

Fabric Extension Study Group Draft minutes 03-822v0
December 8, 2003 1:00 PM – 6:00 PM
Paradise Island, Bahamas

Introductions

John Scheible opened the meeting at 1:00 PM. The group began with introductions. John thanked the host for the facilities

Approval of minutes

The minutes of the October meeting (03-691v0) were approved as written.

Approval of Agenda

The agenda was approved as modified (03-788v1).

Review of Action items

- Bob Nixon to send out a call on the reflector for requirements/problems that this group is attempting to address. (Completed)
- Bob Nixon to gather a list of the requirements/problems that this group is trying to address. (Completed)

5.1 Virtual Fabrics 03-777v0 Silvano Gai

Silvano started with terminology, which he attempted to get agreement on. He proposed definitions of Virtual Fabric, Fabric Routing, and Gateway. It was suggested that address translation should be added to the gateway function. Additionally it was suggested that the term “Route” in Fabric Routing was perhaps overloaded, and should be changed to “move”.

Through a set of slides Silvano described the difference between Fabric Routing and Virtual Fabrics. Both models have deployment currently and are useful. Through these examples there was the contention that the Virtual Fabric and Fabric Routing elements were simply different ways of accomplishing the same function. It was also discussed how Address Spaces were related to each other. This is a very complex differentiation between what devices reside in which address spaces vs. address translation that presents devices in different address spaces. For Virtual Fabrics, Silvano claimed that the address spaces are per Virtual fabric. There was a distinction made between Address Space and Name Space, with significant discussion over what exactly Name Space referred to. Each Virtual Fabric has its own unique address space. Gateways interconnect Virtual Fabric but are not in the middle of a single Virtual Fabric.

Silvano added the following terms to his lexicon: TE_Port (E_Port that supports tagging), TF_Port (an F_Port that supports tagging), and TN_Port (an N_Port that supports tagging). There is a requirement to define that a particular non-tagging port is in

a particular virtual fabric. There needs to be work done to determine how the virtual fabric with which a non-tagging port is determined. One mechanism for determining that there is a miss-configuration is use of broadcast packets. It was re-iterated that a tag-address combination is globally unique within the context of all virtual fabrics that share any hardware. There was an extended discussion between Bob Snively and Silvano about what the format of the structure used for Fabric Routing between tag capable switches. This was intended to facilitate the routing of frames between different Virtual Fabrics. It was also suggested that the VF tag could be used for other applications such as QoS. There was disagreement on the difference contended on slide 7 of Silvano's presentation. Bob contended that while VF and FR are orthogonal concepts, they may use the same structure. Silvano strongly disagreed with this direction.

To implement VF and FR the fabric may have to have 2 different headers. The FR header may have different incarnations. The main purpose of the VF header is to tag the packet in accordance with the VF that it is within. The purpose of the FR header is to carry the address of the source and destination virtual fabric. This was agreed by both Bob and Silvano. The presence of a VF tag is a link attribute. It is present only if the link is capable of carrying multiple traffic for multiple VFs. The VF tag is independent and orthogonal to the presence of FR and Gateways. There was concern raised that we are coming to agreement on principals when we have not determined our customer requirements.

5.2 Trace Route and Pong 03-799v0 Mike O'Donnell

This was presented at the meeting in Oklahoma City the concept of Trace Route and Ping. After feedback they changed the material to "Get Trace Route" and "Pong", as these functions are different than "Trace Route" and "Ping". The proposal is divided into 2 parts, one for FC-GS-5 and one for FC-SW-4.

Mike started with a quick overview of GTR. It was asked whether GTR needed information about Virtual Fabrics. This is not required because this is done within a single virtual fabric. This may be extended as Virtual Fabrics are defined. There is a token in the ACC response to the GTR request for the use of the Async GTR response.

Mike then moved on to the description of Pong. It was mentioned that this needs to be directed to the Configuration Server. This is buried in the text prior to table 7, but will be made more apparent. There is a mechanism to respond before the Pong has completed indicating that a later response will be sent to prevent a timeout waiting for the Async response. This is the same as the GTR in the use of a token and an Async response. It was suggested that it would be useful to indicate whether an Async response is acceptable or whether the originator can only accept the synchronous response.

The GR SW_ILS is used for inter-switch communication to send the GTR to the domain controller of the ingress switch. This same SW_ILS is used between the switches in the route to accumulate the information about the route. All responses are sent to the original

recipient switch of the GTR request. There is a need to modify the definition of R_A_TOV in FC-SW-4. Mike will add a ladder diagram to the introductory information that he will include from the original PowerPoint presentation.

Mike then described Pong Devices (PD) SW_ILS. During the presentation it was decided to change the name to FC_Ping Devices and the ELS from the management device to FC_Ping. Additionally Trace Route was changed to FC_Trace Route.

Mike will update the proposal and bring to this group in February as well as at the FC-SW and FC-GS meeting.

5.3 Customer Input 03-809v0 Bob Nixon

Bob at the direction of this group at the last meeting requested and assembled inputs of customer requirements. The submitted material is in clause 3 and the integrated material is in clause 2. Silvano Gai requested to add 2 definitions: “Fabric – Legacy Fabric as defined in FC-FS and FC-SW-3”, and “FabricGroup – A collection of fabrics that share either communication capability or physical infrastructure.” Bob Nixon captured the modifications that were generated during the course of the meeting.

Bob Nixon requested that all switch vendors propose their modifications to the last paragraph of 2.2.1.

Action Items

Old Action Items

None

New Action Items

None

Next meetings

The next meeting will be in February at Dana Pt, CA during the next T11 plenary week, and 6 hours are requested.

Meeting adjourned at 5:25 PM.

Project status:

The project is still in the Study Group stage.

Attendance

| Name | Company |
|--------------------|------------------------------------|
| Roy Elsbernd | AGILENT |
| Fabio Maino | ANDIAMO SYSTEMS |
| Neil Wanamaker | APPLIED MICRO CIRCUITS CORPORATION |
| William R. Martin | BROCADE COMMUNICATIONS |
| Edward McClanahan | BROCADE COMMUNICATIONS |
| Robert Snively | BROCADE COMMUNICATIONS |
| Steven L. Wilson | BROCADE COMMUNICATIONS |
| Ramesh Sivakolundu | CISCO SYSTEMS |
| Silvano Gai | CISCO SYSTEMS, INC. |
| Harry V. Paul | CNT CORPORATION |
| John Tyndall | CROSSROADS |
| David Black | EMC |
| Ron Stern | EMC |
| Bob Nixon | EMULEX |
| Ralph Weber | ENDL TEXAS |
| Kent Lindell | GENERAL DYNAMICS ADV. INFO. SYS. |
| John Scheible | IBM AUSTIN |
| Michael O'Donnell | MCDATA |
| William Wang | PMC-SIERRA |
| Ed McGlaughlin | QLOGIC |
| Craig Carlson | QLOGIC CORP. |
| Paul Suhler | SEAGATE TECHNOLOGY |
| Horst Truestedt | TRUE FOCUS, INC |
| Maurice Gleeson | XANCOM LTD |
| Rich Ramos | XYRATEX |