TO: T10 Membership
FROM: Paul A. Suhler, Seagate Removable Storage Solutions, LLC
DATE: 21 March 2003
SUBJECT: T10/03-077r4, ADI Bridging Proposal

Revision 4:

- Merged both forms of bridging into one.
- Reorganized AER IU. AER IU from automation is named "Automation AER IU" and the one from DTD is "DTD AER IU."
- Added four-byte First Burst Length field to SCSI Request IU for data-in commands.
- Not included:
    - Restricting non-LUN 0 device servers from supporting transport-related mode and log pages.
    - Renaming automation AER IU to DTD AER IU.

Revision 3:

- Revised ADC model to change concept of partitioned SMC device server into two independent device servers with different sets of mandatory commands.
- Banned splitting the Bridging Login IU across multiple frames.
- Clarified that the bridging manager may initiate commands to remote device server while operating in passthrough bridging mode, e.g., to refresh cached data.
- To support caching of inquiry, etc. data by the data transfer device, an asynchronous notification IU from the automation to the drive is now defined.

Revision 2:

- In Bridging Login IU, changed LU Index back to LUN.
- Added padding field to SCSI Request IU.

Revision 1:

- The feature is renamed to "bridging" with "passthrough" and "hosted" variants.
- ADC text no longer refers to ADT.
- Added a model section for ADT. This refers to ADC.
- Changed "frame" references to "information unit."
- 4.2.x: Sentence with "typically" was rewritten as an example.
- "Surrogate mode" field in mode page was changed to an enabled bit plus one bit for each type of bridging. To determine what types are supported, automation can request the changeable mode page. (The alternative is trial and error setting of bits.) Only one type bit shall be set.
- Put back in text saying that if bridging is disabled, then that LUN is not reported.
- "Process login" is changed to "Bridging Login." This makes the intent of the IU apparent to a reader of the ADT standard.
- Moved the ACCEPT bit to byte 0, bit 7 in bridging login and logout IUs.
- Added a LOGICAL UNIT INDEX field to the login IU just in case we need to support bridging for multiple LUs in the automation device. Other text still refers specifically to the SMC device server. Do we want to make it more general at this time? Future changes to support multiple bridged LUs of types other than SMC will not require IU changes, just rewording of the standard.

Revision 0:

Specification of ADI Bridging Operation requires changes to:

- ADC model to describe bridging operation
- ADC automation drive descriptor to select different bridging modes
- ADT link services clause to add bridging login and logout
- ADT SCSI Request frame payload to add an I_T nexus identifier

FCP-2 doesn't have an entry in the definitions clause for process login, so I haven't provided one here for bridging login.

Thanks to Paul Entzel for providing large parts of this text. I wordsmithed some of his text, so blame me for any problems.

The following is to be added to the ADC model section:

4.2.x ADI bridging

4.2.x.1 ADI bridging introduction

The data transfer device may optionally support ADI bridging for the automation device. When this operation is enabled, the data transfer device reports a logical unit to its primary interface ports that implements an SMC device server, and the automation device reports a logical unit to its ADT port that implements an SMC device server. In the process of executing a SCSI command or task management request, the data transfer device's SMC device server (called the "local device server") may invoke execution of a command or request by the automation device's SMC device server (called the "remote device server"). The entity within the local device server that performs invocation of commands or requests on the remote device server is called the bridging manager.

The effect is that some or all commands and requests addressed to the local device server are passed to the remote device server through the ADT port. This can be used, for example, in low-cost automation devices that do not have separate primary interface ports.

4.2.x.2 Local device server operation

The local device server shall support commands as required by the SCSI Medium Changer device type. The remote device canot necessarily implement the full set of commands required by the SCSI Medium Changer device type, as no indication of the original initiator port of the SCSI Request is passed between the data transfer and automation devices. Thus, the local device server shall service commands and task management functions that require knowledge of the originating initiator port. Effectively, the data transfer device acts as a protocol bridge.

If any of the following commands are supported, they shall be serviced by the local device server and not passed through to the remote device server:

a) RESERVE(6) and RESERVE(10)
b) RELEASE(6) and RELEASE(10)
c) PERSISTENT RESERVE IN
d) PERSISTENT RESERVE OUT
e) REPORT LUNS
f) REQUEST SENSE

The local device server shall also perform the following actions:

a) Check for reservation conflicts on all commands routed to the remote device server. Return RESERVATION CONFLICT on all commands that violate a reservation condition and not pass them through to the SMC device.
b) Manage UNIT ATTENTION conditions generated for multiple initiators. Before passing a command through to the remote device server, pending UNIT ATTENTION conditions shall be checked. If the local device server detects that a UNIT ATTENTION condition is pending for an initiator port when a new command is received from it, the local device server shall return Check Condition for the command without passing it to the remote device server.
c) When the primary interface uses contingent allegiance, save sense data on a per initiator port basis.

4.2.x.3 Caching SMC data

In some implementations, the bridging manager may cache some data from the remote device server. For instance, it may save the inquiry data from the remote device server and return it to any initiator port that requests it. The ADT standard provides means for the remote device server to notify the local device server when data that may be cached has changed. The local device server should discontinue using cached data from the remote device server when it is informed of a possible change until it has been refreshed.

The following change is for ADC clause 6.2.2.3.3, Medium Changer descriptor parameters:

- The descriptive text following Table 29 – Medium changer descriptor is changed as follows:

If the ENABLE field is set to one it indicates that the Logical Unit is reported and supported. Commands received for this logical unit ~~shall~~ may be passed on to the automation device using the ADT interface. When it is set to zero, the logical unit is not reported to a REPORT LUNS command and does not respond to commands.

If the ENABLE field is changed from one to zero, then the local device server shall implicitly abort all commands issued to the automation device and shall report a status of CHECK CONDITION with a sense key of COMMAND ABORTED and an additional sense code of I_T NEXUS LOSS OCCURRED for each command. All remaining device servers in the data transfer device shall report a change in the logical unit inventory, as specified in SPC-2, to any application clients connected through a primary interface port.

The following changes are for ADT:

- 7.1.2, Table 8 – addition of first burst length field for data-in commands, to facilitate throttling data-in transfers:

Byte \ Bit	7	6	5	4	3	2	1	0
0	(MSB)			LUN				
1								(LSB)
2	TASK MANAGEMENT FUNCTION							
3	Reserved							
4				CDB				
19								
20	(MSB)			BUFFER ALLOCATION LENGTH				
23								(LSB
24	(MSB)			FIRST BURST LENGTH				
27								(LSB)

When the SCSI Request IU specifies a data-in command, the FIRST BURST LENGTH field shall contain a value indicating the amount of buffer space prepared in the initiator for the first SCSI Data IU and requesting the transfer from the target of an IU of that length. When the SCSI Request IU specifies a non-data or data-out command, the FIRST BURST LENGTH field shall contain zero.

- 7.1.3 SCSI Response IU – add new response code value to Table 11:

Value	Description
06h	Command complete with GOOD status and sense data valid
07h - FFh	Reserved

The response code value of "Command complete with GOOD status and sense data valid" shall be sent by the automation device when bridging is enabled and a command completes with a SCSI status of GOOD that will generate a unit attention to initiator ports other than the one that initiated the command. In this case, the SCSI AUTOSENSE DATA field shall contain the sense data to be reported to those other initiator ports.

- 7.1.4 SCSI Transfer Ready IU description – generalized for both data-in and data-out:

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

In ADT, two types of AER information units are defined, one for transmission by the automation device, the other for transmission by the DTD:

7.2.4 DTD AER information unit

DTD Asynchronous Event Report IUs may optionally be ~~used~~ sent by the data transfer device to report that an event has occurred that may be of interest to the automation device. Only a Data Transfer device may initiate DTD AER IUs. The payload of ~~an~~ a DTD AER IU shall contain the VHF Polling data as defined in ADC.

7.2.6 Automation DTD AER information unit

Automation Asynchronous Event Report IUs may optionally be sent by the automation device to report that an event has occurred that may be of interest to the DTD. The payload of an Automation AER IU shall contain the payload specified in Table x below. It shall be transmitted when bridging operation is enabled and when a change has occurred to data in a mode page, log page, or vital product data page. It may also be transmitted when a media access command being executed by the automation device has failed, regardless of whether bridging is enabled.

Table x -- Automation AER information unit

Bit / Byte	7	6	5	4	3	2	1	0
0	Reserved							FAIL
1	Reserved		ADLE	BUASDV	READY	RSC	IDC	MPC
2	BROADCAST UNIT ATTENTION ASC							
3	BROADCAST UNIT ATTENTION ASCQ							

Editor's Note: The FAIL bit and its descriptive text from T10/03-087r1 is included here, even though it is not relevant to bridging.

The FAIL field is set to one to notify the DTD of the failure of an operation being performed by the automation device. For example, if a MOVE MEDIUM command has failed because of an inability to load a medium into the DTD, then an Automation AER IU with this field set to one can

be interpreted by the DTD as a signal to report a load failure to initators connected through a primary interface port.

The fields in byte 1 are collectively known as the bridging status byte and are used to notify the bridging manager in the DTD of events in the remote SMC device server that may require changing cached SMC data. If bridging is not enabled, then these fields shall be zero.

The Mode Page Changed (MPC) field shall be set to one when bridging operation is active and the contents of any mode page reported by the remote device server has been changed since the last execution of a MODE SENSE or MODE SELECT command.

The Inquiry Data Changed (IDC) field shall be set to one when bridging operation is active and the contents of the standard inquiry data or of any vital product data page reported by the remote device server has changed since the last execution of an INQUIRY command.

The Ready Status Changed (RSC) field shall be set to one when bridging operation is active and the SCSI Status that will be reported by the remote device server on the next command is GOOD and the status reported on the previous command was any value other than GOOD. It shall also be set to one when the status that will be reported by the remote device server on the next command is any value other than GOOD and the status reported on the previous command was GOOD.

The READY field shall be set to one when the the SCSI Status that will be reported by the remote device server on the next command is GOOD. It shall be set to zero when the the SCSI Status that will be reported by the remote device server on the next command is any status other than GOOD.

When the BROADCAST UNIT ATTENTION SENSE DATA VALID (BUASDV) field is set to one, the BROADCAST UNIT ATTENTION ASC and BROADCAST UNIT ATTENTION ASCQ fields shall contain sense data to be reported to all initiators by the local SMC device server. When the BUASDV field is set to zero, these fields shall be ignored.

The automation device may set the AUTOMATION-DETECTED LOAD ERROR (ADLE) field to one when it abandons attempts to load a medium into the data transfer device because of errors in loading.

The data transfer device's bridging manager is not required to request any or all of the changed pages reported. It may request pages not reported as changed and may request standard inquiry data.