v1 into FC-BB-5.
(Carry)

080604-7 FC-BB-5 editor to add a sentence qualifying figure 28 to the effect that a lossless Ethernet bridging element, if present, may be connected to multiple lossless Ethernet MACs.
(Carry)

080604-8 FC-BB-5 editor to add a requirement that the response to a FIP message use the same VLAN as the request.
(Carry)

080604-11 FC-BB-5 editor to add text recommending the use of independent VLAN learning if separate VLANs are used to isolate Virtual Fabrics.
(Carry)

080806-1 Bob Nixon to publish T11/08-250v3 reflecting T11/08-250v3 and the agreements at other work groups.

080806-2 FC-BB-5 editor to incorporate the FC-BB-5 changes proposed in T11/08-250v3 into FC-BB-5.

080806-3 Joe Pelissier to revise T11/08-264v3 to reflect the agreements at the FC-BB-5 ad hoc work group regular meeting 6 August 2008.
(Completed by posting of revised T11/08-264v3)

080806-4 Joe Pelissier to provide source for T11/08-264v3 to FC-BB-5 editor.
(Completed by email)

080806-5 FC-BB-5 editor to incorporate T11/08-264v3 into FC-BB-5.

080806-6 Roger Hathorn to define a mapping from Ethernet physical link statistics to at minimum some of the counters in the Link Status Block.

080806-7 FC-BB-5 chair to schedule a full day meeting near Minneapolis MN on September 25.

10 Meeting Schedule

10.1 Interim request

A straw poll was taken of whether an interim meeting is desirable on September 25 in the Minneapolis area. The results were 12 favoring and 6 opposing.

12 people would attend such a meeting.

A full day meeting will be requested near Minneapolis MN on September 25.

ACTION FC-BB-5 chair to schedule a full day meeting near Minneapolis MN on September 25.

10.2 Next Plenary Week

Request 6 hours at the T11 plenary week hosted by EMC in Providence RI, 6-9 October 2008.

11 Adjournment

Bill Martin (Emulex) moved and Bob Kembel (Solution Technology) seconded to adjourn. Approved unanimously.

The regular meeting was adjourned at 3:55 PM PDT on 6 August 2008.

12 Status of Open Proposals

Document Title	Number	Disposition	Author
Adapting the FC_Port definition to VN_Ports	T11/08-250	Close. Approved version was T11/08-250v3.	Nixon (Emulex)
Increasing FCoE Robustness using FIP Snooping	T11/08-264	Close. Approved version was T11/08-264v3.	Pelissier (Cisco)
FCoE: MIB Module Tutorial	T11/08-458	Close. Presented version was T11/08-458v0. This presentation was tutorial, and was accepted in principle.	Pelissier (Cisco)
FCoE: MIB Module Proposal	T11/08-427	Carry. Additional review time needed. Version presented was T11/08-427v0. Additional review time needed.	Pelissier (Cisco)
FIP Keep Alive	T11/08-434	Carry. Version presented was T11/08-434v0. Expect discussion to continue with a revision.	Ghanwani (Brocade)
Multiple Fabric support via FIP	T11/08-450	Deferred. Version presented was T11/08-450v0. Expect discussion at next meeting.	Peterson (Brocade)
Open issues	T11/08-451	Close. Version presented was T11/08-451v0. Expect editor to incorporate results of discussion.	Peterson (Brocade)
Ethernet Link Status	T11/08-466	Close. Presented version was T11/08-466v0. This presentation was tutorial, and was accepted in principle.	Thaler (Broadcom)
SPMA and FPMA and FLOGI and FDISC	T11/08-470	Close. Presented version was T11/08-470v0. This presentation was tutorial, and was accepted in principle.	Gai (Cisco)
Followup on FIP Keepalive	T11/08-471	Close. Presented version was T11/08-471v0. This presentation was tutorial, and was accepted in principle. Expect it to be incorporated in T11/08-434.	Ghanwani (Brocade)

13 Attendance

Organization	Representative
BLADE NETWORK TECHNOLOGIES	Chetan Yaliwal
BROADCOM	Pat Thaler
BROADCOM	Uri Elzur
BROCADE	Anoop Ghanwani
BROCADE	David Peterson
BROCADE	John Hufferd
CISCO	Joe Pelissier
CISCO SYSTEMS	Fabrizio Corno
CISCO SYSTEMS	Silvano Gai
CISCO SYSTEMS, INC.	Keith McCloghrie
DELL	Gaurav Chawla
EMC	David Black
EMC	Erik Smith
EMULEX	Bob Nixon
EMULEX	Tuan Nguyen
EMULEX	William R. Martin
ENDL TEXAS	Ralph Weber
FINISAR	Paul Gentieu
IBM	Louis Ricci
IBM	Roger Hathorn
IBM	Scott Carlson
INTEL CORPORATION	Gary Tsao
LSI CORP.	John Lohmeyer
MELLANOX TECHNOLOGIES	Diego Crupnicoff
MICROSOFT CORPORATION	Robert Griswold
NETAPP	Frederick Knight
PANDUIT CORPORATION	Robert Elliot
PMC-SIERRA	Brian L'Ecuyer
PMC SIERRA	Gregg Goyins
QLOGIC CORP	Alan Spalding
QLOGIC CORP.	Craig W. Carlson
SOLUTION TECHNOLOGY	Robert Kembel
SPIRENT COMMUNICATIONS	John McLendon

Organization	Representative
SUN MICROSYSTEMS	Doug Meyers
SUN MICROSYSTEMS	Matt Gaffney
SUN MICROSYSTEMS	Michael Roy
SYMANTEC	Roger Cummings
VMWARE	Lawrence Lamers