

Draft Minutes

T11.3 FC-BB-6 ad hoc work group regular meeting

8 June 2011 - 9 AM to 4 PM PDT

Coeur d'Alene ID

The FC-BB-6 ad hoc work group of the Fibre Channel Protocol (T11.3) Task Group held a regular meeting at Coeur d'Alene ID on 8 June 2011, hosted by FCIA and Chris Lyon. Attendance was 26 people from 18 organizations, and one emeritus member. Attendance 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 absence of chairperson Claudio DeSanti (Cisco), Craig Carlson (QLogic) opened the regular meeting Wednesday, 8 June 2011 **on time** at 9:00 AM PDT. He thanked our hosts, FCIA and Chris Lyon, and led a round of introductions.

2 Meeting Policy

2.1 Attendance and Membership

In absence of chairperson Claudio DeSanti (Cisco), The acting chair 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).

The acting chair 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, it requires nonmembers to identify themselves as such. Nonmembers that expect they may participate in the activities of T11 regularly were encouraged to become members.

No person identified himself as a member of an organization that is not a member of T11.

2.2 Patents

The acting chair 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 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.

2.3 Antitrust

The acting chair 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-6 ad hoc work group regular meeting 8 June 2011 has been posted as T11/11-229v0.

Erik Smith (EMC) advised that the presentation "FIP Clear Virtual Link Reason Codes" has been updated to T11/11-209v1.

Landon Noll (Cisco) moved and Lou Ricci (IBM) seconded to accept T11/11-229v0 with the changes noted above as the agenda for this regular meeting. Approved by acclamation.

3.2 Review of Minutes

Minutes for the FC-BB-6 ad hoc work group regular meeting 6 April 2011 have been posted as T11/11-141v0.

Bill Martin (Emulex) moved and Fred Knight (NetApp) seconded to accept T11/11-141v0 as the minutes of the FC-BB-6 ad hoc work group meeting on 6 April 2011. Approved by acclamation.

4 Review of Old Action Items

There were no action items carried into this meeting.

5 Old Business

No old business was presented.

6 Scheduled Business

6.1 FIP Clear Virtual Link Reason Codes

6.1.1 FIP Clear Virtual Link Reason Codes

T11/11-064v2

Smith (EMC)

This presentation describes a method of associating reason codes with FIP Clear Virtual LKinks messages. It is posed to be desirable in support of failure diagnosis.

6.1.2 FIP Clear Virtual Link Reason Codes (proposed text changes)

T11/11-209v1

Smith (EMC)

This document gives detailed text to implement the proposal introduced in 6.1.1.

Erik Smith (EMC) moved and Dave Peterson (Brocade) seconded to incorporate T11/11-209v1 into FC-BB-6. The motion passed by acclamation.

ACTION FC-BB-6 editor to incorporate T11/11-209v1 into FC-BB-6.

6.2 Controlling Switch Redundancy Protocol

6.2.1 Redundancy Considerations - Presentation

T11/11-227v0

DeSanti (Cisco)

In absense of Claudio, Joe Pelissier (Cisco) presented.

The presentation introduced a method of automating recovery from double failures in a distributed FCF (Controlling FCF pair plus dependent FDFs).

As a means of limiting the scope of the effort, certain double failures that are not automatically recovered in native FC SANs are posed to be out of scope for resolution in distributed FCFs.

It was asked that the consequences of unrecovered multiple failures be specified.

6.2.2 Redundancy Considerations

T11/11-224v0

DeSanti (Cisco)

This is a first pass at detailed text for the feature introduced in 6.2.1. It was not reviewed in any detail, but was indicated as a source for more detailed definitions of the terms and acronyms in the introductory presentation.

6.3 Distributed FCF

6.3.1 Distributed FCF

T11/11-026v2

DeSanti (Cisco)

This document was derived from an earlier document to isolate text concerning Distributed Switch definitions for FCoE. This is intended to be incorporated into FC-BB-6.

6.3.2 VA_Port Protocols

T11/11-225v0

DeSanti (Cisco)

This document was derived from an earlier document to isolate text concerning the protocols to maintain a Distributed Switch. This may be appropriate material for FC-SW-x if Distributed Switches are agreed to be .of possible reelvance beyond FCoE.

6.4 IBM Comments on Distributed FCF proposal 11-026v1

6.4.1 Multilevel FDFs

T11/11-221v0

Hathorn (IBM)

The presentation poses a shortcoming in an earlier presentation, and posed corrections to the earlier document. The writer of the earlier presentation has already accepted the correction in principle.

The members present offered friendly ridicule to the use of the wording "didn't not"

6.4.2 Distributed FCF Fabric Login

T11/11-222v0

Hathorn (IBM)

The presentation offered several design points that should be considered in the detailed design for Distributed FCF protocols, and that may be of particular relevance to the protocols in Distributed FCF using multi-level FDF connectivity.

There was extended discussion of a proposed interlock of Fabric Login response with completion of zoning update distribution induced by that Fabric Login. This issue was the potential delay of the Fabric Login response.

The writer of the original Distributed FCF proposal has agreed with the proposed changes.

6.5 FC-BB-6: Virtual Domain Merge

T11/11-214v0

Johnson (Brocade)

The presentation raised a number of questions that did not seem clearly answered in the specification and proposals to date.

The question of continued communication within a portion of a distributed FCF that has become isolated from its controlling FCF was raised at the prior meeting, and while value was determined, issues were also raised. The question was raised again and actively discussed. More issues were raised. The sense of the group was strongly but not unanimously to discontinue communication.

ACTION Craig Carlson to advise Claudio DeSanti that a proposal should detail how to terminate communication within a portion of a distributed FCF that has become isolated from its controlling FCF.

ACTION Dave Peterson to initiate a proposal to clarify the behavior in the event of a lost link between FDFs that is restored quickly in comparison with the time it takes to coordinate with the controlling FCF.

ACTION Dave Peterson to initiate a proposal to require that each virtual domain shall have a unique switch name.

ACTION Dave Peterson to initiate a proposal to resolve the cases where two independent controlling FCFs are both configured and linked to the same FDF, including what happens in the event of other link failures.

ACTION Dave Peterson to initiate a proposal to require that the switch name for a virtual domain shall persist over all transitions from primary to secondary controlling FCFs.

6.6 Changes to the Fabric Configuration Server Model

T11/11-220v0

Crandall (Brocade)

John reviewed work in progress on his work in progress for a new model for Fabric management. This work was reviewed in more detail in the FC-GS-7 meeting earlier this week.

6.7 FC-BB-6: Controlling Switch Redundancy

T11/11-215v0

Peterson (Brocade)

The presentation proposed that a redundancy group of controlling switches needs to allow for more than two switches.

A protocol was outlined for building a virtual domain of possibly more than two controlling switches and retaining its connectivity. It is based on FSPF.

It was explained that full connectivity was not required among the controlling switch group. The protocol is modeled on a Fabric build in an arbitrarily linked physical Fabric, where the complete protocol does not pass across all possible links.

Routing the protocol through an attached native FC Fabric is not prohibited; it would occur if the FSPF shortest path was via the native FC Fabric.

6.8 FC-BB-6: SW_ILS TLV format(s) T11/11-216v0 Peterson (Brocade)

This presentation proposed that future work on SW_ILSs presume the payload will be a common header followed by parameters structured as a TLV list. "Future work" includes any new SW_ILSs, and may include transition/replacement of current SW_ILSs.

It was suggested to extend the type field from 8 bits to 16 bits.

The presenter observed that, as this project is devining several new SW_ILSs, it would be opportune to begin using a TLV structure.

6.9 FC-BB-6: Virtual Domain text T11/11-217v0 Peterson (Brocade)

This presentation proposes text for virtual domain initialization and maintenance, using SW_ILSs in TLV format.

6.10 BB6 Requirements T11/11-218v0 Ayandeh (HP)

(Note, the presentation given had two more slides than the posted version. The presenter committed to post the new version.)

The presentation observed the increasing capability and scale of level 2 networks (i.e., Ethernet LANs). It then proposed delegation of all FC routing determination and operation to Ethernet bridges, and maintain only Fabric Services in FCFs.

It was observed that FC specifies, implements, and presumes FC zoning, Ethernet implements but does not standardize ACL behavior. The presenter posed that standardized ACL behavior was unnecessary.

The presenter argued that the current Dist-FCF proposal would likely end up being deployed as single-vendor; as E-ports are NOT interoperable. Otherwise it would be missing vendor proprietary features that are dear to customers, and will have higher TCO. He was promptly and firmly contradicted.

Several in the group present observed that this proposal appeared to be a reincarnation of very early discussions of FCF-less routing, most of the proponents of which had early been convinced that it was difficult to meet the expectations of the SAN community.

At least one member expressed sympathy that routing via FCFs introduced choke points that direct L2 routing would avoid.

The acting chair suggested that arguments about the lack of market potential of Dist-FCF and single vs. multi-vendor were of less relevance to T11 than to FCIA. He then allowed the presenter to continue.

7 Unscheduled Business

No unscheduled business was presented.

8 Review of Action Items

- 110608-1 FC-BB-6 editor to incorporate T11/11-209v1 into FC-BB-6.
- 110608-2 Craig Carlson to advise Claudio DeSanti that a proposal should detail how to terminate communication within a portion of a distributed FCF that has become isolated from its controlling FCF.
- 110608-3 Dave Peterson to initiate a proposal to clarify the behavior in the event of a lost link between FDFs that is restored quickly in comparison with the time it takes to coordinate with the controlling FCF.
- 110608-4 Dave Peterson to initiate a proposal to require that each virtual domain shall have a unique switch name.
- 110608-5 Dave Peterson to initiate a proposal to resolve the cases where two independent controlling FCFs are both configured and linked to the same FDF, including what happens in the event of other link failures.
- 110608-6 Dave Peterson to initiate a proposal to require that the switch name for a virtual domain shall persist over all transitions from primary to secondary controlling FCFs.

9 Meeting Schedule

Request 8 hours at the T11 plenary week hosted by FCIA in Edmondton AB Canada, 1-5 August 2011.

10 Adjournment

Joihn Crandall (Brocade) moved and Bill Martin (Emulex) seconded to adjourn. Approved by acclamation.

The regular meeting was adjourned at 5:04 PM PDT on 8 June 2011.

11 Status of Open Proposals

Document Title	Number	Disposition	Author
FC-BB-6 FDF Requirements	T11/10-343	Deferred. Carry for further development. Most recent version presented was T11/10-343v0.	Hathorn (IBM)
Distributed FCF Specification	T11/11-026	Carry for further development. Version available was T11/11-026v2	DeSanti (Cisco)
FIP Clear Virtual Link Reason Codes	T11/11-064	Close, tutorial for accepted proposal T11/11-209v1. Version presented was T11/11-064v2.	Smith (EMC)
Distributed FCF Presentation	T11/11-070	Carry for further development. Version posted was T11/11-070v1	DeSanti (Cisco)
Clear Virtual Link Reason Codes (proposed text changes)	T11/11-209	Close. Accepted T11/11-209v1.	Smith (EMC)

Document Title	Number	Disposition	Author
FC-BB-6: Virtual Domain Merge	T11/11-214	Close, actions were captured to resolve issues it raised. Version presented was T11/11-214v0	Johnson (Brocade)
FC-BB-6: Controlling Switch Redundancy	T11/11-215	Carry for further development. Version posted was T11/11-215v0	Peterson (Brocade)
FC-BB-6: SW_ILS TLV format(s)	T11/11-216	Carry for further development. Version posted was T11/11-216v0	Peterson (Brocade)
FC-BB-6: Virtual Domain text	T11/11-217	Carry for further development. Version posted was T11/11-217v0	Peterson (Brocade)
BB6 Requirements	T11/11-218	Close. Presenter was advised that the nature of the proposal was appropriate to an FCIA audience, but not to T11. Version presented was T11/11-218v0	Ayandeh (HP)
Changes to the Fabric Configuration Server Model	T11/11-220	Close, tutorial in this work group. Version presented was T11/11-220v0	Crandall (Brocade)
Multilevel FDFs	T11/11-221	Close. Recommendations will be incorporated into a future version of T11/11-026. Version presented was T11/11-221v0	Hathorn (IBM)
Distributed FCF Fabric Login	T11/11-222	Close. Recommendations will be incorporated into a future version of T11/11-026. Version presented was T11/11-222v0	Hathorn (IBM)
Redundancy Considerations	T11/11-224	Carry for further development. Version available was T11/11-224v0	DeSanti (Cisco)
Redundancy Considerations - Presentation	T11/11-227	Carry for further development. Version presented was T11/11-227v0	DeSanti (Cisco)
VA_Port Protocols	T11/11-225	Carry for further development. Version available was T11/11-225v0	DeSanti (Cisco)

12 Attendance

Organization	Representative
BROADCOM	Pat Thaler
BROCADE	David Peterson
BROCADE	John Crandall
BROCADE	Steven L. Wilson
CISCO	Joe Pelissier
CISCO SYSTEMS	Landon Noll
EMC	Erik Smith
Emeritus	Robert Kembel
EMULEX	Bob Nixon
EMULEX	William R. Martin
ENDL TEXAS	Ralph Weber
FUJITSU AMERICA, INC.	Sandy Wilson
HEWLETT-PACKARD	Siamack Ayandeh
HEWLETT-PACKARD COMPANY	Nadaraha (Nava) Navaruparajah
HUFFERD ENTERPRISES	John Hufferd
IBM	Louis Ricci
IBM	Roger Hathorn
IBM	Scott Carlson
INTEL CORPORATION	Prafulla Deuskar
JDSU	Jason Rusch
JUNIPER	Joseph White
LSI CORP.	John Lohmeyer
MELLANOX TECHNOLOGIES	Diego Crupnicoff
NETAPP	Frederick Knight
ORACLE	Roger Dickerson
QLOGIC	Ed McGlaughlin
QLOGIC CORP.	Craig W. Carlson