

To: T10 Technical Committee  
 From: Rob Elliott, HP (elliott@hp.com)  
 Date: 25 February 2004  
 Subject: 04-072r0 SAM-3 SPC-3 Unit attention for status delivery problems

### Revision history

Revision 0 (25 February 2004) First revision

### Related documents

sam3r11 - SCSI Architecture Model - 3 revision 11  
 spc3r17 - SCSI Primary Commands - 3 revision 17  
 sbc2r11 - SCSI Block Commands - 2 revision 11  
 ssc2r09 - SCSI Stream Commands - 2 revision 9  
 ses2r06 - SCSI Enclosure Services - 2 revision 6  
 sas1r03 - Serial Attached SCSI 1.1 revision 3

### Overview

If a device server is unable to deliver the Status byte to the initiator port (e.g. the Status byte keeps getting parity errors in SPI, the RESPONSE frames get repeated CRC errors in SAS, or the Resource Recovery Timer kicks in in FCP), it eventually just terminates the command and proceeds as if nothing had happened. The fact that an error occurred just disappears. There is no standard mechanism to report a unit attention or log the error.

A similar problem existed a few years ago: an initiator port did not know that its tasks had been aborted by another initiator port using LOGICAL UNIT RESET or CLEAR TASK SET. That was solved by the TASK ABORTED status, added late in SAM-2. It's possible that the TASK ABORTED gets lost, however.

If any status cannot be delivered, the device server should be allowed to create a unit attention and log information about the command that failed. If that initiator successfully communicates later, it will find out about its lost command. If the initiator port has broken, other initiator ports (e.g. in the cluster) can determine the status of its commands.

A new additional sense code is proposed named COMMAND COMPLETE DELIVERY PROBLEM. It may be sent by the device server if it decides it cannot deliver status or sense data. It's still protocol- and vendor-specific how many retry attempts are done and/or what timers expire to prompt this decision.

The log page records:

- a) I\_T\_L\_Q information
  - A) I: initiator port using the TransportID format (variable length)
  - B) T: target port using the relative target port identifier format (2 bytes)
  - C) L: implied by which logical unit is recording the log page and setting the unit attention
  - D) Q: task tag (protocol-specific width; allow for 64 bits)
- b) service response that was not delivered (encoded in one byte)
- c) status that was not delivered (one byte per SAM-3)
- d) sense data that was not delivered (variable length per SPC-3)

The log page is structured like the Last n Deferred Errors log page. It is vendor-specific how many errors the logical unit records (if any at all). It is also be up to the vendor to choose how many fields to fill in - if the vendor only wants to log the tag, that is better than nothing.

---

Editor's Note 1: the same general issue exists for problems delivering Task Management Function Executed results (e.g. RESPONSE frames for task management functions). This is not as important as the command issue, since there are no ORDERED attributes for task management functions and any software concerned about ordering has to limit itself to one at a time. Thus, they are not covered by this proposal. If wanted, the log page could be expanded to return the "association" (if any) in the task tag field, a larger set of service responses ( FUNCTION COMPLETE, FUNCTION SUCCEEDED, FUNCTION REJECTED, INCORRECT LOGICAL UNIT

NUMBER, and SERVICE DELIVERY OR TARGET FAILURE), no status, and no sense data.

---

## **Suggested changes to SAM-3**

### **5.3 Status**

#### **5.3.1 Status codes**

The status codes are specified in table 21. Status shall be sent from the device server to the application client whenever a command ends with a service response of TASK COMPLETE or LINKED COMMAND COMPLETE..

...

### **5.4 SCSI transport protocol services in support of Execute Command**

#### **5.4.1 Overview**

The SCSI transport protocol services that support the Execute Command procedure call are described in 5.4.

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 5.4.2. The SCSI transport protocol services that support the data transfers associated with processing a SCSI command are described in 5.4.3.

#### **5.4.2 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 Response (from device server):

Send Command Complete (IN ( I\_T\_L\_Q Nexus, [Sense Data], [Sense Data Length], Status, Service Response, [Delivery Result](#) ))

Input Arguments:

I\_T\_L\_Q Nexus: The I\_T\_L\_Q nexus identifying the task (see 4.12).

Sense Data: If present, this argument instructs the SCSI target port to return sense data to the SCSI initiator port (see 5.9.6).

Sense Data Length: The length in bytes of the sense data to be returned to the SCSI initiator port.

Status: Command completion status (see 5.1).

Service Response: Possible service response information for the command (see 5.1).

Output Arguments:

Delivery Result: an encoded value representing one of the following:

DELIVERY SUCCESSFUL: The command complete information was delivered successfully.

DELIVERY FAILURE: An unrecoverable service delivery subsystem error occurred while attempting to deliver the command complete information.

If the target port returns a Delivery Result set to DELIVERY FAILURE, the device server:

- a) shall terminate the command; and
- b) may establish a unit attention condition for all initiator ports with an additional sense code set to COMMAND COMPLETE DELIVERY PROBLEM and log the error in the Command Complete Delivery Problems log page (see SPC-3).

---



---

Editor's Note 2: a unit attention to this initiator port is helpful if it ever successfully communicates again. Is a unit attention for the others a good idea or unnecessary distraction? The log page is shared by all initiator ports, so they'll all see reports of the problems.

---



---

...

## 5.7 Aborting tasks

### 5.7.1 Mechanisms that cause tasks to be aborted

A task is aborted when an event or SCSI initiator device action causes termination of the task prior to its successful completion.

The following events cause a task or several tasks to be aborted:

- a) The return of an Execute Command service response of SERVICE DELIVERY OR TARGET FAILURE as described in 5.1;
- b) A logical unit reset (see 6.3.3);
- c) A hard reset (see 6.3.2);
- d) A power on condition;
- e) [Failure to deliver command complete information \(see 5.4.2\)](#); or
- f) SCSI transport protocol specific events.

...

## Suggested changes to SPC-3

### 7.2 Log parameters

#### 7.2.1 Log page structure and page codes for all device types

...

The page code assignments for the log pages are listed in table 176.

**Table 1 — Log page codes**

Value	Description
...	...
<a href="#">15h</a>	<a href="#">Command Complete Delivery Problems</a>
...	...
<a href="#">11h - 14h</a>	<a href="#">Reserved</a>
<del>14h</del> <a href="#">16h</a> - 17h	Reserved

---



---

Editor's Note 3: Also add to Annex E for device types D T M E V F

---



---

#### 7.2.x Command Complete Delivery Problems log page

The Command Complete Delivery Problems log page is used to report errors that have occurred when returning command complete information (e.g., SCSI status and sense data) to a SCSI initiator port. These errors may generate unit attention conditions with an additional sense code of COMMAND COMPLETE DELIVERY PROBLEM (see SAM-3).

The log page format is defined in 7.2.1.

Table 2 defines the format for a Command Complete Delivery Problems log parameter.

**Table 2 — Command Complete Delivery Problems log parameter format**

Byte\Bit	7	6	5	4	3	2	1	0	
0	(MSB)	PARAMETER CODE							(LSB)
1									
2		PARAMETER CONTROL							
3		PARAMETER LENGTH (z - 3)							
4		Reserved							
5									
6		UNDELIVERED SERVICE RESPONSE							
7		UNDELIVERED STATUS							
8	(MSB)	TASK TAG							(LSB)
15									
16	(MSB)	RELATIVE TARGET PORT IDENTIFIER							(LSB)
17									
18	(MSB)	INITIATOR PORT TRANSPORTID LENGTH (x - 19)							(LSB)
19									
20		INITIATOR PORT TRANSPORTID							
x									
x + 1	(MSB)	UNDELIVERED SENSE DATA LENGTH (z - y)							(LSB)
y									
y + 1		UNDELIVERED SENSE DATA, if any							
z									

The PARAMETER CODE field indicates the relative time at which the error occurred. A higher parameter code indicates that the error event occurred later in time.

When the last supported parameter code is used by a Command Complete Delivery Problem log parameter, the recording on this log page of all subsequent error information shall cease until one or more of the list parameters with the highest parameter codes have been reinitialized. If the RLEC bit of the Control mode page (see 7.4.6) is set to one, the device server shall return CHECK CONDITION status with the sense key set to RECOVERED ERROR and the additional sense code set to LOG LIST CODES EXHAUSTED.

The PARAMETER CONTROL byte is defined in 7.2.1.

The PARAMETER LENGTH field is set to the length of the log parameter minus three.

The UNDELIVERED SERVICE RESPONSE field contains the service response (see SAM-3 and table 3) which could not be delivered.

**Table 3 — Service response**

Value	Description
00h	TASK COMPLETE
01h	LINKED COMMAND COMPLETE
02h	SERVICE DELIVERY OR TARGET FAILURE
All others	Reserved

The UNDELIVERED STATUS field contains the SCSI status (see SAM-3) which could not be delivered.

The TASK TAG field contains the task tag (see SAM-3) of the command for which command complete information could not be delivered (i.e., the Q of the I\_T\_L\_Q nexus). Task tag definitions are protocol-specific.

The RELATIVE TARGET PORT IDENTIFIER field contains the relative port identifier (see SAM-3) of the target port through which command complete information could not be delivered (i.e., the T of the I\_T\_L\_Q nexus).

The INITIATOR PORT TRANSPORTID LENGTH field contains the length of the INITIATOR PORT TRANSPORTID field.

The INITIATOR PORT TRANSPORTID field contains the TransportID (see 7.5.4) of the initiator port to which the command complete information could not be delivered (i.e., the I of the I\_T\_L\_Q nexus).

The UNDELIVERED SENSE DATA LENGTH field contains the length of the UNDELIVERED SENSE DATA field.

The UNDELIVERED SENSE DATA field contains the sense data (see 4.5), if any, which could not be delivered.

---



---

**Editor's Note 4:** If task management function reporting is wanted, the log page would return the association (if supported by the transport protocol) in the task tag field, return service responses of FUNCTION COMPLETE, FUNCTION SUCCEEDED, FUNCTION REJECTED, and INCORRECT LOGICAL UNIT NUMBER instead of TASK COMPLETE and LINKED COMMAND COMPLETE, and not use the status and sense data fields. All logical units would probably have to set a unit attention on failure of an I\_T based function (thankfully none exist right now).

---



---

### **Suggested changes to SBC-2, SSC-2, SES-2 [and maybe MMC-4, SMC-2, OSD, and ADC]**

Since the philosophy of T10 is that SPC-3 just provides a selection of log pages but cannot claim that any apply to all command sets, this log page has to be added to each command set standard to the list of supported log pages.

### **Suggested changes to SAS-1.1**

The SAS target transport layer state machine should report failure to send a RESPONSE frame to the SCSI application layer with a Delivery Result argument (matching the protocol service) so it can create a unit attention condition.

Since the transport layer is undergoing changes right now in 03-245 (Terminate Data Transfer protocol service), specific changes are not yet proposed. Affected sections will include:

- a) 9.2.6.3 ST\_T state machine
- b) 10.2.1.4 Send Command Complete transport protocol service, and
- c) 10.2.3 Device server error handling