

Draft Minutes

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

4 June 2008 - 9 AM to 6 PM PDT

Coeur d'Alene ID

The FC-BB-5 ad hoc work group of the Fibre Channel Protocol (T11.3) Task Group held a regular meeting at Coeur d'Alene ID on 4 June 2008, hosted by FCIA. Attendance was 41 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

Claudio DeSanti opened the regular meeting Wednesday, 4 June 2008 at 9:02 AM PDT. He thanked our host organization, FCIA, and led a round of introductions.

2 Meeting Policy

2.1 Attendance and Membership

Claudio DeSanti 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).

Claudio DeSanti 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

Claudio DeSanti indicated that among the rules and policies under which this working group operates are the INCITS intellectual property policies. He requested persons wishing to make statements or ask questions relevant to this policy not to pursue discussion at this work group meeting, but instead to do so with an officer of T11.3 or T11.

The INCITS intellectual property policies are specified in INCITS Policies and Guidelines (<http://www.incits.org/rd1/in071686.pdf>) subclause 8.4.

A summary of the INCITS intellectual property policies as applicable to the work of this meeting is given in document T11/06-739v0, which is available on www.t11.org.

2.3 Antitrust

Claudio DeSanti 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 4 June 2008 has been posted as T11/08-302v0.

It was agreed to reorder the agenda to place related proposals in sequence.

It was agreed to add a discussion of FC-BB-4 timer clarifications (T11/08-324v0) by Dave Peterson.

Silvano Gai (Cisco) moved and Ed McGlaughlin (QLogic) seconded to accept T11/08-302v0 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 16 April 2008 have been posted as T11/08-248v0.

Minutes for the FC-BB-5 ad hoc work group interim meeting 8 May 2008 have been posted as T11/08-265v0.

Landon Noll (Cisco) moved and Bob Nixon (Emulex) seconded to accept T11/08-248v0 as the minutes of the FC-BB-5 ad hoc work group meeting on 16 April 2008 and to accept T11/08-265v0 as the minutes of the FC-BB-5 ad hoc work group interim meeting 8 May 2008. Approved unanimously.

4 Review of Old Action Items

- A1 Further clarification to be provided on the need for and requirements for timeouts relating to advertisements and solicitations: Cisco, Brocade, QLogic and Emulex.
(Completed by T11/08-263v1))
- A2 Landon Noll (Cisco) to lead a group of interested parties, via the email exploder, to work up a set of criteria for FCoE threat models for consideration by the FC-BB-5 ad hoc.
(Completed at interim meeting 8 May 2008)

- A3 Dave Peterson to integrate the current FC-BB-5 text and new texts in T11/08-172v2 and 08-209v1 in time for the interim meeting.
(Completed by T11/08-258v1)
- A4 FC-BB-5 editor to capture in the specification that 0EFC00h is the default FC-MAP value.
(Completed by T11/08-258v1)
- 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.
(Opened 8 May 2008)(Carry)

5 Old Business

No old business was presented.

6 Scheduled Business

6.1 Additional FIP Functions

T11/08-263v1

DeSanti (Cisco)

The presentation summarized the current operation of FIP discovery, and then proposed some additional FIP protocol options to deal with certain loss of connectivity cases. FCoE may use a "network cloud" to provide communication between an end device and an FCF, where native FC uses point-to-point physical links. This means FCoE can not rely on the physical signalling methods used for native FC links. The presentation offers new timers and FIP protocols to supplement the physical methods. As defined, these new elements deal with a granularity of individual VN_Port/VF_Port Virtual Links, and they are modeled as activities of the FCoE Controllers (the control plane), not the individual VN_Ports (the data plane).

The new protocol elements in combination with additional behavior in FIP-aware intermediate switches that are not FCFs, may assist in maintaining and removing Access Control List entries.

A member observed that if Server-Provided MAC Addressing is used, what is now defined as a Virtual Link may carry multiple VN_Ports, and the names of the new protocol elements are inconsistent with their effect, which is then at a finer granularity than Virtual Links. The presenter had already discussed this issue with a major proponent of the SPMA method, with whom he had agreed to use the current naming until a better choice of names becomes obvious.

A member observed that sending a FIP Clear in response to a FIP Keep Alive must be carefully limited to cases of known previously assigned addresses that have not been reassigned. In other cases, it may be a denial of service attack.

A member observed that this will result in much slower fault detection than native FC. This was granted, with the further observations that an FCoE installation that highly valued very quick fault recovery could directly connect end devices to FCFs. As an alternative, the Ethernet cloud may be installed with high redundancy and quick failover.

A member observed that it will be important for FC-BB-5 to define the Ethernet physical link conditions that should be treated as equivalent to native FC link events by FC upper levels (e.g., RSCN).

ACTION Pat Thaler to define a mapping from Ethernet physical link conditions to FCoE Virtual Link events.

The presentation pointed out that discovery processes that are not time-critical may rely on accumulating periodic advertisements rather than soliciting them actively.

It was agreed the randomization delay for responses to multicasts should be 100 milliseconds, not configurable.

It was agreed to specify the default advertisement period to 8000 milliseconds, and to remain open to supported proposals to change it.

Brocade indicated that it will bring an additional technical contribution that may affect ACL management in intermediate FIP-aware switches. This contribution is not available today, but is intended for early subsequent publication.

The presentation added a partitioning of the descriptor code space between critical descriptors (ignore the whole message that contains an unknown critical descriptor), and noncritical descriptors (ignore just the descriptor if it is unknown). A few in the noncritical range are reserved as vendor specific. A definition was also provided for an entirely vendor specific FIP message. This may include standard descriptors used in nonstandard ways.

There was a good deal of contention over the details of the vendor specific FIP message.

Ralph Weber (ENDL Texas) moved and David Black (EMC) seconded to recast the T10 Vendor_ID in vendor specific FIP messages as a required critical descriptor that occurs first in the descriptor list, and that the descriptor list may contain both standard and vendor specific descriptors, all of which are standard descriptor format. The motion passed 13 Favoring, 0 Opposing, 1 Abstaining.

There was a good deal of contention over the details of the vendor specific FIP message.

Claudio DeSanti (Cisco) moved and Dave Peterson (Brocade) seconded to accept the FIP protocol specified in T11/08-263v1 as the FIP protocol to be specified in FC-BB-5, as updated to include the agreements at the FC-BB-5 ad hoc work group regular meeting 4 June 2008. The motion passed 14 Favoring, 0 Opposing, 3 Abstaining.

ACTION 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.

ACTION FC-BB-5 editor to incorporate the protocol described in T11/08-263v1 into FC-BB-5.

6.2 FC-BB-4 timer clarifications T11/08-324v0 Peterson (Brocade)

Dave pointed out text in FC-BB-4 that was difficult to interpret. He suggested improvements.

Dave Peterson (Brocade) moved and Bob Nixon (Emulex) seconded to incorporate T11/08-324v1 into FC-BB-5, reflecting T11/08-324v0 and the agreements at the FC-BB-5 ad hoc work group regular meeting 4 June 2008. The motion passed unanimously.

ACTION 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.

ACTION FC-BB-5 editor to incorporate T11/08-324v1 into FC-BB-5.

6.3 Clause 7 proposed text T11/08-258v1 Peterson (Brocade)

Dave reviewed the current combined proposal for the FCoE clause in FC-BB-5. It includes all material agreed prior to this meeting. Dave reviewed certain questions that he encountered while merging the material from different sources.

He questioned the need to show a "lossless Ethernet bridging element" in the functional model (figure 28) of a VE_Port/VF_Port. Those who wished it to be present wanted to emphasize that it may optionally be included in an FCF-capable device. Those who questioned the need pointed out that it was not the subject of any normative text in the FCoE specifications proposed to date.

Silvano Gai (Cisco) moved and Bob Nixon (Emulex) seconded that figure 28 be left unchanged, and a sentence be added to the effect that a lossless Ethernet bridging element, if present, may be connected to multiple lossless Ethernet MACs. The motion passed unanimously/16 Favoring, 0 Opposing, 1 Abstaining.

ACTION 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.

It was agreed that ENodes should transmit multicast discovery solicitations only in support of necessary Fabric Login, not periodically.

It was understood some terminology clarification is pending.

It was understood that the use of FIP for ELP has not been entirely resolved.

It was agreed that the 802.1Q tag is optional, though it will in general be shown in diagrams.

ACTION FC-BB-5 editor to add a requirement that the response to a FIP message use the same VLAN as the request.

It was agreed that the FLOGI descriptor could have a fixed length (i.e., it was not subject to optional extensions or ESP effects).

Dave Peterson (Brocade) moved and Joe Pelissier (Cisco) seconded to incorporate T11/08-258v2 into FC-BB-5, reflecting T11/08-258v1, and the agreements at the FC-BB-5 ad hoc work group regular meeting 4 June 2008. The motion passed unanimously.

ACTION 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.

ACTION FC-BB-5 editor to incorporate T11/08-258v2 into FC-BB-5.

6.4 ... Robustness using FIP Snooping ...

T11/08-264v2

Pelissier (Cisco)

This is proposed to be an informative annex describing how FIP messages are intended to be used by a FIP-aware Ethernet switch to maintain ACLs. It is limited to Fabric-Provided MAC Address configurations, only because Server-Provided MAC Addresses are outside the scope of interest of the presenter, not as an intentional restriction on eventual content of the resulting annex or other annexes.

The proposal does not yet include the work approved in T11/08-263v1.

The only missing piece from a fully automated ACL maintenance design that Joe identified is a way to securely identify FCFs. This proposal presumes administrative specification of the switch ports to which only FCFs will be

attached. This may be specified as an ACL that rejects FIP advertisements from all ports, plus ACLs that permit FIP advertisements from the dedicated FCF ports.

It was advised that if an installation uses separate VLANs to implement separate Virtual Fabrics, it should configure its switches for independent VLAN MAC learning.

ACTION FC-BB-5 editor to add text recommending the use of independent VLAN MAC learning if separate VLANs are used to isolate Virtual Fabrics.

Joe then deferred to the presentation of T11/08-282v1 by Roger Hathorn.

6.5 ACL annex proposal Review and Comment T11/08-282v1 Hathorn (IBM)

Roger presented commentary on proposal T11/08-264, which was introduced at the interim meeting 8 May 2008. He advised that T11/08-264v2, presented earlier in this meeting, resolves some but not all of the issues he identified.

One issue for which he offered several alternative resolutions was determining ports to which FCFs may be attached. The issue was recognized, but the need for protocol to control the time period in which FCF learning is permitted was contested. Some possible value was seen in the suggestion of specifying a management plane operation (e.g., a MIB) to determine FCF-permitting ports.

Roger then offered to review in detail the comments he and Joe had exchanged privately. The offer was respectfully refused by some present. The refusal was respectfully ignored.

It became clear that the membership was not of a single understanding of the expected rules for the relationships among VLANs, VFs, and FC-MAP values.

Roger and Joe then offered to work together with other SPMA representatives to extend the ACL recommendations in T11/08-264 to cover SPMA as well.

6.6 FIP Snooping T11/08-283v0 Ghanwani (Brocade)

Anoop presented from a slightly modified unpublished version of the posted proposal.

ACTION 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.

Anoop presented examples of current applications of “snooping” by transparent Ethernet bridges. He then similarly examined the proposed FIP snooping.

It became obvious that the “hard shell, soft core” model for ACLs (i.e., all ACLs are set at ports of traffic entry, not interswitch ports) was necessary to allow Ethernet transparent path failover to work for FCoE. Internal ACLs, if absolutely necessary, would cause any internal link failure to lead to Virtual Link failures and Relogins for end devices routed via the failed link.

TRILL was identified as a circumstance that might lead to a desire to use internal ACLs.

6.7 Group addresses for FCoE T11/08-334v0 Thaler (Broadcom)

Pat advised that Broadcom intends to donate three group addresses to meet the needs of FIP.

Claudio DeSanti (Cisco) moved and Pat Thaler (Broadcom) seconded to accept the donations of Ethernet group addresses from Broadcom as designated in T11/08-334v0. The motion passed unanimously.

ACTION FC-BB-5 editor to incorporate the FIP group address assignments as designated in T11/08-334v0.

6.8 Discussion of the FCoE Threat Model nodoc Noll (Cisco), Black (EMC)

The presentation summarized progress toward developing an FCoE threat model. It has not yet been posted, but will be. The presentation proposed a set of security objectives relevant to the threat model.

ACTION David Black to publish the FCoE threat model discussion presented at the FC-BB-5 ad hoc work group regular meeting 4 June 2008.

7 Unscheduled Business

No unscheduled business was presented.

8 Project Schedule

Taken with a good dose of optimism, it is not broken yet.

Milestone	Expected date
Last Technical Input	June 2008
T11 letter Ballot	August 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.
(Opened 8 May 2008)(Carry)

080604-2 Pat Thaler to define a mapping from Ethernet physical link conditions to FCoE Virtual Link events.

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.

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

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.

080604-6 FC-BB-5 editor to incorporate T11/08-324v1 into FC-BB-5.

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

- 080604-8 FC-BB-5 editor to add a requirement that the response to a FIP message use the same VLAN as the request.
- 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.
- 080604-10 FC-BB-5 editor to incorporate T11/08-258v2 into FC-BB-5.
- 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.
- 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.
- 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.

10 Meeting Schedule

10.1 Next Plenary Week

Request 7 hours at the T11 plenary week hosted by Microsoft in Seattle WA, 4-7 August 2008.

11 Adjournment

It was moved by Bob Nixon (Emulex) and seconded by Silvano Gai (Cisco) to adjourn. Approved unanimously.

The regular meeting was adjourned at 6:03 PM PDT on 4 June 2008.

12 Attendance

Organization	Representative
BLADE NETWORK TECHNOLOGIES	Chetan Yaliwal
BLADE NETWORK TECHNOLOGIES	Tienwei (Tim) Chao
BROADCOM	Pat Thaler
BROADCOM	Uri Elzur
BROCADE	Anoop Ghanwani
BROCADE	Ezio Valdevit
BROCADE	John Hufferd
BROCADE	Robert Snively
CISCO SYSTEMS	Joe Pelissier
CISCO SYSTEMS	Fabrizio Corno
CISCO SYSTEMS	Landon Noll
CISCO SYSTEMS	Silvano Gai
CISCO SYSTEMS	Claudio DeSanti
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
FINISAR	Tim Beyers
IBM	Louis Ricci
IBM	Roger Hathorn
IBM	Scott Carlson
INTEL	Luke Chang
INTEL	Gary Tsao
LSI	John Lohmeyer
MELLANOX TECHNOLOGIES	Diego Crupnicoff
MICROSOFT	Robert Griswold
NETAPP	Frederick Knight
QLOGIC	Ed McGlaughlin
QLOGIC	Alan Spalding

Organization	Representative
QLOGIC.	Craig W. Carlson
SOLUTION TECHNOLOGY	Robert Kembel
SPIRENT COMMUNICATIONS	John McLendon
SUN MICROSYSTEMS	Matt Gaffney
SYMANTEC	Roger Cummings
TRUE FOCUS	Horst Truestedt
UNIVERSITY OF NEW HAMPSHIRE INTEROPERABILITY LAB	Mikkel Hagen
VMWARE	Lawrence Lamers