

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 7 November 2004
 Subject: 04-367r0 SAM-4 SAS-1.1 Task priority cleanup

Revision history

Revision 0 (7 November 2004) First revision

Related documents

04-099r1 SAS-1.1 FCP-3 SRP-2 Per-command Priority Tagging (George Penokie, IBM)
 sas1r06 - Serial Attached SCSI 1.1 revision 6
 sam3r14 - SCSI Architecture Model - 3 revision 13

Overview

04-099r1 added a TASK PRIORITY field to the COMMAND frame of SAS-1.1, but did not describe how to add the optional [Task Priority] argument to the protocol services. It also neglected to add the argument to the state machines.

If the TASK PRIORITY field is set to 0h in the frame, is a Task Priority argument of 0h passed, or is no argument passed? SAM-3 discusses a value of 0h, implying that it is intended to be or could be passed.

If the task attribute is not SIMPLE, is the Task Priority argument passed or not?

As proposed, 04-099r1 requests that a column be added to the TASK ATTRIBUTE field table, marking the task priority code as Reserved for all but the SIMPLE task attribute. That means SAS could use the field for something other than the Task Priority argument in protocol services in those cases, and would need to define the Task Priority argument like this: "If the task attribute is SIMPLE, then the Task Priority argument is from the TASK PRIORITY field. If the task attribute is not SIMPLE, then the Task Priority argument is not present." It would be simpler to always seed it from the TASK PRIORITY field and let SAM-4 define any interactions between it and the task attribute.

Editorial changes are proposed in SAM-4 to follow this convention:

- a) task priority: a 4-bit value that is optionally part of the I_T_L_Q nexus. From either the Task Priority argument, the SET PRIORITY command, or the Control Extension mode page.
 - A) optional value missing means task has vendor-specific priority
 - B) 0h means the task has vendor-specific priority
 - C) 1h means the task has highest priority
 - D) Fh means lowest priority.
- b) priority: a 4-bit value that may be assigned to an I_T_L nexus by a SET PRIORITY command or the Command Extension mode page
- c) Task Priority argument: optional argument in Send SCSI Command and SCSI Command Received that chooses the task priority.

Suggested changes to SAM-4

3.1.123 task priority: The relative scheduling importance of a task having a SIMPLE task attribute in relation to other tasks [having SIMPLE task attributes](#) already in the task set. See 8.7.

4.11 Tasks and task tags

A task is represented by an I_T_L_Q nexus (see 4.12) and is composed of:

- a) A definition of the work to be performed by the logical unit in the form of a command or a group of linked commands;
- b) A task attribute (see 8.6) that allows the application client to specify processing relationships between various tasks in the task set; and
- c) Optionally, a task priority (see 8.7).

8.7 Task priority

A ~~The~~ task priority ~~argument that if set to a value other than zero~~ specifies the relative scheduling importance of a task having a SIMPLE task attribute in relation to other tasks having SIMPLE task attributes already in the task set. If the task has a task attribute other than SIMPLE, the task priority is not used. Task priority is a 4-bit field (i.e., has values 0h through Fh). A task with either no task priority or a task priority set to 0h has a vendor-specific level of scheduling importance. A task with a task priority set to 1h has the highest scheduling importance, with increasing task priority values indicating lower decreasing scheduling importance. A task with a task priority set to Fh has the lowest scheduling importance.

If the ~~task priority~~ Task Priority argument is set to zero or is not contained within the SCSI Command Received SCSI transport protocol service indication (see 5.4.2) and a priority has been assigned to the I_T_L nexus, the device server shall use a priority assigned to the I_T_L nexus may be used that priority as the task priority. A priority may be assigned to an I_T_L nexus by a SET PRIORITY command (see SPC-3) or by the INITIAL PRIORITY field in the Control Extension mode page (see SPC-3). If no priority has been assigned to the I_T_L nexus using the SET PRIORITY command and the logical unit does not support the INITIAL PRIORITY field in the Control Extension mode page ~~the priority assigned to the task is vendor specific,~~ the device server shall set the task priority to 0h (i.e., vendor specific) or the task shall have no task priority.

A task manager may use task priority to determine an ordering to process tasks with the SIMPLE task attribute within the task set. A difference in task priority between tasks may not override other scheduling considerations (e.g., different times to access different logical block addresses) or vendor specific scheduling considerations. However, processing of a collection of tasks with different task priorities should cause the subset of tasks with the higher ~~priority~~ task priorities to return status sooner in aggregate than the same subset would if the same collection of tasks were submitted under the same conditions but with all task priorities being equal.

For a task that processes linked commands, the task priority shall be that specified for the first command in the series of linked commands. The task priority specified for the second and subsequent commands shall be ignored.

~~The size of the task priority argument shall be four bits.~~

Suggested changes to SAS-1.1 (as proposed to be modified by 04-099r1)

9.2.2.1 COMMAND information unit

...

The TASK PRIORITY field specifies the relative scheduling of the task containing this command in relation to other tasks already in the task set, if the tasks have SIMPLE task attributes (see SAM-3).

The TASK ATTRIBUTE field is defined in table 1.

Table 1 — TASK ATTRIBUTE field

Code	Task attribute	Task priority code	Description
000b	SIMPLE	Any	Requests that the task be managed according to the rules for a simple task attribute and for task priority (see SAM-3).
001b	HEAD OF QUEUE	Reserved	Requests that the task be managed according to the rules for a head of queue task attribute (see SAM-3).
010b	ORDERED	Reserved	Requests that the task be managed according to the rules for an ordered task attribute (see SAM-3).
011b	Reserved		
100b	ACA	Reserved	Requests that the task be managed according to the rules for an automatic contingent allegiance task attribute (see SAM-3).
101b-111b	Reserved		

9.2.6.2.2 ST_IFR (initiator frame router) state machine

...

If this state machine receives a Send SCSI Command transport protocol service request then this state machine shall send a Request (Send Command) message to the ST_ITS1:Initiator_Start state in an ST_ITS state machine that does not have an active task. A Request (Send Command) message shall include the following arguments:

- a) the connection rate;
- b) the initiator connection tag;
- c) the destination SAS address;
- d) the tag;
- e) the logical unit number;
- f) the enable first burst value;
- g) the task attribute;
- h) [the task priority](#);
- i) the additional CDB length;
- j) the CDB; and
- k) the additional CDB bytes.

...

9.2.6.2.3.4 ST_ISF3:Prepare_Command state**9.2.6.2.4.3.1 State description**

This state shall construct a COMMAND frame. This state shall include the following values received with the Request (Send Command) message in the frame:

- a) frame type;
- b) hashed destination SAS address;
- c) hashed source SAS address;
- d) logical unit number;
- e) tag;
- f) task attribute;
- g) [task priority](#);
- h) additional CDB length;
- i) CDB; and
- j) additional CDB bytes; and
- k) number of fill bytes.

...

9.2.6.3.2 ST_TFR (target frame router) state machine

...

If the frame type is COMMAND, the items checked in the frame are correct, and this state machine receives an ACK Transmitted confirmation, then this state machine shall send a SCSI Command Received transport protocol service indication to the SCSI application layer. The indication shall include:

- a) the source SAS address;
- b) the tag;
- c) the logical unit number;
- d) the task attribute;
- e) [the task priority](#);
- f) the CDB; and
- g) any additional CDB bytes.

...

10.2.1.1 SCSI transport protocol services overview

...

An application client requests the processing of a SCSI command by invoking SCSI transport protocol services, the collective operation of which is conceptually modeled in the following remote procedure call (see SAM-3):

Service response = Execute Command (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Data-In Buffer Size], [Data-Out Buffer], [Data-Out Buffer Size], [\[Task Priority\]](#)), OUT ([Data-In Buffer], [Sense Data], [Sense Data Length], Status))

10.2.1.2 Send SCSI Command transport protocol service

An application client uses the Send SCSI Command transport protocol service request to request that an SSP initiator port transmit a COMMAND frame.

Send SCSI Command (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Data-In Buffer Size], [Data-Out Buffer], [Data-Out Buffer Size], [\[Task Priority\]](#), [Command Reference Number], [First Burst Enabled]))

Table 2 shows how the arguments to the Send SCSI Command transport protocol service are used.

Table 2 — Send SCSI Command transport protocol service arguments

Argument	SAS SSP implementation
I_T_L_Q nexus	I_T_L_Q nexus, where: <ol style="list-style-type: none"> I_T used to select a connection; L used to set the LOGICAL UNIT NUMBER field in the COMMAND frame header; and Q used to set the TAG field in the COMMAND frame header.
CDB	Used to set the CDB field in the COMMAND frame.
Task Attribute	Used to set the TASK ATTRIBUTE field in the COMMAND frame.
[Data-In Buffer Size]	Maximum of 2^{32}
[Data-Out Buffer]	Internal to the SSP initiator port.
[Data-Out Buffer Size]	Maximum of 2^{32}
[Task Priority]	Used to set the TASK PRIORITY field in the COMMAND frame.
[First Burst Enabled]	Used to set the ENABLE FIRST BURST field in the COMMAND frame and to cause the SSP initiator port to transmit the number of bytes indicated by the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see 10.2.6.1.5) for the SCSI target port without waiting for an XFER_RDY frame.

10.2.1.3 SCSI Command Received transport protocol service

An SSP target port uses the SCSI Command Received transport protocol service indication to notify a device server that it has received a COMMAND frame.

SCSI Command Received (IN (I_T_L_Q Nexus, CDB, Task Attribute, [\[Task Priority\]](#), [Command Reference Number]))

Table 3 shows how the arguments to the SCSI Command Received transport protocol service are determined.

Table 3 — SCSI Command Received transport protocol service arguments

Argument	SAS SSP implementation
I_T_L_Q nexus	I_T_L_Q nexus, where: a) I_T indicated by the connection; b) L indicated by the LOGICAL UNIT NUMBER field in the COMMAND frame header; and c) Q indicated by the TAG field in the COMMAND frame header.
CDB	From the CDB field in the COMMAND frame.
Task Attribute	From the TASK ATTRIBUTE field in the COMMAND frame.
[Task Priority]	From the TASK PRIORITY field in the COMMAND frame.
[Command Reference Number]	Ignored
[First Burst Enabled]	Indicates that first burst data is being delivered based on the ENABLE FIRST BURST field in the COMMAND frame and the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see 10.2.6.1.5).

If an SSP target port calls SCSI Command Received () with a TAG already in use by another SCSI command (i.e., an overlapped command) or by a SCSI task management function, the device server responses are defined in SAM-3.