

Date: February 24, 2004

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli), David Chambliss (IBM Almaden Research Center)

Subject: Per-Command Priority Tagging

1 Overview

The following proposed wording represents changes to FCP-3, SAS 1.1, SAM-3, and SRP-2 to enable the transmission of priority information on a per-command basis.

It also request the removal of the priority field from any OSD command that contains one.

This proposal standardizes the interface by which device servers can offer differentiated quality of service to different applications associated with the same initiator. Examples of its use would include offering lower priority on IO associated with background destage writes within a storage controller or on IO associated with background applications, so that response time may be reduced for those IO operations that directly affect the responsiveness offered to the end user.

The method defined in this proposal to accomplish this involves changes to the protocol standards to accommodate an extension to the task attribute field to allow different priorities to be assigned to simple task attributes. There should be no changes required to SAM-3 as a simple task with a priority is still architecturally handled the same as a simple task.

2 FCP-3, SAS 1.1, and SRP-2 additions

The changes below will be applied to the FCP-3, SAS 1.1, and SRP-2 standards. Only the section numbers are different.

Table 1 — COMMAND information unit

Byte\Bit	7	6	5	4	3	2	1	0
0	(MSB)							
7	LOGICAL UNIT NUMBER (LSB)							
8	Reserved							
9	Reserved	PRIORITY				TASK ATTRIBUTE		
10	Reserved							
11	ADDITIONAL CDB LENGTH (n dwords)						Reserved	
12	CDB							
27								
28	ADDITIONAL CDB BYTES							
27+n×4								

[The PRIORITY field specifies the relative scheduling of this task in relation to other tasks already in the task set for processing by the device server \(see SAM-3\).](#)

The TASK ATTRIBUTE field is defined in table 2.

Table 2 — TASK ATTRIBUTE field

Task Attribute Code	Task attribute	Priority Code	Description
000b	SIMPLE	0h	Requests that the task be managed according to the rules for a simple task attribute (see SAM-3).
		1h-Fh	Requests that the task be managed according to the rules for a simple task attribute (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		

3 Additions to SPC

3.1 Reservations

3.1.1 Persistent Reservations overview

Table 3 — SPC commands that are allowed in the presence of various reservations (part 1 of 3)

Command	Addressed LU has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From I_T nexus not registered	
	Write Excl	Excl Access		Write Excl RR	Excl Access – RR
ACCESS CONTROL IN	Allowed	Allowed	Allowed	Allowed	Allowed
ACCESS CONTROL OUT	Allowed	Allowed	Allowed	Allowed	Allowed
CHANGE ALIASES	Conflict	Conflict	Allowed	Conflict	Conflict
EXTENDED COPY	Conflict	Conflict	Allowed	Conflict	Conflict
INQUIRY	Allowed	Allowed	Allowed	Allowed	Allowed
LOG SELECT	Conflict	Conflict	Allowed	Conflict	Conflict
Key: LU =Logical Unit, Excl =Exclusive, RR =Registrants Only or All Registrants, <> Not Equal					
^a Exceptions to the behavior of the RESERVE and RELEASE commands described in SPC-2 are defined in 5.6.2.					

Table 3 — SPC commands that are allowed in the presence of various reservations (part 2 of 3)

Command	Addressed LU has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From I_T nexus not registered	
	Write Excl	Excl Access		Write Excl RR	Excl Access – RR
LOG SENSE	Allowed	Allowed	Allowed	Allowed	Allowed
MODE SELECT(6)/ MODE SELECT(10)	Conflict	Conflict	Allowed	Conflict	Conflict
MODE SENSE(6)/ MODE SENSE(10)	Conflict	Conflict	Allowed	Conflict	Conflict
PERSISTENT RESERVE IN	Allowed	Allowed	Allowed	Allowed	Allowed
PERSISTENT RESERVE OUT	see table 34				
PREVENT ALLOW MEDIUM REMOVAL (Prevent=0)	Allowed	Allowed	Allowed	Allowed	Allowed
PREVENT ALLOW MEDIUM REMOVAL (Prevent<>0)	Conflict	Conflict	Allowed	Conflict	Conflict
READ ATTRIBUTE	Conflict	Conflict	Allowed	Conflict	Conflict
READ BUFFER	Conflict	Conflict	Allowed	Conflict	Conflict
READ MEDIA SERIAL NUMBER	Allowed	Allowed	Allowed	Allowed	Allowed
RECEIVE COPY RESULTS	Conflict	Conflict	Allowed	Conflict	Conflict
RECEIVE DIAGNOSTIC RESULTS	Conflict	Conflict	Allowed	Conflict	Conflict
RELEASE(6)/ RELEASE(10)	As defined in SPC-2 ^a				
REPORT ALIASES	Allowed	Allowed	Allowed	Allowed	Allowed
REPORT DEVICE IDENTIFIER	Allowed	Allowed	Allowed	Allowed	Allowed
REPORT I_T NEXUS PRIORITY	Allowed	Allowed	Allowed	Allowed	Allowed
REPORT LUNS	Allowed	Allowed	Allowed	Allowed	Allowed
REPORT SUPPORTED OPERATION CODES	Conflict	Conflict	Allowed	Conflict	Conflict
REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS	Conflict	Conflict	Allowed	Conflict	Conflict
REPORT TARGET PORT GROUPS	Allowed	Allowed	Allowed	Allowed	Allowed
REQUEST SENSE	Allowed	Allowed	Allowed	Allowed	Allowed
RESERVE(6)/ RESERVE(10)	As defined in SPC-2 ^a				
SEND DIAGNOSTIC	Conflict	Conflict	Allowed	Conflict	Conflict
SET DEVICE IDENTIFIER	Conflict	Conflict	Allowed	Conflict	Conflict
SET I_T NEXUS PRIORITY	Conflict	Conflict	Allowed	Conflict	Conflict
Key: LU =Logical Unit, Excl =Exclusive, RR =Registrants Only or All Registrants, <> Not Equal					
^a Exceptions to the behavior of the RESERVE and RELEASE commands described in SPC-2 are defined in 5.6.2.					

Table 3 — SPC commands that are allowed in the presence of various reservations (part 3 of 3)

Command	Addressed LU has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From I_T nexus not registered	
	Write Excl	Excl Access		Write Excl RR	Excl Access – RR
SET TARGET PORT GROUPS	Conflict	Conflict	Allowed	Conflict	Conflict
TEST UNIT READY	Allowed	Allowed	Allowed	Allowed	Allowed
WRITE ATTRIBUTE	Conflict	Conflict	Allowed	Conflict	Conflict
WRITE BUFFER	Conflict	Conflict	Allowed	Conflict	Conflict
Key: LU =Logical Unit, Excl =Exclusive, RR =Registrants Only or All Registrants, <> Not Equal					
^a Exceptions to the behavior of the RESERVE and RELEASE commands described in SPC-2 are defined in 5.6.2.					

3.2 REPORT I_T NEXUS PRIORITY command (new section)

The REPORT I_T NEXUS PRIORITY command (see table 4) requests the priority that has been assigned to the I_T nexus associated with the received task.

The REPORT I_T NEXUS PRIORITY command is a service action of the MAINTENANCE IN command. Additional MAINTENANCE IN service actions are defined in SCC-2 and in this standard. The MAINTENANCE IN service actions defined in SCC-2 apply only to logical units that return a device type of 0Ch or the sccs bit equal to one in their standard INQUIRY data (see 6.4.2).

Table 4 — REPORT I_T NEXUS PRIORITY command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (A3h)							
1	Reserved			SERVICE ACTION (xxh)				
2	REPORT I_T NEXUS PRIORITY		Reserved					
3	Reserved							
5								
6	(MSB)							
9	ALLOCATION LENGTH (4h or larger)							(LSB)
10	Reserved							
11	CONTROL							

The REPORT I_T NEXUS PRIORITY field (see table 5) specifies the I_T nexus and the location of the priority value to be reported.

Table 5 — REPORT I_T NEXUS PRIORITY field

Code	Description
00b	The priority for the I_T nexus associated with this command shall be reported and only the I_T nexus descriptor for this command returned (see table 6).
01b	The priority for all I_T nexus that have been assigned with a SET I_T NEXUS PRIORITY command shall be reported in the REPORT I_T NEXUS PRIORITY parameter data (see table 6).
10b	Reserved
11b	Reserved

The ALLOCATION LENGTH field indicates the number of bytes that have been allocated for the returned parameter data. An allocation length that is not sufficient to contain the entire parameter list shall not be considered an error. If the complete list is required, the application client should send a new REPORT i_t NEXUS PRIORITY command with allocation length large enough to contain the entire list.

The format of the parameter data returned by the REPORT I_T NEXUS PRIORITY command is shown in table 6.

Table 6 — REPORT I_T NEXUS PRIORITY parameter data

Bit Byte	7	6	5	4	3	2	1	0
0	I_T NEXUS PARAMETER DATA LENGTH (m - 4)							
3								
	I_T nexus descriptor list							
4	I_T nexus descriptor (first)							
n								
	⋮							
	I_T nexus descriptor (last)							
m								

The I_T NEXUS PARAMETER DATA LENGTH field specifies the number of bytes of parameter data that follow.

The I_T nexus descriptor format is defined in table 7.

Table 7 — I_T NEXUS descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	I_T NEXUS DESCRIPTOR DATA LENGTH (m - 2)							
1								
2	Reserved				CURRENT PRIORITY			
3	Reserved							
4	Identification descriptor							
n								
n+1	TransportID							
m								

The CURRENT PRIORITY field contains the priority (see SAM-3) assigned to the I_T nexus defined by this descriptor.

The I_T NEXUS DESCRIPTOR DATA LENGTH field specifies the number of bytes of descriptor data that follow.

The format of an Identification descriptor is specified in table 274xx.

The format of a TransportID is specified in 7.5.4.

3.3 SET I_T NEXUS PRIORITY command (new section)

The SET I_T NEXUS PRIORITY command (see table 8) requests that a priority be set to the specified value. The priority set by this command shall remain in effect until one of the following occurs:

- a) another SET I_T NEXUS PRIORITY command is received;
- b) a Control Extension mode page (0Ah/01h) is received with the INITIAL PRIORITY field set to a value the is different than the current INITIAL PRIORITY field (see
- c) hard reset event; or
- d) power off.

The priority set by this command shall not be effected by an I_T nexus loss.

The SET I_T NEXUS PRIORITY command is a service action of the MAINTENANCE OUT command. Additional MAINTENANCE OUT service actions are defined in SCC-2 and in this standard. The MAINTENANCE OUT service actions defined only in SCC-2 apply only to logical units that return a device type of 0Ch or the sccs bit equal to one in their standard INQUIRY data

Table 8 — SET I_T NEXUS PRIORITY command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (A4h)							
1	Reserved			SERVICE ACTION (xxh)				
2	I_T NEXUS TO SET		Reserved					
3	Reserved							
4	Reserved							
5	Reserved							
6	(MSB)							
7	PARAMETER LIST LENGTH							
8								
9								
10	Reserved							
11	CONTROL							

The I_T NEXUS TO SET field (see table 9) specifies the I_T nexus and the location of the priority value to be assigned to that I_T nexus.

Table 9 — I_T NEXUS TO SET field

Code	Description
00b	The priority for the I_T nexus associated with this command shall be set to the value contained in the SET PRIORITY field.
01b	The priority for the I_T nexus identified in the I_T nexus parameter list (see table 10) shall be set to the value contained in the SET PRIORITY field. The device server shall terminate the command with CHECK CONDITION status if the parameter list length results in the truncation of the I_T nexus parameter list. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code shall be set to PARAMETER LIST LENGTH ERROR. On successful completion of a SET I_T NEXUS PRIORITY command a unit attention shall be generated for the initiator port identified in the I_T nexus parameter list. When reporting the unit attention condition the additional sense code shall be set to PRIORITY CHANGED.
10b	The priority value specified in the INITIAL PRIORITY field of the Control Extension mode page (see 3.3.1) shall be used for all initiators regardless of any prior priority. On successful completion of a SET I_T NEXUS PRIORITY command a unit attention shall be generated for all other initiator ports. When reporting the unit attention condition the additional sense code shall be set to PRIORITY CHANGED.
11b	Reserved

Editor's Note 1: The PRIORITY CHANGED is a new ASCQ. I would suggest it be coded as 0Fh 12h.

The PARAMETER LIST LENGTH field specifies the length in bytes of the I_T nexus parameter list that shall be contained in the Data-Out Buffer. A parameter list length of zero indicates that the Data-Out Buffer shall be empty. This condition shall not be considered as an error.

Table 10 — SET I_T NEXUS PRIORITY commands I_T nexus parameter list

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1								
2	Reserved				SET PRIORITY			
3	Reserved							
4	Identification descriptor							
n								
n+1	TransportID							
m								

The SET PRIORITY field indicates the priority to be assigned. The value in the SET PRIORITY field shall be returned in subsequent REPORT I_T NEXUS PRIORITY commands.

The command shall be terminated with a CHECK CONDITION status and the sense key set to ILLEGAL REQUEST if the value in the I_T NEXUS PARAMETER DATA LENGTH field results in the truncation of the I_T nexus parameter list.

The format of an Identification descriptor is specified in table 274xx. If the I_T NEXUS TO SET field is set to zero this field shall be ignored.

The format of a TransportID is specified in 7.5.4. If the I_T NEXUS TO SET field is set to zero this field shall be ignored.

3.3.1 Control Extension mode page

The Control Extension mode page (see table 223) is a subpage of the Control mode page (see 7.4.6) provides controls over SCSI features that are applicable to all device types. The mode page policy (see 6.7) for this subpage shall be shared. If a field in this mode subpage 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.

Table 11 — Control Extension mode page

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (1b)	PAGE CODE (0Ah)					
1	SUBPAGE CODE (01h)							
2	(MSB)	PAGE LENGTH (1Ch)						
3	(LSB)							
4	Reserved							IALUAE
5	Reserved				INITIAL PRIORITY			
6	Reserved							
31								

The PS bit, SPF bit, PAGE CODE field, SUBPAGE CODE field, and PAGE LENGTH field are described in 7.4.5.

An implicit asymmetric logical unit access enabled (IALUAE) bit set to one specifies that implicit asymmetric logical unit access state changes (see 5.8.6) are allowed. An IALUAE bit set to zero specifies that implicit asymmetric logical unit access state changes be disallowed and indicates that implicit asymmetric logical unit access state changes are disallowed or not supported.

The INITIAL PRIORITY field specifies the priority to be used on commands received in any I_T nexus where the priority has not been modified by a SET I_T NEXUS PRIORITY command (see x.x.x) since the last power on or hard reset. If a MODE SELECT command specifies an initial priority value that is different than the current initial priority then the device server shall set all the priorities to the new priority and shall generate an unit attention for all initiator ports. When reporting the unit attention condition the additional sense code shall be set to PRIORITY CHANGED.

4 Additions to SAM-3

4.1 The Execute Command procedure call

An application client requests the processing of a SCSI command by invoking the SCSI transport protocol services described in 4.4, the collective operation of which is conceptually modeled in the following procedure call:

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

Input Arguments:

I_T_L_Q Nexus: The I_T_L_Q nexus identifying the task (see 4.12).

CDB: Command descriptor block (see 5.2).

Task Attribute: A value specifying one of the task attributes defined in 8.6. SCSI transport protocols may or may not provide the ability to specify a different task attribute for each task (see 8.6.1). For a task that processes linked commands, the Task Attribute shall be that specified for the first command in the sequence of linked commands. The Task Attribute specified for the second and subsequent commands shall be ignored.

Data-In Buffer Size: The number of bytes available for data transfers to the Data-In Buffer (see 5.4.3).

Data-Out Buffer: A buffer containing command specific information to be sent to the logical unit, such as data or parameter lists needed to process the command. The buffer size is indicated by the Data-Out Buffer Size argument. The content of the Data-Out Buffer shall not change during the lifetime of the command (see 5.5) as viewed by the application client.

Data-Out Buffer Size: The number of bytes available for data transfers from the Data-Out Buffer (see 5.4.3).

Command Reference Number (CRN): When this argument is used, all sequential commands of an I_T_L nexus shall include a CRN argument that is incremented by one. The CRN shall be set to one for each I_T_L nexus involving the SCSI port after the SCSI port receives a hard reset or detects I_T nexus loss. The CRN shall be set to one after it reaches the maximum CRN value supported by the protocol. The CRN value zero shall be reserved for use as defined by the SCSI transport protocol. It is not an error for the application client to provide this argument when CRN is not supported by the SCSI transport protocol or logical unit.

Priority: [The priority assigned to the task. For specific requirements on the Priority argument see 5.2](#)

Output Arguments:

Data-In Buffer: A buffer to contain command specific information returned by the logical unit by the time of command completion. The **Execute Command** procedure call shall not return a status of GOOD, CONDITION MET, INTERMEDIATE, or INTERMEDIATE-CONDITION MET unless the buffer contents are valid. The application client shall not assume that the buffer contents are valid unless the command completes with a status of GOOD, CONDITION MET, INTERMEDIATE, or INTERMEDIATE-CONDITION MET. While some valid data may be present for other values of status, the application client should rely on additional information from the logical unit, such as sense data, to determine the state of the buffer contents. If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be undefined.

Sense Data: A buffer containing sense data returned in the same I_T_L_Q nexus transaction (see 3.1.46) as a CHECK CONDITION status (see 5.9.6). The buffer length is indicated by the Sense Data Length argument. If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be undefined.

Sense Data Length: The length in bytes of the Sense Data.

Status: A one-byte field containing command completion status (see 5.3). If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be undefined.

Service Response assumes one of the following values:

- TASK COMPLETE:** A logical unit response indicating that the task has ended. The Status argument shall have one of the values specified in 5.3 other than INTERMEDIATE or INTERMEDIATE-CONDITION MET.
- LINKED COMMAND COMPLETE:** Logical unit responses indicating that the task has not ended and that a linked command has completed successfully. As specified in 5.3, the Status argument shall have a value of INTERMEDIATE or INTERMEDIATE-CONDITION MET.
- SERVICE DELIVERY OR TARGET FAILURE:** The command has been ended due to a service delivery failure (see 3.1.113) or SCSI target device malfunction. All output parameters are invalid.

4.2 Command descriptor block (CDB)

4.3 Status

4.4 SCSI transport protocol services in support of Execute Command

4.4.1 Overview

The SCSI transport protocol services that support the **Execute Command** procedure call are described in 4.4. Two groups of SCSI transport protocol services are described. The SCSI transport protocol services that support the request and confirmation for the **Execute Command** procedure call are described in 4.4.2. The SCSI transport protocol services that support the data transfers associated with processing a SCSI command are described in 5.4.3.

4.4.2 Execute Command request/confirmation SCSI transport protocol services

All SCSI transport protocol standards shall define the SCSI transport protocol specific requirements for implementing the **Send SCSI Command** SCSI transport protocol service request and the **Command Complete Received** confirmation. Support for the **SCSI Command Received** indication and **Send Command Complete** response by a SCSI transport protocol standard is optional. All SCSI I/O systems shall implement these SCSI transport protocols as defined in the applicable SCSI transport protocol specification.

SCSI Transport Protocol Service Request:

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

Input Arguments:

I_T_L_Q Nexus: The I_T_L_Q nexus identifying the task (see 4.12).

CDB: Command descriptor block (see 5.2).

Task Attribute: A value specifying one of the task attributes defined in 8.6. For specific requirements on the Task Attribute argument see 4.1.

Data-In Buffer Size: The number of bytes available for data transfers to the Data-In Buffer (see 5.4.3).

Data-Out Buffer: A buffer containing command specific information to be sent to the logical unit, such as data or parameter lists needed to process the command (see 4.1). The content of the Data-Out Buffer shall not change during the lifetime of the command (see 5.5) as viewed by the application client.

Data-Out Buffer Size: The number of bytes available for data transfers from the Data-Out Buffer (see 5.4.3).

Command Reference Number (CRN): When this argument is used, all sequential commands of an I_T_L nexus shall include a CRN argument that is incremented by one (see 4.1).

Priority: [The priority assigned to the task. For specific requirements on the Priority argument see 5.2](#)

First Burst Enabled: An argument specifying that a SCSI transport protocol specific number of bytes from the Data-Out Buffer shall be delivered to the logical unit without waiting for the device server to invoke the **Receive Data-Out** SCSI transport protocol service.

SCSI Transport Protocol Service Indication:

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

Input Arguments:

I_T_L_Q Nexus: The I_T_L_Q nexus identifying the task (see 4.12).

CDB: Command descriptor block (see 5.2).

Task Attribute: A value specifying one of the task attributes defined in 8.6. For specific requirements on the Task Attribute argument see 4.1.

Command Reference Number (CRN): When this argument is used, all sequential commands of an I_T_L nexus shall include a CRN argument that is incremented by one (see 4.1).

Priority: [The priority assigned to the task. For specific requirements on the Priority argument see 5.2](#)

First Burst Enabled: An argument specifying that a SCSI transport protocol specific number of bytes from the Data-Out Buffer are being delivered to the logical unit without waiting for the device server to invoke the **Receive Data-Out** SCSI transport protocol service.

5 Task Set Management

5.1 Introduction to task set management

Clause 5 describes some of the controls application clients have over task set management behaviors (see 8.3). Clause 5 also specifies task set management requirements in terms of:

- a) [Priority \(see 5.2\)](#)
- b) Task states (see 8.5);
- a) Task attributes (see 8.6);
- b) The events that cause transitions between task states (see 8.4 and 8.5); and
- c) A map of task state transitions (see 8.7).

Clause 5 concludes with several task set management examples (see 8.8).

Task behavior, as specified in clause 5, refers to the functioning of a task as observed by an application client, including the results of command processing and interactions with other tasks.

The requirements for task set management only apply to a task after it has been entered into a task set. A task shall be entered into a task set unless:

- a) A condition exists that causes that task to be completed with a status of BUSY, RESERVATION CONFLICT, TASK SET FULL, or ACA ACTIVE;
- b) Detection of an overlapped command (see 5.9.3) causes that task to be completed with a CHECK CONDITION status; or
- c) SCSI transport protocol specific errors cause that task to be completed with a status other than GOOD.

5.2 Priority

The priority 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 already in the task set. Priority 1h is the highest priority, with increasing priority values indicating lower scheduling importance.

If the priority is set to zero or not used a priority code assigned to the I T nexus may be used by the task manager to determine an ordering to process tasks with the SIMPLE task attribute in addition to its vendor specific ordering rules. A priority may be assigned to an I T nexus by a SET I T NEXUS 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 nexus using the SET I T NEXUS PRIORITY command (see SPC-3) and the logical unit does not support the INITIAL PRIORITY field of the Control Extension mode page the priority assigned to the task is vendor specific.

A difference in priority between tasks does not necessarily override other scheduling considerations (e.g., different times to access different logical block addresses). However, processing of a collection of tasks with different priorities should cause the subset of tasks with the higher priority 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 priorities equal.

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