To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 14 October 2003
Subject: 03-351r0 SAM-3 SPC-3 Extended Task Attributes mode page

Revision history
Revision 0 (14 October 2003) First revision

Related documents
sam3r09 - SCSI Architecture Model - 3 revision 9
spc3r15 - SCSI Primary Commands - 3 revision 15

Overview

There are currently four task attributes defined in SAM-3:

a) SIMPLE;
b) ORDERED;
c) HEAD OF QUEUE; and
d) ACA.

Problems:

1. There is no way for an application client to tell which task attributes are supported. Some logical units may prefer not to support ORDERED, for example. This proposal lets software read a mode page to determine which task attributes are supported.

2. When asymmetric logical unit access is used (the SET/REPORT TARGET PORT GROUPS commands), applications can access a logical unit through more than one target port. There is no ordering relationship guaranteed between the I_T nexuses to the different target ports, even if the I represents the same initiator port. Multipathing software typically sits below the application in the initiator device, choosing which initiator port to use and which target port to address (i.e., choosing the I_T nexus). If the application tries to use

ORDERED tags, there is no good way for the multipathing software to preserve ordering across separate I_T nexuses. The commands may arrive at the logical unit in a different order than the application intended.

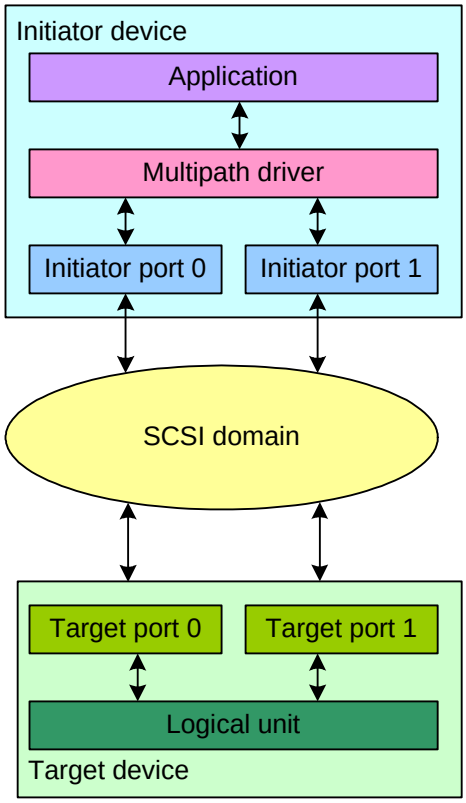

Application sends: Cmd1, Cmd2, Cmd3, Cmd4

Assuming multipath software just uses round-robin and always has Initiator port 0 send to Target port 0:

Initiator port 0 sends: Cmd1, Cmd3
Initiator port 1 sends: Cmd2, Cmd4

Target port 0 receives:
Cmd1, Cmd3 in ordered SCSI transports or
Cmd3, Cmd1 in unordered SCSI transports

Target port 1 receives:
Cmd2, Cmd4 in ordered SCSI transports or
Cmd4, Cmd2 in unordered SCSI transports

Logical unit receives:
Cmd1, Cmd2, Cmd3, Cmd4 or
Cmd1, Cmd3, Cmd2, Cmd4 or
Cmd2, Cmd4, Cmd1, Cmd3
in ordered SCSI transports

or any permutation in unordered SCSI transports

Figure 1 — Asymmetric logical unit access

This proposal lets multipathing software disable support for ORDERED tags via the mode page so the higher-level application knows not to use them. It could either write to the mode page and disable the ORDERED tag, or it could intercept and modify reads to the mode page and make it appear that the ORDERED tag is disabled.

Applications that do rely on ORDERED tags, like backup applications to tape drives, should not be used over multipathing software. This helps them determine if they are in the wrong environment.

3. In SBP-3, there is no task attribute sent with each command. A logical unit is set to treat all commands as SIMPLE or all commands as ORDERED. The proposed mode page lets SCSI software determine which policy is in place.

Suggested changes to SAM-3

5.9.5 Task attribute exception conditions

~~When an attempt to create a task with an invalid task attribute is detected, the command shall be terminated~~ If a device server receives a command with a task attribute that is not supported or is not valid (e.g., an ACA task attribute when an ACA condition does not exist), it shall terminate the command with CHECK CONDITION status, sense key of ILLEGAL REQUEST and additional sense code of INVALID MESSAGE ERROR. If a device server receives a command with a task attribute other than SIMPLE that is not enabled, it may terminate the command with CHECK CONDITION status, sense key of ILLEGAL REQUEST and additional sense code of INVALID MESSAGE ERROR. Device servers shall not reject tasks with SIMPLE task attributes. Task attribute support should be reported with the Extended Task Attribute subpage (see SPC-3).

NOTE 1 (really 12) - The use of the INVALID MESSAGE ERROR additional sense code is based on its similar usage in previous versions of this standard. The use of the INVALID MESSAGE ERROR additional sense code is not to be interpreted as a description of how the task attributes are represented by any given SCSI transport protocol.

8.2 Implicit head of queue

A command standard (see 3.1.15) may define tasks that may be processed by the task manager as if the task's task attribute is HEAD OF QUEUE without regard to the actual task attribute received with the task.

8.3 Controlling task set management

The Control mode page (see SPC-2) contains fields that specify particular task set management behaviors. The standard INQUIRY data CmdQue bit (see SPC-2) indicates support for tagged tasks (command queuing). One specific combination of task set management behaviors is identified as the basic task management model. Support for the basic task management model is indicated by values returned in the CMDQUE and BQUE bits in the standard INQUIRY data (see SPC-2). The basic task management model requires the following task set management behaviors:

Editor's Note 1: above paragraph is reworded by 03-271 but its changes are not listed. The only change in this section requested by this proposal is as follows:

---

a) ~~The only task attribute supported shall be SIMPLE~~ only the SIMPLE task attribute is supported, or only the ORDERED task attribute is supported;

...

8.6 Task attributes

8.6.0 Task attribute overview

Each task shall have one of the task attributes listed in table 1. Transport protocols should provide the capability to specify a unique task attribute for each task.

Table 1 — Task attributes [new table]

Task attribute	Description	Reference
SIMPLE	Task may be processed in any order.	8.6.1
ORDERED	Task is processed in the order received.	8.6.2
HEAD OF QUEUE	Task is processed next.	8.6.3
ACA	Task allowed during an ACA condition.	8.6.4

8.6.1 Simple task

A task having the SIMPLE attribute shall be accepted into the task set in the dormant task state. The task shall not enter the enabled task state until all older head of queue and older ordered tasks in the task set have ended (see 8.4).

The QUEUE ALGORITHM MODIFIER field in the Control mode page (see SPC-2) provides additional constraints on task completion order for tasks with the SIMPLE attribute.

8.6.2 Ordered task

A task having the ORDERED attribute shall be accepted into the task set in the dormant task state. The task shall not enter the enabled task state until all older tasks in the task set have ended (see 8.4).

8.6.3 Head of queue task

A task having the HEAD OF QUEUE attribute shall be accepted into the task set in the enabled task state.

8.6.4 ACA task

A task having the ACA attribute shall be accepted into the task set in the enabled task state. There shall be no more than one ACA task per task set (see 5.9.2.1).

Suggested changes to SPC-3

Possible approaches:

a) create a new Extended Task Attributes mode page
b) create a new REPORT SUPPORTED TASK ATTRIBUTES command and a matching SET SUPPORTED TASK ATTRIBUTES command
c) make the Control mode page longer

Only the first option is proposed.

7.4.8 Extended mode page

7.4.8.1 Extended mode page overview

The Extended mode page (see table 231) provides a means to specify subpages that are defined for all device types. Subpage code 00h is reserved, therefore all Extended mode pages use the sub_page format.

Table 2 — Extended mode page

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	SPF (1b)	PAGE CODE (15h)					
1	SUBPAGE CODE (01h)							
2	(MSB)	PAGE LENGTH (n - 3)						
3								(LSB)
4	Subpage specific mode parameters							
n								

7.4.8.2 Extended Task Attributes subpage [new]

The Extended Control subpage (see table 231) provides controls over SCSI task attributes (see SAM-3). The mode page policy for this subpage should be per-initiator or per-I_T nexus.

Table 3 — Extended Task Attributes subpage

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	SPF (1b)	PAGE CODE (15h)					
1	SUBPAGE CODE (01h)							
2	(MSB)	PAGE LENGTH (12)						
3								(LSB)
4	Reserved				ACASUP	HEADSUP	ORDSUP	SIMPSUP
5	Reserved				ACAEN	HEADEN	ORDEN	SIMPEN
6	Reserved							
15								

An ACA supported (ACASUP) bit set to one indicates the ACA task attribute is supported by the logical unit. An ACASUP bit set to zero indicates the ACA task attribute is not supported. The ACASUP bit shall not be changeable.

A HEAD OF QUEUE supported (HEADSUP) bit set to one indicates the HEAD OF QUEUE task attribute is supported by the logical unit. A HEADSUP bit set to zero indicates the HEAD OF QUEUE task attribute is not supported. The HEADSUP bit shall not be changeable.

An ORDERED supported (ORDSUP) bit set to one indicates the ORDERED task attribute is supported by the logical unit. An ORDSUP bit set to zero indicates the ORDERED task attribute is not supported. The ORDSUP bit shall not be changeable.

A SIMPLE supported (SIMPSUP) bit set to one indicates the SIMPLE task attribute is supported by the logical unit. An SIMPSUP bit set to zero indicates the SIMPLE task attribute is not supported. The SIMPSUP bit shall not be changeable.

An ACA enabled (ACAEN) bit set to one indicates the ACA task attribute is enabled. An ACAEN bit set to zero indicates the ACA task attribute is disabled.

A HEAD OF QUEUE enabled (HEADEN) bit set to one indicates the HEAD OF QUEUE task attribute is enabled. A HEADEN bit set to zero indicates the HEAD OF QUEUE task attribute is disabled.

An ORDERED enabled (ORDEN) bit set to one indicates the ORDERED task attribute is enabled. An ORDEN bit set to zero indicates the ORDERED task attribute is disabled.

A SIMPLE enabled (SIMPEN) bit set to one indicates the SIMPLE task attribute is enabled. An SIMPEN bit set to zero indicates the SIMPLE task attribute is disabled. The SIMPEN bit shall not be changeable.

NOTE 2 Multipathing application software may use the enable bits to disable ORDERED task attributes when accessing a logical unit through more than one target port group in an active asymmetric access state (see 5.8).

If an application client attempts to set an enable bit to one when the corresponding supported bit is set to zero, the device server shall reject the MODE SELECT command with CHECK CONDITION status with a sense key of ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN PARAMETER LIST.

Application clients should not use a task attribute unless it is reported as supported and enabled.

NOTE 3 The logical unit rejects tasks with unsupported task attributes, and may reject tasks with disabled task attributes, with a CHECK CONDITION status, sense key of ILLEGAL REQUEST and additional sense code of INVALID MESSAGE ERROR (see SAM-3).