

FIBRE CHANNEL TEST SUITE
Fabric Build Procedure

Technical Document



Last Updated: October 12, 2001

*Fibre Channel Consortium
InterOperability Laboratory
Research Computing Center
University of New Hampshire*

*121 Technology Drive Suite 2
Durham, NH 03824-3525
Phone: (603) 862-0701
Fax: (603) 862-4181
<http://www.iol.unh.edu>*

*The University of New Hampshire
InterOperability Laboratory*

MODIFICATION RECORD

- | | |
|---------------------|--|
| 1. August 19, 2001 | Initial version |
| 2. October 12, 2001 | Fixed Test Setup
Edited Test Procedures (Ryan Long) |
| 3. October 12, 2001 | Editorial changes (Stephen Schaeffer) |

*The University of New Hampshire
InterOperability Laboratory*

INTRODUCTION

Overview

The University of New Hampshire's InterOperability Laboratory (IOL) is an institution designed to improve the interoperability of standards based products by providing an environment where a product can be tested against other implementations of a standard. This Fabric Build procedure has been developed to isolate communication issues within a Fabric environment. This Test Suite is a pure interoperability procedure in that the conditions created during the testing process are based on the devices used and not on a testing station. It should be clear that successful completion of this test process does not guarantee that the tested device will operate with other Fibre Channel devices. However, it does provide a good opportunity to observe the behavior of an Initiator, Target, and Switch in a multi-vendor, Fibre Channel Fabric. This test procedure primarily focuses on the implementation of the FC-FLA, FC-GS-3, and FC-SW-2 standards. The results of this test procedure will include if and why communication is possible across the Fabric, as well as performance and reliability.

Test Label

The Label associated with each test follows that of hierarchical domain names, with subgroups being separated by periods. In domain name format, the more specified identifiers are on the left; the higher level entries are on the right. For example, the "fc-ph_version", which is one of the FL_Port features tests, is identified using the following label:

fabric_build_procedure.interop.fc

Purpose

The purpose is a short statement describing what the test attempts to achieve. The test is written at the functional level.

References

The references section lists cross-references to the Fibre Channel standards and other documentation that might be helpful in understanding and evaluating the test and results.

Resource Requirements

The requirements section specifies the software, hardware, and test equipment that will be needed to perform the test. The items contained in this section are special test devices, software that must reside on the DUT, or other facilities that may not be available on all devices.

Last Modification

This specifies the date of the last modification to this test.

Discussion

The discussion covers the assumptions made in the design or implementation of the test as well as known limitations. Other items specific to the test are covered here. It also contains additional information that may be helpful in understanding the operation of the test.

Test Setup

The setup section describes in detail the configuration of the test environment and includes a block diagram for clarification as well as information such as the interconnection of devices, what monitoring equipment should capture, what the generation equipment should send, and any other configuration information vital to carrying out the test. Small changes in the configuration are included in the procedure section.

Procedure

The procedure section of the test description contains the step-by-step instructions for carrying out the test. It provides a cookbook approach to testing, and will often be interspersed with observable results.

The University of New Hampshire
InterOperability Laboratory

Observable Results

The observable results section lists observables that can be examined by the tester to verify that the DUT is operating properly. When multiple values are possible for an observable, this section provides a short discussion on how to interpret them.

Possible Problems


This section provides some clues to look for if the test does not yield the expected results.

REFERENCES

The following documents are referenced in this text:

- T11/Project 1305-D/Rev 5.3 – Fibre Channel – Switch Fabric 2 – (FC-SW-2)
- T11/Project 1235-DT/Rev 2.7 – Fibre Channel – Fabric Loop Attachment (FC-FLA)
- T11/Project 1356-D/Rev 6.42 – Fibre Channel – Generic Services 3 (FC-GS-3)

Test Setup for Fabric Build



*The University of New Hampshire
InterOperability Laboratory*

Test Label: fabric_build_procedure.interop.fc

Purpose: To isolate issues which affect communication throughout a multi-vendor Fabric.

Resource Requirements:

- Monitoring station that can capture Fibre Channel frames at a speed of 2 GHz.
- Monitoring station that can capture traces for 10 Ports.
- Two Switches, each with at least four F\L_Ports and one E_Port.

Last Modification: August 19, 2001

Discussion: This procedure is intended to create a Fabric with multiple Initiators and Targets attached to multiple Switches. The Fabric is constructed by attaching “pairs” of Nx_Port devices to a Switch one by one. A pair will be one Initiator and one Target assigned to each other at random. One pair will be local (attached to the same Switch) and the other two pairs will be remote (attached to separate Switches).

All devices and pairs will be given a label before the procedure begins. This label will remain assigned to that device or pair throughout the procedure. If the device requesting the testing procedure is an Initiator or a Target, than that device will always be assigned to pair 1. An example of a label is shown here:

Pair 1: including Target 1 and Initiator 1.

There are three types of tests that will be run during this procedure. They are: Inter Switch Link Tests, Event Tests, and Performance Tests. The ISL test verifies that each Switch initializes properly during Fabric Configuration. The Pair tests verify whether or not the attached devices form an interoperable Fabric. The Performance tests analyze the speed and reliability of the Fabric.

These tests are not pass or fail tests. They are purely observational.

Procedure: Label each device and pair. Unless otherwise specified, each trace should capture the data flow between all devices which are attached to the Fabric. The tests should be run in this order:

- 1) **E Port to E Port:**
 - a) Test 1

- 2) **Pair Tests:**
 - a) Pair 1 – Test 1
 - b) Pair 1 – Test 2
 - c) Pair 2 – Test 1
 - d) Pair 2 – Test 2
 - e) Pair 3 – Test 1
 - f) Pair 3 – Test 2
 - g) Pair 1 – Test 3

- 3) **Performance Tests:**
 - a) Pair 1 – Test 1
 - b) Pair 1 – Test 2
 - Pair 2 – Test 2

The University of New Hampshire InterOperability Laboratory

E Port to E Port:

Initial Setup: Switch 1 and Switch 2 are powered off. Each Switch has an E_Port connection through the analyzer. There are no Nx_Ports connected to either Switch.

Test 1:

Purpose: To confirm the interoperability between two Switches.

References:

- FC-SW2 – Chapter 6
- FC-SW2 – Chapter 7

Procedure: Power on each Switch.

Observables: Confirm that:

- 1) Each Switch Port achieves word synchronization.
- 2) Each Port comes up as E_Port operational. To confirm this, verify that:
 - a) Link initialization is attempted and completed.
 - b) Exchange Link Parameter (ELP) Switch_Interlink_Service (SW_ILS) is sent from one Switch to another.
 - c) If the Switches do not transmit ELP at the same time, both Ports agree on the operating conditions and SW_ACC is sent from the responder of the ELP SW_ILS to the originator.
 - d) If both Switches transmit ELP at the same time, the Switch with a lower Switch_Name transmits SW_RJT and the Switch with a higher Switch_Name transmits SW_ACC.
 - e) No FLOGI's are transmitted during this exchange.
 - f) The originator of the ELP SW_ILS performs the Link Reset Protocol.
 - g) The Link Reset Protocol is successful.
 - h) If both Ports support Exchange Switch Capabilities (ESC), the ELP originator shall become the ESC originator and the ELP responder shall become the ESC responder. Verify that:
 - i. The originator transmits ESC SW_ILS.
 - ii. If the receiving Switch does not support any of the protocols listed in the ESC SW_ILS, it responds with a SW_RJT and a reason code of "Unable to perform command request"
 - iii. If the receiving Switch does support one of the protocols listed, it chooses a single protocol from the list presented in the ESC SW_ILS and responds with this protocol in the payload of the ESC SW_ACC.
- 3) A Principal Switch is selected. To confirm this, verify that:
 - a) If Build Fabric (BF) is generated first,
 - i. SW_ACC is sent back to the originator.
 - ii. If Reconfigure Fabric (RCF) is sent by the BF SW_ILS originator.
 - iii. SW_ACC is sent back to the originator.
 - iv. F_S_TOV has expired.
 - b) Exchange Fabric Parameters (EFP) is completed.
- 4) Address Distribution occurs. To confirm this, verify that:
 - a) The principal Switch granted itself a Domain_ID and generated an EFP SW_ILS request that includes the principal Switches Switch Priority and Switch Name in the Domain_ID List.
 - b) The non-principal Switch transmits an SW_ACC to the EFP SW_ILS request.
 - c) The principal Switch transmits a Domain Identifier Assigned (DIA) request.
 - d) The non-principal Switch transmits an SW_ACC to the DIA SW_ILS request.
 - e) The non-principal Switch transmits a Request Domain_ID (RDI).
 - f) The principal Switch grants the non-principal Switch a Domain_ID in its SW_ACC to the RDI SW_ILS.
 - g) The principal Switch transmits an EFP SW_ILS request with a Domain_ID List that includes both the principal Switch and the non-principal Switch.

The University of New Hampshire
InterOperability Laboratory

- h) The non-principal Switch transmits an SW_ACC to the received EFP SW_ILS request with a Domain_ID List that includes both the principal Switch and the non-principal Switch.
- 5) The link is maintained. To confirm this, verify that:
 - a) Hello (HLO) frames are transmitted every hello interval.
- 6) The link database is maintained. To confirm this, verify that:
 - a) Link State Updates (LSU)'s are transmitted consistently at an appropriate time interval determined by each port.

The University of New Hampshire InterOperability Laboratory

Pair Testing:

Initial Setup: Switch 1 and Switch 2 have an interoperable E_Port connection. There are no powered on Nx_Ports attached to the Fabric.

Test 1:

Purpose: To add a Target to the Fabric.

References:

- FC-FLA – Chapter 5
- FC-FLA – Chapter 7
- FC-FLA – Annex A
- FC-GS3 – Chapter 5

Procedure: Power on Target and connect it to its corresponding Switch through the analyzer.

Observables: Confirm that:

- 1) The Target properly initializes with the attached Fx_Port. To confirm this, verify that:
 - a) The Fx_Port and Nx_Port achieve word synchronization.
 - b) If the Target is an N_Port, Link Initialization is completed.
 - c) If the Target is an NL_Port, Loop Initialization is completed.
 - d) The Target transmits FLOGI to the Fabric Controller.
 - e) The Fabric Controller transmits FLOGI ACC to the Target.
 - f) The Target transmits PLOGI to the Directory Server.
 - g) The Directory Server transmits PLOGI ACC to the Target.
 - h) The Target transmits the Register FC-4 Type Common Fibre Channel Service (RFC-4) to the Directory Server.
 - i) The Directory Server transmits an ACC to that RFC-4 request to the Target.
- 2) Address Resolution occurs. To do this, confirm that:
 - a) The local Switch transmits Registered State Change Notification (SW_RSCN) to the remote Switch.
 - b) The remote Switch transmits SW_ACC to that SW_RSCN.
 - c) Each Switch transmits RSCN to all devices registered for this service by transmitting State Change Registration (SCR) during the initialization phase.
 - d) If an Initiator queries the remote Switch for the nature of the state change, it forwards that request to the local Switch.
 - e) If the remote Switch queries the local the Switch for a new Address Map, the local Switch supports the request and responds to the query with correct information.
 - f) Each Switch responds to every Nx_Port who queried for the nature of the State Change.
- 3) If there are other initiators attached to the Fabric, verify that each Initiator completes PLOGI with the affected Nx_Port.
- 4) Note the difference in time from the point where the Target was removed from the fabric until the last PLOGI was sent.

*The University of New Hampshire
InterOperability Laboratory*

Test 2:

Purpose: To add an Initiator to the Fabric.

References:

- FC-FLA – Chapter 5
- FC-FLA – Chapter 7
- FC-FLA – Annex A
- FC-GS3 – Chapter 5

Procedure: Power on Initiator and connect it to Switch.

Observables: Confirm that:

- 1) The Initiator properly initializes with the attached Fx_Port. To confirm this, verify that:
 - a) The Fx_Port and Initiator achieve word synchronization.
 - b) If the Initiator is an N_Port, Link Initialization is completed.
 - c) If the Initiator is an NL_Port, Loop Initialization is completed.
 - d) The Initiator transmits FLOGI to the Fabric Controller.
 - e) The Fabric Controller transmits FLOGI ACC to the Initiator.
 - f) The Initiator transmits PLOGI to the Directory Server.
 - g) The Directory Server transmits PLOGI ACC to the Initiator.
 - h) The Initiator transmits RFC-4 to the Directory Server.
 - i) The Directory Server transmits an RFC-4 ACC to the Initiator.
 - j) The Initiator transmits an SCR request to the Fabric Controller.
 - k) The Fabric Controller responds by transmitting an SCR ACC to the Initiator.
- 2) Address Resolution occurs. To do this, confirm that:
 - a) The local Switch transmits SW_RSCN to the remote Switch.
 - b) The remote Switch transmits SW_ACC to that SW_RSCN.
 - c) Each Switch transmits RSCN to all devices registered for this service by transmitting SCR during the initialization phase.
 - d) If an Initiator queries the remote Switch for the nature of the state change, it forwards that request to the local Switch.
 - e) If the remote Switch queries the local Switch for a new Address Map, the local Switch supports the request and responds to the query with correct information.
 - f) Each Switch responds to every Nx_Port who queried for the nature of the State Change.
- 3) If there are other initiators attached to the Fabric, verify that each Initiator completes PLOGI with the affected Nx_Port.
- 4) Note the amount of time it takes for the Fabric to recover from an event until all PLOGI's have been completed.

Test 3:

Purpose: To verify that an event does not cause an application to be terminated.

Procedure: With one pair, start an application that can format a disk. Proceed to format the Target. Remove and then re-attach the Fibre connecting that same Target to the analyzer.

Observables: Confirm that the application is not terminated. To do this, confirm that:

- 1) After the application is completed, the Initiator can write data to the Target.

The University of New Hampshire InterOperability Laboratory

Performance Tests:

Initial Setup: Switch 1 and Switch 2 have an interoperable connection. All pairs are attached to the Fabric.

Test 1:

Purpose: To observe the current operating conditions on the Fabric.

Procedure: Start an application that can generate and record data rates. Record the following I/O rates between the designated Initiator and Target pair:

1. Total I/Os per second.
2. Total MBs per second.
3. Average I/O response time.
4. Maximum I/O response time.
5. % CPU usage.
6. Current Operating System.
7. Application Description.

Test 2:

Purpose: To observe the impact of an event on I/O across the Fabric.

References:

- FC-FLA – Chapter 5
- FC-FLA – Chapter 7
- FC-FLA – Annex A
- FC-GS3 – Chapter 5

Procedure: This is a multi-pair procedure. This test requires two pairs, Pair A (Initiator A and Target A) and Pair B (Initiator B and Target B). The DUT, if it is an Nx_Port, should be a part of Pair B. Start transmitting I/O across the Fabric using Pair A. With Pair B, unplug and then reattach the DUT. Record the changes in the categories below.

1. Total I/Os per second.
4. Total MBs per second.
5. Average I/O response time.
6. Maximum I/O response time.
7. % CPU usage.
8. Current Operating System.
9. Application Description.

Observables:

- 1) Observe the impact of the event in terms of data rates.
- 2) The unplugged Nx_Port properly initializes with the attached Fx_Port. To confirm this, verify that:
 - a) The Fx_Port and Nx_Port achieve word synchronization.
 - b) If the Nx_Port is an N_Port, Link Initialization is completed.
 - c) If the Nx_Port is an NL_Port, Loop Initialization is completed.
 - d) The Nx_Port transmits FLOGI to the Fabric Controller.
 - e) The Fabric Controller transmits FLOGI ACC to the Nx_Port.
 - f) The Nx_Port transmits PLOGI to the Directory Server.
 - g) The Directory Server transmits PLOGI ACC to the Nx_Port.
 - h) The Nx_Port transmits RFC-4 to the Directory Server.
 - k) The Directory Server transmits an RFC-4 ACC to the Nx_Port.
 - l) If the Nx_Port is an initiator, it transmits an SCR request to the Fabric Controller.
 - i.) The Fabric Controller responds by transmitting an SCR ACC to the Initiator.