

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 17 March 2007
 Subject: 07-144r0 FCP-4 QUERY UNIT ATTENTION task management function

Revision history

Revision 0 (17 March 2007) First revision, split off from 07-067r0 after the March 2007 FCP WG.

Related documents

sam4r08 - SCSI Architecture Model - 4 (SAM-4) revision 8
 sas2r08 - Serial Attached SCSI - 2 (SAS-2) revision 8
 fcp4r00 - Fibre Channel Protocol - 4 (FCP-4) revision 0
 07-066 - SAM-4 SAS-2 QUERY TASK SET task management function (Rob Elliott, HP)
 07-067 - SAM-4 SAS-2 QUERY UNIT ATTENTION task management function (Rob Elliott, HP)
 07-072 - FCP-4 QUERY TASK task management function (Rob Elliott, HP)
 07-143 - FCP-4 QUERY TASK SET task management function (Rob Elliott, HP)

Overview

After a logical unit establishes a unit attention condition, commands that report the unit attention condition like TEST UNIT READY also clear the unit attention condition (unless the unit attention interlock feature is enabled). It is desirable to have a way to determine if a unit attention condition exists without clearing it. Although a command could be defined to do this, a task management function is a better fit. A new QUERY UNIT ATTENTION task management function is proposed, returning:

- a) FUNCTION SUCCEEDED if there is a unit attention condition pending; and
- b) FUNCTION COMPLETED if there is not a unit attention condition pending.

QUERY UNIT ATTENTION is useful for layered software stacks in initiators when the lower layer needs to determine if a unit attention condition exists (e.g., in SAS, after receiving a Broadcast (Asynchronous Event) or Broadcast (SES)). If the lower layer sends a full-fledged SCSI command which draws out the unit attention condition, the automatic clearing nature prevents the upper layer software from learning about the unit attention condition. There is no way for the lower layer to pass along that knowledge - each command that the upper layer has outstanding will receive its own status, and the unit attention cannot be returned on top of another status value.

Since SAS-2 and FCP-4 both have 3 reserved bytes in their RESPONSE frames, the additional sense code (ASC/ASCQ) of the highest priority unit attention condition is also returned along with the response code. This consumes 2 bytes, leaving one more still reserved.

Changes are proposed for FCP-4. SAM-4 and SAS-2 changes are proposed in 07-067.

Suggested changes to FCP-4 (in places, the base text assumes 07-143r1 has been incorporated)

4.9 Task management

An application client requests a task management function to control explicitly the processing of one or more FCP I/O operations (see 9.2.2.5).

The ABORT TASK task management function is mapped to the FC-FS-2 ABTS basic link service while the other task management functions are mapped into control bits (see table 20) in the FCP_CMND IU. Task management functions that use the FCP_CMND IU are transmitted as the first IU in a new Exchange. A task management function that uses the FCP_CMND IU ends with an FCP_RSP IU that indicates the completion status of the function. If the addressed logical unit is not supported or is not available (e.g., not connected or not configured) the FCP_CMND IU:

- a) should end with an FCP_RSP IU completion status of 09h (i.e., Task Management function incorrect logical unit number) (see table 24); and
- b) may end with an FCP_RSP IU completion status of 00h (i.e., Task Management function complete) (see table 24).

The FCP_CDB field in FCP_CMND IUs that perform task management functions is ignored.

The QUERY TASK task management function is not supported.

The task management function mappings are specified in table 3.

Table 3 — Task management functions, SAM-3 to FCP-4

SAM-3 function	FCP-4 equivalent
ABORT TASK	FCP recovery abort ^a
ABORT TASK SET	FCP_CMND ABORT TASK SET
CLEAR TASK SET	FCP_CMND CLEAR TASK SET
CLEAR ACA	FCP_CMND CLEAR ACA
LOGICAL UNIT RESET	FCP_CMND LOGICAL UNIT RESET
QUERY UNIT ATTENTION	FCP_CMND QUERY UNIT ATTENTION
^a FC-FS-2 basic link services are used to perform the ABORT TASK function.	

FC-FS-2 basic link services and FC-LS extended link services are used to perform the ABORT TASK task management function, to recover Exchange resources, and to re-establish other initial conditions.

The ABORT TASK task management function causes the device server to abort the specified task using the recovery abort protocol, if the task exists. The action is defined in SAM-3. The ABORT TASK task management function is performed by the initiator FCP_Port (i.e., Exchange Originator) using the recovery abort (see 12.3).

The specified Exchange shall be terminated by the initiator FCP_Port using the recovery abort. To be compliant with FC-FS-2, the ABORT TASK task management function may not immediately release all Exchange resources, since a Recovery_Qualifier may be established to allow for the management of information that may already have been delivered to the fabric.

In addition to recovering Exchange resources that may have been left unavailable while processing task management functions, recovery abort may be used to recover Exchange resources left in an undefined state by any of the task abort events defined in SAM-3 or by any similar events

9.2 FCP_CMND IU

9.2.1 Overview and format of FCP_CMND IU

9.2.2 FCP_CMND IU field descriptions

9.2.2.5 TASK MANAGEMENT FLAGS field

The TASK MANAGEMENT FLAGS field specifies the task management function to be performed, if any.

Task management functions shall be requested by the initiator FCP_Port (Exchange Originator) using a new Exchange. If the TASK MANAGEMENT FLAGS field is set to a nonzero value, the FCP_CDB field, the FCP_DL field, the TASK ATTRIBUTE field, the RDDATA bit, and the WRDATA bit shall be ignored and the FCP_BIDIRECTIONAL_READ_DL field shall not be included in the FCP_CMND IU payload. If the TASK MANAGEMENT FLAGS field is set to a reserved value, the target FCP_Port shall return an FCP_RSP IU containing the RSP_CODE field set to 02h (i.e., FCP_CMND fields invalid).

The clearing actions performed by task management functions are specified in table 5.

The format of the TASK MANAGEMENT FLAGS field is specified in table 20.

Table 20 — TASK MANAGEMENT FLAGS field

Code	Task management function ^a	Support
00h	None ^b	Mandatory
01h	QUERY TASK SET per 07-143	Optional
02h	ABORT TASK SET	Mandatory
04h	CLEAR TASK SET	Mandatory
08h	QUERY UNIT ATTENTION	Optional
10h	LOGICAL UNIT RESET	Mandatory
20h	Obsolete	
40h	CLEAR ACA	See ^c
80h	Obsolete	
All others	Reserved	
^a The ABORT TASK management function is specified in 4.9. ^b The FCP_CDB field is honored instead. ^c The CLEAR ACA task management function is mandatory in the Fibre Channel Protocol if the FCP device sets the NORMACA bit to one in the standard INQUIRY data (see SPC-4) and . It shall not be sent to a logical unit with a NORMACA bit equal to zero in the standard INQUIRY data.		

The QUERY TASK SET task management function is defined in SAM-4.

The ABORT TASK SET task management function is defined in SAM-4. The ABORT TASK SET task management function resets internal states of the target FCP_Port as shown in 4.10. Exchange resources may be cleared by a recovery abort sequence (see 12.3) generated by the initiator FCP_Port that sent the ABORT TASK SET task management function for each task known to the initiator FCP_Port.

The CLEAR TASK SET task management function is defined in SAM-4. The CLEAR TASK SET task management function resets internal states of the target FCP_Port as shown in 4.10. Exchange resources to be cleared may be cleared by one or more of the following mechanisms:

- a recovery abort sequence (see 12.3) may be generated by the initiator FCP_Port that sent the CLEAR TASK SET for each task known to that initiator FCP_Port;
- a task, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the CLEAR TASK SET is ended in the logical unit. The initiator FCP_Port for that task shall determine by a timeout that the task did not finish. Subsequent retries fail because the task resources have been cleared in the logical unit, so the initiator FCP_Port shall clear the Exchange resources with a recovery abort sequence. See 12.3; or
- a task for an initiator FCP_Port other than the initiator FCP_Port that sent the CLEAR TASK SET may be completed by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and the additional sense code set as specified in SAM-3.

NOTE 1 - SAM-3 has defined the TASK ABORTED status for tasks terminated by a CLEAR TASK SET task management function if the Control mode page indicates that the TASK ABORTED status is supported.

| [The QUERY UNIT ATTENTION task management function is defined in SAM-4.](#)

The LOGICAL UNIT RESET task management function is defined in SAM-4.

The LOGICAL UNIT RESET task management function resets the internal states of the target FCP_Port and logical unit as shown in 4.10. Exchange resources to be cleared may be cleared by the following mechanisms:

- a) a recovery abort sequence (see 12.3) may be generated by the initiator FCP_Port that sent the LOGICAL UNIT RESET task management function for each task in the logical unit known to that initiator FCP_Port;
- b) a task, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET task management function is ended in the logical unit. The initiator FCP_Port for that task shall determine by a timeout that the task did not finish. Subsequent retries fail because the task resources have been cleared in the logical unit, so the initiator FCP_Port shall clear the Exchange resources with a recovery abort sequence. See 12.3; or
- c) a task for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET task management function may be completed by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and the additional sense code set as specified in SAM-3.

NOTE 3 - SAM-3 has defined the TASK ABORTED status for tasks terminated by a LOGICAL UNIT RESET task management function if the Control mode page (see SPC-4) indicates that the TASK ABORTED status is supported.

The **CLEAR ACA** task management function is defined in SAM-4. When the task manager clears the ACA condition, any task within that task set may be completed subject to the rules for task management specified by SAM-3. If there is no ACA condition present, the CLEAR ACA task management function shall be accepted and the FCP_RSP IU shall contain a RSP_CODE field set to 00h (i.e., Task Management function complete).

The use of the ACA bit in the CDB control field and the implementation of ACA is described in SAM-3.

Depending on the mode page parameters that have been established (see SPC-3), additional FCP I/O operations may have to be aborted by the recovery abort as part of the process of clearing the automatic contingent allegiance.

9.2.2.8 FCP_CDB field

The FCP_CDB field contains the CDB to be sent to the addressed logical unit. The maximum CDB length is 16 bytes unless the ADDITIONAL_FCP_CDB_LENGTH field has specified that there is an ADDITIONAL_FCP_CDB field. The FCP_CDB field shall be ignored if the task management flags field is set to a nonzero value.

The CDB format is defined by SAM-3 and the contents of the CDB are defined in the SCSI command standards.

Bytes between the end of a CDB and the end of the FCP_CDB field or, if applicable, the ADDITIONAL_FCP_CDB field shall be reserved.

9.2.2.9 ADDITIONAL_FCP_CDB field

The ADDITIONAL_FCP_CDB field contains any CDB bytes beyond those contained within the 16 byte FCP_CDB field.

The ADDITIONAL_FCP_CDB field shall not be present if the task management flags field is set to a nonzero value. The contents of the field shall be those bytes of an extended CDB beyond the first 16 bytes of the CDB as defined in the SCSI command standards.

9.5 FCP_RSP IU

9.5.16 FCP_RSP_INFO field

The FCP_RSP_INFO field contains information describing only the protocol failures detected during the processing of an FCP I/O operation. If none of the specified protocol failures have occurred, the FCP_RSP_INFO field shall not be included in the FCP_RSP IU and the FCP_RSP_LEN_VALID bit shall be zero. The FCP_RSP_INFO does not contain link error information, since FC-FS-2 provides the mechanisms for presenting such errors. The FCP_RSP_INFO field does not contain SCSI logical unit error information, since that is contained in the FCP_SNS_INFO field as described in 9.5.17. The FCP_RSP_INFO field shall contain valid

information if the target FCP_Port detects any of the conditions indicated by an FCP FCP_RSP_CODE. The format of the FCP_RSP_INFO field is specified in table 23.

Table 23 — FCP_RSP_INFO field format

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved							
2	ADDITIONAL_RSP_INFO							
3	RSP_CODE							
4	Reserved (if any)							
7	Reserved							

The ADDITIONAL_RSP_INFO field contains additional response information for certain task management functions (e.g., QUERY UNIT ATTENTION) as defined in SAM-4. If the task management function does not define additional response information or the logical unit does not support additional response information, the target FCP_Port shall set the ADDITIONAL_RSP_INFO field to 000000h.

The RSP_CODE field is defined in table 24.

Table 24 — RSP_CODE field

Code	Description
00h	Task Management function complete
01h	FCP_DATA length different than FCP_BURST_LEN
02h	FCP_CMND fields invalid
03h	FCP_DATA parameter mismatch with FCP_DATA_RO
04h ^a	Task Management function rejected
05h	Task Management function failed
08h ^a	Task Management function succeeded
09h ^a	Task Management function incorrect logical unit number
06h - 07h 0Ah - FFh	Reserved
^a Only valid when responding to task management functions	

The completion status of the task management function is indicated by the RSP_CODE field. If the Exchange is aborted before the FCP_RSP IU is returned, the completion status is unknown. If the RSP_CODE field is set to 05h (i.e., Task Management function failed), the state of the logical unit is unknown.

Activities started by a task management function may continue after the FCP_RSP IU for the task management has been delivered.