

Gigabit Ethernet Consortium

Clause 31 Flow Control Test Suite

Technical Document



Last Updated: May 20, 2003, 10:29 AM

Ethernet Testing Services

Gigabit Ethernet Consortium

*InterOperability Laboratory
Research Computing Center
University of New Hampshire*

121 Technology Drive, Suite 2

Durham, NH 03824

Phone: (603) 862-0166

Fax: (603) 862-0898

<http://www.iol.unh.edu/consortiums/ge>

MODIFICATION RECORD

- May 20, 2003 Organized and combined all Tests into one document

The University of New Hampshire
InterOperability Lab

ACKNOWLEDGMENTS

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

Rupert Dance	University of New Hampshire
Adam Healey	University of New Hampshire
Mike Henninger	University of New Hampshire
Gerard Nadeau	University of New Hampshire
Robert Noseworthy	University of New Hampshire
Barry Reinhold	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 4 Media Access Control (MAC) based products. The tests do not determine if a product conforms to the IEEE 802.3, nor are they purely interoperability tests. Rather, they provide one method to isolate problems within a MAC device. Successful completion of all tests contained in this suite does not guarantee that the tested device will operate with other 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 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 University of New Hampshire
InterOperability Lab

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.

Test #31.1.1 - Receive PAUSE frame with zero pause_time

Purpose: To verify that the device under test (DUT) operates normally when it receives a PAUSE frame with pause_time set to zero.

References:

- IEEE 802.3z, 1998 Edition - subclause 37.2.1.4 and 37.2.4.2
- IEEE 802.3x/D3.2 February 9, 1997 - subclauses 31.3 and 31B.3.3
- Appendix 31A - Flow Control modes
- Appendix 31B - Generating packet streams
- Appendix 31C - PAUSE frame format

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 24, 1998

Discussion: MAC Control sublayers implementing the PAUSE function incorporate a timer called pause_timer. When pause_timer is non-zero, the MAC Control sublayer inhibits the transmission of frames by the MAC client. The timer counts down, in real time, to zero at which point the transmission of frames by MAC client may continue.

When a MAC Control frame indicating the PAUSE function, henceforth referred to as a PAUSE frame, is received, the MAC Control sublayer checks the destination address field. If it is the unique address of the station or the multicast address reserved for PAUSE transactions (01-80-C2-00-00-01), pause_timer is set to pause_time pause_quanta where pause_quanta is defined to be 512 bit times. For example, if pause_time was 10, then pause_timer would be set to 5,120 bit times (5.12µs for Gigabit Ethernet). Otherwise, the MAC Control sublayer waits for the next frame to be received.

Zero is a valid value for pause_time. If a PAUSE frame is received with pause_time set to zero, it enables the recipient to send MAC Client frames if pause_timer was non-zero. If pause_timer was already zero, such a frame would have no effect.

The reception of a PAUSE frame does not interfere with the recipient's transmission of MAC Control frames. Only the transmission of MAC Client frames is inhibited.

Test Setup: Using the appropriate medium (e.g. balanced copper, multi-mode fiber), connect the transmitter of the DUT to the receiver of the testing station and the receiver of the DUT to the transmitter of the testing station.

The University of New Hampshire
InterOperability Lab

Procedure:

1. Configure the DUT to operate in any of its available flow control modes. This may be achieved through auto-negotiation or manual configuration of the DUT (refer to appendix 31A).
2. Force the DUT to transmit a stream of M frames (refer to appendix 31B). The output of the DUT is monitored and the gap between consecutive frames is observed.
3. Force the DUT to transmit a stream of M frames.
4. While the DUT transmits its packet stream, instruct the testing station to transmit a PAUSE frame to the multicast address 01-80-C2-00-00-01 with `pause_time` set to zero. The output of the DUT is observed and, if accessible, the appropriate counters are read.
5. Repeat steps 3 and 4 replacing the well-known multicast address with the address of the DUT.
6. Repeat steps 1 through 5 for all available flow control configurations.

Observable results:

- a. The DUT shall transmit all M frames.
- b. The gap between consecutive frames shall not be altered by the reception of a PAUSE frame with `pause_time` set to zero.
- c. The management entity, if present, shall increment all appropriate counters.

Possible Problems:

- The gap between consecutive frames may not be constant. This could make it difficult for the effect of the PAUSE frame to be assessed.

Test #31.1.2 - Receive PAUSE frame with non-zero pause_time

Purpose: To verify that the device under test (DUT) inhibits the transmission of MAC client frames when a PAUSE frame with a non-zero pause_time is received.

References:

- IEEE 802.3z, 1998 Edition - subclause 37.2.1.4 and 37.2.4.2
- IEEE 802.3x/D3.2 February 9, 1997 - subclauses 31.3 and 31B.3.3
- Appendix 31A - Flow Control modes
- Appendix 31B - Generating packet streams
- Appendix 31C - PAUSE frame format

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 24, 1998

Discussion: MAC Control sublayers implementing the PAUSE function incorporate a timer called pause_timer. When pause_timer is non-zero, the MAC Control sublayer inhibits the transmission of frames by the MAC client. The timer counts down, in real time, to zero at which point the transmission of frames by MAC client may continue.

When a MAC Control frame indicating the PAUSE function, henceforth referred to as a PAUSE frame, is received, the MAC Control sublayer checks the destination address field. If it is the unique address of the station or the multicast address reserved for PAUSE transactions (01-80-C2-00-00-01), pause_timer is set to pause_time pause_quanta where pause_quanta is defined to be 512 bit times. For example, if pause_time was 10, then pause_timer would be set to 5,120 bit times (5.12□s for Gigabit Ethernet). Otherwise, the MAC Control sublayer waits for the next frame to be received.

Note that the reception of a PAUSE frame does not interfere with the recipient's transmission of MAC Control frames. Only the transmission of MAC Client frames is inhibited.

For effective flow control, it is necessary to place an upper bound on the time a device may transmit following the reception of a PAUSE frame with a non-zero pause_time. Since the reception of a PAUSE frame cannot interfere with a transmission in progress, the device is constrained to not begin the transmission of a new frame (assertion of TX_EN at the GMII) later than one pause_quanta following the reception of the PAUSE frame (de-assertion of RX_DV at the GMII). If this time is measured at the MDI rather than the GMII, the specification is increased to pause_quanta + 64 bit times.

Test Setup: Using the appropriate medium (e.g. balanced copper, multi-mode fiber), connect the transmitter of the DUT to the receiver of the testing station and the receiver of the DUT to the transmitter of the testing station.

The University of New Hampshire
InterOperability Lab

Procedure:

1. Configure the DUT to operate in any of its available flow control modes. This may be achieved through auto-negotiation or manual configuration of the DUT (refer to appendix 31A).
2. Force the DUT to transmit a stream of M frames (refer to appendix 31B). The output of the DUT is monitored and the gap between consecutive frames is observed.
3. Force the DUT to transmit a stream of M frames. While the DUT transmits its packet stream, instruct the testing station to transmit a PAUSE frame to the multicast address 01-80-C2-00-00-01 with `pause_time` much greater than the gap observed in step 2. The output of the DUT is observed and, if accessible, the appropriate counters are read.
4. Repeat steps 3 and 4 replacing the well-known multicast address with the unicast address of the DUT.
5. Repeat steps 1 through 5 for all available flow control modes.

Observable results:

PAUSE receive is enabled:

- a. The DUT shall transmit all M frames.
- b. The DUT shall inhibit the transmission of frames for `pause_time`. The reception of a PAUSE frame shall not interfere with a transmission in progress. However, the DUT shall not begin a new transmission any later than `pause_quanta` bit times (measured at the GMII) following the reception of the PAUSE frame.
- c. The management entity, if present, shall increment all appropriate counters.

PAUSE receive is disabled:

- a. The DUT shall transmit all M frames.
- b. The gap between consecutive frames shall not be altered by the reception of the PAUSE frame.
- c. The management entity, if present, shall increment all appropriate counters.

Possible Problems:

- The gap between consecutive frames may not be constant. This could make it difficult for the effect of the PAUSE frame to be assessed.

Test #31.1.3 - Resume Transmission

Purpose: To verify that the device under test (DUT) resumes the transmission of MAC Client frames when `pause_timer` is non-zero and it receives a PAUSE frame with `pause_time` set to zero.

References:

- IEEE 802.3z, 1998 Edition - subclause 37.2.1.4 and 37.2.4.2
- IEEE 802.3x/D3.2 February 9, 1997 - subclauses 31.3 and 31B.3.3
- Appendix 31A - Flow Control modes
- Appendix 31B - Generating packet streams
- Appendix 31C - PAUSE frame format

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 24, 1998

Discussion: MAC Control sublayers implementing the PAUSE function incorporate a timer called `pause_timer`. When `pause_timer` is non-zero, the MAC Control sublayer inhibits the transmission of frames by the MAC client. The timer counts down, in real time, to zero at which point the transmission of frames by MAC client may continue.

When a MAC Control frame indicating the PAUSE function, henceforth referred to as a PAUSE frame, is received, the MAC Control sublayer checks the destination address field. If it is the unique address of the station or the multicast address reserved for PAUSE transactions (01-80-C2-00-00-01), `pause_timer` is set to `pause_time pause_quanta` where `pause_quanta` is defined to be 512 bit times. For example, if `pause_time` was 10, then `pause_timer` would be set to 5,120 bit times (5.12us for Gigabit Ethernet). Otherwise, the MAC Control sublayer waits for the next frame to be received.

Zero is a valid value for `pause_time`. If a PAUSE frame is received with `pause_time` set to zero, it enables the recipient to send MAC Client frames if `pause_timer` was non-zero. If `pause_timer` was already zero, such a frame would have no effect.

Test Setup: Using the appropriate medium (e.g. balanced copper, multi-mode fiber), connect the transmitter of the DUT to the receiver of the testing station and the receiver of the DUT to the transmitter of the testing station.

Procedure:

1. Configure the DUT to operate in any flow control mode for which PAUSE receive is enabled. This may be achieved through auto-negotiation or manual configuration of the DUT (refer to appendix 31A).
2. Force the DUT to transmit a stream of M frames (refer to appendix 31B). The output of the DUT is monitored and the gap between consecutive frames is observed.
3. Force the DUT to transmit a stream of M frames (refer to appendix 31B).

The University of New Hampshire
InterOperability Lab

4. While the DUT transmits its packet stream, instruct the testing station to transmit a PAUSE frame to the multicast address 01-80-C2-00-00-01 with `pause_time` set to 0xFFFF. The output of the DUT is observed.
5. Before `pause_timer` expires, instruct the testing station to transmit a PAUSE frame to the multicast address 01-80-C2-00-00-01 with `pause_time` set zero. The output of the DUT is observed and, if accessible, the appropriate counters are read.
6. Repeat steps 1 through 6 for all available flow control modes in which PAUSE receive is enabled.

Observable results:

- a. The DUT shall transmit all M frames.
- b. The reception of the PAUSE frame with non-zero `pause_time` shall not interfere with a transmission in progress. However, the DUT shall not begin a new transmission any later than `pause_quanta` bit times (measured at the GMII) following the reception of the PAUSE frame.
- c. Following the reception of the PAUSE frame with `pause_timer` set to zero, the DUT will transmit any remaining frames.
- d. The management entity, if present, shall increment all appropriate counters.

Possible Problems:

- As the interpacket gap between the data frames approaches 0xFFFF, it becomes increasingly more difficult to assess the affects of the PAUSE frames. If the PAUSE frame with `pause_time` set to zero is transmitted too early, the entire PAUSE transaction will be indistinguishable from normal inter_packet gap. If the frame is transmitted too late, it will be unclear if the frame enabled transmissions due to the reception of a PAUSE frame with `pause_time` equal to zero or if `pause_timer` simply expired. Note that the time it takes for the DUT to enable the transmission of MAC Client frames is unspecified.

Test #31.1.4 - MAC Control Frame Reception and Handling

Purpose: To verify that the device under test (DUT) rejects invalid MAC Control PAUSE frames.

References:

- IEEE 802.1D, 1993 Edition - subclause 3.12.6
- IEEE 802.3z, 1998 Edition - subclause 37.2.1.4 and 37.2.4.2
- IEEE 802.3x/D3.2 February 9, 1997 - subclause 31.4
- IEEE 802.3, 1996 Edition - subclauses 3.2.3, 3.2.6, 3.2.8 and 3.4
- Appendix 31A - Flow Control modes
- Appendix 31B - Generating packet streams
- Appendix 31C - PAUSE frame format
- Appendix 31D – MAC Control / PAUSE Receive State Diagram

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 24, 1998

Discussion: The MAC Control sublayer repeatedly calls the ReceiveFrame function of the underlying MAC. When the MAC returns a frame, the MAC Control sublayer parses the frame and determines if it is a MAC Control frame, that is, if the two-octet lengthOrType field is 88-08. If the lengthOrType field is not 88-08, the MAC Control sublayer passes the frame up to the MAC Client.

If a MAC Control frame is received, the MAC Control sublayer continues to parse the frame in order to find the opcode. The opcode is a two-octet field that indicates the MAC Control function. If MAC Control sublayer incorporates this function, it acts in a function-dependent manner. Otherwise, it waits for the next frame to be received. An opcode of 00-01 indicates the PAUSE function.

When a MAC Control frame indicating the PAUSE function, henceforth referred to as a PAUSE frame, is received, the MAC Control sublayer checks the destination address field. If it is the unique address of the station or the multicast address reserved for PAUSE transactions (01-80-C2-00-00-01), pause_timer is set to pause_time pause_quanta where pause_quanta is defined to be 512 bit times.

Thus the MAC Control sublayer will only act on PAUSE frames that have a lengthOrType field equal to 88-08, an opcode of 00-01, and the unique address of the station or the multicast address reserved for PAUSE transactions.

Table 1: Description of frames to be used in this test

Label	Description
Test Frame 1	Properly encapsulated MAC Control PAUSE frame with destination address that is not the well-known multicast address (01-80-C2-00-00-01) or unicast address of the DUT and pause_time equal to FF-FF.
Test Frame 2	Properly encapsulated MAC frame with lengthOrType not equal to 88-08, with

The University of New Hampshire
InterOperability Lab

	the destination address that is the well-known multicast address (01-80-C2-00-00-01), the PAUSE opcode and pause_time equal to FF-FF.
Test Frame 3	Properly encapsulated MAC Control frame with the opcode not equal to 00-01, with the destination address that is the well-known multicast address (01-80-C2-00-00-01), and pause_time equal to FF-FF.

Test Setup: Using the appropriate medium (e.g. balanced copper, multi-mode fiber), connect the transmitter of the DUT to the receiver of the testing station and the receiver of the DUT to the transmitter of the testing station.

Procedure:

1. Configure the DUT to operate in any available flow control mode. This may be achieved through auto-negotiation or manual configuration of the DUT (refer to appendix 31A).
2. Instruct the DUT to transmit a stream of M request frames (refer to appendix 31B).
3. While the DUT transmits its packet stream, instruct the testing station to transmit test frame 1 from table 1. The output of the DUT is observed and, if accessible, the appropriate counters are read.
4. Repeat step 3 for all test frames described in table 1.
5. Repeat steps 1 through 4 for all available flow control modes.

Observable results:

- a. The DUT shall transmit all M frames.
- b. The gap between consecutive frames shall not be altered by the reception of an invalid PAUSE frame.
- c. The management entity, if present, shall increment all appropriate counters.

Possible Problems:

- The gap between consecutive frames may not be constant. This could make it difficult for the effect of the PAUSE frame to be assessed.

Test #31.1.5 - Receive MAC Control PAUSE frames of incorrect size

Purpose: To determine how the device under test (DUT) handles PAUSE frames that are the incorrect size.

References:

- IEEE 802.3z, 1998 Edition - subclause 37.2.1.4 and 37.2.4.2
- IEEE 802.3x/D3.1 December 16, 1996 - subclause 31.5.1.
- Appendix 31A - Flow Control modes
- Appendix 31B - Generating packet streams
- Appendix 31C - PAUSE frame format

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 24, 1998

Discussion: The MAC will return any received frame at least minFrameSize in length with the appropriate destination address when the ReceiveFrame function is called. By definition, MAC Control frames are fixed to be minFrameSize in length. The MAC Control sublayer may either discard a MAC Control frame longer than minFrameSize or truncate the frame to minFrameSize before parsing it.

Test Setup: Using the appropriate medium (e.g. balanced copper, multi-mode fiber), connect the transmitter of the DUT to the receiver of the testing station and the receiver of the DUT to the transmitter of the testing station.

Procedure:

1. Configure the DUT to operate in any flow control mode for which PAUSE receive is enabled. This may be achieved through auto-negotiation or manual configuration of the DUT (refer to appendix 31A).
2. Force the DUT to transmit a stream of M frames (refer to appendix 31B).
3. While the DUT transmits its packet stream, instruct the testing station to transmit a PAUSE frame N bytes in length to the multicast address 01-80-C2-00-00-01 with pause_time set to 0xFFFF. The output of the DUT is observed and, if accessible, the appropriate counters are read.
4. Repeat steps 1 through 3 for several values of N such as N = 24, 63, 64, 65, 1518, 1519, 1522, 1523 and 2052.
5. Repeat steps 1 through 3 for all available flow control modes in which PAUSE receive is enabled.

Observable results:

- a. The DUT shall transmit all M frames.

The University of New Hampshire
InterOperability Lab

- b. All frames less than minFrameSize and greater than maxFrameSize should be discarded. If the DUT inhibits the transmission of MAC Client frames, then we assume the MAC Control sublayer has truncated the frames greater than 64 bytes to minFrameSize and parsed it. Otherwise, we assume the MAC control sublayer has discarded the frame.

Possible Problems:

- The gap between consecutive frames may not be constant. This could make it difficult for the effect of the PAUSE frame to be assessed.

Test #31.2.1 - PAUSE frame transmission

Purpose: To verify that the device under test (DUT) transmits properly encapsulated PAUSE frames.

References:

- IEEE 802.3z/D4.1 February 5, 1998 - subclause 37.2.1.4 and 37.2.4.2
- IEEE 802.3x/D3.1 December 16, 1996 - subclause 31.4
- Appendix 31A - Flow Control Modes
- Appendix 31C - MAC Control PAUSE Frames.

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: February 19, 1998

Discussion: The MAC Control sublayer may issue a PAUSE frame at any time with any pause_time value necessary. The destination address is required to be the multicast address reserved for PAUSE transactions. PAUSE transactions may only occur between two devices on a full-duplex network. The use of a well-known multicast address allow the devices to issue PAUSE frames without having to first discover the address of the intended recipient. The use of a multicast address also prevents 802.1D-conformant bridges from forwarding PAUSE frames to the rest of the network.

In this test, we will attempt to force the DUT to issue a PAUSE frames by sending it echo request frames at line rate. This is done in an attempt to place the receive buffers in danger of overflowing. The number of consecutive echo request frames required to cause the DUT to issue a PAUSE frame is used as a reference point for the remainder of the test. When the DUT is configured so that PAUSE transmit is enabled, it should issue a PAUSE frame when sent more than this number of consecutive echo request frames. When PAUSE transmit is disabled, the DUT should not issue the PAUSE frame.

Test Setup: Using the appropriate medium (e.g. balanced copper, multi-mode fiber), connect the transmitter of the DUT to the receiver of the testing station and the receiver of the DUT to the transmitter of the testing station.

Procedure:

1. Configure the DUT to operate in any flow control mode for which PAUSE transmit is enabled. This may be achieved through auto-negotiation or manual configuration of the DUT (refer to appendix 31A).
2. Instruct the testing station to transmit a stream of frames each separated by the minimum inter-packet gap. The output of the DUT is observed. Stop when the DUT issues a PAUSE frame. Let R be the number frames it took for the DUT to generate the PAUSE frame. If accessible, the appropriate counters are read.
3. Repeat step 2 several times in order to find an upper bound on R. Let the upper bound of R equal R_{MAX} .

The University of New Hampshire
InterOperability Lab

4. Repeat steps 1 through 3 for all available flow control modes in which PAUSE transmit is enabled.
5. Configure the DUT to operate in any flow control mode for which PAUSE transmit is disabled.
6. Instruct the testing to send a stream of at least $2R_{MAX}$ frames. The output of the DUT is observed and, if accessible, the appropriate counters are read.
7. Repeat step 6 for all available flow control modes in which PAUSE transmit is disabled.

Observable results:

PAUSE transmit is enabled:

- a. Given that the DUT issued a PAUSE frame, that frame shall meet all of the requirements outlined in appendix 31C.
- b. The management entity, if present, shall increment all appropriate counters.

PAUSE transmit is disabled:

- a. The DUT shall not issue a PAUSE frame during the reception of the at least $2R_{MAX}$ frames.
- b. The management entity, if present, shall increment all appropriate counters.

Possible Problems:

- It is conceivable that the DUT may not need to issue a PAUSE frame even in the face of line rate traffic.

Appendix 31A - Flow Control modes

Purpose: An explanation of various flow control modes and how they might be achieved.

References:

- IEEE 802.3z/D3.1 July 18, 1997 - subclause 37.2.1, 37.2.1.3 and 37.2.5.2.
- IEEE 802.3x/D3.2 February 9, 1997 - subclauses 31.3 and 31B.3.3.

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 17, 1997

Discussion: The device under test may be capable of supporting multiple flow control modes. These modes are symmetric flow control (SFC), asymmetric flow control (AFC), and no flow control at all. SFC implies that PAUSE frames may flow in either direction. AFC implies that PAUSE frames may only flow in one direction, whether that direction is towards the local device or away from the local device. No flow control implies that PAUSE frames are not allowed. Essentially, the flow control mode dictates whether or not a device may transmit PAUSE frames and whether or not the device will honor PAUSE requests.

Each test in the flow control test suite should be performed for each flow control mode the DUT supports. Flow control modes may be configured via auto-negotiation or through manual configuration.

If auto-negotiation is used, bits D7 and D8 in the auto-negotiation base page transmitted by the DUT will determine what flow control modes are available. The testing station will send specific values in bits D7 and D8 of its own base page to enable each available mode. Table 1 is provided as guide for the construction of base pages to be sent by the testing station. Note that bit D7 corresponds to the PAUSE bit and bit D8 corresponds to the ASM_DIR bit. Also note that this process can be used to verify that the DUT implements the priority resolution function properly.

Given the advertised abilities of the DUT, a limited number of the flow control modes may be configured through auto-negotiation. Manual configuration may lend more flexibility but this flexibility and procedure used to perform such configuration is device dependent.

*The University of New Hampshire
InterOperability Lab*

Table 1: Guide for configuration of flow control mode through auto-negotiation

DUT Target Configuration	DUT advertises...		Testing station advertises...	
	PAUSE	ASM DIR	PAUSE	ASM DIR
Disable PAUSE Transmit Disable PAUSE Receive	0	0	X	X
	0	1	0	X
	0	1	1	0
	1	0	0	X
	1	1	0	0
Disable PAUSE Transmit Enable PAUSE Receive	0	0	—	—
	0	1	—	—
	1	0	—	—
	1	1	0	1
Enable PAUSE Transmit Disable PAUSE Receive	0	0	—	—
	0	1	1	1
	1	0	—	—
	1	1	—	—
Enable PAUSE Transmit Enable PAUSE Receive	0	0	—	—
	0	1		
	1	0	1	X
	1	1	1	X

NOTE — "X" indicates value may be either 0 or 1, "—" indicates target configuration cannot be achieved.

Appendix 31B - Generating packet streams

Purpose: To describe procedures that may be used to cause the device under test (DUT) to transmit a stream of M frames.

References:

- IEEE 802.3x/D3.2 February 9, 1997.

Resource Requirements:

- A testing station capable of transmitting and receiving arbitrary sequences of 10-bit code_groups using the signaling method of clause 38 or clause 39. The testing station must implement or be able to emulate the auto-negotiation process described in clause 37.

Last Modification: September 17, 1997

Discussion: In order to observe the DUT's behavior when it receives a PAUSE frame it must be transmitting frames. The following procedures can be used to cause the DUT to transmit a stream of M frames.

Procedure 1

1. Instruct the testing station to transmit M request frames each separated by a minimum inter-packet gap. This should cause the DUT to transmit M replies.

Possible Problems

1. The replies received by the testing station may not be separated by a minimum inter-packet gap.
2. If the DUT cannot process the echo requests as quickly as they are being received, the receive buffer of the DUT may overflow resulting in the loss of frames. If so configured, the DUT may also issue a PAUSE frame. In some cases, the value of M can be adjusted so that this will not happen.

Appendix 31C - PAUSE frame format

Purpose: A description and explanation of each field of a MAC Control PAUSE frame.

References:

- IEEE 802.3x/D3.2 February 9, 1997 - subclauses 31.4, ANNEX 31A and ANNEX 31B.

Last Modification: September 17, 1997

Discussion: MAC Control PAUSE frames may only be transmitted by stations that are configured for full-duplex operation and by definition only two stations can be on a full-duplex network. A MAC Control frame comprises MAC Client-data specified in clause 3. The data is encapsulated by the MAC with preamble and Start-of-Frame delimiter and FCS. MAC Control frames can be distinguished from other MAC frames by their Length/Type field.

The MAC Control frame consist of destination (DA) and source (SA) addresses, a Length/Type field, a MAC Control opcode, a MAC control parameters field and a reserved field which is transmitted as zeros.

7 OCTETS	PREAMBLE
1 OCTET	SFD
6 OCTETS	DESTINATION ADDRESS
6 OCTETS	SOURCE ADDRESS
2 OCTETS	LENGTH/TYPE (88-08)
2 OCTETS	OPCODE (00-01)
2 OCTETS	PARAMETERS (pause_time)
42 OCTETS	RESERVED
4 OCTETS	FCS

Figure 1: MAC Control Packet (MAC Control Frame in **bold**)

Preamble and SFD

The CSMA/CD MAC adds the preamble and start-of-frame delimiter (SFD) to the MAC Control frame.

Destination Address (DA)

The destination address for a MAC Control frame is 01-80-C2-00-00-01. Stations implementing PAUSE will enable reception of frames with this globally-assigned multicast address. The use of the well-known multicast address relieves the MAC Control sublayer and its client from having to know the other station's MAC address. PAUSE frames may only be transmitted by stations that are configured for full-duplex operation and by definition only two stations can be on a full-duplex network. If a station decides to transmit a PAUSE frame, IEEE 802.1D-conformant bridges will not forward the frame to the rest of the network.

Source Address (SA)

The source address field contains the 48 bit address of the station sending the PAUSE frame.

Length/Type

The Length/Type field identifies the frame as a MAC Control frame. The 2-octet field contains the hexadecimal value: 88-08. This is a type interpretation (see 3.2.6) and has been assigned for MAC Control of CSMA/CD LANs.

Opcode

The opcode field identifies the MAC Control frame as a PAUSE frame. The 2-octet field contains the hexadecimal value: 00-01. Since all other values are reserved (not defined), this is the only valid opcode.

Parameters

The parameters field contains MAC Control opcode-specific parameters. The parameters field will contain an integral number of octets. Only one parameter is necessary for PAUSE frames however future MAC Control frames may contain zero, one or many parameters. The parameter for PAUSE frames is called a request_operand. For each PAUSE code the request_operand contains a pause_time.

The variable pause_time is a 2 octet, unsigned integer which contains the length of time the receiving station is requested to inhibit transmission of data. The pause_time is measured in units of pause_quanta, equal to 512 bit times. If a station transmits a PAUSE frame with the request_operand pause_time set to 1000 and the receiving station is accepting PAUSE frames then the receiving station should inhibit transmission of frames for (pause_time × pause_quanta) or (1000 × 512) = 512,000 bit times (512µs for Gigabit Ethernet). The range of pause_time is 0 to 65535 pause_quanta. Thus the maximum time a station can request to inhibit transmission with one PAUSE frame is (65535 × 512) = 33,553,920 bit times (33.554ms for Gigabit Ethernet).

Reserved

The reserved field is used when the parameters field does not fill the fixed length of the MAC Control frame. The size of the reserved field is determined by the size of the MAC Control parameters field and the minimum frame size. For Gigabit Ethernet the minFrameSize is 64 bytes. This means that the length of the reserved field for a MAC Control PAUSE frame is 42 bytes. The reserved field is transmitted as zeros.

Frame Check Sequence (FCS)

The CSMA/CD MAC appends a 32-bit cyclic redundancy check (CRC) to the MAC Control frame. The CRC covers the destination address through the reserved field.