Draft Minutes
T11.3 FABEXP Study Group Meeting
4 October 2004 - 3 PM to 7 PM
Austin TX

An Study Group Meeting of work group FABEXP of INCITS Technical Committee T11 Task Group T11.3 was held at Austin TX on 4 October 2004, hosted by Crossroads. Attendance at this Study Group Meeting was 21 people from 18 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 remarks and introductions — Dave Peterson

Chairperson Dave Peterson opened the meeting Monday, 4 October 2004 at 3:01 PM. He thanked our host company, Crossroads, and led a round of introductions.

1.1 Antitrust

Dave Peterson indicated that among the rules and policies under which this working group operates is the INCITS Antitrust Guidelines. Any member of the meeting is responsible for objecting if he believes discussion in the meeting violates those guidelines.

1.2 Patents

Dave Peterson indicated that among the rules and policies under which this working group operates is the ANSI patent policy. He directed persons wishing to make statements or ask questions relevant to this policy to do so at the T11.3 or T11 plenary meeting.

2 Attendance and Membership

Dave Peterson stated that all persons present are considered members of this meeting and may vote on questions, limited to one vote per company present. 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).

3 Approval of Agenda — 04-636v1 — All

It was moved by Dave Peterson and seconded by Bob Nixon to accept 04-636v1 as the agenda for this meeting. Approved unanimously.

4 Review of Minutes and Action Items — All

4.1 Approval of Minutes of Prior Meetings — 04-553v0

It was moved by Bob Nixon and seconded by Bill Martin to accept 04-553v0 as the minutes of the FABEXP Study Group Meeting meeting on 2 August 2004. Approved unanimously.

4.2 Review of Old Action Items

There were none.

5 Old Business

5.1 Considerations on broadcast ***04-653v0*** ***DeSanti/Cisco***

This presntation examines several options for treatment of broadcast in a routed fabric.

The preferred solution in the presentation in fact does not route generic broadcasts, but solves the needs of IP routing by using IP routers that may coreside with fabric routers. The issue of what to do with broadcasts and multicasts remains, unless not routing them is a sufficient resolution.

Of note, FC-SB-x does not use broadcast.

5.2 Fabric Routing Unification ***04-645v0*** ***Pelissier/McData***

Joe introduced his presentation as an expansion on ideas raised by Bob Snively at the last meeting.

It was agreed that the term for a device that contains the fabric routing function will be “Fabric Router”.

It was agreed that a strong requirement on receivers to support both encapsulated and unencapsulated FR headers was necessary. The scope of such a requirement for transmitters was not resolved.

The issue was raised of devices that sit on a link totally transparently but still expect to interpret frame headers (see FC-BB-3). This appears to require FC-BB-3 expert consultation.

ISSUE: How would the presence of intervening switches between two fabric routers be determined in the presense of transparent bridges?

This proposal would move both the source and destination address proxy translation to the ingress router. An architectural issue was raised with this, in that it causes observers of frames in transit to require proxy information from the destination to identify the source.

ISSUE: Some, but not all, proposed address proxying methods require an observer between the entry and exit routers to obtain information from the entry or exit routers to determine the local source and destination addresses of a routed frame. It is not yet clear whether this disadvantage outweighs other advantages of these methods.

ISSUE: How is ELS payload translation handled when both source and destination translation is done in the routher at one end of the routing path?

It was agreed that the presense of a bit indicating expiration time invalid should not imply that a compliant device need not maintain accurate time.

5.3 Fabric Routing and Encap Headers ***04-527v1*** ***Peterson/CNT***

Dave’s proposal detailed the fabric routing header. His intention is that for most purposes, where it conflicts with 04-645, this proposal should subordinate to 04-645.

A discussion began as to whether the fabric IDs used for fabric routing are the same as the fabric IDs used for VSAN.

It was agreed that the fabric IDs used by Virtual Fabrics are completely separate from the fabric IDs used for Fabric Routing.

His main issue is to use NTP rather than FC clock timing for the fabric routing frame expiration time.

Agreed to use NTP rather than FC clock server timing for the fabric routing frame expiration time.

5.4 Transport Fabric Model ***04-663v0*** ***Wilson/Brocade***

04-662v0 is a specification. 04-663v0 is a supporting presentation. It proposes a reference model within which to resolve many ancillary questions raised by FR (e.g., RSCN registration and delivery). The model envisions end fabrics connected via a single transport fabric. All routers would be considered at least logically to directly communicate. This was immediately questioned, as it appears to prevent a fully meshed set of fabrics.

One possible resolution is to presume a meshed set of fabrics will derive a simpified “transport fabric” to simplify operational issues, while retaining the full generality of the routing discovery mechanism. This appeared to be acceptable, if it was presumed that multiple paths through the transport fabric were still possible.

The case of devices attached to switches in the transport fabric was still unclear: is this permitted?

A proposed resolution to this is to consider the collection of routers as an independent “overlay” fabric, possibly interconnected using “classic” fabrics (each of which may also be an end fabric), but running its own FSPF. This did not resolve the perceived issues of all present.

Steve will pursue this further.

5.5 Fabric Expansion MIBS ***04-642v0*** ***Kipp/McData***

Several management interfaces are recognized. This proposal uses MIBs as the model for the information to be provided consistently through all the interfaces. Two MIBs are contemplated: one for VSAN and one for Fabric Routing. The first is obviously much closer to maturity. The work of defining the MIBs is beginning in the T11.5 Management meetings, to which additional participants are enthusiatically invited.

5.6 Fabric ID Aliasing ***04-667v0*** ***Pelissier/McData***

The presenter posed that this presentation has been rendered moot by previous discussion and withdrew it.

6 New Business

None was presented.

7 Review of Schedule

Claudio DeSanti proposed that this study group be terminated by proposing a workgroup to define a Fabric Routing standard. This was debated, though most of the debaters were playing devil’s advocates, and personally agreed.

ACTION: Claudio DeSanti to publish a project proposal for a Fabric Routing standard.

It is possible that the next meeting may be the last.

8 Review of Actions and Issues — Bob Nixon

8.1 Action Items

AI 1 Claudio DeSanti to publish a project proposal for a Fabric Routing standard. (Opened 4 October 2004)

AI 2 next Action Item.

8.2 Open Issues

Issue 1 How would the presence of intervening switches between two fabric routers be determined in the presense of transparent bridges?

Issue 2 How is ELS payload translation handled when both source and destination translation is done in the router at one end of the routing path?

Issue 3 Some, but not all, proposed address proxying methods require an observer between the entry and exit routers to obtain information from the entry or exit routers to determine the local source and destination addresses of a routed frame. It is not yet clear whether this disadvantage outweighs other advantages of these methods.

Issue 4 Next open issue

9 Meeting Schedule — Dave Peterson

Request 4 hours at the T11 Plenary Week the week of December 6, 2004 in Indian Wells, CA.

10 Adjournment

The meeting was adjourned at 6:40 PM on 4 October 2004.

11 Actions on Proposals at This Meeting

Document Title	Number	Disposition
Considerations on broadcast	04-653v0	Close. Presenter to pursue at his discretion
Fabric Routing Unification	04-645v0	Carry. Expect revision.
Fabric Routing and Encap Headers	04-527v1	Close. Salient points to be incorporated in 04-645
Transport Fabric Model	04-663v0	Carry. Expect revision.
Fabric Expansion MIBS	04-642v0	Close. Informative only.
Fabric ID Aliasing	04-667v0	Close. Withdrawn by presenter.

12 Attendance

Name	Organization Represented
Haluk Aytac	AMCC
Steven L. Wilson	BROCADE COMMUNICATIONS
Claudio DeSanti	CISCO SYSTEMS, INC.
Silvano Gai	CISCO SYSTEMS, INC.
Harry V. Paul	CNT CORPORATION
David Peterson	COMPUTER NETWORK TECHNOLOGY
David Black	EMC
Ken Hirata	EMULEX
Bob Nixon	EMULEX
Ralph Weber	ENDL TEXAS
Neil Wanamaker	FINISAR CORP.
Vinod Bhat	HEWLETT PACKARD
Robert Dugan	IBM POUGHKEEPSIE
George Penokie	IBM TIVOLI SYSTEMS
Jeff Whitt	LSI LOGIC
Scott Kipp	MCDATA
Carl Zeitler	NEWISYS, INC.
William R. Martin	SIERRA LOGIC
Robert Kembel	SOLUTION TECHNOLOGY
Horst Truestedt	TRUE FOCUS, INC
Rich Ramos	XYRATEX