

CONGRUENT SOFTWARE, INC.
 3998 Whittle Avenue
 Oakland, CA 94602
 (510) 531-5472
 (510) 531-2942 FAX

FROM: Peter Johansson
 TO: X3T10 Plenary
 DATE: September 18, 1996
 RE: New Task Management Models for SAM-2

At the May X3T10 SCSI-3 Working Group meeting in Ft. Lauderdale Florida, Jim McGrath and I introduced and discussed proposed simplifications to the SAM task management (queuing) model. The Working Group was asked for a straw poll to endorse the concepts and the straw poll passed 17:1. This proposal is the result of additional work to refine these ideas and present them in a concrete form for approval by the SCSI-3 Working Group.

These ideas were discussed in the SCSI-3 Working Group in Colorado Springs and received substantial support. Subsequent discussion in Natick resulted in a proposal that the working group believes is ready to be brought before the plenary.

The essence of the proposal is to modify SAM-2 to permit two levels of task management complexity: basic¹ and full. The salient features of each of these schemes is summarized in the table below.

Model	Task attributes	Task management functions	Control byte	Control mode page		
			NACA	Queue algorithm modifier	QErr	DQue
Basic	SIMPLE	ABORT TASK ABORT TASK SET TARGET RESET	0 (SCSI-2 contingent allegiance)	1 (unrestricted reordering)	1	0
Full	SIMPLE ORDERED HEAD OF QUEUE ACA	ABORT TASK ABORT TASK SET CLEAR ACA CLEAR TASK SET TARGET RESET TERMINATE TASK	—	—	—	—

Since SAM-2 may adopt a style or organization that differs from SAM, the proposals below that suggest a way to express the basic task management model may not be appropriate in the forthcoming document. If the X3T10 Plenary endorses the technical details of this proposal, I suggest that I coordinate with both the SAM-2 and SPC-2 editors to arrive at the appropriate language to express the changes.

Section 6 in SAM, Task Management Functions, establishes optional and mandatory requirements for task management function implementation by SCSI targets. I suggest that only the definition of clear task set requires modification:

CLEAR TASK SET (Logical Unit Identifier ||) – Abort all tasks in the specified task set. This function shall be supported by all logical units that support tagged tasks (see object definition 7) that implement

¹ In Ft. Lauderdale, the new model was referred to as "simple task management." Since this name causes some confusion with the SIMPLE task attribute, I suggest that the model be renamed to "basic."

the full task management model. This function may be supported by logical units that do not support tagged tasks.

Section 7 in SAM, Task Set Management, seems to be the appropriate place to distinguish between basic and full queuing models. I suggest that the following paragraphs be added at the end of the first clause in section 7, just before the header for 7.1:

SAM specifies two models for task set management: basic and full. The basic task management model is a subset of the full model and differs from it in five particulars:

- The only task attribute supported is SIMPLE;
- The device server may reorder the actual execution sequence of tasks in any manner. Any data integrity exposures related to task sequence order shall be explicitly handled by the application client through the selection of appropriate commands;
- All the blocked tasks in the task set shall be aborted after an ACA condition is cleared;
- It is not possible to disable tagged queuing; and
- Support for the CLEAR TASK SET task management function is a logical unit option.

Once SAM-2 is modified to permit the two different task management models, a mechanism is needed for devices to report which model they implement. Consensus emerged in the SCSI-3 Working Group in Colorado Springs that INQUIRY data is the most suitable location to report task management model variants.

The following text and tables are extracted from SPC Revision 9b and show (underlined) the addition of a new bit, BQue, to indicate support for the basic task management model. The addition of the BQue bit also necessitates a change in the definition of the CmdQue bit.

Table 19 – Standard INQUIRY data format

Bit Byte	7	6	5	4	3	2	1	0	
0	Peripheral qualifier			Peripheral device type					
1	RMB	Reserved							
2	ISO / IEC version		ECMA version			ANSI version			
3	AERC	TrmTsk	NormACA	Reserved	Response data format				
4	Additional length (n - 4)								
5	Reserved								
6	<u>BQue</u>	EncServ	VS	MultiP	MChngr	ACKREQQ	Addr32	Addr16	
7	RelAdr	WBus32	WBus16	Sync	Linked	TranDis	CmdQue	VS	
8	(MSB)	Other INQUIRY data as specified by SPC-2							
n								(LSB)	

The basic queuing (BQue) bit shall be zero if the CmdQue bit is one.

When the CmdQue bit is zero, the BQue bit shall have the following meaning. A BQue bit of zero indicates that the device does not support tagged tasks (command queuing) for this logical unit. A value of one indicates that the device supports, for this logical unit, the basic task management model defined by SAM.

A device that reports a BQue bit of one that also supports the control mode page for the logical unit shall report a value of one for each of the queue algorithm modifier field and the QErr and DQue bits. These fields shall be unchangeable.

A command queuing (CmdQue) bit of one indicates that the device supports tagged tasks (command queuing) for this logical unit (see SAM). A value of zero indicates the device server ~~may~~ does not support tagged tasks for this logical unit.

Because SBP-2 makes use of the basic task management model and because disk drive vendors wish to use the basic task management model in the parallel SCSI environment, I propose that these changes be promptly incorporated in SAM-2 and SPC-2.