

To: INCITS T10 Membership
From: Paul Entzel, Quantum
Date: 11 November 2002
Document: T10/02-329r2
Subject: Proposed frame format for ADT

1 Related Documents

T10/02-233r0 ADT Frame Format Notes (Paul Suhler)
T10/02-274r0 ADI ADT Frame Format Proposal (Rod Wideman)

2 Revision History

Revision 0:

Initial revision. Posted to the T10 web site on 26 August 2002 and presented at the ADI working group meeting in Minneapolis on 10 September 2002.

Revision 1:

Changes based on discussion at the ADI working group meeting in Minneapolis. See the minutes of the meeting for more explanation (T10/02-328r0):

- Change the Exchange_ID field to be 1 byte instead of 2.
- Add a Port Logout Link Layer frame type and an explanation of the action to be taken by a port that receives one.
- Add an ACK status code value to indicate a frame is aborted because the port is logged out.
- Add clarification to the description of ACK status "Out of Resources".
- Clarify baud rate field in the last paragraph of section 7.4.
- Add text indicating that an Automation device may initiate a Port Login exchange of it's own when it receives a new Port Login exchange.
- Changed the name of the Resume frame to NOP.
- Add timeout value for ACK frame.
- Add error recovery procedure.
- Add Abort Task and Target Reset TMF.
- Clarify the Buffer Allocation Length field description in the SCSI Command Frame.
- Change all use of "xfer ready" to "transfer ready".
- Fixed UA bit in references in AER frame description.

Revision 2:

Changes based on discussion at the ADI working group meeting in Huntington Beach. See the minutes of the meeting for more explanation (T10/02-411r0):

- Changes to the frame header X_Origin, Exchange_ID, and Frame Number fields.
- In section 10.1, add text clarifying that a sender is not required to terminate a frame upon receiving a Port Login or Port Logout frame.
- Change "out-of-order" to "non-sequential" with respect to frame numbers.
- Add text indicating that an NAK frame indicating frame sequence error must be sent with the expected frame number after receiving a frame with an incorrect frame number.
- Change time-out calculation to use the word "period" instead of "rate".
- Add an ACK frame type with zero length payload, change current ACK to NAK.
- Add a NAK status code of "Awaiting Initiate Recovery" for use when sending NAK for non-sequential frame numbers.
- Change AER to only affect the VHF frame. The notification and control frames will use the VHF structure, with the control frame sending mask for field that should enable a notification when changed.

- Remove all polling frame types except VHF.
- Replace all instances of “Initiate Retry” with “Initiate Recovery”.

3 General

This proposal builds upon the documents listed above to create a more complete frame format. Several of the concepts that are described in the proposal are slightly different approaches than these other proposals, while others are simply building on top of them.

No attempt has been made to lay out this document in the same manner as the finished ADT document is to be organized. If this document is approved for inclusion in ADT, the editor will need to re-organize the concepts from this proposal into the clauses of ADT.

4 Assumptions

Several assumptions were used that had not been clearly described in the other documents:

1. The physical interface will be a point-to-point connection.
2. The recipient shall acknowledge every frame except ACK and NAK frames.
3. Because ACK frames will be used so often, they should be kept short.
4. Like other serial transport layers used by T10, there will be several different protocols used on this one. I have identified four so far:
 - A link layer protocol that includes the ACK frame mentioned earlier as well as a protocol for port login. The link protocol should also include a method to negotiate link parameters, such as baud rate and maximum frame size.
 - A protocol that allows for encapsulation of SCSI information. The protocol must be capable of supporting simultaneous command threads initiated by both sides.
 - A protocol that allows for quick access to the polling frame defined in T10/02-097r1 without most of the overhead associated with processing SCSI commands.
 - Provisions for supporting vendor specific protocols that can be used to access diagnostic features in the devices.
5. It is acceptable to have multiple layers of header that are dependent on the protocol the frame is hosting. The basic frame header should include everything that is required to validate and route the frame. Other information that is common to frames of a particular protocol can be included in a second layer of header within the frame’s payload.

5 Basic Frame Format

Document T10/02-274 describes a basic frame format in it section labeled “ADI ADT Frame Format”. This proposal is based on that format.

6 ADT Frame Header

This proposal is based around the concept of layered headers, with the ADT frame header only containing the information needed to validate and route the frame to the proper protocol handler. Table 1 defines the ADT Frame Header.

Table 1 - ADT Frame Header

Bit	7	6	5	4	3	2	1	0
Byte								
0	Reserved	PROTOCOL			PAYLOAD TYPE			
1	X_ORIGIN	EXCHANGE ID			Reserved	FRAME NUMBER		
2 - 3	PAYLOAD SIZE							

The first byte in the header is a set of bit fields collectively referred to as the Frame Type byte.

The X_ORIGIN bit shall be set to zero if the Automation device originates the exchange. The X_ORIGIN bit shall be set to one if the Data Transfer Device originates the exchange. This bit remains constant for all frames associated with a given exchange.

The PROTOCOL field indicates the protocol that is carried in the payload. Table 2 defines the values for the PROTOCOL field.

Table 2 - PROTOCOL field values

Protocol	Description
0	Link Services
1	SCSI
2	ADC Fast Access
3	Vendor Specific
4 – 7	Reserved

The PAYLOAD TYPE field specifies the type of data that can be found in the payload of the frame. See the individual protocol sections for a description of the values in this field.

The FRAME NUMBER field is a continuously incrementing number assigned by the transmitting port that uniquely identifies a frame from other frames sent by that port over a small period of time. It ranges from 0 to 7, and repeats. ACK frames return the FRAME NUMBER field value of the frame that they are acknowledging. The FRAME NUMBER field of a NAK frame shall contain the next expected frame number. A transmitting port shall assign all other frame types the next frame number in sequence, independent of the traffic the port is receiving.

The EXCHANGE ID field contains the identifier used to distinguish frames that are part of the same exchange. Some sequences require more than one frame to complete, often involving frames originating in both devices. All frames that are associated in these sequences shall have the same exchange ID. An originator of an exchange may not re-use an exchange ID value until all frames associated with that exchange have been acknowledged, or until an Abort Exchange frame has been acknowledged.

The PAYLOAD SIZE field contains a count of byte in the Payload area of the frame. This count does not include the SOF, EOF, ADT Frame Header, Checksum, or Escape bytes within the payload.

7 Link Service Frames

7.1 Link Service frame overview

Either port may initiate link Service frames. Link Service frames are used to manage the transport layer. Table 3 defines the values for Frame Type field in Link Service protocol frames.

Table 3 - Link Service frame types

Payload Type	Description
0h	ACK (acknowledge)
1h	NAK (negative acknowledge)
2h	Port Login
3h	Port Logout
4h	Pause
5h	NOP
6h	Initiate Recovery
7h - Fh	Reserved

7.2 ACK

ACK frames are sent by the transport layer to indicate that the port has received a frame without error. An ACK frame shall be sent for every frame that is received without error, except ACK and NAK frames.

7.3 NAK

NAK frames are sent by the transport layer to indicate that the port has detected an error during the reception of a frame. The Payload of a NAK frame contains 1 byte indicating the status. The FRAME NUMBER field in the ADT Header of the NAK frame shall be the next expected frame. Table 4 lists the status values.

Table 4 - NAK Frame status code values

Status	Description
00h	Reserved
01h	Bad checksum
02h	Over-length (more bytes received than Payload Size indicate)
03h	Under-length (less bytes received than Payload Size indicate)
04h	Framing error
06h	Hardware over-run
07h	Non-sequential frame numbers
08h	Awaiting Initiate Recovery frame
09h – 7Fh	Reserved
80h	Unsupported protocol
81h	Out of resources, retry later. The receiving port has run out of buffers to store the frame.
82h	Aborted, login in progress
83h	Invalid or illegal Pause frame received
84h	Illegal operation in Special state
85h	Rejected, port is logged out
86h	Maximum ACK offset exceeded
87h	Maximum frame size exceeded
88h	Unsupported frame type
89h – 0FFh	Reserved

7.4 Port Login

Port Login frames are used to establish link parameters. The login process is a negotiation between the ports that shall result in the determination of a set of operating parameters that are acceptable to both ports. A port shall discard any frame type other than Link Service frames until the login process has completed. The login process consists of the following steps:

- 1) One of the ports on the link sends a Port Login frame containing parameters that are acceptable to that port. The PC field shall be one and the ACCEPT field shall be zero. Any frames received by the port shall be discarded and a NAK shall be sent contain a Status field of “Aborted, login in progress”.
- 2) Upon receiving a Port Login frame, all open exchanges shall be aborted and an ACK frame is sent. If the parameters in the Port Login frame are acceptable, a Port Login frame is sent back with the PC bit set to zero and the ACCEPT bit set to one. Skip to step 4. If one or more parameter is unacceptable, the values are adjusted down and a Port Login frame is sent with the adjusted values. The PC bit shall be set to one and the ACCEPT bit set to zero.
- 3) Upon receiving a Port Login frame, the login exchange originator shall send an ACK frame. If the parameters in the Port Login frame are acceptable, a Port Login frame is sent back with the PC bit set to zero and the ACCEPT bit set to one. Skip to step 4. If the values are unacceptable, they are adjusted down and a Port Login is sent with the adjusted values. The PC bit shall be set to one and the ACCEPT bit set to zero. Skip to step 2.
- 4) Once negotiations have reached a point where a port receives a Port Login frame with acceptable parameters, it shall set the ACCEPT bit to one and the PC bit to zero and send it back. Upon receiving a Port Login frame with the ACCEPT bit set to one, a port shall send an ACK frame. If the port has not yet sent a Port Login frame with the ACCEPT bit set to one, it shall do so in response to this frame. The login parameters shall take affect after both ports have acknowledged a Port Login frame with identical parameters that have the ACCEPT bit set to one and the PC bit set to zero.

An automation device that receives a frame indicating a new port login exchange that has initiated a port login exchange that is not yet complete, shall discard the frame. If an automation device receives a frame indicating a new port login exchange that is not already participating in a port login negotiation may discard the frame and initiate a new port login exchange. A Data Transfer device that receives a Port Login frame shall abort all open exchanges other than the exchange associated with the Port Login frame.

Editors Note: Add a state transition table here so this stuff makes some sense.

Table 5 defines the payload of the Port Login Frame.

Table 5 - Port Login frame payload contents

Bit	7	6	5	4	3	2	1	0
Byte								
0	PC	ACCEPT	Reserved					SPECIAL
1	MAXIMUM ACK OFFSET							
2 - 3	MAXIMUM PAYLOAD SIZE							
4 - 5	BAUD RATE							

The parameters changed (PC) bit is set to one on the first frame of a negotiation sequence, and for all subsequence frames if a negotiated parameter in the frame is different than was last received. The PC bit is set to zero if the Port Login frame is identical to the last one received, with the exception of the PC and Accept bits.

The ACCEPT bit is set to zero on the first frame of a negotiation sequence and all subsequent Port Login frames sent by a port until the frame it is sending matches the last one it received, except for the PC and ACCEPT bits.

The SPECIAL bit is set to one if the port is negotiating parameters for a special operation. Each device may place limitations on the operations that can be performed while operating in special mode. The intention of this mode is to improve transfer rates for special operations such as Firmware Updates and remote diagnostics. Devices may offer higher baud rates, larger frames, or larger offsets in this mode.

The MAXIMUM ACK OFFSET field indicates the number of frames that may be sent to the port without receiving ACK frames in response. The offset count is incremented for each frame sent by a port and decremented for each ACK frame received. Link Service frames are not counted in the offset. A value of zero indicates the port is disabled for all but Link Service traffic.

The MAXIMUM PAYLOAD SIZE value indicates the maximum number of bytes in the payload of a frame that the port can accommodate. A port must be capable of supporting a frame payload size of at least 270 bytes to support this transport layer.

The BAUD RATE field indicates the speed that the port is negotiating the physical interface shall run. The BAUD RATE field contains the desired Baud rate divided by 100. All ports shall default to operating at 9600 Baud at power-up and following error conditions that require re-establishment of the link.

7.5 Port Logout

Only automation may send this frame. Upon receiving a Port Logout frame, the device shall:

- a) Abort all open exchanges.
- b) Disable surrogate mode.
- c) Disable Asynchronous Event Reporting.
- d) Disable initiating Port Login exchanges.

Knowledge of the logged out state may be volatile, so a power cycle or other hard reset condition in the logged out device may cause the port to become active again and attempt to log in to the automation port.

The Port Logout frame has no data in its payload.

7.6 Pause

A Pause Link Service frame may be sent by either device to temporarily stop traffic on the link. When a port receives a Pause frame, it shall acknowledge the frame and then temporarily discontinue sending any more frames on the link. Once in the paused state, receipt of any valid frame other than an ACK frame shall place the port back into active state. The paused state is volatile, so a power cycle or other hard reset condition in the paused device may cause the port to become active again. A Pause frame shall not be sent until the ports have successfully negotiated the link parameters with a Port Login exchange. The paused state only affects the sending of frames, a port must always be capable of receiving frames unless it has placed the opposite port into paused state.

The Pause frame has no data in its payload.

Editors note: We need to include link state definitions in the model section.

7.7 NOP

A NOP frame may be sent by a port to re-activate the other device's port after it has been paused.

7.8 Initiate Recovery

An initiate Recovery frame shall be sent by a port when it detects an error has occurred with a frame that it sent. The Frame Number field in the ADT Frame Header shall contain the frame number of the frame in error. The X_ORGIN and EXCHANGE ID fields are undefined and shall be ignored. There is no payload in an Initiate Recovery frame. See Section 10 for a full explanation of the error recover process.

8 SCSI Frames

8.1 SCSI frame overview

SCSI frames contain information required to implement the SCSI protocol. Each SCSI frame shall include a payload header at the start of the Payload to carry other information that is required to perform the selected function. All SCSI commands use the simple queue model described in SAM-2. The X_ORGIN bit in the ADT Frame Header implies the Initiator and Target identities. . The EXCHANGE ID value from the ADT Frame Header takes on the roll of the Queue Tag from SAM-2. The LUN is included in the SCSI Command frame payload header.

SCSI frames support the Payload Type values defined in Table 6.

Table 6 – SCSI protocol Frame Types

Payload Type	Description
0h	Command
1h	Response
2h	Transfer Ready
3h	Data
4h- Fh	Reserved

8.2 SCSI Command frame

The SCSI Command frame payload shall contain information described in Table 7

Table 7 - SCSI Command frame payload contents

Bit	7	6	5	4	3	2	1	0
Byte								
0 – 1	LUN							
2	TASK MANAGEMENT FLAGS							
3	Reserved							
4 - 19	CDB							
20 – 23	BUFFER ALLOCATION LENGTH							

The LUN field indicates the Logical Unit Number that the command or task management function shall be routed to within the target device. If the logical unit does not exist within the device, the device shall return a Check Condition status with a Sense Key of Illegal Request and addition sense code of Invalid Logical Unit.

The TASK MANAGEMENT FLAGS field indicated the type of action that is to be performed by the target logical unit or device. The values for this field are defined in Table 8.

Table 8 – Task Management Flags values

TMF	Description
00h	Frame contains a SCSI Command
01h	Abort Task
02h	Abort Task Set
03h – 07h	Reserved
08h	Logical Unit Reset
09h – 1Fh	Reserved
20h	Target Reset
21h - FFh	Reserved

If the TASK MANAGEMENT FLAGS field contain 00h, the CDB field contains a SCSI Command Descriptor Block. Otherwise, the target shall ignore the CDB field, and the task management function indicated by the TASK MANAGEMENT FLAGS field shall be executed. See SAM-2 for a definition of the task management functions provided.

The BUFFER ALLOCATION LENGTH field contains a count of the greatest number of data bytes to be transferred to or from the application client data buffer by the SCSI CDB. The BUFFER ALLOCATION LENGTH field is the data buffer size defined by SAM-2. A Buffer Allocation Length value of zero indicates that no data transfer is expected and that no SCSI Transfer Ready or SCSI Data frames shall be transferred for this exchange.

8.3 SCSI Response frame

A SCSI Response frame shall be returned to the Exchange Initiator for every SCSI Command frame that is sent. Table 9 defines the payload of a SCSI Response frame.

Table 9 – SCSI Response frame contents

Bit	7	6	5	4	3	2	1	0
Byte								
0	RESPONSE CODE							
1	SCSI STATUS							
2 - 3	SENSE LENGTH							
4 - N	SCSI AUTOSENSE DATA							

The RESPONSE CODE field indicates the results of the operation as an extension to the SCSI Status. Table 10 defines the values for this field.

Table 10 – Response code values

Response Code	Description
00h	No failure or Task Management function complete
01h	More data transferred than requested
02h	Invalid field in SCSI Command frame
03h	Incorrect relative offset in data frame
04h	Task Management function not supported
05h	Task Management function failed
06h - FFh	Reserved

The SCSI STATUS field contains SCSI Status as defined in SAM-2. This is only valid if the RESPONSE CODE field is set to 00h and the SCSI Command frame for the Exchange was a SCSI command containing a CDB.

The SENSE LENGTH field indicates how many bytes of sense data can be found in the frame. This field shall be set to 0 if the response code is not 00h, and no sense data shall be included in the frame. If the response code is 00h and the SCSI STATUS field contains Check Condition, autosense data shall be included in the frame as defined in SPC-2 and the SENSE LENGTH field shall be set to indicate how much sense data is included.

8.4 SCSI Transfer Ready frame

A SCSI Transfer Ready frame shall be sent by a target of an exchange to inform the Exchange initiator that it is ready to receive data associated with the command. The target port may request all of the data associated with a command with a single SCSI Transfer Ready, or it may use multiple SCSI Transfer Ready frames within the exchange context to request the data a little bit at a time. The contents of the SCSI Transfer Ready frame payload are described in Table 11.

Table 11 – SCSI Transfer Ready frame payload contents

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	BUFFER OFFSET							
4 - 7	BURST LENGTH							

The BUFFER OFFSET field indicates the offset from the beginning of the buffer associated with the first byte sent shall be sent. Data shall not be requested out of order. This field can be used to recover from an error detected in transmission by allowing the target device to request re-transmission of the previous burst of data.

The BURST LENGTH field indicates the size of the buffer that has been allocated to receive data within the target device. The initiator shall transmit Burst Length bytes of data using a SCSI Data frame or sequence of frames to satisfy the request.

8.5 SCSI Data frame

The SCSI Data frame is used to send data associated with SCSI Data In and Data Out operations. Table 12 describes the contents of a SCSI Data frame.

Table 12 – SCSI Data frame payload contents

Bit	7	6	5	4	3	2	1	0
Byte								
0 – 3	BUFFER OFFSET							
4 – 7	DATA LENGTH							
8 – N	DATA							

The BUFFER OFFSET field indicates the offset from the beginning of the buffer associated with the first byte sent shall be sent. Data shall not be sent out of order, however, a port may re-send a data burst should it detect an error in transmission.

The DATA LENGTH field indicates the number of bytes of data included in this frame.

The DATA field contains data.

9 ADC Fast Access Frames

9.1 ADC Fast Access overview

This protocol is intended to provide a feature set beyond what is provided by SCSI to both take advantage of the features of the transport layer and work around the slower speed of it. The ADC Fast Access protocol provides:

- A simple method for accessing the Very High Frequency Polling frame define in ADC,
- a method for device to report asynchronous activity, and
- a method to control these asynchronous reports.

SCSI frames support the Payload Type codes listed in Table 13.

Table 13 – ADC Fast Access frame types

Payload Type	Description
0h	Request for VHF Polling Frame
1h	VHF Polling Frame
2h	AER Frame
3h	AER Control Frame
4h – Fh	Reserved

9.2 Request for VHF Polling Frame

Only Automation devices may initiate a Request for Polling Frame. This frame has no payload.

9.3 VHF Polling Frame

Only Data Transfer devices may initiate a VHF Polling frame. The payload of this frame is described in ADC.

9.4 AER Frame

Asynchronous Event Report frames may optionally be used to report that an event has occurred that may be of interest. Only a Data Transfer device may initiate AER frames. The payload of an AER frame shall contain the VHF Polling data as defined in ADC.

9.5 AER Control Frame

The AER Control frame may optionally be sent by an Automation device to a Data Transfer device to enable or disable AER reporting. The payload of this frame shall contain a VHF Polling data structure, with the bits set to one for each field that the device shall report a change. Multiple bit fields shall have either all of the bits of the field set to 1 or all of the bits in the field set to 0.

Data Transfer devices that do not support AER shall send a NAK frame in response with the STATUS field set to Unsupported frame type.

Data Transfer devices that support AER shall respond to the receipt of a AER Control Frame by sending an AER Control Frame back to the Automation Device with the same X_ORIGIN and EXCHANGE ID values. The payload of the frame shall contain a VHF Polling Frame data structure. Each field that has been enabled for AER notification and is supported by the device shall contain all 1 bits. Each field that has been either disabled for AER notification or is not supported for AER notification by the device shall contain 0. The default setting for all AER events in a Data Transfer device shall be zero.

10 Error Recovery Sequences

10.1 Error detection

Errors in the transport layer may be detected by either the sender of a frame, receiver of a frame, or by both.

10.1.1 Error Detection by the frame sender

10.1.1.1 ACK frame timeout

The sender of a frame, other than an ACK frame, shall timeout the resulting acknowledge. It shall be considered an error condition if a corresponding ACK frame is not received within the timeout period. The time-out period shall start after the EOF of the frame is sent. When operating with a ACK offset greater than 1, a port may start the time-out period for a frame that has completed transmission after the ACK frame for a previously sent frame has been received.

Ports shall not terminate transmission of a frame to acknowledge frames that it has received except in the case of receiving a Port Login or Port Logout frame. A port that receives a Port Login or Port Logout frame may terminate a transmission in progress, or it may complete the transmission. Once a port completes transmission of a frame, it shall acknowledge all frames that it has received before starting transmission of any other frame type, except Port Login, Port Logout, or Pause frames. The minimum ACK timeout period shall be calculated using the formula in Figure 10-1.

Figure 10-1 Minimum ACK timeout period

$$\text{Timeout}_{\text{ACK}} = (\text{Period} * \text{Size}_{\text{Max}} * 2) + (\text{Period} * (\text{Offset}_{\text{Max}} * \text{Size}_{\text{ACK}} * 2)) + 0.1 \text{ seconds}$$

Where:

- $\text{Timeout}_{\text{ACK}}$ is the minimum timeout period in seconds.
 Period is the time per byte calculated as $(10 / \text{Baud Rate})$ and is expressed in seconds per byte.
 Size_{Max} is the Maximum Payload Size negotiated with the Port Login process, plus SOF, EOF, ADT Header, and checksum bytes.
 $\text{Offset}_{\text{Max}}$ is the maximum ACK offset negotiated with the Port Login process.
 Size_{ACK} is the size in bytes of the ACK frame including SOF, EOF, and checksum bytes.

For example, at 9600 Baud with a negotiated Maximum Payload Size of 1024 and Maximum ACK Offset of 2, the minimum timeout period would be approximately 2.28 seconds.

10.1.1.2 NAK frame

It shall also be considered an error condition if a port receives a NAK frame.

10.1.2 Error detection by the frame receiver

The port that receives a frame shall detect and report the following link level errors:

- a) Checksum, over-length, under-length, or improperly formatted frames.
- b) Unsupported protocol or frame type values.
- c) Frames with protocol other than Link Services when logged out.
- d) Frames with non-sequential Frame Numbers.

When a port detects an error on a frame it receives it shall send a NAK frame to the other port with appropriate status so that the port that sent the frame in error can initiate recovery steps. The Frame Number field of the NAK frame shall be set to the of the next expected frame number.

10.2 Error recovery for transmission problems

10.2.1 Differentiating transmission problems

The status codes for NAK frames are grouped into 2 categories:

- a) Link level problems (01h through 7Fh)
- b) Formatting or ULP problems (80h through FFh)

Transmission problems are those that are reported with the link level status codes. A timeout on an ACK frame also qualifies as a transmission problem. Some of the problems in the formatting or ULP category may also be recovered using techniques from this clause, when a retry is called for.

10.2.2 Error recovery for Port Login frames

If a transmission error is detected on a Port Login frame, the recovery process is to restart the negotiation process. This is accomplished by initiating a Port Login frame with a frame number of one and a new exchange ID value. Values in the payload of the frame should be set to the default values for the port.

If an error is detected on a Port Login frame, or a Port Login sequence is being performed as part of the recovery process for error on other frames, the port's operating parameters shall be set to the default values. The default operating parameters for the port are as follows:

- a) The baud rate shall be set to 9600.
- b) The ACK offset shall be set to 1.
- c) The Maximum Payload size shall be 16 bytes.

These values shall remain in effect until the negotiation process is complete, at which time the negotiated values shall take effect.

10.2.3 Error recovery for other frame types

After detecting a transmission error condition has occurred with a frame that it sent, a port shall initiate an error recovery process. A port that detects an error on a frame that it sent must retry sending the frame at least once and no more than 4 times. The frame retry sequence is:

- 1) The port that sent the frame in error sends an Initiate Recovery frame. The Initiate Recovery frame contains the frame number that was detected in error in the FRAME NUMBER field.
- 2) After sending the Initiate Recovery frame, the port in error waits for an ACK frame for that frame. No other frames shall be sent by that port except ACK frames for frames it receives until an ACK frame is received for the Initiate Recovery frame, a timeout occurs on the ACK frame, or a Port Login frame is received.
- 3) If a ACK frame is received for the Initiate Recovery frame, the error port shall resume normal operation by re-sending the frame in error and all frames sent after it before the error was detected, with the exception of ACK frames.
- 4) If no ACK frame is received for the Initiate Recovery frame before the ACK timeout, or a NAK frame is received indicating an error on the Initiate Recovery frame, and the Initiate Recovery frame has not been retried, the port in error shall re-send the Initiate Recovery frame.
- 5) If the Initiate Recovery frame has been sent twice with no ACK frame returned, or a NAK frame indicating an error on the Initiate Recovery frame, the port in error shall abort all open exchanges, set the operating parameters of the interface to default settings, and initiate a Port Login exchange.

When a port receives an Initiate Recovery frame it is an indication that the other port has detected an error in transmission and is attempting to recover from the error. The following steps shall be taken by the receiving port to accommodate the recovery process.

- a) An ACK frame shall be sent to acknowledge receipt of the Initiate Recovery frame.
- b) The FRAME NUMBER field in the Initiate Recovery frame shall be compared to the next expected frame number. If the frame numbers match, the port shall proceed as normal.
- c) If the frame number does not match, this is an indication that an ACK frame was lost in transmission. Frames that are received by the port shall be acknowledged and discarded until the frame that matches the next expected frame number is received. Normal operation shall continue once the frame received matches the frame expected.