July 7, 2000

To: T10 Technical Committee

From: Ron Roberts, Adaptec, Inc.

Subj: Additional Expander Proposal

This proposal contains corrections and additions as well as modified section of SDVR02. Discussions with various users of Expanders have led to many of the changes within this proposal. Many of the statements included are to provide better understanding of the function outlined in the SDV document.

Overview of operation

The Expander shall monitor Bus Arbitration to identify the Initiator attempting to obtain control of Bus. The Expander retains the Initiator ID for use during other functions.

The sequence of proceeding through the process of determining the addresses and the topology of the entire sub-system are as follows:

INQUIRY command to all SCSI devices
Address Discovery
Topology Discovery

The expanders shall have monitored any arbitration between the initiator and target and determined the transfer mode and width. It is required that the Initiator negotiate for Asynchronous Data Phase to the targets before starting Topology Discovery. It is recommended that it negotiate for a wide data bus, however the target may negotiate for a narrow bus. The expanders shall use the negotiated width when inserting Data In Phase information as described below.

Expanders shall respond to a selection with 5 or more bits asserted (5 bits in 7:0 and 1 to 7 bits in 15:8). The 5 bits in 7:0 ensure that neither the legacy nor extended address devices will consider the Expander Selection valid. The value in bits 7:0 of this selection is called the Expanded Signature and denoted as XSIG. The value in bits 15:8, denoted as XID, are called the Expanded ID, a binary number from 00h through FFh. This type of addressing restricts the Expanded Address devices to be selectable targets that cannot disconnect and reconnect.

Broadcast Selection

Broadcast Selection is a means by which the Initiator shall address the Expanders and does not require a response from the Expander. A Broadcast Selection shall be valid for all SCSI Expanders of the Signature Type (XSIG), unless that Expander does not support that function, e.g. Hardwired Expanders will ignore any broadcasts relative to automatic addressing.

Unique Selection

A Unique Selection provides the Initiator with the means to select each Expander individually. Expanded ID (XID) 80h through EFh) requires a response (BSY) from the Expander and a subsequent SCSI Command. Hardwired and Auto-Addressable Expanders are assigned within different address ranges to allow mixing within the same SCSI I/O Subsystem.

Topology Discovery Phase

Using the INQUIRY command, the Initiator determines the SCSI devices attached to the I/O sub-system. The Initiator uses the Topology Discovery Phase to discover the following:

- The number of Expanders in the path to each SCSI Device
- The address of the expanders in the path to the SCSI Device
- Any un-addressed Expanders in the path to the SCSI Device
- and the Expander Characteristics

To begin the discovery of the I/O sub-system the Initiator shall arbitrate for the SCSI Bus, and begin a Selection phase. During arbitration and selection the Expander shall monitor the SCSI Bus to determine Initiator ID and the Broadcast Select (FAA7h). The Expanders shall save the ID of the Initiator and begin to monitor the SCSI Bus for a WRITE/READ BUFFER Command issued from that Initiator.

After the Broadcast Select, the Initiator will select each device, sending a WRITE/READ BUFFER Command. All devices that support this command shall return a Status In Response. The content of the status is not relevant. The first expander in the path back from the device will detect the incoming Status In Phase and capture the status byte. It then changes the outgoing phase to Data In and inserts its ID and expander characteristics. Each expander remaining in the path back to the Initiator will simply pass this Data In Phase transfer. When the first expander has inserted all of its information, it propagates the Status In Phase which is detected by the next in-line expander which inserts its ID and expander characteristics. Eventually the Status In reaches the initiator and the discovery of the path to the device is complete. If any expander does not have an address, it will insert the NULL Address (XID, XSIG=FFA7h). If a WRITE/READ BUFFER command is issued to any device from a Initiator that has not issued a Broadcast Select, the WRITE/READ BUFFER Status In phase is passed untouched. After issuing the WRITE/READ BUFFER Command to all devices, the Initiator analyzes the Topology Map to determine which expanders, are not addressed.

Table 1 - Topology Discovery Expander Entry

Bit Word	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Expander ID (XID)								Expander Signature (XSIG)							
1	See Table 13 on page 12, byte 3								See Table 13 on page 12, byte 2							
2	See Table 13 on page 12, byte 5								See Table 13 on page 12, byte 4							
3	Reserved															
4																
5																
6																
7																

After all SCSI Devices have been issued the WRITE/READ BUFFER and inserted the ID and characteristics, the full Expander/Device topology of the I/O Subsystem will be known. The number of inserted Data In Phase bytes shall determine the number of in-line expanders. The Initiator Software must check for illegal configurations such as too many in-line expanders. This inserted Data In Phase data is used to build a Topology Map. The Topology Map is defined by a structure where each possible Device ID has a set of up to n expander entries (nx16 bytes per Device ID) used to contain the Expander path (total map size is nx256 bytes for the 16 possible SCSI Devices, where n is the maximum number of in-line expanders allowed). The nx16 bytes per Device ID allows up to n in-line expanders to any SCSI Device. The first entry in the set defines the Expander to which the device is physically attached. The

second describes the next in-line expander and so on. An Expander Entry in the set with the first word of FF00h indicates the end of the set. If there are no Expanders in-line to the device the first Expander entry of the set is FF00h. If there are n in-line expanders there are no FF00h entries, and the termination is assumed after the eighth entry. If there are more than 8 in-line expanders the first word of the eighth entry is FFFFh, and this is considered an error. The Topology Discovery by Initiator 7 of Figure 1 (assuming all expanders are hardwired as shown) would yield the steps/responses as shown in Table 2. The resultant Topology Map is built from the Expander ID Segment and Path received in the WRITE/READ BUFFER response. Note that any expander could have been addressable and either already addressed or not addressed. The un-addressed expander would have inserted FFA7h into the first word.

Expander Address Phase

After the Topology Discovery described in Section 0.3 on page 1 the Initiator will have a map showing the addressed and un-addressed expanders. If all expanders are addressed no further action is required, but if any are not addressed the Initiator must execute the Expander Address Phase.

The Initiator arbitrates for the SCSI Bus and issues the Broadcast Select Start Expander Address (FEA7h). During arbitration and selection the Expander shall monitor the SCSI Bus to determine Initiator ID. The selection tells all Auto-Addressable Expanders to monitor the SCSI Bus for the WRITE/READ BUFFER Command from the identified Initiator and Set Expander Address Command (FCA7h) from that Initiator. Any other Initiator may issue the Broadcast Select Start Expander Address but the expanders will only accept it from the first Initiator to issue it. The Initiator may either keep the existing addresses or it may reset all of the addresses using the Broadcast Select Reset Expander Address. Only the first Initiator that issues the Start Expander Address Broadcast Select will have the Reset Expander Address Broadcast Select accepted by the expanders.

After the Initiator has issued the Start Expander Address it issues the WRITE/READ BUFFER to the first SCSI Device. The un-addressed expanders that are in-line to the device monitor this command from the sensitized Initiator and are then armed for the Set Expander Address Command (EFA7h). The first un-addressed in-line expander will see a Status In Phase, know it is the last in-line expander to the device, and temporarily take the XID of EFh and insert a single byte or word (depending on the negotiated width) of zeroes. Any other un-addressed in-line expander will see the Data In Phase before the Status In Phase and know it is not the last device and pass all further phases and information untouched. This handling of the response to the WRITE/READ BUFFER Command allows the expanders to know who is the last in-line expander, that expander begin the one to assert BSY for the subsequent Set Expander Address Select. The Initiator then issues the Set Expander Address Select. The last armed in-line Expander will respond by asserting BSY, both back to the initiator and on to the far-side segment(s). The last in-line expander then transitions to the Data Out Phase and asserts REQ. REQ is passed through the in-line expanders back to the initiator. The initiator responds with ACK with the Expander ID on the data bus. The first armed un-addressed in-line expander will capture the ACK and data use it as its ID and negate REQ. ACK is not passed down the line. The initiator will respond by negating ACK. The first in-line expander will then re-assert REQ since the far-side REQ is still asserted. When the initiator responds with ACK and the next ID, the now addressed first in-line expander will pass the ACK and data. The next armed un-addressed in-line expanders capture the ACK and data to set its ID. Eventually the ACK will reach the last in-line expander driving BSY. It will negate REQ that is propagated back to the initiator who negates ACK. The negated ACK propagates back to the last in-line expander that then negates BSY. Refer to Figure 1 on page 4 and Figure 2 on page 5 for an example configuration and associated timing for the Set Expander Address Select Command.

If a second Initiator had issued a Start Expander Address Broadcast Select it could also try to issue the WRITE/READ BUFFER Command and/or the Set Expander Address Select Command. The responses it gets for these actions will alert it that it another Initiator has already started the process. The WRITE/READ BUFFER Command used to arm the in-line devices will return only Status In information since the in-line expanders are not sensitized to that Initiator. The Software shall only issue this command to a device that is reached through un-addressed in-line expanders, and if it does not get a Data In Phase before the Status In Phase it knows that it does not have the permission to set addresses. This Initiator must then go back to the Topology Discovery Phase to discover the new addresses. The Initiator should wait approximately 5 microseconds times the number of addressable expanders before entering the Topology Discover Phase to allow the other Initiator time to address all of the expanders. If there are still un-addressed expanders in the Topology Discovery, the Initiator waits and retries the discovery. The Initiator should implement a retry timer that allows it to try another Expander Address phase or reset the SCSI Bus if the other Initiator never seems to finish the addressing.

The following further actions are implemented in the expanders:

- The first issue of Start Expander Address shall turn OFF all SCSI Bus Reset Isolation modes within all Expanders. It does not affect the Segment Isolation Mode, and the Start Expander Address is not propagated through the isolated expander.
- The Far-Side Reset Function is disabled during Topology Discovery and Expander Addressing.
- The Topology Discovery Phase and Expander Address Phases are aborted in the event of a SCSI Bus Reset.

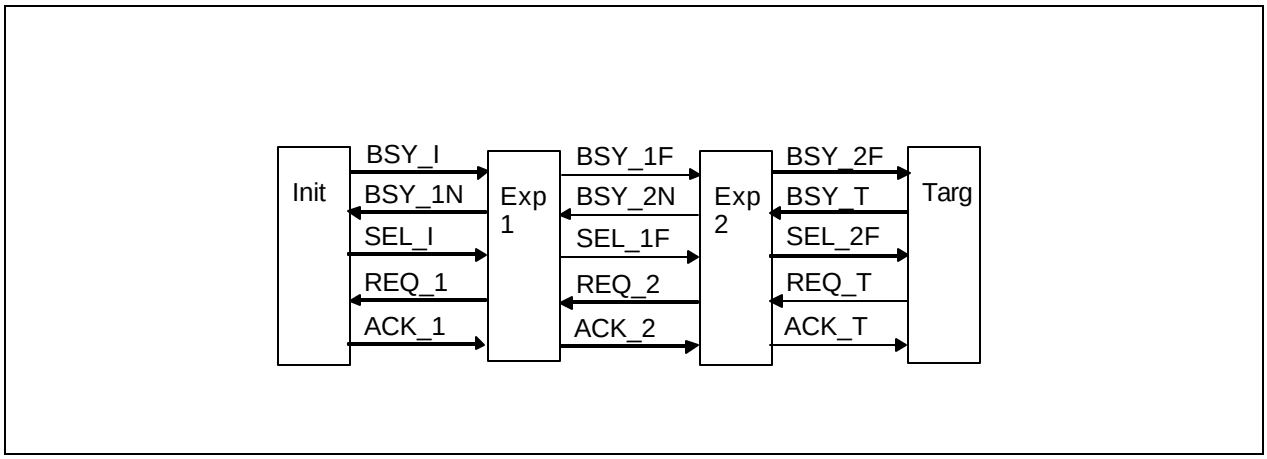

Figure 1 - Set Expander Address Example Configuration

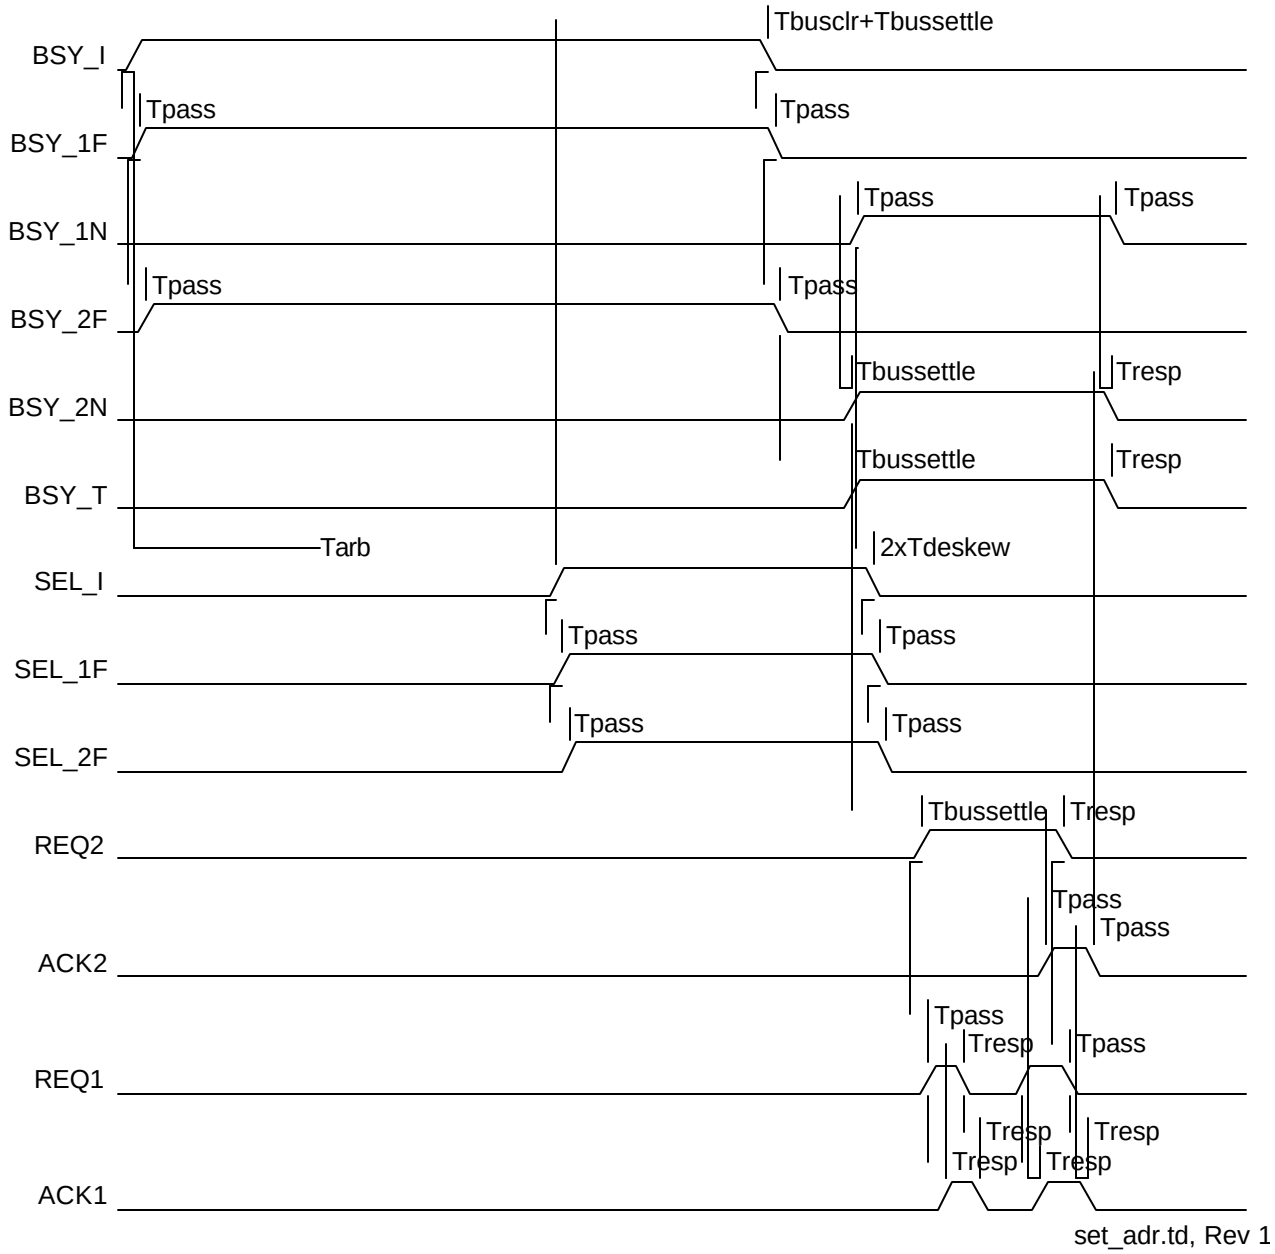

Figure 2 - Set Expander Address Select Command Timing

The Initiator that is allowed to set the Expander Addresses provides the addresses for all of the expanders. It does this by examining the Topology Map and determining which addresses are used. Some addresses may be used by hardwired devices and some by addressed expanders. From this information it can generate a pool of free ID's which it uses to send in the Data Out Phase of the Set Expander Address Select Commands. If the Initiator issues the Reset Expander Address Broadcast Select it must generate the free ID pool based only on the number of hardwired expanders in the Topology Map.

When any change is made in the topology, a new Topology Discovery Phase is required. The Expanders must also monitor the SCSI Bus for SCSI negotiation phases so it can set the Nexus transfer modes. This is true for both Hardwired and Auto-Addressable Expanders.

Expander Communication

Communication to the Expander is accomplished with standard SCSI Asynchronous Command/Status/Data protocol. This is used for both discovery as well as functional operation.

The Expander is selected using the Expander Selection Phase. Once the Expander is selected it is responsive to the following SCSI SPC-2 commands:

- INQUIRY - Responds with Expander Characteristics.
- MODE SELECT - Allows programmable options in the Expander.
- MODE SENSE - Responds with Expander Characteristics and defines changeable parameters.
- REQUEST SENSE - Method to determine errors on CHECK CONDITION.
- WRITE BUFFER - Used for Domain Validation and Code Download (if required)
- READ BUFFER - Used for Domain Validation

Once a connection is established, all data phase communications are asynchronous wide transfer mode. After power-up or SCSI Bus reset the default transfer mode is Asynchronous Narrow. If the Initiator attempts any negotiation (WDTR, SDTR, PPR) the SCSI Expander will respond and negotiate for Asynchronous Wide, but will negotiate down to Asynchronous Narrow.

INQUIRY Command

The Expanders support the INQUIRY Command as described below.

Table 2 - INQUIRY Command

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (12h)							
1	Reserved						CmdDt	EVPD
2	Page or Operation Code							
3	Reserved							
4	Allocation Length							
5	Control							

CmdDt The Command Support Data function is not supported by the SCSI Expander, and if set to 1 will result in a CHECK CONDITION.

EVPD The Enable Valid Product Data function is not supported by the SCSI Expander, and if set to 1 will result in a CHECK CONDITION.

Control The Control Byte is always 0, and any non-zero value will cause a CHECK CONDITION.

Page or Operation Code Not used, this field is ignored.

Allocation Length The Allocation Length must be set to 3Ah to receive all of the INQUIRY data defined in Table 6 on page 7. Any value less than 3Ah will cause a truncation of the INQUIRY Data followed by a CHECK CONDITION. Any value greater than 3Ah will transfer only the 58 bytes of INQUIRY Data and will not result in CHECK CONDITION.

Control The control byte is always 00h, and a non-zero value will result in a CHECK CONDITION.

The Expanders supports the Standard INQUIRY Data page shown in Table 6.

Table 3 - Standard INQUIRY Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier			Peripheral Device Type				
1	RMB	Reserved						
2	Version							
3	AERC	Obsolete	NormACA	HiSup	Response Data Format			
4	Additional Length (n-4)							
5	SCCS	Reserved						
6	BQue	EncServ	VS	MultiP	MChngr	Obsolete	Obsolete	Addr16
7	RelAdr	Obsolete	WBus16	Sync	Linked	TranDis	CmdQue	VS
8	Vendor Identification (8 bytes)							
15								
16	Product Identification (8 bytes)							
31								
32	Product Revision Level (4 bytes)							
35								
36	Vendor Specific							
55								
56	Reserved			FRC/SM	Clocking		QAS	IUS
57	Reserved							

Peripheral Qualifier This 3-bit field is always 000b specifying that the SCSI Expander is currently connected.

Peripheral Device Type This 5-bit field is always 10h for Bridging Expanders.

RMB This 1-bit field is always 0.

Version This 1-byte field is always 00h (TBD - will a version be defined for SPI-4?)

AERC This 1-bit field is always 0.

Obsolete All obsolete bits are always 0.

NormACA This 1-bit field is always 0.

HiSup This 1-bit field is always 0.

Response Data Format This field is always 010b.

Additional Length This field is always 53d (35h).

SCCS This 1-bit field is always 0.

BQue This 1-bit field is always 0.

EncServ This 1-bit field is always 0 (until an Enclosure Services Device is integrated)

VS This 1-bit field is always 0.

MultiP This 1-bit field is always 0.

MChngr This 1-bit field is always 0.

Addr16 This 1-bit field is always 1.

RelAdr This 1-bit field is always 0.

WBus16 This 1-bit field is always 1.

Sync This 1-bit field is programmable with the MODE SELECT Command and defaults to 1.

Linked This 1-bit field is always 0.

TranDis This 1-bit field is always 0.

CmdQue This 1-bit field is always 0.

VS This 1-bit field is always 0.

Vendor Identification This 8-byte field is always “ADAPTEC”, left aligned with trailing 00h.

Product Identification This 8-byte field is always “AIC-732”, left aligned with trailing zeroes.

Product Revision Level This 4-byte field defines the revision level of the SCSI Expander and starts with 00000000h.

Vendor Specific This 20-byte field contains the 14 bytes of Expander Characteristics, the remaining 6 bytes are unused and are always 00h.

WPC This 1-bit field is a 1 if the SCSI Expander supports SPI-4 Driver Write Pre-compensation (this is the SPI-4 Write Pre-compensation Enable Bit which may change per the SCSI Committee). This field is programmable by the MODE SELECT Command.

RCVEQ This 1-bit field is a 1 if the SCSI Expander supports SPI-4 Receiver Equalization (this is the SPI-4 Receiver Equalization Enable Bit which may change per the SCSI Committee). This field is programmable by the MODE SELECT Command.

FRC/SM This 1-bit field is a 1 for a SPI-4 SCSI Expander (this is the SPI-4 Free Running Clock and Skew Management Enable Bit, effectively the SPI-4 select mode, which may change per the SCSI Committee).

Clocking This 2-bit field defines the clocking support for the SCSI Expander (00b = ST Only; 01b = DT Only; 10b = Reserved; 11b = ST and DT).

QAS This 1-bit field defines whether the expander supports QAS Message snooping and passing of the QAS Message when not in-line to the IT Nexus.

IUS This 1-bit field is always 0.

DBIAS DRIVER BIAS CANCELLATION: This bit is a 1 if the SCSI Expander supports Driver Bias Cancellation.

RBIAS RECEIVER BIAS CANCELLATION: This bit is a 1 if the SCSI Expander supports Receiver Bias Cancellation.

MODE SELECT Command

The SCSI Expander supports the MODE SELECT(10) Command as described below. The Mode Pages are described in Section 1.3 on page 10.

Table 4 - MODE SELECT(10) Command

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (55h)							
1	Reserved			PF	Reserved			SP
2	Reserved							
3	Reserved							
4	Reserved							
5	Reserved							
6	Reserved							
7	Parameter List Length[15:8]							
8	Parameter List Length[7:0]							
5	Control							

PF The Page Format bit is used to specify vendor-specific (PF=0) or standard (PF=1) Mode Pages.

SP The Save Pages bit is always set to 0. The SCSI Expander does not save pages to non-volatile memory and will only keep the Current page. The Default Page is used after any SCSI Bus Reset or when the SCSI Expander is powered on. A value of 1 in this bit will result in a CHECK CONDITION.

Parameter List Length This 2-byte field is set to define the number of bytes that will be transferred in the Data Phase of this command. A value of 00h is valid and indicates no Parameter Pages will be transferred.

Control The Control Byte is always 0, and any non-zero value will result in a check condition.

In a multi-initiator environment, the SCSI Expander will accept a MODE SELECT Command that changes parameters from any initiator. Since the SCSI Expander is not normally in the functional communication domain, it does not generate a Unit Attention Condition since it has no method to report the condition (i.e. Message In Phase to a selection). Therefore, it is the responsibility of the initiator(s) to ensure compatible operation.

MODE SENSE Command

The SCSI Expander supports the MODE SENSE(10) Command as described below. The Mode Pages are described in Section 1.3 on page 10.

Table 5 - MODE SENSE(10) Command

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (5Ah)							
1	Reserved				DBD	Reserved		
2	PC		Page Code					
3	Reserved							
4	Reserved							
5	Reserved							
6	Reserved							
7	Allocation Length[15:8]							
8	Allocation Length[7:0]							
9	Control							

DBD This bit is ignored by the SCSI Expander since Block Descriptors are not required.

PC All Page Control values are supported (00 = Current Values; 01 = Changeable Values; 10 = Default Values; 11 = Saved Values). The SCSI Expander makes no distinction between the Current and Saved Values.

Page Code This 6-bit field is set with the Page Code of the page(s) to be returned. Refer to Section 1.3 on page 10 for valid Page Codes.

Allocation Length This 2-byte field is set to a value that allows transfer of all Mode Pages. If this size does not allow for all information to be transferred a CHECK CONDITION is received. If the size is greater than that required, only the required amount is transferred and CHECK CONDITION is not signaled.

Control The Control Byte is always 0, and any non-zero value will cause a CHECK CONDITION.

MODE Parameters

The Expanders support the following Mode Parameters for the MODE SENSE(10) and MODE SELECT(10) commands.

Table 6 - Mode Parameter List

Bit Byte	7	6	5	4	3	2	1	0
0 - 7	Mode Parameter Header							
8 - n	Parameter Pages							

Table 7 - Mode Parameter Header(10)

Bit Byte	7	6	5	4	3	2	1	0
0	Mode Data Length[15:8]							
1	Mode Data Length[7:0]							
2	Medium Type							
3	Device Specific Parameter							
4	Reserved							
5	Reserved							
6	Block Descriptor Length [15:8] (00h)							
7	Block Descriptor Length [7:0] (00h)							

Mode Data Length This field varies depending on which page(s) will be transferred. This is the length of the page(s) plus 6 bytes of this header (Mode Data Length is not included in the count). Refer to the descriptions of each page defined in Section 1.3.2 on page 12 for the appropriate Mode Data Length.

Medium Type This field is always 00h.

Device Specific Parameter This field is always 00h.

Block Descriptor Length This field is always 0000h since the SCSI Expander does not require description of the blocks of a medium.

Table 8 - Mode Page Format

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Reserve d	Page Code					
1	Page Length (n-1)							
2 - n	Mode Parameters							

PS This read-only bit is always 0.

Page Code This field defines the Page Code. For the MODE SELECT Command this is the page to be modified with the data in bytes 2-n. For the MODE SENSE Command this defines which page(s) will be returned in bytes 2-n.

Page Length (n-1) This field defines the number of bytes for the Mode Parameters in bytes 2-n. If this field is not configured for a MODE SELECT Command to be the same as returned for a MODE SENSE Command a CHECK CONDITION occurs (Sense Key ILLEGAL REQUEST, Sense Code INVALID FIELD IN PARAMETER LIST).

Mode Parameters These bytes define the Mode Parameters that will be modified with the MODE SELECT Command and provide the parameters for the MODE SENSE Command. Refer to Section 1.3.2 on page 12.

1.3.2

MODE Page Codes

The supported are defined in Mode Page Codes in Table 12. Any code not supported will result in CHECK CONDITION if specified in the MODE SELECT (Sense Key ILLEGAL REQUEST, Additional Sense Code INVALID FIELD IN PARAMETER LIST) and MODE SENSE (Sense Key ILLEGAL REQUEST, Additional Sense Code INVALID FIELD IN CDB) commands.

Table 9 - Supported Mode Page Codes

Page Code	Page Name	Section
00h	Vendor Specific	Vendor Specific
01h	SCSI Expander Capabilities Page	Section 1.3.2.1 on page 12
02h	Disconnect-Reconnect	Not Supported
03h - 08h	Device Specific	Not Supported
09h	Obsolete	Not Supported
0Ah	Control Mode Page	Not Supported
0Bh - 17h	Device Specific	Not Supported
18h	Protocol Specific LUN Page	Not Supported
19h	Protocol Specific Port Page	Not Supported
1Ah	Power Condition Page	Not Supported
1Bh	Device Specific	Not Supported
1Ch	Informational Exceptions Control Page	Not Supported
1Dh - 3Eh	Device Specific	Not Supported
3Fh	Return All Pages (valid only for the MODE SENSE Command)	Section 1.3.2.2 on page 14

Page Code 01h (SCSI Expander Capabilities Page)

The Expander Capabilities page is used to define and set capabilities of the SCSI Expander.

Table 10 - SCSI Expander Capabilities Page

Bit Byte	7	6	5	4	3	2	1	0
0	PS (0)	Reserve d	Page Code (01h)					
1	Page Length (09h)							
2	Reserved			FRC/SM	Clocking		Reserved	
3	Reserved		RstIso	IsoSeg	RBIAS	DBIAS	WPC	AAF
4	Near	CTLA	BSYA	OffA	GRstA	RstA	SigA[1:0]	
5	Reserve d	CTLB	BSYB	OffB	GRstB	RstB	SigB[1:0]	
6 - 10	Reserved							

PS This read-only field is always 0.

Page Code This field is 01h to define the SCSI Expander Capabilities Page.

Page Length This field is always 09h.

FRC/SM FREE-RUNNING CLOCK WITH SKEW MANAGEMENT: This read-only field defines SPI-4 protocol is supported. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = 1; Changeable = 0; Default = 0; Saved = 1. The Reset State is (X,X,X). See Section 3 on page 15.

Clocking SCSI BUS CLOCKING: This read-only field defines the SCSI Bus clocking protocol. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = 11; Changeable = 00; Default = 11; Saved = 11. The Reset State is (X,X,X). See Section 3 on page 15.

RstIso RESET ISOLATION: This read-only bit defines whether the expander is in the Reset Isolation Mode. See Section 4 on page 15. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 0; Default = 0; Saved = Current Setting. The Reset State is (0,X,X). See Section 3 on page 15. Additionally, this bit is reset when a Start Topology Discovery or Start Expander Addressing Broadcast Select is issued from either side.

IsoSeg ISOLOATE SEGMENTS: This bit defines whether the expander is in the Isolation Mode. See Section 5 on page 15. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1; Default = 1; Saved = Current Setting. The Reset State is (0,X,X). See Section 3 on page 15.

Near NEAR-SIDE SELECT: This read-only bit defines which side of the SCSI Expander is the near-side, where 0 = A and 1 = B. This bit will always read the current side for any MODE SENSE Command Page Control setting. It could be the same for two initiators (both on the same side) or it could be different (on opposite sides). For the MODE SENSE Command this field returns the following for the Page Control settings: Current = 0 (Side A) or 1 (Side B); Changeable = 0; Default = 0 (Side A) or 1 (Side B); Saved = 0 (Side A) or 1 (Side B). The Reset State is (X,X,X). See Section 3 on page 15.

BSYn BSY SIGNAL STATE Side n: This read-only bit reflects the current state of BSY on Side n. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current State of BSY; Changeable = 0; Default = Current State of BSY; Saved = Current State of BSY. The Reset State is (X,X,X). See Section 3 on page 15.

CTLn CONTROL SIGNAL STATE Side n: This read-only bit reflects whether any of the signals SEL, C/D, I/O, MSG, ATN, REQ, ACK, SD or SDP are asserted on Side n. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current State of BSY; Changeable = 0; Default = Current State of BSY; Saved = Current State of BSY. The Reset State is (X,X,X). See Section 3 on page 15.

RBIAS RECEIVER BIAS CANCELLATION: This bit defines whether SCSI Receiver Bias Cancellation is supported and enabled. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1; Default = 1; Saved = Current Setting. The Reset State is (1,X,X). See Section 3 on page 15.

DBIAS DRIVER BIAS CANCELLATION: This bit defines whether SCSI Driver Bias Cancellation is supported and enabled. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1; Default = 1; Saved = Current Setting. The Reset State is (1,X,X). See Section 3 on page 15.

WPC WRITE PRE-COMPENSATION: This bit defines whether SPI-4 Write Pre- compensation is supported and enabled. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1; Default = 1; Saved = Current Setting. The Reset State is (1,X,X). See Section 3 on page 15. If this bit is set (enabled) the Write Pre-compensation turned ON or OFF depending on information snooped in the PPR Message. If this bit is reset (disabled) the expander will never attempt Write Pre-compensation.

AAF ADJUSTABLE ACTIVE FILTER: This bit defines whether SPI-4 Adjustable Active Filtering is supported and enabled. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1 (0 if not supported); Default = 1 if supported and 0 if not supported; Saved = Current Setting. The Reset State is (1,X,X if supported and 0,0,0 if not supported). See Section 3 on page 15. If this bit is set (enabled) the

AAF is turned ON or OFF depending on information snooped in the PPR Message. If this bit is reset (disabled) the expander will never attempt input filtering.

Rstn n-SIDE RESET CONTROL: This field defines the SCSI Bus Reset Propagation capability from Side n to Side m, where 0 = Do not Propagate and 1 = Propagate. Refer to Section 4 on page 15. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1; Default = 1; Saved = Current Setting. The Reset State is (0,X,X). See Section 3 on page 15.

GRstn GENERATE n-SIDE RESET: This field is written with a 1 to generate a SCSI Bus Reset on Side n. Refer to Section 4 on page 15. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current State of Side n RST Signal; Changeable = 1; Default = 0; Saved = Current State of Side n RST Signal. The Reset State is (0,0,0). See Section 3 on page 15.

Offn n-SIDE OFF: This field defines the on/off capability of the n-side SCSI I/O, where 0 = ON; 1 = OFF. Refer to Section 5 on page 15. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Current Setting; Changeable = 1; Default = 0; Saved = Current Setting. The Reset State is (0,X,X). See Section 3 on page 15.

Sign[1:0] SIGNAL MODE n-SIDE: This field defines the signaling mode of the n-Side SCSI Bus Segment after a reset or power-up, where 00 = Single-Ended; 01 = LVD; 10 = HVD; 11 = reserved. If the n-side comes up HVD the n-Side is OFF and cannot be turned on. For the MODE SENSE Command this field returns the following for the Page Control settings: Current = Detected Setting; Changeable = 00b; Default = Detected Setting; Saved = Detected Setting. The Reset State is (X,X). See Section 3 on page 15.

1.3.2.2 Page Code 3Fh

A 3Fh Page Code selected in a MODE SELECT Command will result in CHECK CONDITION (Sense Key ILLLEGAL REQUEST, Additional Sense Code INVALID FIELD IN PARAMETER LIST).

A 3Fh Page Code selected in a MODE SENSE Command will result in the transfer of Page Code 01h, 19h, and 00h (in that order). If the Allocation Length of the MODE SENSE Command is less than that required to transfer all parameters, the command results in CHECK CONDITION (Sense Key ILLEGAL REQUEST, Additional Sense Code INVALID FIELD IN CDB).

REQUEST SENSE Command

The REQUEST SENSE Command is supported to enable reporting of error status for the commands supported by the expander.

WRITE BUFFER Command

The WRITE BUFFER Command is supported to enable downloading information and for providing fault isolation during Domain Validation. This command shall be functional for all speeds supported by the expander and can be margined by using the Margin Control Message the MODE SELECT Command.

READ BUFFER Command

The READ BUFFER Command is supported to enable fault isolation during Domain Validation. This command shall be functional for all speeds supported by the expander and can be margined by using the Margin Control Message the MODE SELECT Command.

Expander Negotiation

The SCSI Expander shall negotiate with the initiator for the Expander Communication. The default after Power-On or SCSI Bus reset is Asynchronous Narrow. The negotiation with the expander only affects

the data phases for the communication with the expander. WARNING: The initiator must negotiate for asynchronous data phases, narrow or wide, before issuing the Set Expander Address Select Command.

The pass-through mode and speed is snooped by the expander as negotiation occurs between the initiator and target when the expander is in-line to the IT Nexus. The speed negotiated with the target is applicable to all pass-through data as well as capture and modify data used in the Topology Discovery and Expander Addressing. The expanders capture the Status In Phase of the WRITE/READ BUFFER Command and insert the Data In Phase, that Data In Phase being the same negotiated with the target. WARNING: The initiator must negotiate for asynchronous data phases, wide or narrow, before issuing the WRITE/READ BUFFER Command for the Topology Discovery and Expander Addressing Phases.

I/O Subsystem Reset

The SCSI I/O Subsystem has three reset conditions:

- Power-On Reset
- SCSI Reset
- Expander Reset

The Expanders require a Power-On reset to initialize certain parameters, allowing them to become fully operational. SCSI Reset is the SCSI Bus RST# signal. The Expander Reset is a Broadcast Select.

The Reset State of the Expander Capabilities is defined in the definition for each function in the form (a,b,c), where a is the state from Power-On Reset, b is the state from SCSI Reset, and c is the state from Expander Reset. A value of X for a, b, or c signifies that bit is not affected by that reset.

Reset Isolation

The SCSI I/O Subsystem can be configured to isolate the reset of the bus segments. The Isolation Mode is enabled with the MODE SELECT Command, Expander Characteristics Mode Page.

When the Expander is in this mode, any assertion of RST on a bus segment will not be propagated to the other side. The Expander enters the Reset Isolation Mode which does not allow signals from the other side to propagate to the reset side. If the hang condition on the bus was on the far-side segment, the disable of signal propagation will allow the reset segment to be usable.

For example, if a device isolated from a Initiator by several in-line expanders has hung the bus with BSY asserted all segments are hung. The Initiator system detects the hang and issues a RST on its immediate segment. All Expanders and devices on that segment are reset, but not the other segments. The Expanders isolate the hang condition from the immediate segment, allowing the Initiator to check the state of the far side using either the INQUIRY or MODE SENSE (Expander Characteristics Page). This activity on the near-side is not propagated to the far side. The Reset Isolation Mode is automatically exited if the SCSI Bus signals on the far side are not asserted. If they are asserted, the far side must be reset using the MODE SELECT Command. The far side reset will automatically exit the Reset Isolation Mode. The path of the hang condition is followed until the offending device is finally reset.

For a multi-initiator environment, the Initiators must provide communication protocol so isolated reset of devices is known. The Initiator software must be able to distinquish which devices have been reset as opposed to normal SCSI Bus non-isolated reset which resets all devices.

Segment Isolation

The SCSI I/O Subsystem can be configured to isolate segments of the bus. This can be used to isolate faulty devices or to allow Spatial Reuse.

The Segment Isolation Mode is selected by the MODE SELECT Command. When isolated, no signals are propagated from one side to the next. The Expander will respond to Expander Selection and commands from either side.

For example, a Initiator on Segment X directs another Initiator on Segment Y to back up a device to tape, both residing on Segment Y. Initiator X isolates the segments by putting the expander in the Segment Isolation Mode. Initiator Y polls the expander until it is in the Segment Isolation mode and then executes the backup on Segment Y. In the meantime, Initiator X can use Segment X but must not attempt an access to any device(s) on Segment Y. When Initiator Y is done, it can turn off Segment Isolation in the Expander and communicate the done status to Initiator X.

Additionally, the Expander can be completely isolated. The MODE SELECT Command can set the OffA or OffB control. When this bit is asserted, nothing can pass through the OFF side. This includes any SCSI Bus signal, and the expander will not respond to Expander Selection from the OFF Side. The Initiator software must be careful not to turn OFF the expander to the near side. If this is done, the expander cannot be enabled again until a power-on reset.

Supported Functions

Faulty Segment Isolation

The Expanders could be used as a switch to disable a particular Segment and any other downstream devices. This allows removing ill-behaved devices until they can be removed or replaced.

The Expander accepts the Expander Control Command with parameters to disable the far-side drivers and receivers.

Domain Validation

The Expanders can be controlled to modify only one side to isolate faulty cables, connections, drivers and receivers down to a particular Expander Segment.

The Expander accepts the Expander Control Command with parameters to manipulate the defined Margining Domain Validation parameters for the near-side and/or far-side. Additionally, it snoops the Margin Control Message when in-line to the IT Nexus and adjusts it parameters accordingly.

The terminators and enclosure chips can be modified in the same manner.

Performance Status

The Expanders Characteristics are received during the Topology Discovery Phase and define the performance.

The Expander also accepts MODE SENSE and MODE SELECT Commands to provide other status and control.

The Expanders will be capable of extending a specific protocol and speed and will not translate from one speed and/or protocol to another (i.e. U160 Packetized to U160 Packetized, not U160 Non-Packetized to any other speed or Packetized, etceteras). This kind of information is required for a useful Topology Map. For example, if a U160 Expander is in-line to a U320 Device from a U320 Initiator, the Initiator must not negotiate for U320. Negotiation will succeed but the subsequent transfers will fail through the U160 Expander.