

10GEC

THE 10 GIGABIT ETHERNET CONSORTIUM

RS Test Suite V1.2a

Technical Document



Last Updated: June 7, 2005 6:30 pm

10 Gigabit Ethernet Consortium

*University of New Hampshire
InterOperability Laboratory*

*121 Technology Drive, Suite 2
Durham, NH 03824*

*Phone: +1-603-862-0166
Fax: +1-603-862-4181*

<http://www.iol.unh.edu/consortiums/10gec/>

MODIFICATION RECORD

- June 7, 2005 Version 1.2a Released
 - Modified procedure for test 46.3.4 for devices implementing an XGXS

- February 20, 2005 Version 1.2 Released
 - Modified Test 46.1.3
 - Modified Test 46.2.1
 - Modified Test 46.2.6
 - Modified Procedures for devices implementing a Clause 49 PCS
 - Fixed minor editorial errors and updated references to IEEE 802.3ae-2002

- September 18, 2002 Version 1.1 Released
 - Fixed errors in tests 46.2.4 and 46.3.3

- June 25, 2002 Version 1.0 Released
 - Minor editorial modifications and updating of test suite to D5.0

- October 21, 2001 Version 0.1 Released

ACKNOWLEDGMENTS

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

David Estes	University of New Hampshire
Eric Lynskey	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 evaluate the functioning of their Clause 46 RS based products. The tests do not determine if a product fully conforms to the IEEE Std. 802.3ae-2002 standard. Additionally, successful completion of all tests contained in this suite does not guarantee that the tested device will operate with other Clause 46 RS capable 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 many 10Gb/s 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 by similar functions and further organized by technology. 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 IEEE 802.3 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

MODIFICATION RECORD	2
ACKNOWLEDGMENTS	3
INTRODUCTION	4
TABLE OF CONTENTS	6
Group1: Transmission	7
Test 46.1.1 – Start control character creation and alignment	8
Test 46.1.2 – Terminate control character creation and alignment.....	10
Test 46.1.3 – Deficit Idle Count.....	11
Group2: Reception	13
Test 46.2.1 – Reception of Start control character.....	14
Test 46.2.2 – Reception of Preamble and SFD.....	16
Test 46.2.3 – Reception of Terminate control character.....	17
Test 46.2.4 – Receive IFG Tolerance	18
Test 46.2.5 – Assertion of DATA_VALID_STATUS	19
Test 46.2.6 – De-assertion of DATA_VALID_STATUS	21
Test 46.2.7 – Reception of /E/ during DATA_VALID_STATUS	23
Group3: Fault	25
Test 46.3.1 – Continuous Reception of Fault Sequences.....	26
Test 46.3.2 – Reception of Identical fault_sequences.....	28
Test 46.3.3 – Reception of non-identical fault sequences	30
Test 46.3.4 – Setting of col_cnt	32

Group1: Transmission

Scope: The following tests cover the transmission of data at the Reconciliation Sublayer.

Purpose: These tests are designed to verify that the device under test reacts properly to receipt of data, both valid and invalid, at the Reconciliation Sublayer.

Test 46.1.1 – Start control character creation and alignment

Purpose: To verify that upon reception of the first byte of preamble from the MAC, the RS replaces the preamble byte with a start control character and aligns it to lane 0.

References:

- [1] IEEE Std. 802.3ae-2002 - subclauses 4.2.5, 46.1.7.1.4, 46.2.2, and 46.2.5

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: When the MAC is requested to send a new frame, it calls the procedure PhysicalSignalEncap. This procedure transmits 7 bytes of preamble and then 1 byte of SFD. The preamble pattern has historically been used to stabilize and synchronize the physical medium. In 10 Gigabit Ethernet, the MAC is required to transmit preamble and SFD but it is not necessary for stabilization or synchronization. After the successful transmission of 7 bytes of preamble, the MAC shall transmit an SFD. The entire preamble and SFD pattern is shown below:

10101010 10101010 10101010 10101010 10101010 10101010 10101010 10101011

For a 10Gb/s MAC/RS implementation, the Start control character (0xFB), which replaces the first byte of preamble, is required to be aligned to lane 0 on the XGMII. Error free transmission will not cause the SFD to be on any lane other than lane 3 in the column following the Start control character.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. Instruct the testing station to transmit a valid 64-byte request frame to the DUT.
2. Capture and observe the transmissions from the DUT.
3. Repeat steps 1 and 2 with frame sizes including but not limited to 65, 66, and 67 bytes in length.

Observable results:

- a. If the DUT implements XAUI or 10GBASE-LX4, the DUT should reply to all valid frames with a Start control character aligned to XGMII lane 0 and an SFD aligned to XGMII lane 3. If the DUT implements 10GBASE-R the Start control character should

be observed in position 0 (with the SFD in position 7 of the same block) and in position 4 (with the SFD in position 3 of the following block).

Possible problems: None

Test 46.1.2 – Terminate control character creation and alignment

Purpose: To verify that the DUT inserts a Terminate control character at the end of any frame it transmits, and that the Terminate control character can be aligned to any of the 4 lanes.

References:

- [1] IEEE Std. 802.3ae-2002- subclauses 46.1.7.1.4, 46.2.4, and 46.2.5

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: FMarch 23, 2005

Discussion: The end of frame delimiter is denoted by the assertion of TXC with the appropriate Terminate control character on any one of the four lanes. The Terminate character is represented by TXC=1 and TXD=FD on the lane containing the character across the XGMII, and by the DATA_COMPLETE PLS_DATA.request parameter from the MAC. Unlike the start control character, the terminate control character is valid on any lane. Frames may be any length from 64 bytes to 1518 bytes, and thus the last byte of the CRC can fall on any lane. When transmitting a frame, the RS inserts a terminate control character after the last byte of the CRC. The inter-packet gap begins with this character.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. Instruct the DUT to transmit a 64-byte frame.
2. Capture and observe the transmissions from the DUT.
3. Repeat steps 1 and 2 with frame sizes including but not limited to 65, 66, and 67 bytes in length.

Observable results:

- a. All frames received by the testing station should contain a valid Terminate control character following the last byte of the CRC.
- b. The Terminate control character should be observed on all four XGMII lanes. Additionally, if the DUT implements 10GBASE-R the Terminate control character must be observed in all 8 block positions.

Possible problems: None

Test 46.1.3 – Deficit Idle Count

Purpose: Verify that the DUT properly implements the Deficit Idle Count.

References:

- [1] IEEE Std. 802.3ae-2002 - subclause 46.3.1.4

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: February 20, 2005

Discussion: When the MAC is ready to transmit a frame, it may be necessary for the RS to modify the length of the inter-frame gap in order to align the Start control character to lane 0. One possible way to do this is for the RS to always insert additional idle characters to align the Start control character. However, this method will reduce the effective data rate for frames with minimum inter-frame spacing. Another method allows the RS to insert and delete idle characters to align the Start to lane 0. When using this method, the RS maintains a Deficit Idle Count (DIC) that represents the cumulative count of idle characters inserted or deleted. The count has a minimum value of zero, and a maximum value of three. The DIC is incremented when idle characters are deleted, and decremented when idle characters are inserted. This means that up to three idle characters can be deleted at a time, thus shrinking the minimum inter-frame gap to 9 idle characters. It should also be noted that the average minimum inter-frame gap remains at 12 idle characters. The following table depicts the proper behavior for a device implementing DIC:

Packet Length Modulo 4	Current DIC = 0		Current DIC = 1		Current DIC = 2		Current = 3	
	IPG Length	New DIC value	IPG Length	New DIC value	IPG Length	New DIC value	IPG Length	New DIC value
n+0	12	0	12	1	12	2	12	3
n+1	11	1	11	2	11	3	15	0
n+2	10	2	10	3	14	0	14	1
n+3	9	3	13	0	13	1	13	2

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. The testing station is instructed to transmit the following frame sequence to the DUT, with all frames separated by minimum IPG:
 - PAUSE frame
 - 64-byte frame
 - 65-byte frame
 - 64-byte frame
 - 65-byte frame
 - 64-byte frame
 - 65-byte frame
 - 64-byte frame
 - 65-byte frame
 - 66-byte frame
 - 66-byte frame
 - 65-byte frame
 - 66-byte frame
 - 66-byte frame
 - 67-byte frame
 - 67-byte frame
 - 67-byte frame
 - 67-byte frame
 - 64-byte frame
2. The testing station is instructed to capture and observe all transmissions from the DUT.

Observable results:

- a. The IPG lengths transmitted from the DUT should match the proper IPG lengths for each DIC and packet length in the table above. The first group of 64 and 65-byte frames can be used to determine the initial DIC value.

Possible Problems: If the DUT does not support DIC, then this test cannot be done. Also, if the DUT does not support PAUSE frames then this test sequence may have to be modified and run in several steps.

Group2: Reception

Scope: The following tests cover the reception of data at the Reconciliation Sublayer.

Purpose: These tests are designed to verify that the device under test reacts properly to transmission of data, both valid and invalid, at the Reconciliation Sublayer

Test 46.2.1 – Reception of Start control character

Purpose: To verify that the DUT only accepts frames with proper Start control character alignment.

References:

- [1] IEEE Std. 802.3ae-2002 - subclauses 4.2.5, 46.1.7.1.4, 46.2.2, and 46.2.5

Resource Requirements:

- A testing station capable of encoding (decoding) 8-bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: February 20, 2005

Discussion: An RS that is transmitting a frame is required to align the start of frame to lane 0. Once the RS has transmitted a frame, there is no valid mechanism to realign the start of frame to any other lane. Therefore, any frame that is received with a start character on a lane other than lane 0 indicates that an error condition is present and the frame should be discarded.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

Part A (constant preamble length, SFD on lanes 0 - 2)

1. The testing station is instructed to transmit a properly encapsulated, valid 64-byte request frame with the Start control character aligned to lane 0.
2. The testing station is instructed to transmit a properly encapsulated, valid 68-byte request frame with the Start control character aligned to lane 1, 6 bytes of preamble, and an SFD.
3. The testing station is instructed to transmit a properly encapsulated, valid 72-byte request frame with the Start control character aligned to lane 0.
4. Capture and observe transmissions from the DUT.
5. Repeat steps 1 through 4 changing the Start control character alignment in step 2 to lanes two and three.

Part B (variable preamble length, SFD on lane 3)

1. The testing station is instructed to transmit a properly encapsulated, valid 64-byte request frame with the Start control character aligned to lane 0.
2. The testing station is instructed to transmit a properly encapsulated, valid 68-byte request frame with the Start control character aligned to lane 1, 5 bytes of preamble, and an SFD aligned to lane 3.
3. The testing station is instructed to transmit a properly encapsulated, valid 72-byte request frame with the Start control character aligned to lane 0.

4. Capture and observe transmissions from the DUT.
5. Repeat steps 1 through 4 changing the Start control character alignment in step 2 to lanes two and three, also shrinking the preamble length to align the SFD to lane 3.

Observable results:

- a. The DUT should discard the 68 byte frame and reply to the 64 and 72 byte frames.
- b. The DUT should discard the 68 byte frame and reply to the 64 and 72 byte frames.

Possible Problems: If the DUT implements a 10GBASE-R PCS only the reception of a Start control character in position 0 and 4 can be verified, as there is no permissible encoding for an /S/ in any position other than 0 or 4. In these cases the test will include 6 bytes of preamble and an SFD in position 3 or 7.

Test 46.2.2 – Reception of Preamble and SFD

Purpose: To determine if the DUT is (in)sensitive to preamble shrinkage or growth.

References:

- [1] IEEE Std. 802.3ae-2002 - subclauses 4.2.9 process BitReceiver and procedure PhysicalSignalDecap, 46.3.3.3

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: February 20, 2005

Discussion: Preamble is not required for clock synchronization in a 10GBASE-X/R/W frame. Therefore the DUT should not be sensitive to its size. When the MAC is receiving a frame, the process BitReceiver first calls the procedure PhysicalSignalDecap. This procedure receives one bit at a time from the physical medium and discards all bits until a valid SFD is detected. At this point the BitReceiver process accepts bits while the receiveDataValid signal is asserted and the frame is not finished. However, a 10 Gb/s MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 3 of the column following the column containing the Start control character. This test seeks to determine the (in)sensitivity of the MAC/RS to reception of preamble.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. The testing station is instructed to transmit a request frame with only a Start Control character and SFD present. The output from the DUT is observed.
2. Step 1 is repeated inserting additional preamble between the Start Control character and SFD until the length of the Start, preamble, and SFD is 16 octets long.
3. For a DUT implementing a 10GBASE-R PCS, steps 1 and 2 are repeated to ensure that the Start control character is tested in positions 0 and 4.

Observable results:

- a. The DUT should reply to all frames containing a properly aligned Start control character in lane 0, 6 bytes of preamble, and an SFD in XGMII lane 3.
- b. The DUT may reply to frames with preamble variations other than those listed in part (a).

Possible Problems: None

Test 46.2.3 – Reception of Terminate control character

Purpose: Verify that the DUT can receive frames containing the Terminate control character in any lane.

References:

- [1] IEEE Std. 802.3ae-2002- subclauses 46.2.4, 46.2.5

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: A DUT is allowed to transmit any size frame from 64-1518 bytes in length. Therefore, it is possible for the frame to terminate on any of the four lanes. A receiver must therefore be insensitive to the Terminate control character being on any lane.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. The testing station is instructed to transmit a properly encapsulated, valid 64-byte request frame.
2. The testing station is instructed to capture and observe all transmissions from the DUT.
3. Repeat steps 1 and 2 with frame sizes including but not limited to 65-67 bytes in length to force the Terminate control character to appear on all of the four XGMII lanes.

Observable results:

- a. The DUT should reply to frames containing the Terminate control character on any of the 4 XGMII lanes. Additionally, if the DUT implements 10GBASE-R the DUT should reply to frames containing the Terminate control character in any of the 8 block positions.

Possible Problems: None

Test 46.2.4 – Receive IFG Tolerance

Purpose: Verify that the DUT can properly receive frames with an IFG between 5 and 12 bytes in length.

References:

- [1] IEEE Std. 802.3ae-2002- subclauses 4.4.2, 46.2.1, 46.3.1.4

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: February 20, 2005

Discussion: It may be necessary for the RS to modify the length of the inter-frame gap in order to align the Start control character to lane 0. One method allows the RS to insert and delete idle characters to align the Start to lane 0. When using this method, the RS maintains a Deficit Idle Count (DIC) that represents the cumulative count of idle characters inserted or deleted. This means that up to three idle characters can be deleted at a time, thus shrinking the minimum inter-frame gap to 9 idle characters. However, due to variable network delays and clock tolerances, it is possible for an additional column of Idle to be deleted, which would shrink the minimum IFG to 5 bytes. Taking both DIC and clock compensation into consideration, the minimum IFG can vary from 5 to 12 bytes under normal network conditions.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. The testing station is instructed to transmit a valid request frame followed by a total of 5 bytes of IFG.
2. The testing station is instructed to transmit a 68-byte request frame followed by a total of 12 bytes of IFG.
3. The testing station is instructed to transmit a 72-byte request frame.
4. The testing station is instructed to capture and observe all transmissions from the DUT.
5. Parts 1 through 4 are repeated increasing the IFG in part 1 up to 12 bytes in length. Additionally, the frame size may need to be increased or decreased to properly transmit the correct amount of IFG.

Observable results:

- a. For every tested value of IFG, the DUT should reply to all three frames.

Possible Problems: None

Test 46.2.5 – Assertion of DATA_VALID_STATUS

Purpose: Verify that the DUT does not accept frames that are preceded by anything but a full column of Idle or a sequence ordered set.

References:

- [1] IEEE Std. 802.3ae-2002 – subclause 46.1.7.5.3

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: February 20, 2005

Discussion: The receiving RS sets the values of the PLS_DATA_VALID.indicate service primitive whenever the DATA_VALID_STATUS parameter changes. The parameter can take on a value of DATA_VALID or DATA_NOT_VALID, depending on the data received. When a start control character is received on lane 0, the RS sets DATA_VALID_STATUS to DATA_VALID if the previous column contained four valid Idle characters, or a Sequence ordered set. Anything else received by the RS in the column preceding the start control character will cause the RS to set DATA_VALID_STATUS to DATA_NOT_VALID.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

Test Cases:

1. A full column of Idle
 2. A sequence ordered set corresponding to Local Fault
 3. A sequence ordered set corresponding to Remote Fault
 4. A sequence ordered set corresponding to a reserved value
 5. A column containing a Terminate control character
 6. A column containing a Start control character
 7. A column containing an Error control character
 8. A column containing Data code groups
1. The testing station is instructed to transmit a valid 64-byte request frame followed by 12 bytes of Idle.
 2. The testing station is instructed to transmit the pattern described in Test Case 1.
 3. The testing station is instructed to transmit a valid 68-byte request frame followed by 12 bytes of Idle, followed by a 72-byte request frame.
 4. The testing station is instructed to capture and observe all transmissions from the DUT.
 5. Repeat steps 1 through 4 for all listed test cases.

Observable results:

- a. The DUT should reply to all three frames in test case 1.
- b. The DUT should reply to all three frames in test case 2.
- c. The DUT should reply to all three frames in test case 3.
- d. The DUT should reply to all three frames in test case 4.
- e. The DUT should reply to only the two 64-byte frames in test case 5.
- f. The DUT should reply to only the two 64-byte frames in test case 6.
- g. The DUT should reply to only the two 64-byte frames in test case 7.
- h. The DUT should reply to only the two 64-byte frames in test case 8.

Possible Problems: If the DUT implements a 10GBASE-R PCS then cases 6, 7, and 8 cannot be tested, as there are no valid transitions for these test cases in the 10GBASE-R PCS Receive State Diagram.

Test 46.2.6 – De-assertion of DATA_VALID_STATUS

Purpose: Verify that the DUT properly accepts and discards the appropriate frames when DATA_VALID_STATUS changes from DATA_VALID to DATA_NOT_VALID.

References:

- [1] IEEE Std. 802.3ae-2002 – subclause 46.1.7.5.3

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: The receiving RS sets the values of the PLS_DATA_VALID.indicate service primitive whenever the DATA_VALID_STATUS parameter changes. This parameter can take on a value of DATA_VALID or DATA_NOT_VALID, depending on the data received. Upon reception of a valid start control character in lane 0 that is preceded by a full column of Idle or a Sequence ordered set, the RS will set DATA_VALID_STATUS to DATA_VALID. As long as no errors occur in the frame, the DUT will set DATA_VALID_STATUS to DATA_NOT_VALID upon the reception of a Terminate control character. When DATA_VALID_STATUS takes on a value of DATA_NOT_VALID for any control character other than a Terminate, the RS will make sure that the MAC detects a CRC error in the frame.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

Part A (DUT sets DATA_NOT_VALID for a control character other than Terminate)

1. Instruct the testing station to transmit a properly formed 64-byte request frame to the DUT followed by minimum IPG.
2. Instruct the testing station to transmit a 68-byte request frame to the DUT that replaces the terminate control character with an idle control character followed by minimum IPG.
3. Instruct the testing station to transmit a properly formed 72-byte request frame to the DUT.
4. Observe all transmissions from the DUT.
5. Repeat steps 1 through 4 replacing the Terminate control character in part 2 with an /O/, and then with an /S/.

Part B (DUT sets DATA_NOT_VALID for a Terminate control character)

1. Instruct the testing station to transmit a properly formed 64-byte request frame to the DUT followed by minimum IPG.
2. Instruct the testing station to transmit a properly terminated 68-byte request frame to the DUT.
3. Transmit S/D/D/D in the XGMII column following the Terminate control character, followed by minimum IPG.
4. Instruct the testing station to transmit a properly formed 72-byte request frame to the DUT.
5. Observe all transmissions from the DUT.
6. Repeat steps 1 through 5 replacing the column in step 3 with D/D/D/D, T/K/K/K, and E/E/E/E. Fault sequences should also be transmitted in step 3 if the DUT implements 10GBASE-R.

Observable results:

- a. The DUT should reply to both the 64-byte and 72-byte frames, discard the 68-byte frame in step 2, and should increment its CRC error counter.
- b. The DUT should reply to all three frames.

Possible Problems: Part A cannot be done for a DUT implementing a 10GBASE-R PCS, as these are no valid transitions in the PCS Receive state machine for these test cases. In Part B with a device implementing 10GBASE-R, test cases D/D/D/D, E/E/E/E, and T/K/K/K cannot be tested; these test cases do not represent valid transitions in the PCS Receive state machine.

Test 46.2.7 – Reception of /E/ during DATA_VALID_STATUS

Purpose: Verify that the DUT properly discards frames that are received with errors and increments its CRC error counters.

References:

- [1] IEEE Std. 802.3ae-2002 – subclauses 46.1.7.5.3, 46.3.3.1

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: During frame reception, if a control character other than a Terminate control character is received, then the RS will make sure that the MAC indicates a CRC error for that frame. In particular, if the received frame contains one or more Error control characters in an otherwise valid frame, the RS must still cause the MAC to detect a CRC error.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

Test Cases:

1. /E/ after /S/, before SFD
 2. /E/ replacing SFD
 3. /E/ after SFD, replacing one octet of an otherwise valid frame
 4. /E/ replacing /T/
1. Instruct the testing station to transmit a properly formed 64-byte request frame to the DUT followed by minimum IPG.
 2. Instruct the testing station to transmit a 68-byte request frame to the DUT that contains the pattern described in Test Case 1, followed by minimum IPG.
 3. Instruct the testing station to transmit a properly formed 72-byte request frame to the DUT.
 4. Observe all transmissions from the DUT.
 5. Repeat steps 1 through 4 for test cases 2 through 5.

Observable results:

- a. The DUT should reply to both the 64-byte and 72-byte frames, discard the 68-byte frame, and increment its CRC error counter.

Possible Problems: For a DUT implementing 10GBASE-R, only a full block of /E/ within the frame can be tested.

Group3: Fault

Scope: The following tests cover the reception and transmission of fault sequences at the Reconciliation Sublayer.

Purpose: These tests are designed to verify that the device under test reacts properly to receipt and transmission of fault sequences, both valid and invalid, at the Reconciliation Sublayer.

Test 46.3.1 – Continuous Reception of Fault Sequences

Purpose: Verify that the DUT properly reacts upon the continuous reception of either Local Fault or Remote Fault ordered_sets.

References:

- [1] IEEE Std. 802.3ae-2002 - subclause 46.3.4

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: February 20, 2005

Discussion: Clause 46 clearly states “Sublayers within the PHY are capable of detecting faults that render a link unreliable for communication. Upon recognition of a fault condition a PHY sublayer indicates Local Fault status on the data path. When this Local Fault status reaches an RS, the RS stops sending MAC data, and continuously generates a Remote Fault status on the transmit data path (possibly truncating a MAC frame being transmitted). When an RS receives Remote Fault status, the RS stops sending MAC data, and continuously generates Idle control characters. When the RS no longer receives fault status messages, it returns to normal operation, sending MAC data.” The following table depicts the defined fault messages:

Lane 0	Lane 1	Lane 2	Lane 3	Description
Sequence	0x00	0x00	0x00	Reserved
Sequence	0x00	0x00	0x01	Local Fault
Sequence	0x00	0x00	0x02	Remote Fault
Sequence	≥0x00	≥0x00	≥0x03	Reserved

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

1. Establish a link between the testing station and the DUT.
2. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
3. Instruct the testing station to continuously transmit Local Fault sequences to the DUT.
4. Observe all transmissions from the DUT.
5. Repeat steps 1 through 4 with the testing station transmitting Remote Fault sequences.
6. Repeat steps 1 through 4 with the testing station transmitting Reserved sequences.

Observable results:

- a. Upon reception of the Local Fault sequences, the DUT should cease transmission of frames and commence continuous transmission of Remote Fault sequences.
- b. Upon reception of the Remote Fault sequences, the DUT should cease transmission of frames and commence continuous transmission of Idle.
- c. Upon reception of the Reserved sequences, the DUT should not cease transmission of frames.

Possible Problems: If the DUT implements XAUI or 10GBASE-LX4, then the testing station will not receive continuous Remote Fault sequences from the DUT as described in part a; consecutive fault_sequences will follow an ||A|| column and be separated by 16-31 ||K|| and ||R|| columns. Also, when transmitting fault sequences to a DUT implementing XAUI or 10GBASE-LX4 from the testing station, they must follow a properly spaced ||A|| column.

Test 46.3.2 – Reception of Identical fault_sequences

Purpose: Determine the number of identical fault_sequences that the DUT needs to receive before acknowledging the reception of Local or Remote fault.

References:

- [1] IEEE Std. 802.3ae-2002 - subclauses 46.3.4, 46.3.4.3

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: The link fault signaling state machine specifies that the RS will recognize a link_fault when 4 fault_sequences containing the same value have been received without receiving any other fault_sequences in a period of 128 columns or less. Upon reception of a fault_sequence, the RS begins a count with the *col_cnt* variable. Upon reception of four Local Fault or four Remote Fault fault_sequences, the DUT sets the link_fault variable to Local Fault or Remote Fault, respectively.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

Case 1. DUT does not implement an XGXS.

1. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
2. Instruct the testing station to send 1 Local Fault ordered_set to the DUT.
3. Observe all transmissions from the DUT.
4. Repeat parts 1 through 3 increasing the number of ordered_sets until the DUT momentarily stops transmitting frames. Let the number of ordered_sets be referred to as *n*.
5. Repeat parts 1 through 4 sending Remote Fault ordered_sets to the DUT.
6. Repeat parts 1 through 4 using *n* identical ordered_sets that are neither Local Fault nor Remote Fault.

Case 2. DUT implements one or more XGXS.

1. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
2. Instruct the testing station to send 1 column of ||A||, followed by 1 Local Fault ordered_set, followed by 16 columns of ||R|| and/or ||K|| to the DUT.

3. Repeat part 2 zero times.
4. Observe all transmissions from the DUT.
5. Repeat parts 1 through 4, increasing the number of times that part 2 is repeated in part 3 until the DUT momentarily stops transmitting frames. Let the number of ordered_sets be referred to as n .
6. Repeat parts 1 through 5 sending Remote Fault ordered_sets to the DUT.
7. Repeat parts 1 through 5 sending n identical ordered_sets that are neither Local Fault nor Remote Fault to the DUT.

Observable results:

Case 1 and Case 2.

- a. The DUT should continuously transmit frames when the value of n is less than 4 and the ordered_sets represent Local Fault fault_sequences. The DUT should momentarily stop transmitting frames and transmit RF when the value of n is 4 or greater.
- b. The DUT should continuously transmit frames when the value of n is less than 4 and the ordered_sets represent Remote Fault fault_sequences. The DUT should momentarily stop transmitting frames when the value of n is 4 or greater.
- c. The DUT continuously transmit frames regardless of the value of n when the ordered_sets represent neither Local Fault nor Remote Fault fault_sequences.

Possible Problems: If the DUT does not properly increment the col_cnt variable, then Case 2 cannot be done, as this test case relies on the DUT to be able to recognize reception of fault_sequences within 128 columns. Also, in some systems it may not be possible to determine if the DUT has entered the fault state upon reception of RF, as RF will only cause the DUT to transmit Idle and may be indistinguishable from a large IPG.

Test 46.3.3 – Reception of non-identical fault sequences

Purpose: Verify that the DUT properly resets the seq_cnt variable to 0 upon reception of non-identical ordered_sets.

References:

- [1] IEEE Std. 802.3ae-2002 - subclauses 46.3.4, 46.3.4.3

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: The link fault signaling state machine specifies that the RS will recognize a link_fault when 4 fault_sequences containing the same value have been received without receiving any other fault_sequences in a period of 128 columns or less. Upon reception of a fault_sequence, the RS begins a count with the col_cnt variable. Upon reception of four Local Fault or four Remote Fault fault_sequences, the DUT sets the link_fault variable to Local Fault or Remote Fault, respectively. However, when four identical fault_sequences are not received, the seq_cnt variable is reset to zero, and the DUT should not set the link_fault variable.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

For this test, the value of n is obtained from test 46.3.2.

Case 1. DUT does not implement an XGXS.

1. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
2. Instruct the testing station to transmit $n-1$ Local Fault ordered_sets followed by m Remote Fault ordered_sets (where m is initially set to 1), followed by 1 Local Fault ordered_set.
3. Observe all transmissions from the DUT.
4. Repeat steps 1 through 3 increasing the value of m in step 2 until $m = n$.
5. Repeat steps 1 through 4 transmitting $n-1$ Remote Fault ordered_sets followed by m Local Fault ordered_sets, and 1 Remote Fault ordered_set, resetting m to 1.
6. Repeat steps 1 through 4 transmitting $n-1$ Local Fault ordered_sets followed by m Reserved ordered_sets, and 1 Local Fault ordered_set, resetting m to 1.
7. Repeat steps 1 through 4 transmitting $n-1$ Remote Fault ordered_sets followed by m Reserved ordered_sets, and 1 Remote Fault ordered_set, resetting m to 1.

8. Repeat steps 1 through 4 replacing step 2 with the following: Instruct the testing station to transmit 1 Local Fault ordered_set followed by 1 Remote Fault ordered_set. Repeat this pattern such that the total number of ordered_sets is $2*n$.
9. Repeat step 8 with 1 Local Fault ordered_set followed by 1 Reserved ordered_set.
10. Repeat step 9 with 1 Remote Fault ordered_set followed by 1 Reserved ordered_set.

Case 2. DUT implements one or more XGXS.

1. Repeat Case 1 but insert a column of ||A|| before every ordered_set, and 31 columns of ||K|| and/or ||R|| between all ||A|| columns.

Observable results:

Case 1 and Case 2.

- a. The DUT should continuously reply to the frames when the value of m is less than the value of n when the testing station is transmitting $n-1$ Local Fault ordered_sets followed by m Remote Fault ordered_sets and 1 Local Fault ordered_set.
- b. The DUT should momentarily stop replying to the frames when the value of m is greater than or equal to the value of n when the testing station is transmitting $n-1$ Local Fault ordered_sets followed by m Remote Fault ordered_sets and 1 Local Fault ordered_set.
- c. The DUT should continuously reply to the frames when the value of m is less than the value of n when the testing station is transmitting $n-1$ Remote Fault ordered_sets followed by m Local Fault ordered_sets and 1 Remote Fault ordered_set.
- d. The DUT should momentarily stop replying to the frames when the value of m is greater than or equal to the value of n when the testing station is transmitting $n-1$ Remote Fault ordered_sets followed by m Local Fault ordered_sets and 1 Remote Fault ordered_set.
- e. The DUT should momentarily stop replying to the frames after receiving n Local Fault ordered_sets in part 6.
- f. The DUT should momentarily stop replying to the frames after receiving n Remote Fault ordered_sets in part 7.
- g. The DUT should continuously reply to the frames in part 8.
- h. The DUT should momentarily stop replying to the frames in part 9.
- i. The DUT should momentarily stop replying to the frames in part 10.

Possible Problems: In some systems it may not be possible to determine if the DUT has entered the fault state upon reception of RF, as RF will only cause the DUT to transmit Idle and may be indistinguishable from a large IPG.

Test 46.3.4 – Setting of col_cnt

Purpose: Verify that the DUT properly uses the col_cnt variable.

References:

- [1] IEEE Std. 802.3ae-2002 - subclauses 46.3.4, 46.3.4.3

Resource Requirements:

- A testing station capable of encoding (decoding) 8 bit octets to (from) 10-bit code groups as specified in clause 48 and sending (receiving) these code groups using the signaling method described in clause 53, or a testing station capable of encoding (decoding) 64-bit words to (from) 66-bit words as specified in clause 49 and optionally clause 50, and sending (receiving) these code groups using the signaling method described in clause 52.

Last Modification: March 23, 2005

Discussion: The link fault signaling state machine specifies that the RS will recognize a link_fault when 4 fault_sequences containing the same value have been received without receiving any other fault_sequences in a period of 128 columns or less. Upon reception of a fault_sequence, the RS begins a count with the col_cnt variable. Upon reception of four local fault or four remote fault_sequences, the DUT sets the link_fault variable to Local Fault or Remote Fault, respectively. However, when four identical fault_sequences are not received within 128, the col_cnt variable is reset to zero, and the DUT should not set the link_fault variable.

Test Setup: Connect the device under test (DUT) to the testing station (transmit to receive, receive to transmit) with the appropriate medium (i.e. multi-mode fiber or single mode fiber).

Procedure:

For this test, the value of n is taken from test 46.3.2.

Case 1. DUT does not implement an XGXS.

Part 1 (The DUT has not yet set link_fault to Local Fault or Remote Fault)

1. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
2. The testing station is instructed to send n Local Fault ordered_sets to the DUT, with m columns of Idle between each ordered_set, where m is initially set to less than 128.
3. Repeat parts 1 through 2 increasing the value of m until the DUT momentarily stops transmitting frames.
4. Repeat parts 1 through 3 sending Remote Fault ordered_sets to the DUT.

Part 2 (The DUT has set link_fault to Local Fault or Remote Fault)

5. The testing station is instructed to send continuous Local Fault ordered_sets to the DUT.
6. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.

7. The testing station is instructed to send m columns of Idle and frames to the DUT, with m initially set to less than 128.
8. Observe all transmissions from the DUT.
9. Repeat parts 5 through 8, incrementing m until the DUT momentarily exits the fault state and begins transmission of frames.
10. Repeat parts 5 through 9 sending Remote Fault ordered_sets in step 5.

Case 2. DUT implements one or more XGXS.

Part 1 (The DUT has not yet set link_fault to Local Fault or Remote Fault)

1. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
2. The testing station is instructed to send a 116-byte frame, 1 Local Fault ordered_set, and 93 columns of Idle to the DUT.
3. The testing station is instructed to send a 116-byte frame, 1 Local Fault ordered_set, and 93 columns of Idle to the DUT.
4. The testing station is instructed to send 2 116-byte frames, each followed by 1 Local Fault ordered_set and 3 columns of Idle to the DUT.
5. Repeat steps 1 through 4, increasing the frame size (by multiples of 4) in step 3 until the DUT does not enter the Fault state.
6. Repeat steps 1 through 5 sending Remote Fault ordered_sets to the DUT.

Part 2 (The DUT has set link_fault to Local Fault or Remote Fault)

7. The testing station is instructed to send continuous Local Fault ordered_sets to the DUT.
8. Instruct the DUT to transmit frames continuously to the testing station with minimum IPG.
9. The testing station is instructed to send a 116-byte frame, 1 Local Fault ordered_set, and 93 columns of Idle to the DUT.
10. The testing station is instructed to send a 116-byte frame, 1 Local Fault ordered_set, and 93 columns of Idle to the DUT.
11. The testing station is instructed to resume continuous transmission of Local Fault ordered_sets to the DUT.
12. Repeat steps 7 through 11, increasing the frame size (by multiples of 4) in step 10 until the DUT exits the Fault state.
13. Repeat steps 7 through 12 sending Remote Fault ordered_sets to the DUT.

Observable results:

Case 1 and Case 2.

- a. In Part 1, when the value of m is below 128, the DUT should momentarily stop transmitting frames.
- b. In Part 1, when the value of m is greater than or equal to 128, the DUT should continuously transmit frames.
- c. In Part 2, when less than 128 columns of Idle are being sent to the DUT, the DUT should not exit the fault state.
- d. In Part 2, when 128 columns of Idle are being sent to the DUT, the DUT should momentarily exit the fault state and begin frame transmission.

Possible Problems:

Part 1

In some systems it may not be possible to determine if the DUT has entered the fault state upon reception of RF, as RF will only cause the DUT to transmit Idle and may be indistinguishable from a large IPG.

Part 2

In some systems it may not be possible to determine if the DUT has exited the fault state after having previously received RF, as the capture may only show Idle if there is a large delay before the DUT will begin frame transmission.