Date: Aug 13, 1997

To: T10 Committee (SCSI)

From: George Penokie, Jeff Williams

Subject: Packetized SCSI

1 Overview

The packetization of SPI will provide may benefits such as:

a) ability to stream multiple CDBs to a target without having to rearbitrate for the bus,
b) reducing of the time the target takes to receive the initial connection information (i.e., Identify, queue tag, CDB, etc.) by approximately 6 microseconds even without streaming,
c) reducing the disconnection time by approximately 1.5 microseconds,
d) allowing autosence, and
e) close to seamless mapping of fibre channel on to SPI.

NOTE 1 - All times assume Fast-40 timings.

All of the above plus backward compatibility is possible if the following proposal is accepted for inclusion into SPI-3.

2 Signal sources

Table 1 indicates the type of SCSI device allowed to source each signal. No attempt is made to show if the source is driving asserted, driving negated, or is released. All SCSI device drivers that are not active sources shall be in the high-impedance state. The RST signal may be asserted by any SCSI device at any time.

Table 1 - Signal sources

SCSI bus phase	P cable signals						Q cable signals		
	A cable signals								
	BSY	SEL	C/D I/O MSG REQ	ACK ATN	DB7-0 DB(P)	DB15-8 DB(P1)	REQQ	ACKQ	DB31-16 DB(P2) DB(P3)
INFORMATION UNIT IN	Targ	None	Targ	Init	Targ	Targ	Targ	Init	Targ
INFORMATION UNIT OUT	Targ	None	Targ	Init	Init	Init	Targ	Init	Init

All: The signal shall be driven by all SCSI devices that are actively arbitrating.
S ID: A unique data bit (the SCSI ID) shall be driven by each SCSI device that is actively arbitrating; the other data bits shall be released (i.e., not driven) by this SCSI device. The parity bit(s) may be released or driven to the true state, but shall not be driven to the false state during this phase.
I&T: The signal shall be driven by the initiator, target, or both, as specified in the SELECTION phase and RESELECTION phase.
Fort: Initiator or target or neither, depending on the state of the I/O signal and the bus width.
Init: If driven, this signal shall be driven only by the active initiator.
None: The signal shall be released; that is, not driven by any SCSI device. The bias circuitry of the bus terminators pulls the signal to the false state.
Win: The signal shall be driven by the one SCSI device that wins arbitration.
Targ: If the signal is driven, it shall be driven only by the active target.

3 Logical characteristics

3.1 SCSI bus phases

The SCSI architecture includes nine distinct phases:

a) BUS FREE phase,
b) ARBITRATION phase,
c) SELECTION phase,
d) RESELECTION phase,
e) COMMAND phase,
f) DATA phase,
g) INFORMATION UNIT phase,
h) STATUS phase, and
i) MESSAGE phase.

The COMMAND phase, DATA phase, INFORMATION UNIT phase, STATUS phase, and MESSAGE phase are collectively termed the information transfer phases.

The SCSI bus can never be in more than one phase at any given time. In the following descriptions, signals that are not mentioned shall not be asserted.

3.1.0.1 Unexpected bus free

An unexpected bus free occurs when an initiator detects a BUS FREE phase (i.e., the release of BSY) that

is not expected. Initiators only expect a BUS FREE phase to occur after one of the following occurs:

a) after a hard reset is detected;
b) after an ABORT TASK message is successfully received by a target;
c) after an ABORT TASK SET message is successfully received by a target;
d) after an CLEAR TASK SET message is successfully received by a target;
e) after an LOGICAL UNIT RESET message is successfully received by a target;
f) after an TARGET RESET message is successfully received by a target;
g) after an TERMINATE TASK message is successfully received by a target;
h) after a DISCONNECT message is successfully transmitted from a target;
i) after a TASK COMPLETE message is successfully transmitted from a target;
j) after a RELEASE RECOVERY message is successfully received by a target;
k) after the release of the SEL signal after a SELECTION or RESELECTION phase time-out;
l) after the last SPI command information unit is successfully received by a target;
m) after a SPI data information unit is successfully received by or transmitted from a target;
n) after a SPI status information unit is successfully transmitted from a target.

3.1.1 SELECTION phase

The SELECTION phase allows an initiator to select a target for the purpose of initiating some target function (e.g., READ or WRITE command). During the SELECTION phase the I/O signal is negated so that this phase can be distinguished from the RESELECTION phase.

The SCSI device that won the arbitration has both the BSY and SEL signals asserted and has delayed at least a bus clear delay plus a bus settle delay before ending the ARBITRATION phase. The SCSI device that won the arbitration becomes an initiator by not asserting the I/O signal.

The initiator shall set the DATA BUS to a value that is the OR of its SCSI ID bit and the target's SCSI ID bit and it asserts the ATN signal to indicate the target shall not use information unit phases (i.e., indicating that a MESSAGE OUT phase is to follow the SELECTION phase). The initiator shall not asset the ATN signal to indicate the target shall use information unit phases (i.e., indicating that a INFORMATION UNIT OUT phase is to follow the SELECTION phase). The initiator shall then wait at least two deskew delays and release the BSY signal. The initiator shall then wait at least a bus settle delay before looking for a response from the target.

3.1.2 Information transfer phases

The COMMAND, DATA, INFORMATION UNIT, STATUS, and MESSAGE phases are all grouped together as the information transfer phases because they are all used to transfer data or control information via the DATA BUS. The actual content of the information is beyond the scope of this section.

Table 2 - Information transfer phases

Signal			Phase	Direction of transfer	Comment
MSG	C/D	I/O			
0	0	0	DATA OUT	Initiator to target	Data phase
0	0	1	DATA IN	Initiator from target	
0	1	0	COMMAND	Initiator to target	
0	1	1	STATUS	Initiator from target	
1	0	0	INFORMATION UNIT OUT	Initiator to target	Information unit phase
1	0	1	INFORMATION UNIT IN	Initiator from target	
1	1	0	MESSAGE OUT	Initiator to target	Message phase
1	1	1	MESSAGE IN	Initiator from target	
Key: 0 = False; 1 = True					

3.1.2.1 Synchronous data transfer

Synchronous data transfer is optional and is only used in data phases and information unit phases. It shall be used in a data phase or a information unit phase if a synchronous data transfer agreement has been established (see 6.6.21). The agreement specifies the REQx/ACKx offset and the minimum transfer period.

3.1.3 Information unit phase

The data phase is a term that encompasses both the INFORMATION UNIT IN phase and the INFORMATION UNIT OUT phase.

3.1.3.1 INFORMATION UNIT IN phase

The INFORMATION UNIT IN phase allows the target to request that information units be sent to the initiator from the target.

The target shall assert the I/O and MSG signals and negate the C/D signal during the REQx/ACKx handshake(s) of this phase.

3.1.3.2 INFORMATION UNIT IN phase exception condition handling

Editors Note 1 - GOP: Do we want to put any thing here?

3.1.3.3 INFORMATION UNIT OUT phase

The INFORMATION UNIT OUT phase allows the target to request that information units be sent from the initiator to the target.

The target shall negate the C/D and I/O signals and assert the MSG signal during the REQx/ACKx handshake(s) of this phase.

3.1.3.4 INFORMATION UNIT OUT phase exception condition handling

Editors Note 2 - GOP: Do we want to put any thing here?

3.2 SCSI bus phase sequences

The order in which phases are used on the SCSI bus follows a prescribed sequence.

A hard reset can abort any phase and is always followed by the BUS FREE phase. Also any other phase can be followed by the BUS FREE phase but many such instances are error conditions (see 3.1.0.1).

The additional allowable sequences shall be as shown in figure 1 and figure 2.

The normal progression for selection with attention is from the BUS FREE phase to ARBITRATION, from ARBITRATION to SELECTION or RESELECTION, and from SELECTION or RESELECTION to one or more of the information transfer phases (COMMAND, DATA, STATUS, or MESSAGE). The final information transfer phase is normally the MESSAGE IN phase where a DISCONNECT, or COMMAND COMPLETE message is transferred, followed by the BUS FREE phase.

The normal progression for selection without attention is from the BUS FREE phase to ARBITRATION, from ARBITRATION to SELECTION or RESELECTION, and from SELECTION or RESELECTION to one or more information unit phases (INFORMATION UNIT OUT or INFORMATION UNIT IN). The final information unit phase is followed by the BUS FREE phase.

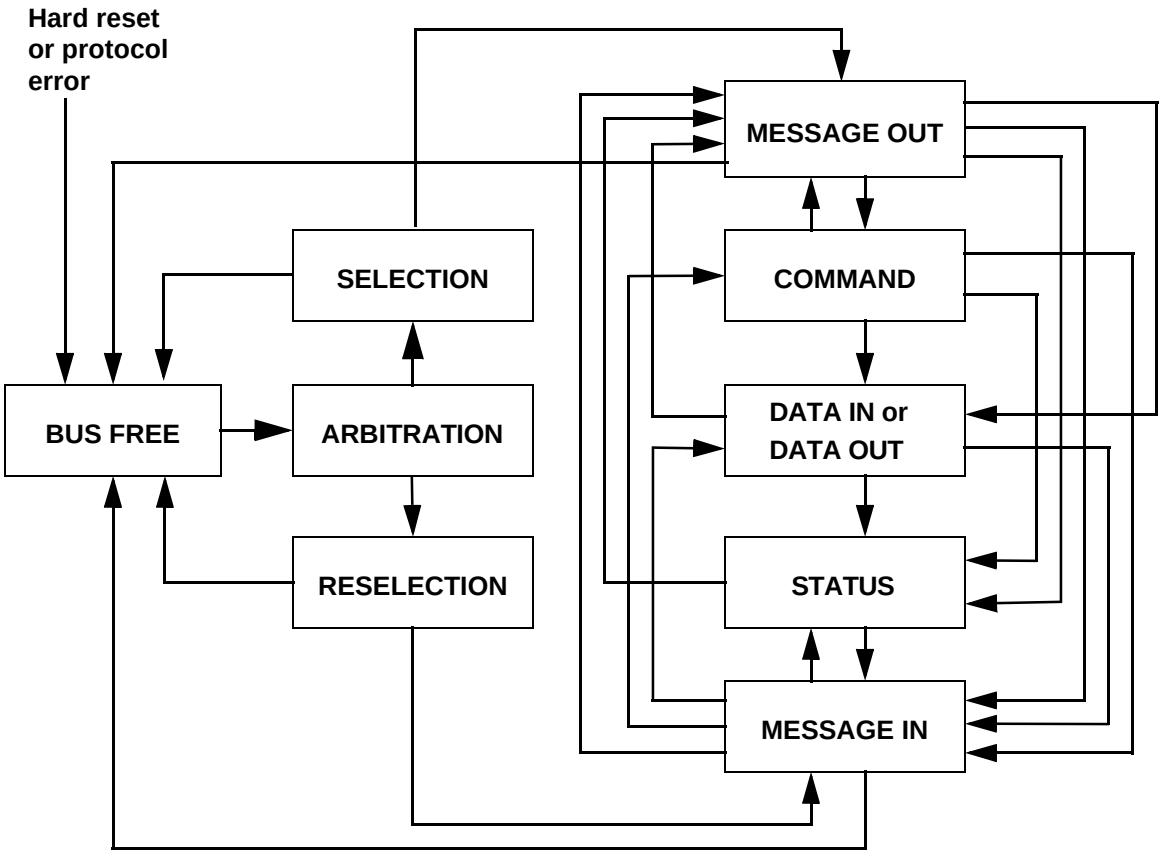

Figure 1 - Phase sequences for selection with attention

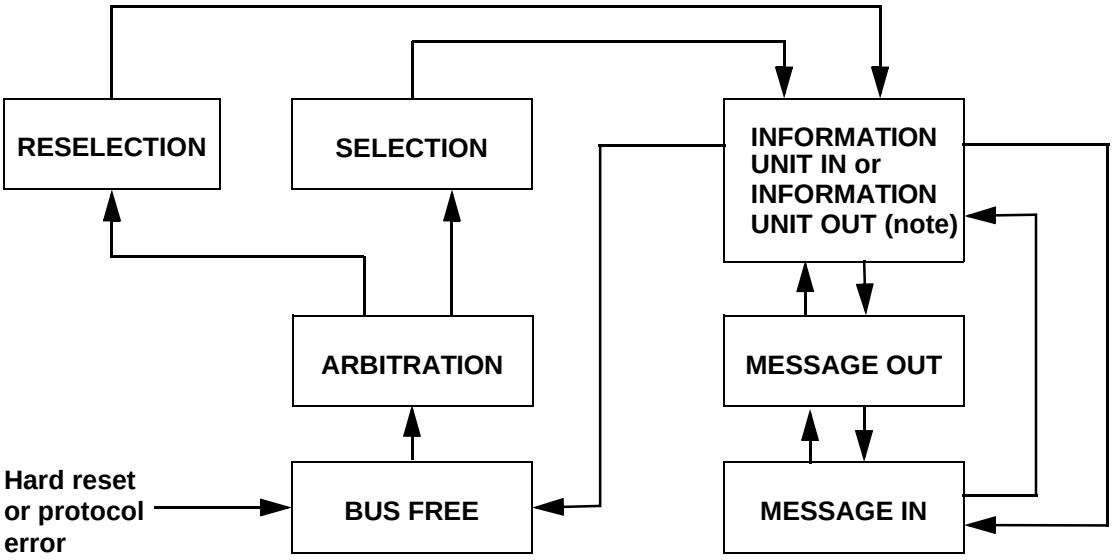

Note: See figure 3 for the sequencing of SPI information units within the INFORMATION UNIT IN and INFORMATION UNIT OUT phases.

Figure 2 - Phase sequences for selection without attention

4 SPI information unit

This clause describes the SPI information units.

4.1 SPI information unit sequences

The information unit phase transfers data in SPI information units. The order in which SPI information units are transfered within the information unit phase follows a prescribed sequence.

The allowable SPI information unit sequences shall be as shown in figure 3.

The normal progression is from one or more SPI L_Q information unit/SPI command information unit pair(s), to one or more SPI L_Q information unit/SPI data information unit pair(s), to a SPI L_Q information unit/SPI status information unit pair.

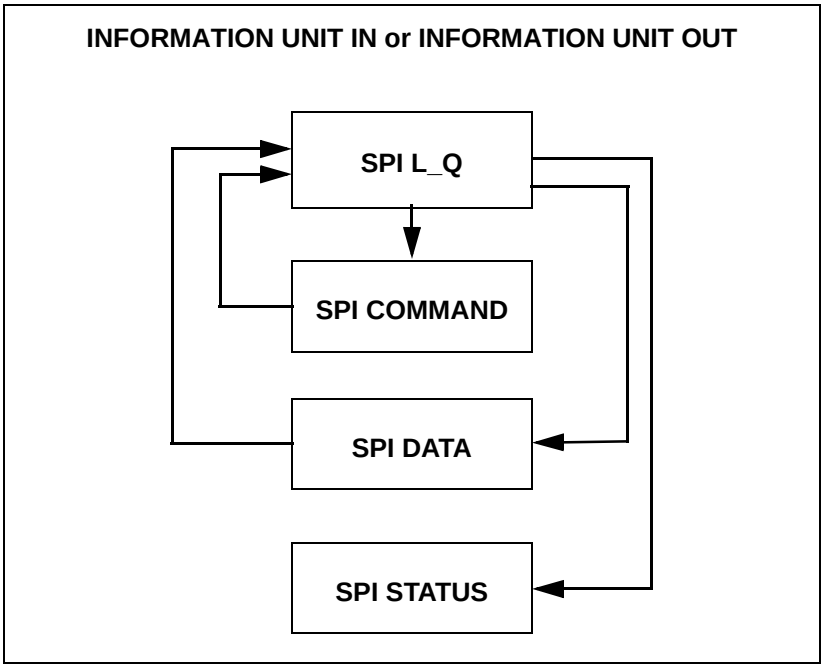

Note: See figure 2 for the sequencing rules between the INFORMATION UINT IN or INFROMATION UNIT OUT phases and the other phases.

Figure 3 - SPI information unit sequences

4.2 SPI command information unit

The SPI command information unit (see table 3) carries CDBs, task attributes, and task management requests to be performed by a device server.

Table 3 - SPI command information unit

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)			LOGICAL UNIT NUMBER				
7								(LSB)
8	RESERVED							
9	RESERVED					TASK ATTRIBUTE		
10	TASK MANAGEMENT FLAGS							
11	RESERVED						RDDATA	WRDATA
12	(MSB)			CDB				
27								(LSB)
28	MSB			DATA LENGTH				
29								
30								
31								LSB
32	MSB			CRC				
33								
34								
35								LSB

The logical unit number field specifies the address of the logical unit of the I_T_L nexus for the current task. The structure of the logical unit number field shall be as defined in SAM-2. If the addressed logical unit does not exist, the task manger shall follow the SCSI rules for selection of invalid logical units as defined in SPC.

The task attribute field is defined in table 4.

TABLE 4 - TASK ATTRIBUTE

Codes	Description
000b	Requests that the task be managed according to the rules for a simple task attribute. (See SAM)
001b	Requests that the task be managed according to the rules for a head of queue task attribute. (See SAM)
010b	Requests that the task be managed according to the rules for a ordered attribute. (See SAM)
011b	Reserved
100b	Requests that the task be managed according to the rules for a automatic contingent allegiance task attribute. (See SAM)
101b	Requests that the task be managed according to the rules for an untagged task. (See SAM).
110b-111b	Reserved

The task management flags field is defined in table 5. If a task management function fails the task manager shall terminate the task with a GOOD status. The protocol failure code shall be set to task management function failed.

TABLE 5 - TASK MANAGEMENT FLAGS

Codes	Description
00h	Indicates no task management requests for the current task.
01h	The task manger shall abort the task as defined in the ABORT TASK message (see xxx).
02h	The task manger shall abort the task set as defined in the ABORT TASK SET message (see xxx).
04h	The task manger shall clear the task set as defined in the CLEAR TASK SET message (see xxx).
08h	The task manager shall perform a hard reset to the selected logical unit as defined in the LOGICAL UNIT RESET message (see xxx).
20h	The task manger shall perform a hard reset as defined in the TARGET RESET message (see xxx).
40h	The task manger shall perform a clear ACA as defined in the CLEAR ACA message (see xxx).
80h	The task manger shall perform a terminate task as defined in the TERMINATE TASK message (see xxx).
others	The task manager shall terminate the task with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to ILLEGAL REQUEST IN SPI INFORMATION UNIT. **Editors Note 3 - GOP: or should it be?** The task manager shall terminate the task with a GOOD status. The protocol failure code shall be set to task management function not supported.

The write data bit (wrdata) and read data bit (rddata) shall be ignored.

The cdb field contains he actual CDB to be interpreted by the addressed logical unit. The maximum CDB length is 16 bytes. The cdb field is not valid and is ignored of any task management flag except the clear ACA flag (40h) is set to one. The contents of the CDB shall be as defined in SAM and SPC.

The data length field shall be ignored.

The crc field shall use the same algorithm as defined in FC-xx.

4.3 SPI L_Q information unit

The SPI L_Q information unit (see table 6) contains L_Q nexus information for the current task, the type of information unit to immediately follow, and the length of information to immediately follow. A SPI L_Q information unit shall immediately proceed all SPI command information units, SPI data information units, and SPI status information units.

Table 6 - SPI L_Q information unit

Bit Byte	7	6	5	4	3	2	1	0
0	TYPE							
1	RESERVED							
2	RESERVED							
3	TAG							
4	(MSB)	LOGICAL UNIT NUMBER						
11								(LSB)
12	MSB	DATA LENGTH						
13								
14								
15								LSB
16	MSB	CRC						
17								
18								
19								LSB

The type field is defined in table 7.

TABLE 7 - TYPE

Codes	Type	Description
00h	Com-mand	Sent by an initiator to indicate a SPI command information unit shall immediately follow this SPI L_Q information unit. Indicates the initiator will not send any more SPI command information units during the current connection. The LOGICAL UNIT NUMBER field shall be ignored and the DATA LENGTH field shall be set to 24h.
01h	Multiple Com-mand	Sent by an initiator to indicate a SPI command information unit shall immediately follow this SPI L_Q information unit. Indicates the initiator will send another SPI L_Q information unit and SPI command information unit during the current connection. The LOGICAL UNIT NUMBER field shall be ignored and the DATA LENGTH field shall be set to 24h.
02h	Data	Sent by a target to indicate a SPI data information unit shall immediately follow this SPI L_Q information unit. A length of zero in the DATA LENGTH field shall not be considered an error.
03h	Status	Sent by a target to indicate a SPI status information unit shall immediately follow this SPI L_Q information unit.
04h-FFh		Reserved

See xxx for a description of the tag field.

The logical unit number field specifies the address of the logical unit of the I_T_L nexus of the current task. The structure of the logical unit number field shall be as defined in SAM-2. If the type is command or multiple command the logical unit number field shall be ignored.

The data length field contains the length in 2-byte words of the following information unit. The data length shall include 4 bytes of CRC (e.g., a 512 byte data transfer would have a data length of 258).

The crc field shall use the same algorithm as defined in FC-xx.

4.4 SPI data information unit

The SPI data information unit (see table 8) contains data.

Table 8 - SPI data information unit

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
				DATA				
n-4								(LSB)
n-3	MSB							
n-2								
				CRC				
n-1								
n								LSB

The data field may contain any type of information (e.g., parameter lists, mode pages, user data, etc.). The data field shall contain an even number of bytes. If the amount of data to transfer has an odd number of bytes the sending SCSI device shall pad the last byte with a 00h.

The crc field shall use the same algorithm as defined in FC-xx.

4.5 SPI status information unit

The SPI status information unit (see table 9) contains the competition status of the task indicated by the proceeding SPI L_Q information unit. After successfully sending this information unit, the target shall go to the BUS FREE phase by releasing the BSY signal. The target shall consider the message transmission to be successful when if detects the negation of ACK for the last byte of the information unit with the ATN signal false.

Table 9 - SPI status information unit

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
7				RESERVED				(LSB)
8				RESERVED				
9				RESERVED				
10				RESERVED	UNDER	OVER	SNSVALID	RSPVALID
11				STATUS				
12	(MSB)							
15				SENSE DATA LIST LENGTH (n-m)				(LSB)
16	(MSB)							
19				PROTOCOL FAILURES LIST LENGTH (m-19)				(LSB)
20	(MSB)							
m				PROTOCOL FAILURES				(LSB)
1+m	(MSB)							
n				SENSE DATA				(LSB)
n+1	MSB							
n+2				CRC				
n+3								
n+4								LSB

The over field and under field shall be set to zero by the target and ignored by the initiator.

A sense data valid bit (snsvalid) of zero indicates the sense data list length is not valid and no sense data is provided. A snsvalid bit of one indicates the sense data list length field specifies the number of bytes in the sense data field.

A protocol failures valid bit (rspvalid) of zero indicates the protocol failures list length is not valid and no protocol failure information is provided. A rspvalid bit of one indicates the protocol failures list length field specifies the number of bytes in the protocol failure field.

If the rspvalid bit is one, the protocol failures list length field contains the length in bytes of the protocol failures field. If there is no protocol failure information the rspvalid bit and the protocol failures list length field shall be set to zero. The protocol failures list length field shall only contain lengths of 0, or 4. Other values or length are reserved for future standardization.

If the snsvalid bit is one, the sense data list length field contains the length in bytes of the sense data field. If there is no sense data the snsvalid bit and the sense data list length field shall be set to zero. The sense data list length field shall only contain even lengths (i.e., 0,2,4, etc.).

The protocol failures field (see table 10) contains information describing the protocol failures detected during the execution of a task. The protocol failures field shall contain valid information if the target detects any of the conditions described by the protocol failure code (see table 11).

Table 10 - PROTOCOL FAILURES **field**

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	RESERVED							
2	RESERVED							
3	PROTOCOL FAILURE CODE							

The protocol failure code field is defined in table 11.

TABLE 11 - PROTOCOL FAILURE CODE

Codes	Description
00h	Indicates no failure or task management function complete.
01h	Reserved
02h	SPI command information unit fields invalid.
03h	Reserved
04h	The task management function not supported.
05h	The task management function failed.
06h-FFh	Reserved

The sense data field contains the information specified by the SCSI-3 Primary Commands Standard for presentation by the REQUEST SENSE command. The proper sense data shall be presented when a SCSI status byte of CHECK CONDITION or COMMAND TERMINATED is presented as specified by the SCSI-3 Primary Commands Standard.

The crc field shall use the same algorithm as defined in FC-xx.

5 Enabling mechanism

5.1 Identify command

A bit is added to Identify data to indicate if a target supports the information unit phases.

5.2 Disconnect/Reconnect mode page

A bit is added to the disconnect/reconnect mode page to enable/disable information unit phases.