

Date: July 13, 2004

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: SAM-3; SPC-3; Allowing only Task Management Functions through ACA

1 Overview

As defined today ACA is difficult to implement on protocols where the target does not have complete control of transfers on the interconnect. This applies to most, if not all, of the serial protocols. The problem usually happens if an ACA occurs and the target has outstanding transfer ready credit(s) to an initiator. If this occurs and an application client sends down a new command with an ACA tag during the ACA condition a dead-lock may occur at the logical unit.

The solution that is being proposed is that a mode page option be added that, when selected, would allow a target to give an ACA ACTIVE status for all commands received from the faulted initiator during the ACA. In this case initiators would only be allowed to send task management functions from the faulted initiator to the logical unit with an ACA.

2 SPC-3 changes

2.0.1 Control mode page

The Control mode page (see table 1) provides controls over SCSI features that are applicable to all device types (e.g., task set management and error logging). If a field in this mode page is changed while there is a task already in the task set, it is vendor specific whether the old or new value of the field applies to that task. The mode page policy (see 6.7) for this mode page shall be shared, per initiator port, or per I_T nexus.

Table 1 — Control mode page

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (0Ah)					
1	PAGE LENGTH (0Ah)							
2	TST			TMF_ONLY	Reserved	D_SENSE	GLTSD	RLEC
3	QUEUE ALGORITHM MODIFIER				Reserved	QERR		Obsolete
4	TAS	RAC	UA_INTLCK_CTRL	SWP	Obsolete			
5	APTG_OWN	Reserved				AUTOLOAD MODE		
6	Obsolete							
7								
8	(MSB)	BUSY TIMEOUT PERIOD						(LSB)
9								
10	(MSB)	EXTENDED SELF-TEST COMPLETION TIME						(LSB)
11								

[The allow task management functions only \(TMF_ONLY\) bit set to zero specifies tasks with a task attribute of ACA may be sent from the faulted initiator port when an ACA condition has been established \(see SAM-3\). A TMF_ONLY bit set to one specifies that all tasks sent from the faulted initiator port when an ACA condition has been established shall be terminated with an ACA ACTIVE status.](#)

3 SAM-3 changes (these changes are already requested in the SAM-3 letter ballot comments)

3.1 Status

3.1.1 Status codes

The status codes are specified in table 21. Status shall be sent from the device server to the application client whenever a command ends with a service response of TASK COMPLETE or LINKED COMMAND COMPLETE.

.....

ACA ACTIVE. This status shall be returned when an ACA exists within a task set and a SCSI initiator port issues a command for that task set when at least one of the following is true:

- a) There is a task with the ACA attribute (see 8.6.5) in the task set;
- b) The SCSI initiator port issuing the command did not cause the ACA condition; ~~or~~
- c) The task created to process the command did not have the ACA attribute and the NACA bit was set to one in the CDB CONTROL byte of the faulting command (see 5.9.1); or
- d) The TMF ONLY (see SPC-3) bit is set to one and the command was issued from the faulted SCSI initiator port.

The SCSI initiator port may reissue the command after the ACA condition has been cleared.

....

3.1.2 Auto contingent allegiance (ACA)

When a command is terminated with a CHECK CONDITION status, the application client may request that device server alter command processing by establishing an ACA condition. When the NACA bit in the CONTROL byte of the CDB for the command is set to one, the device server shall establish an ACA condition as described in 3.1.2.1. Upon establishment of the ACA condition, some tasks other than the task returning the CHECK CONDITION status may be aborted and continued processing of other tasks may be blocked as described in 3.1.2.1.

While the ACA condition is in effect and the TMF ONLY bit (see SPC-3) is set to zero, new tasks received by the logical unit from the faulted initiator port are not allowed to enter the task set unless they have the ACA task attribute (see 8.6.5). One of the results of the ACA task attribute requirement is that commands in-flight when the CHECK CONDITION status occurs are returned unprocessed to the SCSI initiator port with an ACA ACTIVE status. Multiple commands may be sent one at a time using the ACA task attribute to recover from the event that resulted in the ACA condition without clearing the ACA.

While the ACA condition is in effect, and the TMF ONLY bit is set to one no new tasks received by the logical unit from the faulted initiator port are allowed to enter the task set (see 8.6.5).

While the ACA condition is in effect:

- a) New tasks from the faulted initiator port shall be handled as described in 3.1.2.2, and
- b) New tasks from SCSI initiator ports other than the faulted initiator port shall be handled as described in 5.9.2.3.

The methods for clearing an ACA condition are described in 5.9.2.4.

3.1.2.1 Establishing an ACA

.....

3.1.2.2 Handling new tasks from the faulted initiator port when ACA is in effect

Table 2 describes the handling of new tasks from the faulted initiator port when ACA is in effect.

Table 2 — Handling for new tasks from a faulted initiator port during ACA

New Task Properties		ACA Task Present in the Task Set	TMF ONLY Value ^e	Device Server Action	ACA Established If New Task Terminates with a CHECK CONDITION status
Attribute ^a	NACA Value ^b				
ACA	0	No	0	Process the task. ^d	No ^c
	1	No	0		Yes ^c
	n/a	n/a	1	Terminate the task with ACA ACTIVE status.	n/a
	0 or 1	Yes	n/a	Terminate the task with ACA ACTIVE status.	n/a
Any Attribute Except ACA	0 or 1	n/a	n/a	Terminate the task with ACA ACTIVE status.	n/a

^a Task attributes are described in 8.6.

^b The NACA bit is in the CONTROL byte in the CDB (see 5.2).

^c If a task with the ACA attribute terminates with a CHECK CONDITION status, the existing ACA condition shall be cleared and the value of the NACA bit shall control the establishment of a new ACA condition.

^d All the conditions that affect the processing of commands (e.g., reservations) still apply.

^e [The TMF ONLY bit is in the Control mode page \(see SPC-3\).](#)