Fibre Channel Consortium

Point-to-Point Interoperability
Test Suite
Version 1.01

Technical Document

Last Updated: August 15, 2006 1:06PM

Fibre Channel Consortium
InterOperability Laboratory
University of New Hampshire

121 Technology Drive, Suite 2
Durham, NH 03824
Phone: +1-603-862-0701
Fax: +1-603-862-4181
http://www.iol.unh.edu/consortiums/fc

MODIFICATION RECORD

June 22, 2005 – Release Draft v1.0

Matthew Plante – Merged Packet Error Rate Estimation Test Suite and Speed Negotiation Test Suite into one test suite encompassing point to point interoperability tests.

August 15, 2006 – Release Draft v1.01

Michael Davidson – Removed references to Research Computing Center

ACKNOWLEDGEMENTS

The University of New Hampshire would like to acknowledge the efforts of the following individuals in the development of this test suite.

Andy Baldman	University of New Hampshire
Matthew Plante	University of New Hampshire
David Woolf	University of New Hampshire

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 suite of tests has been developed to help implementers identify problems that Fibre Channel devices may have in establishing link and exchanging packets with each other. The tests do not determine if a product conforms to T11 standards. Rather, they provide one method to verify that the two devices can exchange packets and detect and establish a link at the optimal speed between the two devices that make up a link segment.

Note: Successful completion of all tests contained in this suite does not guarantee that the tested device will operate with other compliant devices. However, combined with satisfactory operation in the IOL's interoperability test bed, these tests provide a reasonable level of confidence that the Device Under Test (DUT) will function well in most environments.

Organization of Tests
The tests contained in this document are organized to simplify the identification of information related to a test and to facilitate in the actual testing process. Each test contains an identification section that describes the test and provides cross-reference information. The discussion section covers background information and specifies why the test is to be performed. Tests are grouped in order to reduce setup time in the lab environment. Each test contains the following information:

Test Number
The Test Number associated with each test follows a simple grouping structure. Listed first is the Test Group Number followed by the test's number within the group. This allows for the addition of future tests to the appropriate groups of the test suite without requiring the renumbering of the subsequent tests.

Purpose
The purpose is a brief statement outlining what the test attempts to achieve. The test is written at the functional level.

References
The references section lists cross-references to the T11 standards and other documentation that might be helpful in understanding and evaluating the test and results.

Resource Requirements
The requirements section specifies the hardware, and test equipment that will be needed to perform the test. The items contained in this section are special test devices or other facilities, which 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.

Test Setup
The setup section describes the configuration of the test environment. Small changes in the configuration should be included in the test procedure.

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 may be interspersed with observable results.

Observable Results
The observable results section lists specific items that can be examined by the tester to verify that the DUT is operating properly. When multiple values are possible for an observable result, this section provides a short discussion on how to interpret them. The determination of a pass or fail for a certain test is often based on the successful (or unsuccessful) detection of a certain observable result.

Possible Problems
This section contains a description of known issues with the test procedure, which may affect test results in certain situations.

TABLE OF CONTENTS

GROUP 1: SPEED NEGOTIATION INTEROPERABILITY

Scope: These tests are designed to identify problems that Fibre Channel Standard compliant devices may have in establishing link and exchanging data with each other.

Test #1.1: Speed Negotiation

Purpose: To determine if the DUT establishes the best possible link with a reference set of stations.

References:
[1] T11 Std. FC-FS Rev 1.70, Clause 28

Resource Requirements:
- A reference set of stations that can be used as link partners.
- Link monitoring facilities that are able to determine the signaling being used on the link.
- Local management indicators on the DUT and reference set that indicate the state of the link as perceived by the different stations.
- A channel with known characteristics within allowable margins.

Last Modification: March 2, 2005

Discussion: The ability to detect and establish a link at the optimal speed is dependent on the two devices that make up the link segment, and providing and detecting the signaling method or connection information being passed. The large majority of Fibre Channel products use the T11 standard FC-FS, Clause 28 Speed Negotiation. Some devices may use different proprietary schemes to detect the link partner's speed or do not detect link speed at all. This test procedure addresses three conditions in which link speed detection should work. The first procedure covers the case where the DUT is initialized before the remote station and there is no signal on the DUT's receiver. The second procedure covers the case where the DUT is initialized after the remote station and there is a signal from this remote station on the DUT's receiver. The third procedure covers the final case where the DUT is in an operational state and is connected to a station that is also in an operational state. These three conditions are checked, as there may be different signals on the line during the boot up sequences of the devices that could cause the DUT to detect and establish a link at the wrong speed.

All possible speed configurations between the DUT and the remote station are tested for each condition described above. Please refer to Table 1 below:

Configuration #	DUT's Speed Configuration	Link Partner's Speed Configuration
1	Auto	Auto
2	Auto	Fixed at 4G
3	Auto	Fixed at 2G
4	Auto	Fixed at 1G
5	Fixed at 4G	Auto
6	Fixed at 2G	Auto
7	Fixed at 1G	Auto

Table 1: Possible Speed Configurations

This test is an interoperability test. Failure of this test does not mean that the DUT is non-conformant. It does suggest that a problem in the ability of two devices to work "properly" together exists and further work should be done to isolate the cause of the failure.

Test Setup: Connect the DUT to a link partner via the appropriate media channel.

Procedure:
Part A **:** Speed configuration #1
Case 1: The DUT receives no signal from the link partner during initialization.
1. Power off the DUT and the link partner.
2. Power on the DUT and ensure that the device is initialized and all needed drivers are loaded.
3. Power on the test link partner and verify that it is initialized and all needed drivers are loaded.

4. Check local management information to verify that the link is established at the proper speed and that link auto-negotiation, if supported, negotiated the optimal common values for the two devices.
5. Send the DUT a series of packets and observe whether the packets are accepted or not.

Case 2: The DUT receives signal from the link partner during initialization.

1. Power off the DUT and the link partner.
2. Power on the link partner and ensure that the device is initialized and all needed drivers are loaded.
3. Power on the DUT and verify that it is initialized and all needed drivers are loaded. Check local management information to verify that the link is established at the proper speed and that link auto-negotiation, if supported, negotiated the optimal common values for the two devices.
4. Send the DUT a series of packets and observe whether the packets are accepted or not.

Case 3: The DUT establishes link with a fully powered and operational link partner.

1. Power off the DUT and the link partner.
2. Power both devices back on at the same time and allow them to initialize.
3. Verify that a proper link is established as in Cases 1 and 2.
4. Remove and hold the cable for a few seconds, then reinsert. Repeat five times. Check local management information to verify that the link came up at the proper speed and that link auto-negotiation, if supported, negotiated the optimal common values for the two devices.
5. Send the DUT a series of packets and observe whether the packets are accepted or not.

Part B:

1. Repeat Part A for speed configurations two through seven shown in Table 1.

Observable Results:

a. The DUT and the link partner should establish an HCD link in all cases. Both the DUT and link partner must be examined for indicators of proper link speed and type. Local management may provide information about configuration such as link speed. If management indications are unavailable, please refer to Appendix A for other methods to determine link speed.

Possible Problems: None.

GROUP 2: BIT ERROR RATE ESTIMATION

Scope: The following tests determine if two devices can maintain a BER of 10^{-12} between each other while transmitting a stressing pattern.

Test #2.1: Bit Error Rate Estimation

Purpose: To determine that a Fibre Channel system can facilitate packet exchange with a bit error rate less than or equal to 10^{-12}. This test is not a compliance test.

References:

[1] FC Std. Rev 7.0 - 2004 Subclause 5.1
[2] Appendix B
[3] Appendix C

Resource Requirements:

- A Test Pattern Generator running software capable of generating a specific number of FC frames containing the Compliant Jitter Test Pattern (CJTPAT). This may be a Fibre Channel Target, Initiator, Switch, or Test Tool.
- FC Test Channel with channel response in accordance with current specifications, multiple test lengths, and optionally multiple trace widths.
- A Monitoring Station capable of monitoring traffic sent to the Test Pattern Generator from the DUT that will indicate whether a CRC Error occurred. The Monitoring Station must not effect the characteristics of the Test Channel from the Test Pattern Generator to the DUT.
- Local Error Counters on both the DUT and Test Pattern Generator which will count detected CRC errors.

Last Modification: March 1, 2005

Discussion: This test is designed to verify the ability of a Fibre Channel system to exchange packets between elements within the system. The Fibre Channel system consists of a DUT, a Test Pattern Generator, and a Test Channel. The exchange of packets should produce a packet error ratio that is low enough to meet a desired bit error ratio. The bit error ratio as specified in reference [1] is 10^{-12}. For this test, $3x10^{12}$ bits will be transmitted between the Test Pattern Generator and the DUT. This will ensure that the bit error ratio is less than 10^{-12} with 95% accuracy if no errors are observed. Please note that the derivation of these numbers can be seen in reference [2].

The Test Pattern Generator is instructed to transmit the CJTPAT pattern to the DUT, through the FC Test Channel. Local Error counters on both the Test Pattern Generator and the DUT will be checked to determine if any CRC errors occurring during transmission. The counters can be used to isolate where the CRC errors occurred, either from the DUT or to the DUT. Optionally a Monitoring Station may be used to monitor the channel from the DUT to the Test Pattern Generator for indications of CRC errors that would have occurred on the Test Channel from the Test Pattern Generator to the DUT. This Monitoring Station must not affect the Test Channel from the Test Pattern Generator to the DUT. Using Local Error Counters and information from the Monitoring Station, the number of detected CRC errors occurring during transmission of CJTPAT can be determined. If more than 7 CRC errors are detected during the exchange the bit error rate criterion has not been met and the test fails.

Since a single packet contains many bits, the measurement technique does not really measure the bit error rate. However, due to the large amount of both bits and frames that are being sent, all errors will be treated equally. Some devices may have the ability to count multiple errors in a single frame, while others may only count one error per frame. For this test, both these cases are considered equivalent. This means that all errors observed will be counted towards the maximum allowed number of 7 errors. A device may in theory pass a test with a with a much higher bit error rate than that which is being measured. However, given that any one bit in error will corrupt the packet, multiple errors within a packet do not, in practice, make a difference in the number of packets that must be retransmitted on real links. Thus, a short-term clock deviation that causes 5 bit errors in one packet in a stream of 10^{12} bits will, under most conditions, cause as many packet errors as a device with a bit error rate of 1 in 10^{12}.

Test Setup: The DUT is tested against a Test Pattern Generator. The Test Pattern Generator is the device at the other end of the channel being used for interoperability testing. Test Setups for different device types can be found in reference [3].

Procedure:

1. Connect the DUT, Test Pattern Generator, Test Channel, and optionally the Monitoring Station, as shown in reference [3] for the correct DUT type.
2. Reset all counters that will be used to measure or monitor the exchange of packets between the DUT and the testing station. Configure software as needed.
3. Using the Test Pattern Generator, transmit $3x10^{12}$ observable bits to the link partner. Using 512 byte frames, this will require 585,937,500 frames to be transmitted.
4. Using error counters available on the DUT or Test Pattern Generator, observe the number of bit errors that occurred during the test period. Compare this to any information available from the Monitoring Station, if one was used.
5. Repeat steps 1-4 to test all connections in the system.

Observable results: Using the counters on the Test Pattern Generator, identify the number of packets transmitted. Using the counters on the DUT, identify the number of packets received. The difference between the number of packets transmitted and the number received is the number of lost packets. Knowledge of the Upper Layer Protocol used is important, as transmission of one packet may result in reception of multiple packets (i.e. a single SCSI Command may illicit a response of several data packets and a SCSI Response). The number of lost packets should be examined with other information gathered during the testing process to ensure that the failure is due to bit errors and not resource errors on the DUT or testing stations. In the ideal case all lost packets are identified on the Test Pattern Generator or the DUT as either a CRC error, or some other type of receiver error. If the local information gathered from the DUT is reliable it is often possible to isolate the failure to either the transmitter channel or the receiver channel. It should be noted that an observed BER of greater than 10^{-12} does not necessarily imply two devices are not interoperable. Additionally, an observed BER of greater than or less than 10^{-12} does not imply compliance or non-compliance of either device.

Possible Problems:

- If errors exist within the testing station's generation circuitry, or with the line monitor, then the spirit of this test (BER verification of the FC physical channel) is compromised.
- Verify that the line monitor has little impact on the channel between the test station and expander.
- It is possible that some devices may not be able to properly compensate for clock tolerances with a minimum gap between the frames. It may be necessary to adjust this gap length accordingly. Given that the length of the frames is approximately 512 bytes, a modest increase in the gap should not change the results of the test.

Appendix A: Link Speed Detection Methods

Purpose: To provide alternative methods to verify link speed.

References:

[1] FC-FS.
[2] FC-PI-2.

Last Modification: March 2, 2005 (Version 1.0)

Discussion:

A.1 – Introduction

Devices that support multiple speeds should provide for some indication of the speed established with the device's link partner. Local management indications (either through LEDs or a console interface) provide for the easiest and least invasive means to determine the link speed. However, some devices may not provide for such an interface. This appendix gives examples of how to determine the link speed if local management indications are not provided by the vendor.

A.2 – Optical Links

Optical power splitters can be used in conjunction with multi-speed analyzers to monitor the negotiated link speed. The "analyzer" can be either a high-level protocol analyzer that indicates the link speed, or it can be an optical oscilloscope. If an oscilloscope is used, the length of 1 unit interval (UI) is the indication of the negotiated speed.

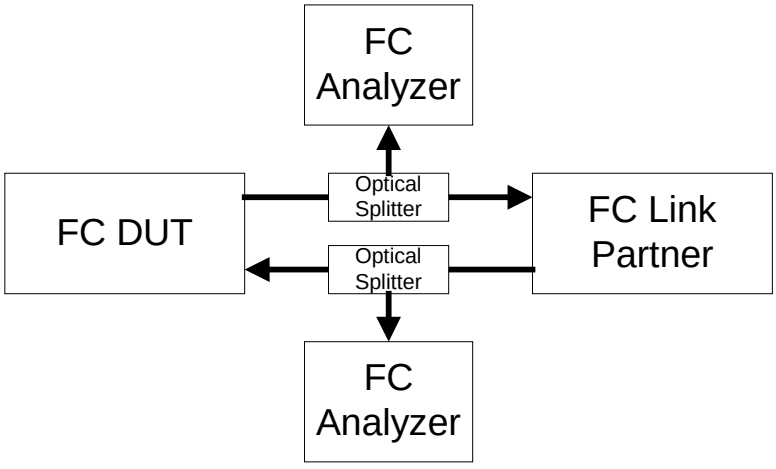

Figure A-1: Method to monitor optical links

A.3 – Electrical Links

If no management indications are available for devices linked over an electrical channel, methods for observing the negotiated speed are limited. An oscilloscope with a high impedance differential probe can be used to observe the transmit and receive lines. Connect the probe to the electrical interconnect between the external connector and the serdes, and measure the length of the UI.

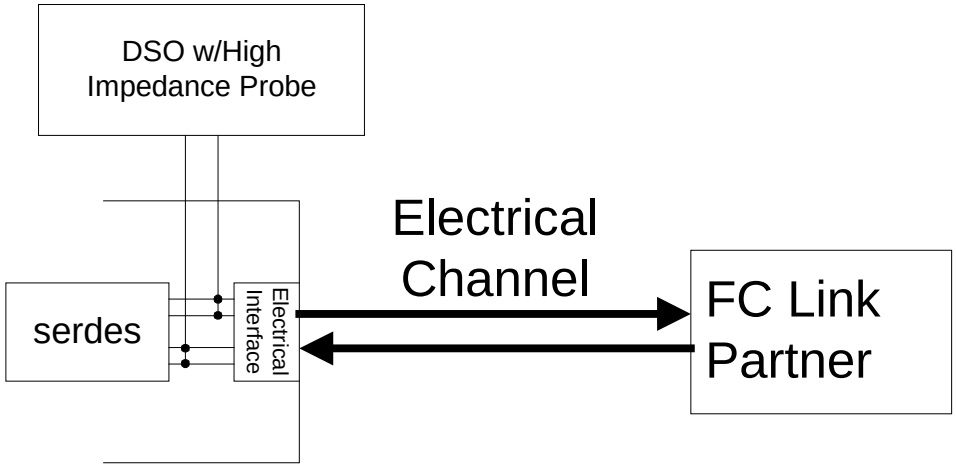

Figure A-2: Method to monitor electrical links

Appendix B: Packet Error Rate Measurement

Purpose: To develop a procedure for bit error rate measurement through the application of statistical methods.

References:

[1] Miller, Irwin and John E. Freund, Probability and Statistics for Engineers (Second Edition), Prentice-Hall, 1977, pp. 194-210, 240-245.

Last Modification: November 4, 2004 (Version 1.0)

Discussion:

B.1 – Introduction

One key performance parameter for all digital communication systems is the bit error rate (BER). The bit error rate is the probability that a given bit will be received in error. The BER may also be interpreted as the average number of errors that would occur in a sequence of n bits.

While the bit error rate concept is quite simple, the measurement of this parameter poses some significant challenges. The first challenge is deciding the number of bits, n, that must be sent in order to make a reliable measurement. For example, if 10 bits were sent and no errors were observed, it would be foolish to conclude that the bit error rate is zero. However, common sense tells us that the more bits that are sent without error, the more reasonable this conclusion becomes. In the interest of keeping the test duration as short as possible, we want to send the smallest number of bits that provides us with an acceptable margin of error.

This brings us to the second challenge of BER measurement. Given that we send n bits, what reasonable statements can be made about the bit error rate based on the number of errors observed? Returning to the previous example, if 10 bits are sent and no errors are observed, it is unreasonable to say that the BER is zero. However, it may be more reasonable to say that the BER is 10^{-1} or better. Furthermore, you are absolutely certain that the bit error rate is not 1.

In this appendix, two statistical methods, hypothesis testing and confidence intervals, are applied to help us answer the questions of how many bits we should be sent and what conclusions can be made from the test results.

B.2 – Statistical Model

A statistical model for the number of errors that will be observed in a sequence of n bits must be developed before we apply the aforementioned statistical methods. For this model, we will assume that every bit received is an independent Bernoulli trial. A Bernoulli trial is a test for which there are only two possible outcomes (i.e. a coin toss). Let us say that p is the probability that a bit error will occur. This implies that the probability that a bit error will not occur is (1-p).

The property of independence implies that the outcome of one Bernoulli trial has no effect on the outcomes of the other Bernoulli trials. While this assumption is not necessarily true for all digital communications systems, it is still used to simplify the analysis.

The number of successful outcomes, k, in n independent Bernoulli trials is taken from a binomial distribution. The binomial distribution is defined in equation B-1.

$$b(k; n, p) = C_{n,k}\, p^k (1-p)^{n-k} \qquad \text{(Equation B-1)}$$

Note that in this case, a successful outcome is a bit error. The coefficient $C_{n,k}$ is referred to as the binomial coefficient or "n-choose-k". It is the number of combinations of k successes in n trials. Returning to coin toss

analogy, there are 3 ways to get 2 heads from 3 coin tosses: (tails, heads, heads), (heads, tails, heads), and (heads, heads, tails). Therefore, $C_{3,2}$ would be 3. A more precise mathematical definition is given in equation B-2.

$$C_{n,k} = \frac{n!}{k!(n-k)!} \quad \text{(Equation B-2)}$$

This model reflects the fact that for a given probability, p, a test in which n bits are sent could yield many possible outcomes. However, some outcomes are more likely than others and this likelihood principle allows us to make conclusions about the BER for a given test result.

B.3 – Hypothesis Test

The statistical method of hypothesis testing will allow us to establish a value of n, the number of bits to be sent, for the BER measurement. Naturally, the test begins with a hypothesis. In this case, we will hypothesize that the probability of a bit error, p, for the system is less than some target BER, P_0. This hypothesis is stated formally in equation B-3.

$$H_0 : p \le P_0 \quad \text{(Equation B-3)}$$

We now construct a test for this hypothesis. In this case, we will take the obvious approach of sending n bits and counting the number errors, k. We will interpret the test results as shown in table B-1.

Table B-1: Acceptance and rejections regions for H_0

Test Result	**Conclusion**
k = 0	H_0 is true
k > 0	H_0 is false

We now acknowledge the possibility that our conclusion is in error. Statisticians define two different categories of error. A type I error is made when the hypothesis is rejected even though it is true. A type II error is made when the hypothesis is accepted even though it is false. The probability of a type I and a type II error are denoted as α and β respectively. Table B-2 defines type I and type II errors in the context of this test.

Table B-2: Definitions of type I and type II errors

Type I Error	k > 0 even though p ≤ BER
Type II Error	k = 0 even though p > BER

A type II error is arguably more serious and we will define n so that the probability of a type II error, β, is acceptable. The probability of a type II error is given in equation B-4.

$$\beta = (1-p)^n < (1-P_0)^n \quad \text{(Equation B-4)}$$

Equation B-4 illustrates that the upper bound on the probability of a type II error is a function of the target bit error rate and n. By solving this equation for n, we can determine the minimum number of bits that need to sent in order to verify that p is less than a given P_0 for a given probability of type II error.

$$n > \frac{\ln(\beta)}{\ln(1-P_0)} \quad \text{(Equation B-5)}$$

Let us now examine the probability of a type I error. The definition of α is given in equation B-6.

$$\alpha = 1-(1-p)^n \le 1-(1-P_0)^n \quad \text{(Equation B-6)}$$

Equation B-6 shows that while we increase n to make β small, we simultaneously raise the upper bound on α. This makes sense since the likelihood of observing a bit error increases with the number of bits that you send, no matter how small bit error rate is. Therefore, while the hypothesis test is very useful in determining a reasonable value for n, we must be very careful in interpreting the results. Specifically, if we send n bits and observe no errors, we are confident that p is less than our target bit error rate (our level of confidence depends on how small we made β). However, if we do observe bit errors, we cannot be quick to assume that the system did not meet the BER target since the probability of a type I error is so large. In the case of k > 0, a confidence interval can be used to help us interpret k.

B.4 – Confidence Interval

The statistical method of confidence intervals will be used to establish a lower bound on the bit error rate given that k > 0. A confidence interval is a range of values that is likely to contain the actual value of some parameter of interest. The interval is derived from the measured value of the parameter, referred to as the point estimate, and the confidence level, (1-α), the probability that the parameter's actual value lies within the interval.

A confidence interval requires a statistical model of the parameter to be bounded. In this case, we use the statistical model for k given in equation B-1. If we were to compute the area under the binomial curve for some interval, we would be computing the probability that k lies within that interval. This concept is shown in figure B-1.

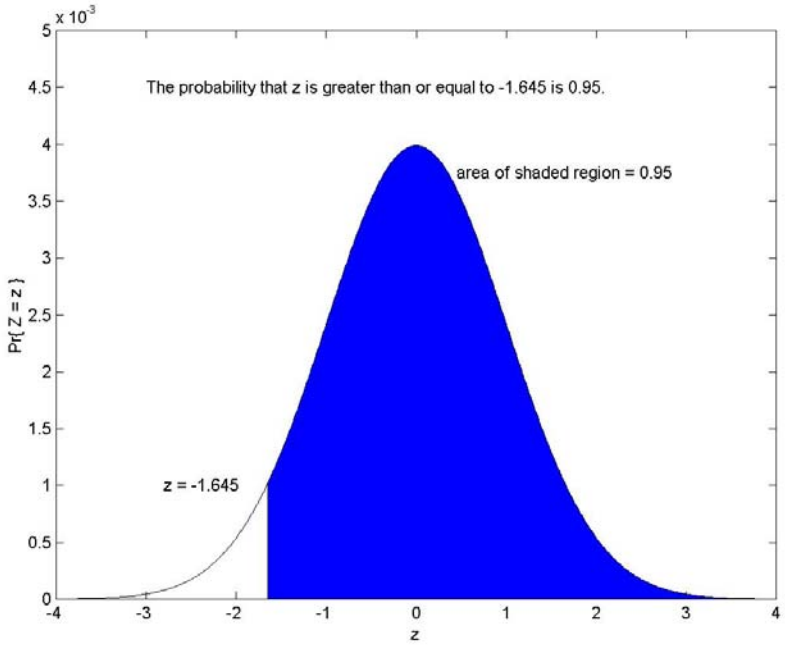

Figure B-1: Computing the probability that z ≥ -1.645 (standard normal distribution).

To compute the area under the binomial curve, we need a value for the parameter p. To compute a confidence interval for k, you assume that k/n, the point estimate for p, is the actual value of p.

Note that figure B-1 illustrates the computation of the lower tolerance bound for k, a special case where the confidence interval is $[k_l, +\infty]$. A lower tolerance bound implies that in a percentage of future tests, the value of k will be greater than k_l. In other words, actual value of k is greater than k_l with probability equal to the confidence level. Therefore, if k_l/n is greater than P_0, we can say that the system does not meet the target bit error rate with probability (1-α). By reducing α, we reduce the probability of making a type I error.

To determine the value of k_l, it is useful to assume that the binomial distribution can be approximated by a normal (Gaussian) distribution when n is large. The mean and variance of this equivalent distribution are the mean and variance of the corresponding binomial distribution (given in equations B-7 and B-8).

$$\mu_K = np \quad \text{(Equation B-7)}$$

$$\sigma_K^2 = np(1-p) \quad \text{(Equation B-8)}$$

Now, let α be the probability that $Z \leq z_\alpha$ where Z is a standard normal random variable. A standard random variable is one whose mean is zero and whose variance is one. The random variable K can be standardized as shown in equation B-9.

$$Z = \frac{K - \mu_K}{\sigma_K} \quad \text{(Equation B-9)}$$

Note that Z is greater than z_α with probability $(1-\alpha)$, the confidence level. We apply this inequality to equation B-9 and solve for K to get equation B-10.

$$K > \mu_K + z_\alpha \sigma_K$$
$$K > np + z_\alpha \sqrt{np(1-p)} \quad \text{(Equation B-10)}$$

As mentioned before, we assume that p is k/n. We can now generate an expression for k_l, the value that K will exceed with probability $(1-\alpha)$. This expression is given in equation B-11.

$$k_l = k + z_\alpha n \sqrt{\frac{(k/n)(1-k/n)}{n}} \quad \text{(Equation B-11)}$$

Finally, we argue that if K exceeds k_l, then the actual value of p must exceed k_l/n. Therefore, we can generate an expression for p_l, the value that p will exceed with probability $(1-\alpha)$, and compare it to the target bit error rate. By applying this comparison (given in equation B-12) the probability of a type I error can be greatly reduced. For example, by setting z_α to –1.645, the probability of a type I error is reduced to 5%.

$$P_0 \geq p_l = \frac{k_l}{n} = \frac{k}{n} + z_\alpha \sqrt{\frac{(k/n)(1-k/n)}{n}} \quad \text{(Equation B-12)}$$

B.5 – Sample Test Construction

We now compress the theory presented in sections B-2 through B-4 into two inequalities that may be used to construct a bit error rate test. First, we take equation B-5 and assume that $\ln(1-P_0)$ is $-P_0$ (valid for P_0 much less than one). The result is equation B-13.

$$n > \frac{-\ln(\beta)}{P_0} \quad \text{(Equation B-13)}$$

Second, we examine equation B-12. Assuming that (1-k/n) is very close to 1 and substituting $-\ln(\beta)/P_0$ for n, we get equation B-14.

$$-\ln(\beta) \geq k + z_\alpha \sqrt{k} \quad \text{(Equation B-14)}$$

The largest value of k that satisfies equation B-14 is k_l. The benefit of these two equations is that a bit error rate test is uniquely defined by β and α and that the test scales with P_0. Table B-3 defines n and k_l in terms of β and α.

Table B-3: n and k_l as a function of β and α.

β	–ln(β)	n	α	z_α	k_l
0.10	2.30	$2.30/P_0$	0.10	–1.29	5
0.10	2.30	$2.30/P_0$	0.05	–1.65	6
0.05	3.00	$3.00/P_0$	0.05	–1.65	7
0.05	3.00	$3.00/P_0$	0.01	–2.33	10
0.01	4.60	$4.60/P_0$	0.05	–1.65	9
0.01	4.60	$4.60/P_0$	0.01	–2.33	13

As an example, let us construct a test to determine if a given system is operating at a bit error rate of 10^{-12} or better. Given that a 5% chance of a type I error is acceptable, the test would take the form of sending $3x10^{12}$ bits and counting the number of errors. If no errors are counted, we are confident that the BER was 10^{-12} or better.

Given that a 5% chance of a type II error is acceptable, we find that k_l is 7. If more than 7 errors are counted, we are confident that the bit error rate is greater than 10^{-12}. However, what if between 1 and 7 errors are counted? These cases may be handled several different ways. One option is to make a statement about the bit error rate (whether it is less than or greater than 10^{-12}) at a lower level of confidence. Another option would be to state that the test result is success since we cannot establish with an acceptable probability of error that the BER is greater than 10^{-12}. Such a statement implies that we failed to meet the burden of proof for the conjecture that the BER exceed 10^{-12}. Of course, the burden of proof could be shifted to the device under test which would imply that any outcome other than k = 0 would correspond to failure (the device under test failed to prove to us that the BER was no more than 10^{-12}). If neither of these solutions are acceptable, it is always an option to perform a more vigorous bit error rate test in order to clarify the result.

B.6 – Packet Error Rate Measurement

It is often easier to measure packet errors than it is to measure bit errors. In these cases, it is helpful to have some linkage between the packet error rate and the bit error rate. To make this linkage, we assume that the bit error rate is low enough and the packet size is small enough so that each packet error contains exactly one bit error.

To complete the linkage, some care must be taken regarding how many packets to send. A bit error is only detectable in the region of the packet that is covered by the cyclic redundancy check (CRC). In the context of Fibre Channel, this region is the first bit after the SOF to the last bit of the CRC. There is no guarantee that errors in the SOF, EOF, and other regions will be detected. Therefore, we must translate n from the number of bits are sent to the number of "observable" bits that are sent. This will increase the test duration since a portion of the time will be spent sending unobservable bits.

For packets of length x bits, at least n/x packets must be sent to perform the equivalent bit error rate test. If no packet errors are observed, the conclusion is that the bit error rate is less than P_0. If more than k_l packet errors are observed, the conclusion is that the bit error rate is greater than P_0.

Note that x is the length of the packet after encoding. In other words, in a 8B10B encoding environment, a 512-byte packet is 5120 bits in length after encoding. Also note that to reinforce the assumption that there is only one bit error per packet error, a test should be run with the shortest possible packets. However, if extremely low bit error rates are to be verified, it may be favorable to use long packets to increase the percentage of observable bits and reduce the test duration.

Appendix C: Test Setups for Specific Device Types

Purpose: To define an optional set of test setups for varying device types and describe the Test Pattern Generator, DUT, Test Channel, and Monitoring Station for each of these setups.

References:

[1] FC-PI-2 Clause 5.1
[2] FC-FS-2
[3] FC-MJSQ A.2.3.3

Last Modification: February 17, 2005 (Version 1.0)

Discussion:

C.1 – Introduction

The ability of a system to maintain the specified bit error rate is an excellent interoperability metric. In this test a Fibre Channel system is built from three parts: a DUT, a Test Channel, and a Test Pattern Generator.

The DUT can be any Fibre Channel device. The Test Channel is the means of connecting the DUT to the Test Pattern Generator. The Test Channel can be optical or electrical, and can be either 'worst-case' or 'ideal'. The Test Pattern Generator can be any Fibre Channel device, or Fibre Channel Test Equipment. If the Test Pattern Generator is a Fibre Channel device, it must be a device that transmits and receives Fibre Channel packets with the DUT. For example, if the DUT were a Target, the Test Pattern generator would be an Initiator.

In order to fully stress the system, the CJTPAT described in FC-MJSQ A.2.3.3 will be used as a stressing pattern in all tests. CJTPAT is hereafter referred to as the stressing pattern.

Since this is a system interoperability test, it is necessary for testing to be performed on all links in the system that involve the DUT. This would mean performing the testing with the DUT acting as the DUT, and as the Test Pattern Generator. The exception to this is Setup 1.

C.2 – Setup 1: DUT Supports Loopback Mode: If the DUT supports loopback mode, Fibre Channel Test Equipment can be used to source the stressing pattern through the Test Channel and DUT, and monitor the transmissions of the DUT to verify that what is transmitted through the loopback of the DUT matches the original transmission of the Test Equipment. In this case the Test Equipment is playing the role of both Test Pattern Generator and Monitoring Station.

Depending on whether the Test Channel used, and the error counting capabilities of the DUT and Test Pattern Generator, a different setup may be used. Setup 1a is the preferred method when testing using either an electrical or optical channel. Since it is difficult to monitor an electrical channel without affecting its characteristics, this is the only setup recommended for testing using an electrical channel. In this setup, Local Error Counters are essential to be able to isolate what side of the channel CRC errors occurred on.

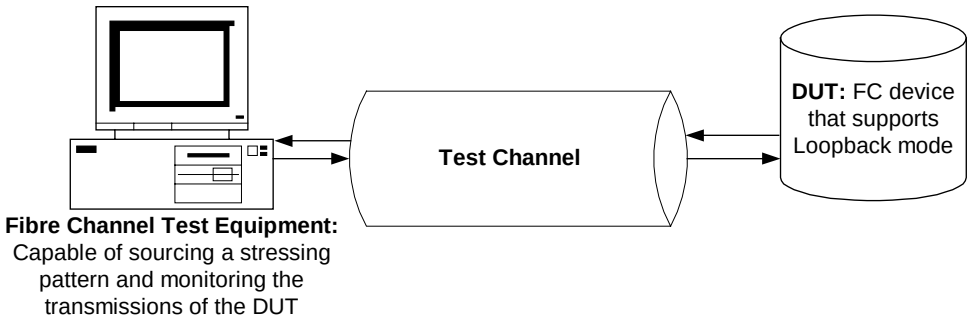

Figure B-1a: Setup 1a – Preferred method for all cases when DUT Supports Loopback Mode.

It is possible to monitor one side of an optical channel without affecting the characteristics of the other side. Setup 1b will work when testing using an optical channel. In this setup, if an ideal channel is used between the DUT and Test Pattern Generator, and a worst case Test Channel is used between the Test Pattern Generator and the DUT, it can be assumed that all CRC errors will occur on worst case Test Channel. Local Error Counters are essential only on the side of the Test Pattern Generator, to be able to isolate what side of the channel CRC errors occurred on. Setup 1b will only work if an optical test channel is used. Setup 1a is preferred for all cases.

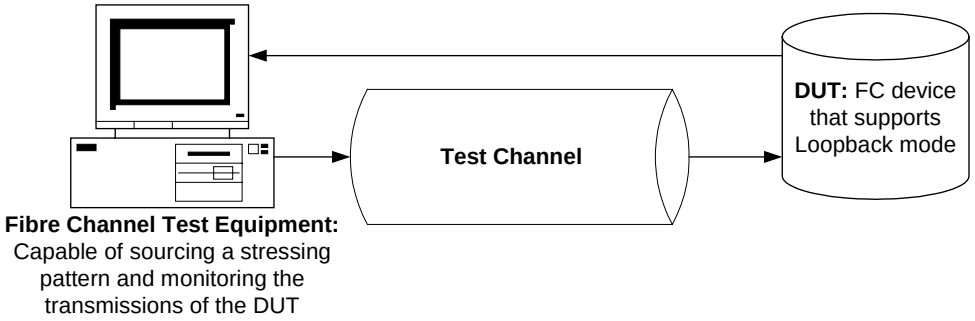

Figure B-1b: Setup 1b - DUT Supports Loopback Mode and an optical test channel is used.

C.3 – Setup 2: DUT is a Target, Test Pattern Generator is an Initiator: The DUT and Test Pattern Generator should be physically connected using the Test Channel. Using a higher layer application, the Test Pattern Generator should source SCSI WRITE commands on the link from the Test Pattern Generator to the DUT. The Data accompanying these WRITE commands should be the stressing pattern. Using Local Error counters on both the Test Pattern Generator and the DUT, the number of detected CRC errors can be determined.

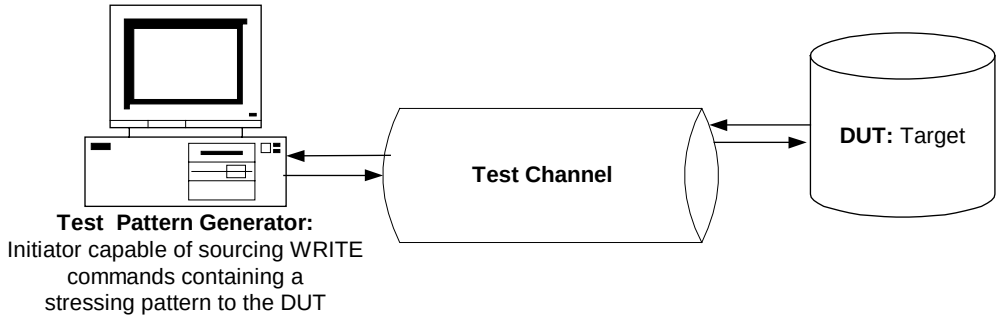

Figure B-2a: Setup 2a - DUT is a Target, Test Pattern Generator is an Initiator.

It is possible to monitor one side of an optical channel without affecting the characteristics of the other side. Setup 2b will work when testing using an optical channel. Local Error Counters are essential only on the side of the Test Pattern Generator, to be able to isolate what side of the channel CRC errors occurred on. Setup 2b will only work if an optical test channel is used. Setup 2a is preferred for all cases. If a Monitoring Station is used , it should monitor the link from the DUT to the Test Pattern Generator for SCSI Response with Status CHECK CONDITION, or for ACC to ABTS or RRQ. These are indicators that the target received a CRC error.

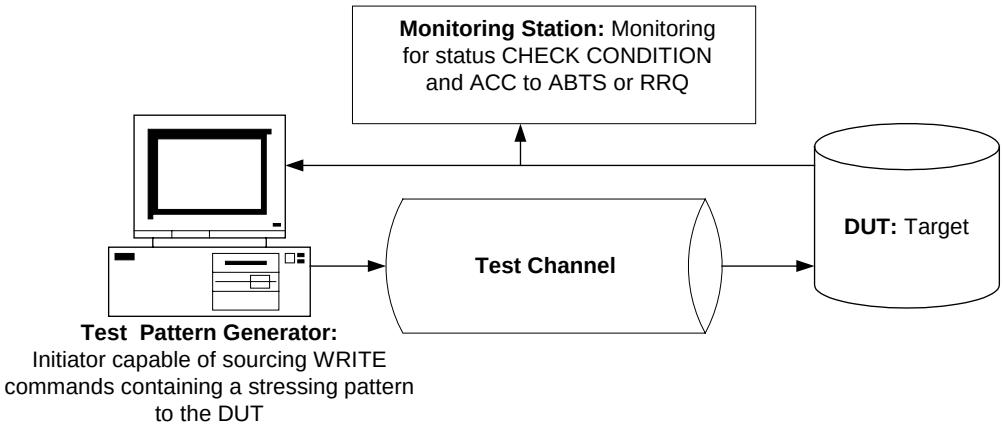

Figure B-2b: Setup 2b - DUT is a Target, Test Pattern Generator is an Initiator, optical Test Channel used.

C.4 – Setup 3: DUT is an Initiator, Test Pattern Generator is a Target: The DUT and Test Pattern Generator should be physically connected using the Test Channel. Using a higher layer application, the DUT should source SCSI WRITE commands on the link from the DUT to the Test Pattern Generator. The Data accompanying these WRITE commands should be the stressing pattern. These WRITE commands will prime the Test Pattern Generator with the stressing pattern. The DUT should then source SCSI READ commands to the Test Pattern Generator for the data that was written to the Test Pattern Generator during the previous WRITE commands. Using Local Error counters on both the Test Pattern Generator and the DUT, the number of detected CRC errors can be determined.

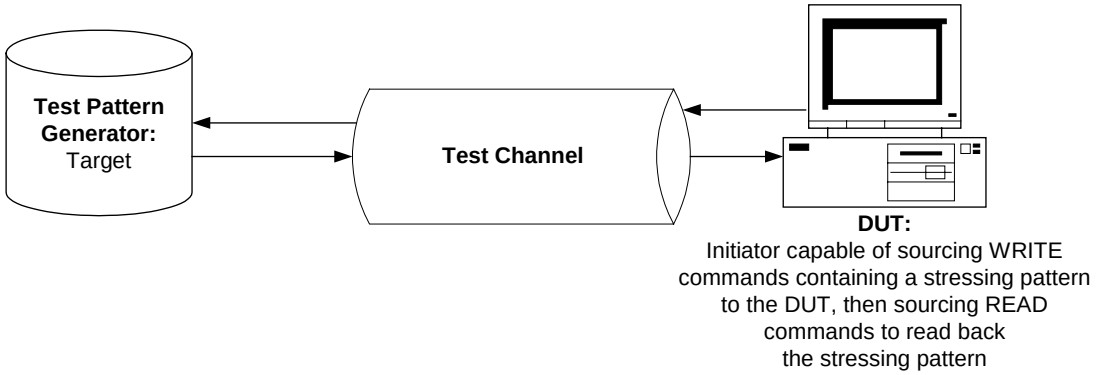

Figure B-3a: Setup 3a DUT is an Initiator, Test Pattern Generator is a Target

It is possible to monitor one side of an optical channel without affecting the characteristics of the other side. Local Error Counters can be used to isolate what side of the channel CRC errors occurred on. Setup 3b will only work if an optical test channel is used. Setup 3a is preferred for all cases. If a Monitoring Station is used, it should monitor the link from the DUT to the Test Pattern Generator for retried SCSI commands.

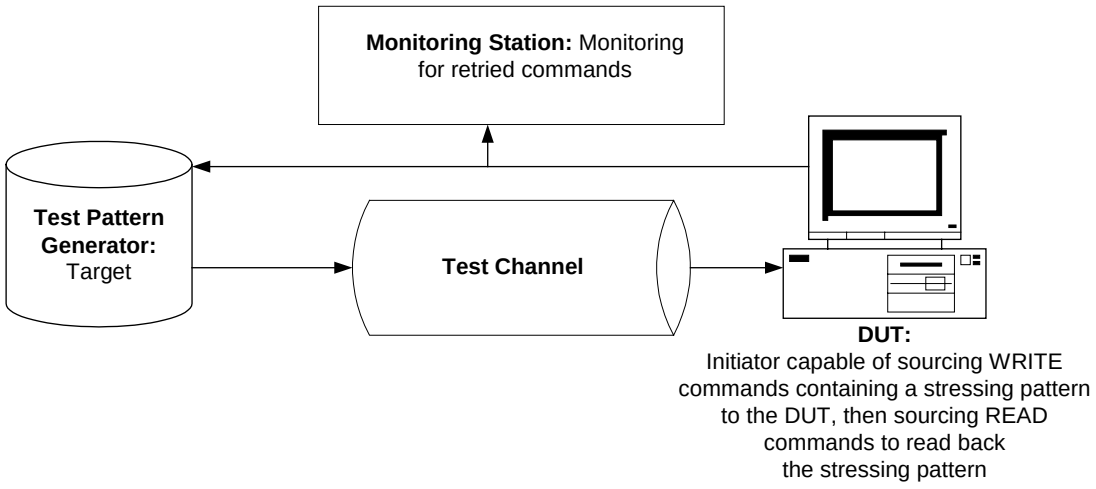

Figure B-3b: Setup 3b DUT is an Initiator, Test Pattern Generator is a Target, optical Test Channel used.

C.5 – Setup 4: DUT is a Switch, Test Pattern Generator is an Initiator: The DUT and Test Pattern Generator should be physically connected using the Test Channel. Using a higher layer application, the DUT should source SCSI WRITE commands on the link from the Test Pattern Generator to a Target attached to the DUT. The Data accompanying these WRITE commands should be the stressing pattern. Using Local Error counters on the Test Pattern Generator, DUT, and the attached Target, the number of detected CRC errors can be determined.

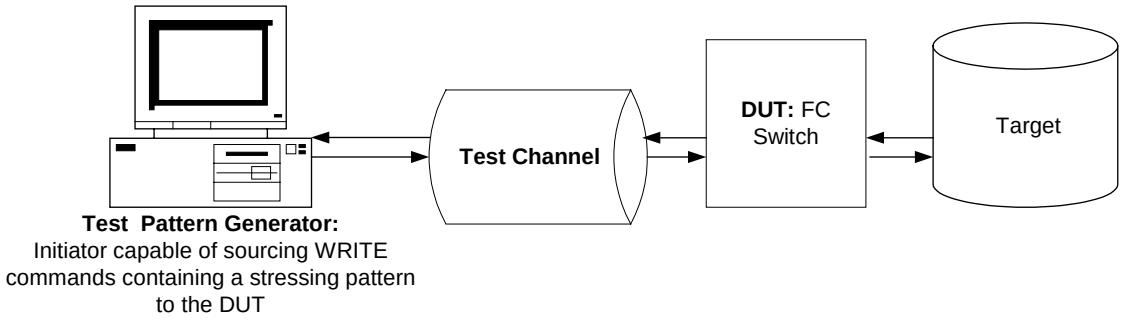

Figure B-4a: Setup 4a DUT is a Switch, Test Pattern Generator is an Initiator

It is possible to monitor one side of an optical channel without affecting the characteristics of the other side. Setup 4b will work when testing using an optical channel. Setup 4a is preferred for all cases. Local Error Counters can be used to isolate what side of the channel CRC errors occurred on. If a Monitoring Station is used , it should monitor the link from the DUT to the Test Pattern Generator for SCSI Response with Status CHECK CONDITION, or for ACC to ABTS or RRQ. These are indicators that the target received a CRC error.

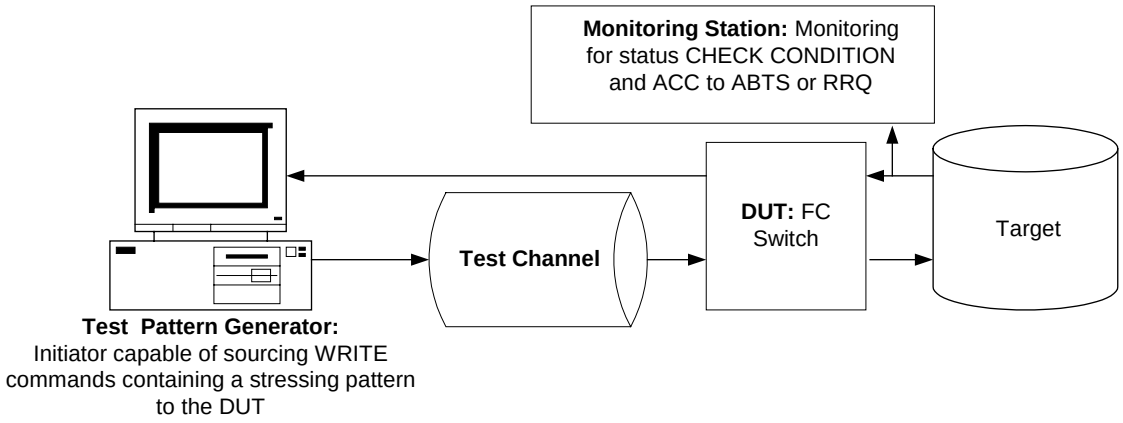

Figure B-4b: Setup 4b DUT is a Switch, Test Pattern Generator is an Initiator , optical test channel is used.

C.6 – Setup 5: DUT is a Switch, Test Pattern Generator is a Target: The DUT and Test Pattern Generator should be physically connected using the Test Channel. Using a higher layer application, the DUT should source SCSI WRITE commands on the link from the DUT to the Test Pattern Generator. The Data accompanying these WRITE commands should be the stressing pattern. These WRITE commands will prime the Test Pattern Generator with the stressing pattern. The DUT should then source SCSI READ commands to the Test Pattern Generator for the data that was written to the Test Pattern Generator during the previous WRITE commands. Using Local Error counters on the Test Pattern Generator, DUT, and the attached Target, the number of detected CRC errors can be determined.

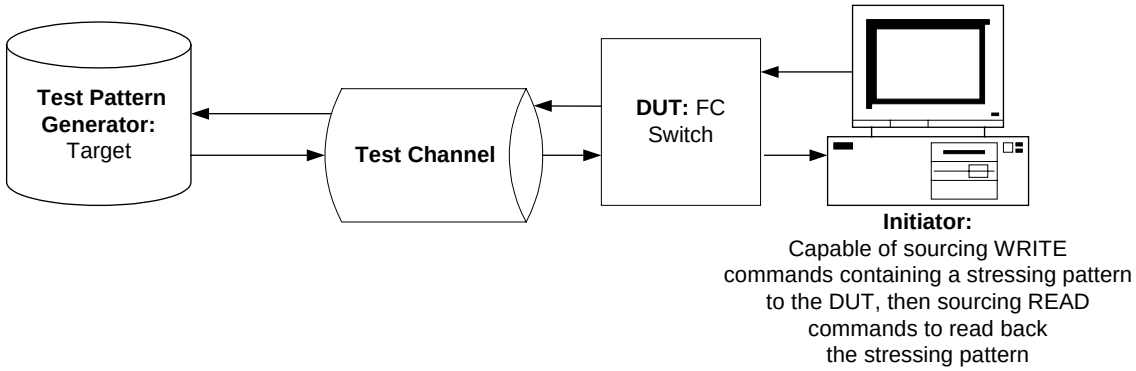

Figure C.5a: Setup 5a DUT is a Switch, Test Pattern Generator is a Target

It is possible to monitor one side of an optical channel without affecting the characteristics of the other side. Setup 5b will work when testing using an optical channel. Setup 5a is preferred for all cases. Local Error Counters can be used to isolate what side of the channel CRC errors occurred on. If a Monitoring Station is used it should monitor the link from the Initiator to the DUT for retried commands.

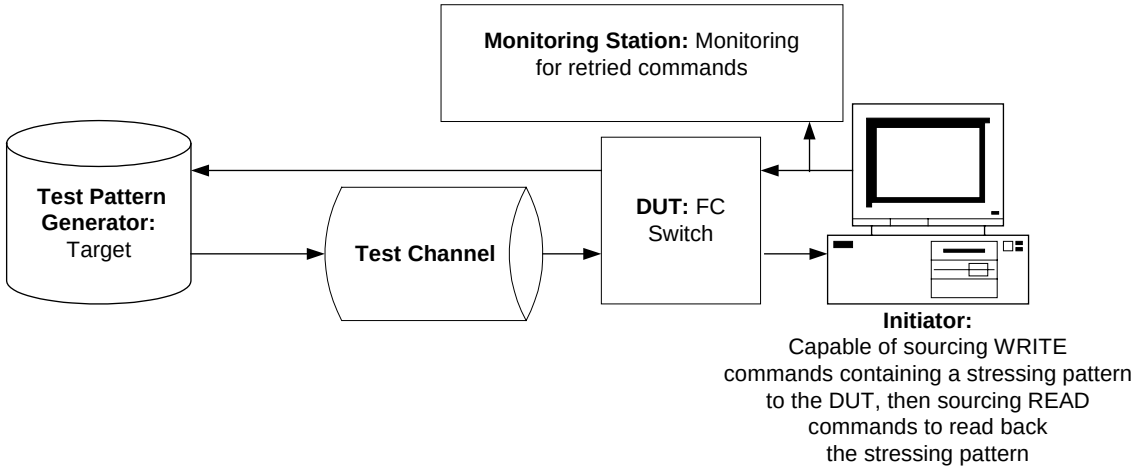

Figure C.5b: Setup 5b DUT is a Switch, Test Pattern Generator is a Target, optical test channel used.