To: T10 Technical Committee
From: Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
Date: 16 April 2002
Subject: T10/02-134r0 SAM-2 SPC-3 SBC-2 SSC-2 Clearing effects of I_T nexus loss

Revision History
Revision 0 (16 April 2002) first revision

Related Documents
sam2r23 – SCSI Architecture Model – 2 revision 23

Command sets:
spc3r05 – SCSI Primary Commands – 3 revision 5
ssc2r08 – SCSI Stream Commands – 2 revision 8
sbc2r05a – SCSI Block Commands – 2 revision 5a

Protocols:
srp-r10 – SCSI RDMA Protocol revision 10
draft-ietf-ips-iscsi-12 – Internet SCSI revision 12
spi4r09 - SCSI Parallel Interface - 4 revision 9

Proposals:
02-078r1 - Reservations & Nexus (Mallikarjun Chadalapaka and Randy Haagens)
02-121r0 - SRP WG Dallas 3/2002 minutes (Cris Simpson)
02-127r1 - LOGOUT signals: Model concerns (Cris Simpson)

Overview
Thanks to Mallikarjun Chadalapaka (HP), Randy Haagens (HP), Cris Simpson (Intel), George Penokie (IBM/Tivoli), Dave Peterson (Cisco), and Ralph Weber (ENDL) for assistance in developing this proposal.

The SCSI architecture model does not describe the effects of logouts, although they have been implemented by numerous protocols. Command sets often indicate an object is cleared based on a “reset”, but don’t mention logouts. Protocol standards are left with large tables listing behavior for commands; these tables are always incomplete and out of date.

This proposal:
a) Defines “I_T nexus loss” in SAM-2.
b) Changes current command set (SPC-3, SBC-2, and SSC-2) references to various forms of reset to just “logical unit reset” and/or “power on”.
c) Adds “I_T nexus loss” to SAM-2 and command sets alongside “logical unit reset” where appropriate.
d) Describes how “I_T nexus loss” should be defined for the iSCSI, SRP, and SPI-5 transport protocols.
e) Defines “channel loss” in SAM-2 (along with its limited effects).
f) Does not define how “channel loss” should be defined for the iSCSI, SRP, and SPI-5 transport protocols.
g) Cleans up relationships of transport attentions in SAM-2.

This table summarizes all the things affected by resets that might be affected by “I_T nexus loss.” Some of the items are biased towards clearing, while others are biased towards preserving. Rather than invert the sense to make them all clearing, the sense in the current standards remains.

Table 1. I_T nexus loss summary

What	Where	When (current phrase(s))	Proposed changes	FCP-2 clearing effects table
Abort background operations (e.g. format, self-test)	nowhere - SAM-2 5.5 should describe	not mentioned anywhere except by self-test log page in SPC-3 which say “reset” affects it	logical unit reset	not mentioned
Deferred errors cleared	nowhere - SAM-2 5.5 should describe	not mentioned	by logical unit resets and by I_T nexus losses of the initiator port that sent the command that spawned the background operation	cleared
Reset CRN value	SAM-2 5.1	The initial, wrap, and reset CRN values shall be one	CRN shall be set to one on logical unit reset or I_T nexus loss	set to one

What	**Where**	**When (current phrase(s))**	**Proposed changes**	**FCP-2 clearing effects table**
Sense data preserved until	SAM-2 5.8.4.1	it is transferred	add logical unit reset or I_T nexus loss	not mentioned
Generate a unit attention	SAM-2 5.8.5	logical unit has been reset or by a power on reset	logical unit reset or I_T nexus loss	not mentioned
Abort all tasks	SAM-2 5.8.7	logical unit reset	(in new I_T nexus loss section) abort for that initiator on I_T nexus loss	aborted
Clear CA or ACA	SAM-2 5.8.7	logical unit reset	(in new I_T nexus loss section) clear for that initiator on I_T nexus loss	cleared
Release all reservations	SAM-2 5.8.7	logical unit reset	don't mention in SAM-2	cleared
Return operating mode to initial conditions, including mode pages	SAM-2 5.8.7	logical unit reset	don't mention in SAM-2	mode pages cleared
Set unit attention	SAM-2 5.8.7	logical unit reset	in 5.8.7, specify LOGICAL UNIT RESET OCCURRED additional sense code. in new I_T nexus loss section, specify additional sense code of I_T NEXUS LOSS OCCURRED	not mentioned
Classic reservations cleared	SPC-3 5.5.2 (model section)	initiator failures and subsequent hard resets, some recovery actions (e.g. hard resets)	logical unit resets	cleared **[CONFLICT]**
Classic reservations cleared	SPC-3 7.26.2 (RESERVE section)	TARGET RESET task management function, LOGICAL UNIT RESET task management function, hard reset condition, power on cycle	logical unit reset	cleared **[CONFLICT]**

What	**Where**	**When (current phrase(s))**	**Proposed changes**	**FCP-2 clearing effects table**
Classic reservations - 3rd party reservation	SPC-3 7.26.3 (RESERVE section)	TARGET RESET task management function, hard reset condition, or power on cycle	logical unit reset	cleared ("for initiator port associated with the reservation" which probably means the reservation holder, not the one granted access) **[CONFLICT]**
Persistent reservations registration and reservation data preserved	SPC-3 5.5.3.1 and 5.5.3.2 (model section)	initiator failures which involve hard resets, TARGET RESET task management function or other global actions any reset	hard resets, logical unit resets or I_T nexus losses	not cleared
Asymmetric logical unit access state preserved	SPC-3 5.7.8	across any reset or across any power cycle	hard reset, logical unit reset or I_T nexus loss	not mentioned
Power condition state set to active	SPC-3 5.8 8.4.11	wakeup or hard reset	(no change)	not mentioned
Alias table associations cleared	SPC-3 7.2.1	any event that resets the logical unit and events designated by the protocol	logical unit reset or I_T nexus loss	not mentioned
INQUIRY data should be available without incurring media access delays	SPC-3 7.4.1	hard reset or power up condition	logical unit reset	not mentioned
Log pages set to saved/default values (shared pages)	SPC-3 7.6	power on or hard reset	logical unit reset or I_T nexus loss of all I_T nexuses	not mentioned
Log pages set to saved/default values (per-initiator port pages) (only the TapeAlert page in SSC-2)	SSC-2 4.2.16.2	on hard reset	logical unit reset or I_T nexus loss	not mentioned

What	Where	When (current phrase(s))	Proposed changes	FCP-2 clearing effects table
Mode pages revert to saved/default (shared pages) [SPC-3 doesn't split out shared/per-initiator yet, much less per-I_T nexus]	SPC-3 7.9.2 7.9.6 9.x (both MODE SENSE command and mode parameters sections)	power on or hard reset condition, power up condition or hard reset condition	logical unit reset or I_T_nexus loss of all I_T nexuses	cleared on LOGO/PLOGI. For both PRLI and PRLO: keep current values if any still logged in; "not specified" once every initiator logs out. **[CONFLICT]**
Mode pages revert to saved/ default (per-initiator port pages)	SPC-3 7.9.2 7.9.6	power on or hard reset condition, power up condition or hard reset condition	logical unit reset or I_T nexus loss of all I_T nexuses involving the initiator port	cleared on LOGO/PLOGI. cleared on PRLI; "not specified" on PRLO **[CONFLICT]**
Mode pages revert to saved/default (per-I_T nexus pages)	SPC-3 7.9.2 7.9.6	power on or hard reset condition, power up condition or hard reset condition	logical unit reset or I_T nexus loss	not mentioned **[CONFLICT]**
PRGENERATION value cleared	SPC-3 7.11.3	part of power on reset process	power on	not cleared
Prevention of medium removal	SPC-3	hard reset condition	logical unit reset	cleared **[CONFLICT]**
EXTENDED COPY COPY STATUS data discarded	SPC-3 7.16.2	hard reset condition	logical unit reset or I_T nexus loss	cleared
EXTENDED COPY RECEIVE DATA data discarded	SPC-3 7.16.3	hard reset condition	logical unit reset or I_T nexus loss	cleared
EXTENDED COPY FAILED SEGMENT DETAILS data discarded	SPC-3 7.16.5	hard reset condition	logical unit reset or I_T nexus loss	cleared
changeable device identifier persists through	SPC-3 7.21	resets, power cycles	logical unit resets, I_T nexus losses	not mentioned
REPORT LUNS data available without incurring media access delays	SPC-3 7.22	hard reset or power up condition	logical unit reset	not mentioned
UNIT ATTENTION definition	SPC-3 7.25	media changed or target has been reset	media changed, logical unit reset, or I_T nexus loss involving that initiator add I_T NEXUS LOSS OCCURRED additional sense code	not mentioned **[CONFLICT]**
Classic reservations [see above]	SPC-3 7.26.2 7.26.3			

What	Where	When (current phrase(s))	Proposed changes	FCP-2 clearing effects table
WRITE BUFFER download microcode Unable to accept due to device condition, so revert to vendor specific condition	SPC-3 7.33.x	power cycle or reset	logical unit reset	not mentioned
WRITE BUFFER download microcode all microcode not written so discard	SPC-3 7.33.x	before reset or power on cycle	logical unit reset or I_T nexus loss	not mentioned **[CONFLICT]**
SEND DIAGNOSTIC invoked background self-test aborted - reasons to report results	SPC-3 8.2.9 (log page)	reset	logical unit reset	not mentioned
Mode parameter block descriptors number of blocks field (capacity setting) retained	SPC-3 8.4.4.2 8.4.4.3	reset events or power cycles	power cycles, hard resets, logical unit resets, and I_T nexus losses	not mentioned
Informational Exceptions Control mode page timer/counter maintained	SPC-3 8.4.10	power cycles and/or resets	power cycles, hard resets, logical unit resets, and I_T nexus losses	not mentioned **[CONFLICT]**
Access Controls left enabled	SPC-3 9.3.1.2	power cycles, logical unit resets, and target resets shall not disable	power cycles, hard resets, logical unit resets, and I_T nexus losses shall not disable	not mentioned
Access Controls: TransportID definition [of an initiator] must persist through	SPC-3 9.3.1.3.2.3	common reset events	hard resets and I_T nexus losses	not mentioned
Access Controls: go to not-enrolled state	SPC-3 9.3.1.5.1.2	power cycles or target resets based on vendor specific criteria	neither	not mentioned
Access Controls: force to pending-enrolled state	SPC-3 9.3.1.5.1.4	any event that causes the access controls coordinator to question (e.g. a process or port logout for FC or a hard bus reset for parallel SCSI)	logical unit reset or I_T nexus loss involving that initiator	yes
Access Controls: validity and proxy access rights of proxy tokens not affected	SPC-3 9.3.1.6.2.1	power cycles and target resets	power cycles, hard resets, logical unit resets, and I_T nexus losses	not mentioned
Access Controls: proxy LUN remain valid until	SPC-3 9.3.1.6.2.2	power cycle or target reset	logical unit reset or I_T nexus loss of the I_T nexus used to create the proxy LUN	not mentioned

What	Where	When (current phrase(s))	Proposed changes	FCP-2 clearing effects table
Access Controls: override lockout timer is set to initial value	SPC-3 9.3.1.8.2.2	reset or power cycle	logical unit reset	not mentioned
Access Control data persists across	SPC-3 9.3.1.12	power cycles and resets	power cycles, hard resets, and logical unit resets	not mentioned
Access Controls previously enrolled initiators to some state	SPC-3 9.3.12	following a power cycle or reset event	following a logical unit reset…following I_T nexus loss…	not mentioned
Mode parameters may also need initialization	SBC-2 4.2.1.5	[after logical unit resets]	after logical unit resets or I_T nexus losses	[see mode pages]
Cause of UNIT ATTENTION	SBC-2 4.2.1.13 4.3.3 4.4.5	logical unit reset or medium change	logical unit reset, I_T nexus loss, or medium change	not mentioned
Retain XOR data across	SBC-2 4.2.3.7	logical unit reset	logical unit reset, I_T nexus loss	retain
Cause of UNIT ATTENTION	SSC-2 4.2.10	logical unit reset or medium change	logical unit reset, I_T nexus loss, or medium change	not mentioned
Write Protect bits set to default state [mode page bits]	SSC-2 4.2.11.3 4.2.11.4 4.2.11.5 4.2.11.6	logical unit reset	[let SPC-3 define mode pages]	[see mode pages]
Unusual per-initiator log page for TapeAlert (see above)	SSC-2 4.2.16.2	[see above]	[see above]	[see above]
SET CAPACITY capacity persists across	SSC-2 7.10	power cycles, logical unit resets	power cycles, hard resets, logical unit resets, I_T nexus losses	not mentioned
Cleaning requirement parameter shall be persistent across	SSC-2 8.2.2	logical unit resets and power cycles	power cycles, hard resets, and logical unit resets, and I_T nexus loss	not mentioned
Device-specific parameters in mode parameter header, mode block descriptor, and data compression mode page retained following	SSC-2 8.3.1	a reset condition	hard reset, logical unit reset, or I_T nexus loss	clear (mode page)
following… report principal density	SSC-2 8.3.1	power on or reset condition	logical unit reset	clear (mode page)
following … retain knowledge of density code	SSC-2 8.3.1	reset condition	logical unit reset	clear (mode page)
Medium Partition mode page field values…shall not be cleared until	SSC-2 8.3.4	logical unit reset	logical unit reset or until I_T nexus loss of all I_T nexuses	clear (mode page)

Background

This is the original list discussed in the SRP WG and on the T10 reflector of some things affected by hard resets that are possibly affected by logouts. References are to SPC-3 revision 4 unless otherwise noted.

Mentioned by the FCP-2 clearing effects table (FCP-2 revision 7 section 4.9):

- **All tasks** aborted
- **Classic reservations** released (5.5.2, 7.26.2, 7.26.3)
- **Shared mode pages** reset to default/saved values (7.9.2, 7.9.6)
- **Non-shared mode pages** reset to default/saved values (7.9.2, 7.9.6)
- **Prevent Allow Medium Removal state** set to "allow" (7.13)
- **Buffered data for EXTENDED COPY** cleared (COPY STATUS data, RECEIVE DATA data, FAILED SEGMENT DETAILS data)(7.16.2, 7.16.3, 7.16.5)
- **Buffered data for XOR commands** cleared (SBC-2 4.2.3.7)
- **AccessID enrollment state** forced to pending-enrolled (future SPC-3)
- Preexisting ACA, unit attention, and deferred error conditions cleared (FCP-2)

Not mentioned by the FCP-2 clearing effects table but affected by resets:

- **Log pages** reset (7.6)
- **Alias list** associations cleared (7.2.1)
- **Generation value for persistent reservations** set to 0 (7.11.3)
- **Download microcode mode for WRITE BUFFER** - partial data cleared (7.33.5-8)
- **Self-test** cancelled (8.2.9)
- **Power condition state** (8.4.11)
- **Interval timer and report count fields** in Informational Exceptions mode page MRIE field are allowed to be cleared (8.4.10) [apparently a typo]

SSC-2 revision 7 section 8.3 prevents clearing of some important mode pages (for tapes), overriding the SPC-3/FCP-2 clearing rules:

> "The device-specific parameters contained in the mode parameter header, mode block descriptor values, and data compression page shall be retained following a reset condition (e.g., Target Reset, SCSI Logical Unit Reset, Fibre Channel Reset LIP or PLOGI)."

SSC-2 discusses these items affected by hard resets:

- **Software, associated, persistent, and permanent write protect bits** in the device configuration mode page cleared (4.2.9.2-5)
- **TapeAlert flags** cleared (4.2.1.4.2)

[These are being changed to “logical unit reset”]

Other command sets like SBC-2 may also clear various items based on resets.

Suggested Changes to SAM-2

3.1.39 hard reset: A ~~target action in response to a reset event in which the target port performs the operations described in~~type of transport attention used to return a device to its default operating conditions (see 5.8.6).

3.1.x I_T nexus loss: A type of transport attention indicating that an I_T nexus no longer exists (see 5.8.x).

3.1.x I_T nexus loss event: A protocol-specific event that triggers an I_T nexus loss (see 5.8.x).

3.1.60 logical unit reset: A logical unit action in response to a logical unit reset event in which the logical unit performs the operations described in 5.8.7.

3.1.61 logical unit reset event: An event that triggers a logical unit reset from a logical unit as described in 5.8.7.

3.1.p power cycle: Power being removed from and later applied to a SCSI device.

3.1.p power on: Power being applied to a SCSI device.

3.1.84 reset event: A ~~protocol specific~~n event that triggers a hard reset ~~from~~ in a SCSI device as described in 5.8.6. Reset events include power on, wakeup, and protocol-specific events.

3.1.x transport attention: A condition in the SCSI protocol layer that affects higher protocol layers. Transport attentions include I_T nexus loss and hard reset (see 5.8.t).

3.1.140 wakeup: A SCSI target port returning from the sleep power condition to the active power condition (see SPC-3).

3.1.141 wakeup event: An event that triggers a wakeup ~~from~~ in a SCSI target port as described in SPC-3.

5 SCSI Command Model

5.1 The Execute Command remote procedure

…

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 initial, wrap, and reset CRN values shall be one.~~ CRN shall be set to one for each I_T_L involving the initiator port after the initiator port receives a hard reset or detects I_T nexus loss. 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 protocol. It is not an error for the application client to provide this argument when CRN is not supported by the SCSI protocol or logical unit.

…

5.5 Task and command lifetimes

This subclause specifies the events delimiting the beginning and end (i.e., lifetime) of a task or tendered SCSI command from the viewpoint of the device server and application client.

The device server shall create a task upon receiving a SCSI Command Received indication unless the command represents a continuation of a linked command as described in 5.1.

The task shall exist until:

a) The device server sends a SCSI protocol service response for the task of TASK COMPLETE; or

b) The task is aborted as described in 5.6.5.8 Command processing considerations and exception conditions.

The application client assumes that the task exists from the time the **Send SCSI Command** SCSI protocol service request is invoked until it receives one of the following target responses:
a) A service response of TASK COMPLETE for that task;
b) Notification of a unit attention condition with one of the following additional sense codes:
A) COMMANDS CLEARED BY ANOTHER INITIATOR (if in reference to the task set containing the task);
[Editor's note: if additional sense code renaming is allowed, make all these changes. If not, just add I_T NEXUS LOSS occurred to the list.]
B) Any additional sense code whose ADDITIONAL SENSE CODE field contains 29h (e.g., ~~POWER ON, RESET, OR BUS DEVICE RESET~~TRANSPORT ATTENTION OCCURRED; POWER ON OCCURRED; ~~SCSI BUS~~HARD RESET OCCURRED; ~~BUS DEVICE RESET FUNCTION~~ LOGICAL UNIT RESET OCCURRED; DEVICE INTERNAL RESET; TRANSCEIVER MODE CHANGED TO SINGLE-ENDED; ~~or~~ TRANSCEIVER MODE CHANGED TO LVD; I_T NEXUS LOSS OCCURRED);
c) A service response of SERVICE DELIVERY OR TARGET FAILURE for the command. In this case, system implementations shall guarantee that the task associated with the failed command has ended;
d) A service response of FUNCTION COMPLETE following an ABORT TASK task management ~~request~~ function directed to the specified task;
e) A service response of FUNCTION COMPLETE following an ABORT TASK SET or a CLEAR TASK SET task management function directed to the task set containing the specified task; or
f) A service response of FUNCTION COMPLETE in response to a LOGICAL UNIT RESET or TARGET RESET task management function.

To the application client, the command is tendered from the time it calls the **Send SCSI Command** SCSI protocol service until one of the above responses or a service response of linked command complete is received.

When a SCSI protocol does not require state synchronization (see 4.6.1), there may be a time skew between the completion of a device server request-response transaction as seen by the application client and device server. As a result, the lifetime of a task or command as it appears to the application client normally is different from the lifetime observed by the device server.

[Editor's note: commands with background operations are:
SBC-2: FORMAT UNIT, PRE-FETCH, START STOP UNIT, SYNCHRONIZE CACHE
SSC-2: ERASE, VERIFY, WRITE FILEMARKS, LOCATE, FORMAT MEDIUM, LOAD UNLOAD, REWIND, SET CAPACITY
SPC-3: SEND DIAGNOSTICS]

Some commands (e.g., SEND DIAGNOSTIC) spawn background operations that operate after the task containing the command is no longer in the task set. These operations shall be aborted by logical unit resets. These operations may generate deferred errors, which are reported in sense data for a completed command (see SPC-3). Deferred errors shall be cleared by logical unit resets and by I_T nexus losses of the initiator port that sent the command that spawned the background operation, if the device server has retained that information.

5.8.4.1 Sense data introduction
Sense data shall be made available by the logical unit in the event a command completes with a CHECK CONDITION status or other conditions. The format, content and conditions under which sense data shall be prepared by the logical unit are specified in this standard, SPC-2, the applicable device command standard and applicable SCSI protocol standard.

Sense data shall be preserved by the logical unit for the initiator until it is transferred by one of the methods listed below or until another task from that initiator is entered into the task set, a logical unit reset occurs, or I_T nexus loss involving the initiator occurs.
…

5.8.5 Unit Attention condition

Each logical unit shall generate a unit attention condition ~~whenever the logical unit has been reset~~on a transport attention condition ~~-~~as described in 5.8.t ~~7 or by a power on reset~~. In addition, a logical unit shall generate a unit attention condition for each initiator whenever one of the following events occurs:

a) A removable medium may have been changed;
b) The mode parameters in effect for this initiator have been changed by another initiator;
c) The version or level of microcode has been changed;
d) Tasks for this initiator were cleared by another initiator;
e) INQUIRY data has been changed;
f) The logical unit inventory has been changed;
g) The mode parameters in effect for the initiator have been restored from non-volatile memory;
h) A change in the condition of a synchronized spindle; or
i) Any other event requiring the attention of the initiator.

Logical units may queue unit attention conditions. After the first unit attention condition is cleared, another unit attention condition may exist (e.g., a power on condition followed by a microcode change condition).

...

5.8.t Transport attentions

Figure 1 shows the relationships between I_T nexus loss, logical unit reset, hard reset, power on, and their associated events.

[Editor's note: need a better term than "state change" for the non-transport attention transport attentions.]

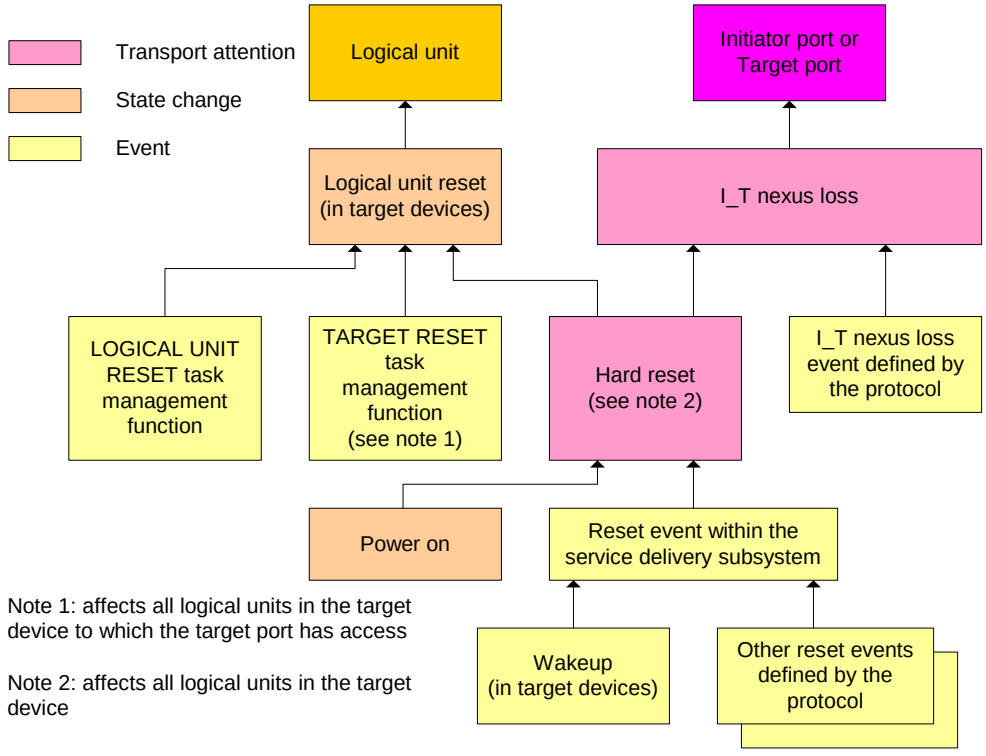

Figure 1. Transport attention relationships

Table xx shows the additional sense code that shall be reported when a unit attention is reported for each type of transport attention that occurs in a logical unit, unless the protocol defines a more

specific additional sense code for the event causing the condition. The conditions are sorted from least-specific to most-specific. The most-specific condition known to the logical unit shall be reported.

Table xx. Additional sense codes for transport attentions

Transport attention or state change	Additional sense code (today's wording)	Additional sense code (if wording changes allowed)
Logical unit cannot distinguish between the events	POWER ON, RESET, OR BUS DEVICE RESET OCCURRED	TRANSPORT ATTENTION OCCURRED
I_T nexus loss	I_T NEXUS LOSS OCCURRED	I_T NEXUS LOSS OCCURRED
Logical unit reset	BUS DEVICE RESET FUNCTION OCCURRED	LOGICAL UNIT RESET OCCURRED
Hard reset	SCSI BUS RESET OCCURRED	HARD RESET OCCURRED
Power on	POWER ON OCCURRED	POWER ON OCCURRED

5.8.p Power on

[Editor's note: define this term here so each protocol document need not include it in its list of protocol-specific reset events.]

When a SCSI device is powered on, it shall be considered a reset event which causes a hard reset.

5.8.6 Hard reset

[Editor's note: hard reset applies to both initiators and targets.]

A hard reset is a transport attention condition for a SCSI device indicating that the device shall be reset. A hard reset is a ~~target~~ SCSI port action in response to a power on or a reset event within the service delivery subsystem. A wakeup event (see 3.1.141) is ~~a~~ one such reset event. The definition of additional reset events is protocol specific. Each SCSI protocol standard that defines reset events shall specify ~~the~~ a target port's action in response to reset events.

~~The~~ A target port's response to a hard reset shall include initiating the equivalent of a logical unit reset for all logical units as described in 5.8.7.

~~While~~ Although the task manager response to task management requests is subject to the presence of access restrictions, as managed by ACCESS CONTROL OUT commands (see SPC-3), a hard reset ~~in response to a reset event within the service delivery subsystem~~ shall be unaffected by access controls.

When an initiator port receives a hard reset, it should terminate all its outstanding Execute Command remote procedure calls with SERVICE DELIVERY OR TARGET FAILURE.

5.8.7 Logical unit reset

A logical unit reset is a transport attention condition for a target device indicating that a logical unit shall be reset. A logical unit reset is:

a) An action in response to a LOGICAL UNIT RESET task management request (see 6.6) ~~or some other logical unit reset event~~; or

b) Part of an action in response to a TARGET RESET task management function (see 6.7); or

c) Part of an action in response to ~~or~~ a hard reset (see 5.8.6).

~~The definition of logical unit reset events is dependent on the SCSI protocol.~~

[Editor's note: Most of this list should be left to the command standards. Mentioning persistent reservations, mode pages, etc. defined outside of SAM-2 is not appropriate here. Mentioning SAM-2 effects like tasks being aborted *is* appropriate here.]

To process a logical unit reset the logical unit shall:

a) Abort all tasks as described in 5.6;
b) Clear a CA (see 5.8.1.6) or ACA (see 5.8.1.7) condition, if one is present;
~~c) Release all reservations established using the reserve/release management method (persistent reservations shall not be affected);~~
~~d) Return the logical unit's operating mode to the appropriate initial conditions, similar to those conditions that would be found following device power-on. The MODE SELECT parameters (see SPC-2) shall be restored to their last saved values if saved values have been established. MODE SELECT parameters for which no saved values have been established shall be returned to their default values;~~
e) Set a unit attention condition (see 5.8.5 and 5.8.t); ~~and~~
f) Initiate a logical unit reset for all dependent logical units (see 4.12);~~.~~ and

[Editor's note: protocol standards should not define any other effects – only command set standards should do so.]
~~In addition to the above, the logical unit shall p~~g) Perform any additional functions required by the applicable command set standards.

5.8.x I_T nexus loss

[Editor's note: this is a new section defining the new term.]
An I_T nexus loss is a transport attention condition indicating that an I_T nexus is no longer available. In target devices, logical unit resets are I_T nexus loss events. In initiator devices, hard resets are I_T nexus loss events. Protocols may define additional I_T nexus loss events.

Each logical unit to which the target port has access shall process an I_T nexus loss.

If the logical unit remembers any state for the initiator port, it shall perform this sequence:

a) Abort all tasks from the initiator port involved in the I_T nexus as described in 5.6;
b) Clear a CA (see 5.8.1.6) or ACA (see 5.8.1.7) condition, if one is present for the initiator port;
c) Set a unit attention condition for the initiator port (see 5.8.5 and 5.8.t); and
d) Perform any additional functions required by the applicable command set standards.

If the logical unit does not remember any state for the initiator port, it shall process the I_T nexus loss as a logical unit reset as described in 5.8.7.

When an initiator port detects an I_T nexus loss, it should terminate all its outstanding Execute Command remote procedure calls to that target port with SERVICE DELIVERY OR TARGET FAILURE.

5.8.x Channel loss

[Editor's note: it may be best to defer this concept until SAM-3.]
Protocols may support multiple channels between an initiator port and target port (i.e., forming the same I_T nexus). Channel loss means one of these channels has disappeared, but the I_T nexus is still accessible through other channels. Loss of all channels is considered an I_T nexus loss (see 5.8.x).

Protocols defining multiple channels include SRP and iSCSI. Protocols not defining multiple channels include SPI-4 and FCP-2.

To process a loss of channel the target port shall:
a) Abort all tasks sent on that channel as described in 5.6.

When an initiator port detects channel loss, it should terminate all its outstanding Execute Command remote procedure calls on that channel with SERVICE DELIVERY OR TARGET FAILURE.

[Editor's note: probably need to add "channel" elsewhere in the architecture, making this a more complicated addition.]

Suggested changes to SPC-3

[Editor's note: Wherever "reset" is listed, change to "logical unit reset" unless it's truly hardware related This parallels changes already made to SAM-2, SBC-2, SSC-2, and SPI-4.]

[Editor's note: Wherever "logical unit reset" is listed, add "I_T nexus loss" if appropriate.]

…

~~**3.1.33 hard reset:** A target response to a reset event or TARGET RESET task management function. A detailed definition of hard reset may be found in SAM-2.~~

3.1.xx. **hard reset**: A port action in response to a reset event in which the port performs the operations described in SCSI Architecture Model-2.

3.1.x I_T nexus loss: A port action in response to an I_T nexus loss event in which the port performs the operations described in SCSI Architecture Model-2.

3.1.x I_T nexus loss event: An event that triggers I_T nexus loss as described in SCSI Architecture Model-2.

3.1.x. **logical unit reset**: A logical unit action in response to a logical unit reset event in which the logical unit performs the operations described in SCSI Architecture Model-2.

3.1.xx. **logical unit reset event**: An event that triggers a logical unit reset from a logical unit as described in SCSI Architecture Model–2.

3.1.xx. **reset event**: An event that triggers a hard reset in a SCSI device as described in SCSI Architecture Model-2. Reset events include power on, wakeup, and other protocol-specific events.

3.1.p power cycle: Power being removed from and later applied to a SCSI device.

3.1.p power on: Power being applied to a SCSI device.

3.1.96 wakeup: A target port returning from the sleep power condition to the active power condition (see 5.8).

3.1.97 wakeup event: An event defined by the protocol that triggers a wakeup from a target port as described in 5.8.

…

5 Model common to all device types

5.5 Reservations

5.5.2 The Reserve/Release management method

[Editor's note: Reserve/Release may be obsoleted in May 2002, making these fixes moot. If they remain, the informal group choose to have classic reservations preserved through I_T nexus losses.]

The reserve/release management method commands, RESERVE(6), RESERVE(10), RELEASE(6), and RELEASE(10) are used among multiple initiators that do not require operations to be protected across initiator failures (and subsequent ~~hard resets~~logical unit resets). The reserve/release reservations management method also allows an application client to provide restricted device access to one additional initiator (a third-party initiator), usually a temporary initiator performing a service for the application client sending the reservation command.

Reservations managed using the reserve/release method do not persist across ~~some recovery actions (e.g., hard resets)~~logical unit resets. ~~When a target performs~~After one of these ~~recovery actions~~events, the application client~~(s)~~ have to rediscover the configuration and re-establish the required reservations. Reserve/release managed reservations are retained by the device server until released or until reset by mechanisms specified in this standard.

…

5.5.3 The Persistent Reservations management method

5.5.3.1 Overview of the Persistent Reservations management method

The persistent reservations management method is the mechanism specified by this standard for use by multiple initiators that require operations to be protected across initiator failures, which usually involve ~~hard resets~~logical unit resets and I_T nexus losses.

Persistent reservations persist across recovery actions, to provide initiators with more detailed control over reservations recovery. Persistent reservations are not reset by ~~the~~ hard resets, logical unit resets or I_T nexus losses~~TARGET RESET task management function or other global actions~~.

The persistent reservation held by a failing initiator port may be preempted by another initiator port as part of the recovery process. Persistent reservations shall be retained by the device server until released, preempted, or cleared by mechanisms specified in this standard. ~~Even though different SCSI protocols that transport SCSI commands handle hard resets differently (e.g., parallel SCSI uses a reset signal, Fibre Channel loops use primitive signals) the persistent reservation shall be preserved.~~ Optionally, persistent reservations may be retained when power to the target is removed.

5.5.3.2 Preserving persistent reservations

…

The device server shall preserve the following information for each registration across any hard reset, logical unit reset~~reset~~ or I_T nexus loss, and if the persist through power loss capability is enabled, across any power cycle:

…

The device server shall preserve the following persistent reservation information across any hard reset, logical unit reset~~reset~~ or I_T nexus loss, and if the persist through power loss capability is enabled, across any power cycle:

…

[Editor's note: This section is not where power cycles are defined as performing hard resets, so remove line a).]

A persistent reservation may also be released by a loss of power, if the persist through power loss capability is not enabled. When the most recent APTPL value received by the device server is zero (see 7.12.3), a power cycle:

~~a) Performs a hard reset;~~

a~~b~~) Releases all persistent reservations; and

b~~c~~) Removes all registered reservation keys (see 5.5.3.4).

…

5.7 Asymmetric logical unit access

5.7.8 Behavior after ~~power cycling or hard resets~~logical unit reset

[Editor's note: I_T nexus loss does not affect target port group states (active/standby).]

For all SCSI target devices that report in the standard INQUIRY data (see 7.4.2) that they support only explicit asymmetric logical unit access (i.e., the ALUA field contains 10b), the target port shall preserve the target port asymmetric access state ~~across~~ during any hard reset, logical unit reset or I_T nexus loss~~ and across any power cycle~~.

…

5.8 Power conditions

[Editor's note: power condition is not affected by I_T nexus loss. Only a wakeup or hard reset causes wakeup, not a logical unit reset.]

…

Table 15

…

sleep Logical units to which the target port has access are not capable of accepting commands routed from the target port. The lowest power consumption, with power applied, occurs in the sleep power condition. The target port requires a wakeup or hard reset to return to the active

power condition. The target port enters the sleep power condition only when all the logical units to which it has access have entered the sleep power condition.

…

If implemented, the target port shall use the optional Power Condition mode page (see 8.4.11) to control the logical unit power conditions after a wakeup or hard reset until a START STOP UNIT command (see SBC-2 or RBC) is received with the POWER CONDITIONS field set to a value other than 0h or 7h.

…

TP1:TP0: A wakeup or hard reset returns the target port to the active power condition and returns each logical unit to the power condition (active, idle, or standby) defined by the saved Power Condition mode page parameters.

…

Table 16

…

LU8:LUx A wakeup or hard reset returns the logical unit to the active, idle, or power condition defined by the saved Power Condition mode page parameters.

…

7.2 CHANGE ALIASES command

7.2.1 CHANGE ALIASES command introduction

[Editor's note: alias tables are cleared by I_T nexus loss]

…

On successful completion of a CHANGE ALIASES command, the device server shall maintain an association of each assigned eight byte alias value to the SCSI device or port designation. These associations shall be cleared ~~under any event that resets the logical unit and events designated by the SCSI protocol~~by a logical unit reset or I_T nexus loss. The device server shall maintain a separate alias list for each initiator.

…

7.4 INQUIRY command

7.4.1 INQUIRY command introduction

[Editor's note: discussion of INQUIRY data availability applies to logical unit resets, not just hard resets or power ups]

…

The INQUIRY data should be returned even though the device server is not ready for other commands. To minimize delays after a ~~hard reset or power-up condition~~logical unit reset, the standard INQUIRY data should be available without incurring any media access delays. If the device server does store some of the INQUIRY data on the media, it may return zeros or ASCII spaces (20h) in those fields until the data is available from the media.

…

[Editor's note: this seems to be discussing the initiator device's hard reset or power-up, not the target device's. So, logical unit reset is not appropriate.]

NOTE 14 - The INQUIRY command may be used by an application client after a hard reset or power~~-up~~ on condition to determine the device types for system configuration.

…

7.6 LOG SENSE command

[Editor's note: shared log pages are affected by I_T nexus loss of all I_T nexuses]

[Editor's note; the TapeAlert log page, a per-initiator port log page, does not follow these rules. A table like in mode pages could be added describing these types of log pages too.]

…

The page control (PC) field defines the type of parameter values to be selected (see 7.5 for the definition of the page control field). The parameter values returned by a LOG SENSE command are determined as follows:

a) The specified parameter values at the last update (i.e., in response to a LOG SELECT or LOG SENSE command or done automatically by the target for cumulative values);

b) The saved values, if saved parameters are implemented and an update has not occurred since the last ~~power-on or hard reset condition~~ logical unit reset or I_T nexus loss of all I_T nexuses involving the initiator port~~and saved parameters are implemented~~; or

c) The default values, if saved values are not available or not implemented anad an update has not occurred since the last ~~power-on or hard reset condition~~ logical unit reset or I_T nexus loss of all I_T nexuses involving the initiator port~~and saved values are not available or not implemented~~.

…

7.7 MODE SELECT(6) command

…

~~If a target supports saved mode pages, it may save only one copy of the mode page for each logical unit and have it apply to all initiators, or it may save separate copies for each initiator for each logical unit. Multiple port implementations may save one copy per logical unit and have it apply to all initiators on all ports or save a separate copy per logical unit for each initiator on each port. If separate copies are saved, the target shall maintain separate current values for each combination of initiator and logical unit that it detects. Mode pages that are common to all initiators are not required to have multiple copies.~~

Logical units may maintain current and saved values of each mode page based on any of the policies listed in Table xx.

Table xx. Mode page policies

Mode page policy	**Logical unit maintains**	**When the mode page shall revert to saved values if supported or default values if saved values are not supported**
shared	One copy of the mode page.	logical unit reset or I_T nexus loss of all I_T nexuses
per-initiator port	Separate copies of the mode page for each initiator port.	logical unit reset or I_T nexus loss of all I_T nexuses involving the initiator port
per-I_T nexus	Separate copies of the mode page for each I_T nexus.	logical unit reset or I_T nexus loss

If an application client sends a MODE SELECT command that changes any parameters applying to other initiators, the device server shall generate a unit attention condition for all initiators except the one that issued the MODE SELECT command (see SAM-2). The device server shall set the additional sense code to MODE PARAMETERS CHANGED.

~~The target may provide for independent sets of parameters for each attached logical unit or for each combination of logical unit and initiator. If independent sets of parameters are implemented, and~~ If a third-party reservation is requested, the device server shall transfer the set of per-initiator port or per-I_T nexus parameters in effect for the initiator port or I_T nexus that sent the RESERVE command to the parameters used for commands from the third-party device (see 7.26.3).

A page format (PF) bit of zero indicates that all parameters after the block descriptors are vendor specific. A PF bit of one indicates that the MODE SELECT parameters following the header and block descriptor(s) are structured as pages of related parameters and are as specified in this standard.

A save pages (SP) bit of zero indicates the device server shall perform the specified MODE SELECT operation, and shall not save any mode pages. If the ~~target~~ logical unit implements no distinction between current and saved mode pages and the SP bit is zero, the command shall be terminated with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code shall be set to INVALID FIELD IN CDB. An SP bit of one indicates that the device server shall perform the specified MODE SELECT operation, and shall save to a nonvolatile vendor specific location all the saveable mode pages including any sent in the Data-Out Buffer. The SP bit is optional~~, even when mode pages are supported by the target~~. Mode

pages that are saved are identified by the parameter saveable (PS) bit that is returned in the page header by the MODE SENSE command (see 8.4). If the PS bit is set to one in the MODE SENSE data then the mode page shall be saveable by issuing a MODE SELECT command with the SP bit set to one. If the ~~target~~ logical unit does not implement saved mode pages and the SP bit is set to one, the command shall be terminated with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code shall be set to INVALID FIELD IN CDB.

…

7.9 MODE SENSE(6) command

7.9.1 MODE SENSE(6) command introduction

…

An application client may request any one or all of the supported mode pages from the device server. If an application client issues a MODE SENSE command with a page code or subpage code value not implemented by the target, the device server shall return CHECK CONDITION status and shall set the sense key to ILLEGAL REQUEST and the additional sense code to INVALID FIELD IN CDB.

…

7.9.2 Current values

[Editor's note: mode pages need not be affected by I_T nexus loss as specified in FCP-2. It doesn't hurt if they are cleared, because software must assume the worst (that some other initiator changed them while it was logged out. The main concern is that device servers may clear per-initiator data when the initiator departs; otherwise how long is it kept? Assume that FC devices, which are required to clear mode pages when all initiators log out, will provide a unit attention indicating logical unit reset occurred rather than the new I_T nexus loss occurred.]

A PC field value of 00b requests that the device server return the current values of the mode parameters. The current values returned are:

a) The current values of the mode parameters established by the last successful MODE SELECT command;

b) The saved values of the mode parameters if a MODE SELECT command has not successfully completed since the ~~last power-on or hard reset condition~~mode parameters were restored to their saved values; or

c) The default values of the mode parameters if a MODE SELECT command has not successfully completed since the ~~, if saved values, are not available or not supported~~mode parameters were restored to their default values.

…

7.9.6 Initial responses

After a ~~power-up condition or hard reset condition~~logical unit reset, the device server shall respond in the following manner:

a) If default values are requested, report the default values;

b) If saved values are requested, report valid restored mode parameters, or restore the mode parameters and report them. If the saved values of the mode parameters are not able to be accessed from the nonvolatile vendor specific location, terminate the command with CHECK CONDITION status and set the sense key set to NOT READY. If saved parameters are not implemented respond as defined in 7.9.5; or

c) If current values are requested and the current values of the mode parameters have not been sent by the application client (via a MODE SELECT command), the device server may return either the default or saved values, as defined above. If current values have been sent, the current values shall be reported.

…

7.11 PERSISTENT RESERVE IN command

7.11.3 PERSISTENT RESERVE IN parameter data for READ KEYS

[Editor's note: PR generation value is not affected by I_T nexus loss]

The Persistent Reservations Generation (PRGENERATION) field shall contain a 32-bit counter maintained by the device server that shall be incremented every time a PERSISTENT RESERVE OUT command requests a REGISTER, a REGISTER AND IGNORE EXISTING KEY, a CLEAR,

a PREEMPT, or a PREEMPT AND ABORT service action. The counter shall not be incremented by a PERSISTENT RESERVE IN command, by a PERSISTENT RESERVE OUT command that performs a RESERVE or RELEASE service action, or by a PERSISTENT RESERVE OUT command that is terminated due to an error or reservation conflict. Regardless of the APTPL bit value, the PRGENERATION value shall be set to zero ~~as part of the power on reset process~~by a power on.

…

7.13 PREVENT ALLOW MEDIUM REMOVAL command

[Editor's note: medium removal permission is not affected by I_T nexus loss]

…

The prevention of medium removal shall begin when any application client issues a PREVENT ALLOW MEDIUM REMOVAL command with a PREVENT field of 01b or 11b (i.e., medium removal prevented). The prevention of medium removal for the logical unit shall terminate:

a) After all initiators with application clients that previously prevented medium removal issue PREVENT ALLOW MEDIUM REMOVAL commands with a PREVENT field of 00b or 10b, and the device server has successfully performed a synchronize cache operation; or

b) Upon a ~~hard reset condition~~logical unit reset.

…

7.16 RECEIVE COPY RESULTS command

7.16.2 COPY STATUS service action

[Editor's note: copy buffers are cleared by I_T nexus loss. This is data associated with a former task, and all tasks get cleared.]

…

After completion of an EXTENDED COPY command, the copy manager shall preserve all data returned by a COPY STATUS service action for a vendor specific period of time. The copy manager shall discard the COPY STATUS data when:

a) A REC~~I~~EIVE COPY RESULTS command with COPY STATUS service action is received from the same initiator with a matching list identifier;

b) When another EXTENDED COPY command is received from the same initiator and the list identifier matches the list identifier associated with the data preserved for the COPY STATUS service action;

c) When the copy manager detects a ~~hard reset condition~~logical unit reset or I_T nexus loss; or

d) When the copy manager requires the resources used to preserve the data.

…

7.16.3 RECEIVE DATA service action

…

Following completion of an EXTENDED COPY command, the copy manager shall preserve all data returned by a REC~~I~~EIVE DATA service action for a vendor specific period of time. The application client should issue a RECEIVE COPY RESULTS command with RECEIVE DATA service action as soon as practical following completion of the EXTENDED COPY command to insure that the data is not discarded by the copy manager. The copy manager shall discard the buffered inline data:

a) After all data held for a specific EXTENDED COPY command has been successfully transferred to the application client;

b) When a REC~~I~~EIVE COPY RESULTS command with RECEIVE DATA service action has been received from the same initiator with a matching list identifier, with the ALLOCATION LENGTH field set to zero;

c) When another EXTENDED COPY command is received from the same initiator and the list identifier matches the list identifier associated with the data preserved for RECEIVE DATA service action;

d) When the copy manager detects a ~~hard reset condition~~logical unit reset or I_T nexus loss; or

e) When the copy manager requires the resources used to preserve the data.

…

7.16.5 FAILED SEGMENT DETAILS service action

…

The application client should issue a RECEIVE COPY RESULTS command with FAILED SEGMENT DETAILS service action immediately following failure of the EXTENDED COPY

command to insure that the information is not discarded by the copy manager. The copy manager shall discard the failed segment details:
a) After all failed segment details held for a specific EXTENDED COPY command have been successfully transferred to the application client;
b) When a RECEIVE COPY RESULTS command with FAILED SEGMENT DETAILS service action has been received from the same initiator with a matching list identifier, with the ALLOCATION LENGTH field set to zero;
c) When another EXTENDED COPY command is received from the same initiator using the same list identifier;
d) When the copy manager detects a ~~hard reset condition~~logical unit reset or I_T nexus loss; or
e) When the copy manager requires the resources used to preserve the data.
…

7.21 REPORT DEVICE IDENTIFIER command

…

[Editor's note: changeable device identifiers are not affected by I_T nexus loss]
The IDENTIFIER field shall contain a vendor specific value. The value reported shall be the last value written by a successful SET DEVICE IDENTIFIER command. The value of the identifier shall be changed only by a SET DEVICE IDENTIFIER command. The identifier value shall persist through ~~resets, power cycles~~logical unit resets, I_T nexus losses, media format operations, and media replacement.
…

7.22 REPORT LUNS command

[Editor's note: similar to INQUIRY data text, REPORT LUNS data should be ready after any logical unit reset, not just any hard reset/power up.]
…
The REPORT LUNS data should be returned even though the device server is not ready for other commands. To minimize delays after a ~~hard reset or power-up condition~~logical unit reset, the default report of the logical unit inventory should be available without incurring any media access delays. The default report of the logical unit inventory shall contain at least LUN 0.
…

7.25 REQUEST SENSE command

[Editor's note: add an additional sense code to tell an initiator that it is logging in again after an I_T nexus loss rather than that it is logging in after a logical unit reset has occurred. Simple software just assumes everything may have been cleared on any ASC of 29h. Advanced software may interpret 29h/07h as indicating that items not cleared by I_T nexus losses alone are still as they were before (e.g. log pages, mode pages, changeable device identifiers, …). One problem left: if it is not really a reset, it probably shouldn't be a 29h code.]

Table 151 - Sense key descriptions

…

6h **UNIT ATTENTION.** Indicates that the removable medium may have been changed ~~or the target has been reset~~, or a logical unit reset occurred, or an I_T nexus loss occurred involving the initiator port. See SAM-2 for more detailed information about the unit attention condition.
…

Table 152 - ASC and ASCQ assignments

…

[Editor's note: renaming the ASCQs may cause confusion, but the current labels don't match current terminology. Perhaps both the new names and the old names could be listed in the database?]
[Editor's note: the 00 label is incomplete; it doesn't mention DEVICE INTERNAL RESET or the TRANSCEIVER MODE CHANGE reset events. Rather than make the label several lines long, a simple name might be better:]

[Editor's note: if major changes are allowed:]
29h 00h D T L PWR S OMC A E B K TRANSPORT ATTENTION OCCURRED
[Editor's note: if only minor changes are allowed:]
29h 00h D T L PWR S OMC A E B K POWER ON, HARD RESET, OR ~~BUS DEVICE~~ LOGICAL UNIT RESET OCCURRED

29h 01h D T L PWR S OMC A E B K POWER ON OCCURRED

29h 02h D T L PWR S OMC A E B K ~~SCSI BUS~~HARD RESET OCCURRED
29h 03h D T L PWR S OMC A E B K ~~BUS DEVICE~~LOGICAL UNIT RESET ~~FUNCTION~~ OCCURRED
29h 04h D T L PWR S OMC A E B K DEVICE INTERNAL RESET
29h 05h D T L PWR S OMC A E B K TRANSCEIVER MODE CHANGED TO SINGLE-ENDED
29h 06h D T L PWR S OMC A E B K TRANSCEIVER MODE CHANGED TO LVD
29h 07h D T L PWR S OMC A E B K I_T NEXUS LOSS OCCURRED
…

7.26 RESERVE(10) command
7.26.2 Logical unit reservation
[Editor's note: as in section 5, a classic reservation is NOT cleared by I_T nexus loss of the reservation holder.]
Logical unit reservations are mandatory if the RESERVE(10) command is implemented. This command shall request that the entire logical unit be reserved for the exclusive use of the initiator until the reservation is:
a) Superseded by another valid RESERVE command from the same initiator;
b) Released by a RELEASE command from the same initiator that made the reservation;
c) Released by a ~~TARGET RESET task management function performed by any initiator~~logical unit reset; or
~~d) Released by a LOGICAL UNIT RESET task management function performed by any initiator;~~
~~e) Released by a hard reset condition; or~~
~~f) Released by a power on cycle.~~

NOTE: Logical units with target ports on some protocols compliant with previous versions of this standard (e.g., FCP-2) may release reservations on I_T nexus loss.
…
7.26.3 Third-party reservation
If the third-party (3RDPTY) bit is zero, then a third-party reservation is not requested. If the 3RDPTY bit is zero then the LONGID bit shall be ignored. If the 3RDPTY bit is one then the device server shall reserve the specified logical unit for the SCSI device specified in the THIRD-PARTY DEVICE ID field. Device ID formats are protocol specific. The device server shall preserve the reservation until it is superseded by another valid RESERVE command from the initiator that made the reservation or until:
a) ~~it is released~~ a RELEASE command from ~~by~~ the same initiator that made the reservation; or
b) ~~,~~ by a ~~TARGET RESET task management function performed by any initiator, a hard reset condition, or by a power on cycle~~logical unit reset.~~.~~

NOTE: Logical units with target ports on some protocols compliant with previous versions of this standard (e.g., FCP-2) may release reservations on I_T nexus loss.

The device server shall ignore any attempt to release the reservation made by any other initiator.
…

7.33 WRITE BUFFER command
7.33.5 Download microcode mode (04h)
[Editor's note: I_T nexus loss clears any partially downloaded microcode data.]
If the logical unit is unable to accept this command because of some device condition, the device server shall terminate each WRITE BUFFER command with this mode (04h) with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

In this mode, vendor specific microcode or control information shall be transferred to the control memory space of the logical unit. After a ~~power-cycle or reset~~logical unit reset, the device operation shall revert to a vendor specific condition. The meanings of the BUFFER ID, BUFFER OFFSET, and PARAMETER LIST LENGTH fields are not specified by this standard and are not required to be zero-filled. When the microcode download has completed successfully the device server shall generate a unit attention condition for all initiators except the one that issued the WRITE

BUFFER command (see SAM-2). The additional sense code shall be set to MICROCODE HAS BEEN CHANGED.

7.33.6 Download microcode and save mode (05h)

If the logical unit is unable to accept this command because of some device condition, the device server shall terminate each WRITE BUFFER command with this mode (05h) with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

In this mode, vendor specific microcode or control information shall be transferred to the logical unit and, if the WRITE BUFFER command is completed successfully, also shall be saved in a nonvolatile memory space (semiconductor, disk, or other). The downloaded code shall then be effective after each ~~power-cycle and reset~~logical unit reset until it is supplanted in another download microcode and save operation or download microcode with offsets and save operation. The meanings of the BUFFER ID, BUFFER OFFSET, and PARAMETER LIST LENGTH fields are not specified by this standard and are not required to be zero-filled. When the download microcode and save command has completed successfully the device server shall generate a unit attention condition (see SAM-2) for all initiators except the one that issued the WRITE BUFFER command. When reporting the unit attention condition, the device server shall set the additional sense code to MICROCODE HAS BEEN CHANGED.

7.33.7 Download microcode with offsets (06h)

In this mode, the application client may split the transfer of the vendor specific microcode or control information over two or more WRITE BUFFER commands. If the logical unit is unable to accept this command because of some device condition, the device server shall terminate each WRITE BUFFER command with this mode (06h) with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

If the last WRITE BUFFER command of a set of one or more commands completes successfully, the microcode or control information shall be transferred to the control memory space of the logical unit. After a ~~power-cycle or reset~~logical unit reset, the device shall revert to a vendor specific condition. In this mode, the Data-Out Buffer contains vendor specific, self-describing microcode or control information.

Since the downloaded microcode or control information may be sent using several commands, when the logical unit detects the last download microcode with offsets and save mode WRITE BUFFER command has been received, the device server shall perform any logical unit required verification of the complete set of downloaded microcode or control information prior to returning GOOD status for the last command. After the last command completes successfully the device server shall generate a unit attention condition (see SAM-2) for all initiators except the one that issued the set of WRITE BUFFER commands. When reporting the unit attention condition, the device server shall set the additional sense code to MICROCODE HAS BEEN CHANGED.

If the complete set of WRITE BUFFER commands required to effect a microcode or control information change (one or more commands) are not received before a ~~reset or power-on cycle~~logical unit reset or I_T nexus loss occurs, the change shall not be effective and the new microcode or control information shall be discarded.

…

7.33.8 Download microcode with offsets and save mode (07h)

In this mode, the initiator may split the transfer of the vendor specific microcode or control information over two or more WRITE BUFFER commands. If the logical unit is unable to accept this command because of some device condition, the device server shall terminate each mode 07h WRITE BUFFER command with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

If the last WRITE BUFFER command of a set of one or more commands completes successfully, the microcode or control information shall be saved in a nonvolatile memory space (e.g., semiconductor, disk, or other). The saved downloaded microcode or control information shall then be effective after each ~~power cycle and reset~~logical unit reset until it is supplanted by another download microcode with save operation or download microcode with offsets and save operation. In this mode, the Data-Out Buffer contains vendor specific, self-describing microcode or control information.

Since the downloaded microcode or control information may be sent using several commands, when the logical unit detects the last download microcode with offsets and save mode WRITE BUFFER command has been received, the device server shall perform any logical unit required verification of the complete set of downloaded microcode or control information prior to returning GOOD status for the last command. After the last command completes successfully the device server shall generate a unit attention condition (see SAM-2) for all initiators except the one that issued the set of WRITE BUFFER commands. When reporting the unit attention condition, the device server shall set the additional sense code to MICROCODE HAS BEEN CHANGED.

If the complete set of WRITE BUFFER commands required to effect a microcode or control information change (one or more commands) are not received before a ~~reset or power on cycle~~logical unit reset or I_T nexus loss occurs, the change shall not be effective and the new microcode or control information shall be discarded.

8.2 Log parameters

8.2.9 Self-Test Results log page

…

[Editor's note: Background self tests are not aborted by I_T nexus losses. They are apparently cleared by logical unit resets since "reset" was mentioned before.]

Table 193 - Self-test results values

…

2h The self-test routine was aborted by an application client using a method other than a SEND DIAGNOSTICS command with the SELF-TEST CODE field set to 100b (e.g., by a task management function, by a ~~reset~~logical unit reset, or by issuing an exception command as defined in 5.4.3).
[Editor's note: which task management functions clear background tasks? SAM-2 doesn't describe this state. These commands invoke background operations:
SBC-2: FORMAT UNIT, PRE-FETCH, START STOP UNIT, SYNCHRONIZE CACHE
SSC-2: ERASE, VERIFY, WRITE FILEMARKS, LOCATE, FORMAT MEDIUM, LOAD UNLOAD, REWIND, SET CAPACITY
SPC-3: SEND DIAGNOSTICS]

…

8.4 Mode parameters

8.4.4 Mode parameter block descriptor formats

8.4.4.2 Direct-access device block descriptor format for LONGLBA=0

[Editor's note: mode pages are not cleared by I_T nexus losses]

…

If the SCSI device doesn't support changing its capacity by changing the NUMBER OF BLOCKS field using the MODE SELECT command, the value in the NUMBER OF BLOCKS field is ignored. If the device supports changing its capacity by changing the NUMBER OF BLOCKS field, then the NUMBER OF BLOCKS field is interpreted as follows:

a) If the number of blocks is set to zero, the device shall retain its current capacity if the block size has not changed. If the number of blocks is set to zero and the block size has changed, the device shall be set to its maximum capacity when the new block size takes effect;

b) If the number of blocks is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through ~~reset events or power cycles~~power cycles, hard resets, logical unit resets, and I_T nexus losses;

c) If the number of blocks field is set to a value greater than the maximum capacity of the device and less than FFFF FFFFh, then the command is terminated with a CHECK CONDITION

status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings;

d) If the number of blocks is set to FFFF FFFFh, the device shall be set to its maximum capacity. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through ~~reset events or power cycles~~power cycles, hard resets, logical unit resets, and I_T nexus losses.

…

8.4.4.3 Long LBA block descriptor format

If the SCSI device doesn't support changing its capacity by changing the NUMBER OF BLOCKS field using the MODE SELECT command, the value in the NUMBER OF BLOCKS field is ignored. If the device supports changing its capacity by changing the NUMBER OF BLOCKS field, then the NUMBER OF BLOCKS field is interpreted as follows:

a) If the number of blocks is set to zero, the device shall retain its current capacity if the block size has not changed. If the number of blocks is set to zero and the block size has changed, the device shall be set to its maximum capacity when the new block size takes effect;

b) If the number of blocks is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through ~~reset events or power cycles~~power cycles, hard resets, logical unit resets, and I_T nexus losses;

c) If the number of blocks field is set to a value greater than the maximum capacity of the device and less than FFFF FFFF FFFF FFFFh, then the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings;

d) If the number of blocks is set to FFFF FFFF FFFF FFFFh, the device shall be set to its maximum capacity. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through ~~reset events or power cycles~~power cycles, hard resets, logical unit resets, and I_T nexus losses.

…

8.4.10 Informational Exceptions Control mode page

[Editor's note: the last sentence is referring to the running timer value and counter, not the fields. Change names a bit to help clarify.]

[Editor's note: the running timer and counter need not be cleared on I_T nexus losses, but they may be.]

The INTERVAL TIME~~R~~ field indicates the period in 100 millisecond increments for reporting that an informational exception condition has occurred. The device server shall not report informational exception conditions more frequently than the time specified by the INTERVAL TIME~~R~~ field and shall report them as soon as possible after the ~~timer~~ interval has elapsed. After the informational exception condition has been reported the interval timer shall be restarted. A value of zero or FFFF FFFFh in the INTERVAL TIME~~R~~ field shall indicate the ~~timer~~ interval time is vendor specific.

The REPORT COUNT field indicates the number of times to report an informational exception condition to the application client. A value of zero in the REPORT COUNT field indicates there is no limit on the number of times the device server reports an informational exception condition.

The maintaining of the ~~INTERVAL TIMER~~interval timer and the ~~REPORT COUNT fields~~ report counter across ~~power cycles and/or resets by the target are~~ power cycles, hard resets, logical unit resets, and I_T nexus losses is vendor specific.

…

8.4.11 Power Condition mode page

[Editor's note: as in chapter 5, power conditions only relate to hard resets not logical unit resets]

The logical unit shall use the Power Condition mode page to control the power conditions after a power on or a hard reset until a START STOP UNIT command is received that sets power conditions.

…

9 Well known logical units
9.3 ACCESS CONTROLS well known logical unit
9.3.1 Access controls model
9.3.1.2 Access controls overview

[Editor's note: Access Controls' being enabled is not affected by I_T nexus loss]

…

Once access controls are enabled, ~~power cycles, logical unit resets, and target resets~~ power cycles, hard resets, logical unit resets, and I_T nexus losses shall not disable them.

…

9.3.1.3 The access control list
9.3.1.3.2 Access identifiers
9.3.1.3.2.3 TransportID access type identifiers

[Editor's note: An identifier or name referred to by a TransportID must not change during I_T nexus losses. Don't include logical unit reset here, since this is about initiators.]

…

The format of the SCSI protocol specific data depends on the value in the PROTOCOL IDENTIFIER field. The SCSI protocol specific data in a TransportID shall only include initiator port identifiers, initiator port names, or initiator device names (see SAM-2) that persist across ~~common reset events~~hard resets and I_T nexus losses~~in the service delivery subsystem~~. TransportID formats specific to SCSI protocols are listed in table 288.

…

9.3.1.5 Enrolling AccessIDs
9.3.1.5.1 Enrollment states
9.3.1.5.1.1 Summary of enrollment states

Initiators enroll an AccessID with an access controls coordinator in order to be allowed access to logical units listed in the ACE having the same AccessID type access identifier. Enrolling an AccessID is accomplished using the ACCESS CONTROL OUT command with ACCESS ID ENROLL service action (see 9.3.3.4). An initiator shall be in one of three states with respect to such an enrollment:

a) **not-enrolled**: The state for an initiator before it sends the first ACCESS CONTROL OUT command with ACCESS ID ENROLL service action to the access controls coordinator. Also the state for an initiator following successful completion of an ACCESS CONTROL OUT command with CANCEL ENROLLMENT service action (see 9.3.3.5);
b) **enrolled**: The state for an initiator following successful completion of an ACCESS CONTROL OUT command with ACCESS ID ENROLL service action; or
c) **pending-enrolled**: The state for an enrolled initiator following:
    A) Events in the service delivery subsystem described in 9.3.1.12; or
    B) Successful completion of an ACCESS CONTROL OUT command with MANAGE ACL service action from any initiator with the flush bit set to one (see 9.3.3.2).

9.3.1.5.1.2 Not-enrolled state

[Editor's note: go to pending-enrolled on I_T nexus loss]

…

The access controls coordinator shall change an initiator from the enrolled or pending-enrolled state to the not-enrolled state in response to the following events:

a) Successful completion of the ACCESS CONTROL OUT command with CANCEL ENROLLMENT service action (see 9.3.3.5);
b) Successful completion of an ACCESS CONTROL OUT command with MANAGE ACL service action (see 9.3.3.2) that replaces the ACL entry for the enrolled AccessID as follows:
    A) If the NOCNCL bit (see 9.3.3.2.2) is set to zero in the ACCESS CONTROL OUT command with MANAGE ACL service action parameter data, the state shall change to not-enrolled; or
    B) If the NOCNCL bit is set to one, the state may change to not-enrolled based on vendor specific criteria~~; or~~

~~c) Power cycles or target resets based on vendor specific criteria (see 9.3.1.12).~~.

[Editor's note: this is just an e.g. so need not be complete]

An enrolled initiator may find itself in the not-enrolled state as a result of actions taken by a third-party (e.g., an ACCESS CONTROL OUT command with MANAGE ACL service action performed by another initiator or a ~~target reset~~logical unit reset). The purpose of placing an enrolled initiator in the not-enrolled state in response to these events is to give the initiator an indication that the ACE defining its logical unit access has changed. One consequence of changes in an ACE is that previous relationships between logical units and LUN values may no longer apply.

…

9.3.1.5.1.4 Pending-enrolled state

The access controls coordinator shall place an initiator in the pending-enrolled state only if that initiator currently is in the enrolled state, and in response to the following:

a) ~~Any event in the service delivery subsystem that causes the access controls coordinator to question whether an initiator in the enrolled state has changed its AccessID (e.g., a process or port logout in Fibre Channel, or a hard bus reset for parallel SCSI)~~a logical unit reset;

aa) I_T nexus loss involving that initiator;

b) Successful completion of an ACCESS CONTROL OUT command with MANAGE ACL service action where the FLUSH bit is set to one in the parameter data; or

~~c) Optionally after a TARGET RESET task management function, as described in 9.3.1.12.~~

…

9.3.1.6 Granting and revoking access rights

9.3.1.6.2 Proxy access

9.3.1.6.2.1 Proxy tokens

…

~~Power cycles and target resets~~Power cycles, hard resets, logical unit resets, and I_T nexus losses shall not affect the validity and proxy access rights of proxy tokens (see 9.3.1.12). A proxy token shall remain valid and retain the same proxy access rights until one of the following occurs:

…

9.3.1.6.2.2 Proxy LUNs

…

Once assigned, a proxy LUN shall remain valid until one of the following occurs:

a) The third party releases the proxy LUN value using the ACCESS CONTROL OUT command with RELEASE PROXY LUN service action (see 9.3.3.12);

~~b) An event in the service delivery subsystem causes the access controls coordinator to question whether the third party initiator that created the proxy LUN value has changed and may no longer be in possession of the proxy token);~~

c) The proxy token is made invalid as described in 9.3.1.6.2.1; or

d) A ~~power cycle or target reset occurs~~ logical unit reset or I_T nexus loss of the I_T nexus used to create the proxy LUN (see 9.3.1.12).

If the third party believes that the invalidation of a proxy LUN value is temporary, it may reissue the ACCESS CONTROL OUT command with ASSIGN PROXY LUN service action in an attempt to re-establish its proxy access rights. The access controls coordinator shall process the request as described in 9.3.1.6.2.1 without reference to any previous assignment of the proxy LUN value.

9.3.1.8 The management identifier key

9.3.1.8.2 Overriding the management identifier key

9.3.1.8.2.2 The override lockout timer

After a logical unit reset ~~reset or power cycle~~, the override lock timer shall be set to the initial override lockout timer value within ten seconds of the non-volatile memory containing the initial override lockout timer value becoming available.

…

9.3.1.12 Access controls information persistence and memory usage requirements

If a SCSI target device supports ~~the~~ access controls, then the SCSI target device shall contain an access controls coordinator that shall maintain the following information in nonvolatile memory:

a) Whether access controls are enabled or disabled; and

b) The access controls data that is described as persistent across power cycles ~~and~~, hard resets, and logical unit resets in table 291 and table 292.

If the access control coordinator's nonvolatile memory is not ready and the access controls coordinator is unable to determine that access controls are disabled, the device servers for all logical units shall terminate all commands except INQUIRY and REQUEST SENSE commands with a CHECK CONDITION status, the sense key shall be set to NOT READY and additional sense data shall be set as described in table 167 (see 7.31).

[Editor's note: since table 291 makes it mandatory to remember enrollment states in NVRAM, the previous choice - a) pending-enrolled if such states are remembered or b) not-enrolled if they are not - is no longer necessary. Just go to pending-enrolled.]
Following a ~~power cycle or~~ logical unit reset ~~event~~, all previously enrolled initiators shall be placed in the pending-enrolled ~~same enrollment~~ state ~~and that state shall be one of the following:~~
~~a) Pending-enrolled (see 9.3.1.5.1.4); or~~
~~b) Not-enrolled (see 9.3.1.5.1.2).~~.

Following I_T nexus loss, a previously enrolled initiator shall be placed in the pending-enrolled state.

The information shown in table 291 shall be maintained by the access controls coordinator.

Table 291 - Mandatory access controls resources
Table 292 - Optional access controls resources
column heading in each table: Persistent ~~A~~across ~~P~~power ~~C~~cycles ~~and Resets~~, hard resets, and logical unit resets

Suggested changes to SBC-2

[Editor's note: Wherever "logical unit reset" is listed, consider whether "I_T nexus loss" is also appropriate.]

4.2.1.5 Initialization

Direct-access block devices may require initialization prior to write or read operations. This initialization is performed by a FORMAT UNIT command. Parameters related to the geometry and performance characteristics may be set with the MODE SELECT command prior to the format operation. Some block devices are initialized by means not specified in this standard. The time when the initialization occurs is specific to the implementation of the direct-access block device.

Block devices using a non-volatile medium may save the parameters and only need to be initialized once. However, some mode parameters may need to be initialized after each logical unit reset. A catastrophic failure of the direct-access block device may require the FORMAT UNIT command to be reissued.

Block devices that use a volatile medium may need to be initialized after each logical unit reset prior to the execution of read or write operations. Mode parameters may also need initialization after logical unit resets or I_T nexus losses.

…

4.2.1.13 Error reporting

If any of the following conditions occur during the execution of a command, the command shall be terminated with CHECK CONDITION status and the sense key shall be set to the appropriate sense key with the appropriate additional sense code for the condition. Some errors may occur after the completion status has already been reported. For such errors, SPC-3 defines a deferred error reporting mechanism. Table 4 illustrates some error conditions and the applicable sense keys. The list does not provide an exhaustive enumeration of all conditions that may cause the CHECK CONDITION status.

Table 4 - Example error conditions

Condition	Sense key
Invalid logical block address	ILLEGAL REQUEST
Unsupported option requested	ILLEGAL REQUEST
Logical unit reset, I_T nexus loss, or medium change since last command from this application client	UNIT ATTENTION
Self diagnostic failed	HARDWARE ERROR
Unrecovered read error	MEDIUM ERROR or HARDWARE ERROR
Recovered read error	RECOVERED ERROR
Overrun or other error that might be resolved by repeating the command	ABORTED COMMAND
Attempt to write on write protected medium	DATA PROTECT

[Editor's note: fix same kind of table in 4.3.3 Table 6, and 4.4.5 Table 8]

4.2.3.7 XOR data retention requirements

The target shall retain XOR data while awaiting retrieval by an XDREAD command until performing one of the following events: a matching XDREAD command, logical unit reset, I_T nexus loss involving the initiator which sent the XDWRITE command, CLEAR TASK SET, ABORT TASK if the task matches the pending XDREAD, or ABORT TASK SET.

5.1.23 START STOP UNIT command

…

If the START STOP UNIT command is issued with the POWER CONDITIONS field set to 1h, 2h, or 3h the block device shall:

a) change power conditions only on receipt of another START STOP UNIT command or a logical unit reset;
b) suspend any Power Condition timers (see SPC-3) that are active on receipt of the START STOP UNIT command until another START STOP UNIT command is received that returns control of the power condition to the block device or a logical unit reset occurs;
c) terminate any command received that requires more power than allowed by the START STOP UNIT command's most recent power condition setting with a CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOW POWER CONDITION ACTIVE.

Suggested changes to SSC-2

[Editor's note: Wherever "logical unit reset" is listed, consider whether "I_T nexus loss" is also appropriate.]

4.2.10 Error reporting

If any of the following conditions occur during the processing of a command or if a deferred error prevented the command from processing, the device server shall return CHECK CONDITION status. The appropriate sense key and additional sense code should be set. Table 1 illustrates some error conditions and the applicable sense keys. Table 1 does not provide an exhaustive enumeration of all conditions that may cause the CHECK CONDITION status.

Table 4 - Error conditions and sense keys

Condition	Sense key
Unsupported option requested.	ILLEGAL REQUEST
Logical unit reset, I_T nexus loss, or medium change since last command from this initiator.	UNIT ATTENTION
…	

4.2.11.3 Software Write Protect for the device server

[Editor's note: let SPC-3 define mode page behavior]

Software Write Protect for the device server controls write protection for the device server. This method of write protect is optionally controlled from the Control mode page (see SPC-3) or the SWP bit in the Device Configuration mode page (see 8.3.3). Either or both methods may be implemented by the device server. If both methods are implemented, each control bit is independently set. Software write protection exists if either bit is non-zero. The state of software write protect for the device server shall not be recorded on media. The value of the SWP bit may be altered by the application client (if the SWP bit is changeable). ~~The state of each control bit shall be set to its default state after a logical unit reset.~~

4.2.11.4 Associated Write Protect

Associated Write Protect controls write protection for the currently mounted volume as long as the current volume is mounted. The associated write protect state is controlled by the ASOCWP bit in the Device Configuration mode page (see 8.3.3). Associated write protection exists if the ASOCWP bit is non-zero. Associated write protection may be altered by the application client (if the ASOCWP bit is changeable) if a volume is mounted. If a volume is de-mounted ~~or after a logical unit reset occurs~~, associated write protection shall be removed.

4.2.11.5 Persistent Write Protect

Persistent Write Protect controls write protection for the currently mounted volume. The persistent write protect state is controlled by the PERSWP bit in the Device Configuration mode page (see 8.3.3). If enabled, persistent write protection shall exist for the mounted volume until disabled by the application client. The state of persistent write protection shall be recorded with the volume. The device server shall report the PERSWP bit as one when a mounted volume is marked with persistent write protection. If a volume is de-mounted ~~or after a logical unit reset occurs~~, the device server shall report the PERSWP bit as zero prior to the mounting of a volume. The means for recording the state of persistent write protect for the volume may be specified in the applicable recording format standard or be vendor-specific.

4.2.11.6 Permanent Write Protect

Permanent Write Protect controls write protection for the currently mounted volume. The permanent write protect state is controlled by the PRMWP bit in the Device Configuration mode page (see 8.3.3). If enabled, permanent write protection shall exist for the mounted volume until disabled by a vendor-specific method. The state of permanent write protection shall be recorded with the volume. The device server shall report the PRMWP bit as one when a mounted volume is marked with permanent write protection. If a volume is de-mounted ~~or after a logical unit reset occurs~~, the device server shall report the PRMWP bit as zero prior to the mounting of a volume. The means for recording the state of permanent write protect for the volume may be specified in the applicable recording format standard or be vendor-specific. Permanent write protection shall

not be removed by a MODE SELECT command using the PRMWP bit. Methods to remove this protection may or may not exist and are vendor-specific.
…

4.2.16.2 TapeAlert log sense format
The TapeAlert interface to the tape drive is based on a LOG SENSE page (2Eh) containing 64 one-byte flags (see table A.1). The specific conditions for any one flag to be set and cleared are vendor-specific. The minimum subset of flags and basic implementation guidelines are described in table 7.

The TapeAlert data is event based and the page control bits in the LOG SENSE command are not applicable and shall be ignored by the device.

[Editor's note: clear per-initiator log pages on I_T nexus loss]
Each flag shall be cleared in the following circumstances:
a) at drive power on;
b) after the TapeAlert log page is read - note in multi-initiator environments the TapeAlert flags shall be cleared on a per-initiator basis such that set flags are still visible to other initiators;
c) when the specified corrective action has been taken (such as using a cleaning cartridge);
d) on ~~hard reset~~logical unit reset or I_T nexus loss; or
e) on LOG SELECT reset.

7.10 SET CAPACITY command
[Editor's note: capacity is not affected by I_T nexus loss]
The SET CAPACITY command (see table 46) sets the available medium for the currently mounted volume to a proportion of the total capacity of that volume. Any excess space shall be unavailable on the volume after successful completion of this command until changed by a new SET CAPACITY command. This change shall persist through power cycles, hard resets, logical unit resets, I_T nexus losses, and unloading or reloading of the volume. Other vendor-specific actions such as physical erasure may reset the total capacity of the volume. The method for recording the available capacity and other marks needed to manage the resulting capacity for volume interchange may be specified in a recording format standard or may be vendor-specific.

8.2.2 Sequential-Access Device log page
…
A non-zero value of the cleaning required parameter specifies a condition requiring cleaning has been detected and a subsequent cleaning cycle has not been completed. The cleaning required parameter shall be persistent across power cycles, hard resets, logical unit resets ~~and power cycles~~, and I_T nexus loss.

8.3.1 Mode parameters overview
This subclause defines the descriptors and pages for mode parameters used with sequential-access devices.

The mode parameter list, including the mode parameter header and mode block descriptor, are described in SPC-3.

The device-specific parameters contained in the mode parameter header, mode block descriptor values, and Data Compression mode page shall be retained following ~~a reset condition~~hard reset, logical unit reset, or I_T nexus loss ~~(e.g., Target Reset, SCSI Logical Unit Reset, Fibre Channel Reset LIP or PLOGI)~~.

NOTE 44 This is to facilitate continued operation for applications such as backup/restore following a reset event.

…
For the MODE SENSE command, the DENSITY CODE field reflects the current operating density of the logical unit. If a current operating density has not been selected, either because no medium is

mounted or because the density of the installed medium has not been determined, the DENSITY CODE field should be set to the principal density code value (see 3.1.45). For some logical units, the principal density code value returned in response to a MODE SENSE command may change dynamically to match the most recently detected density. The DENSITY CODE value returned in response to a MODE SENSE command shall be determined as follows:

a) following a ~~power on or reset condition occuring~~logical unit reset occurring while not ready, the device server shall report the principal density;

b) following a unit attention condition for a not-ready-to-ready transition, the device server shall:
   A) report the principal density if no attempt has been made by the logical unit to determine the density;
   B) report the principal density if the logical unit is unable to automatically determine the density from the medium; or
   C) report the current medium density if the logical unit has determined the density from the medium.

c) following a successful read operation at or after beginning-of-medium, the device server shall report a density code value reflecting the recorded density of the medium. For some implementations, the logical unit may automatically determine this value from the medium. For devices not capable of automatic density determination, the principal density is reported if the density code value is not provided by the preceding MODE SELECT command;

d) following an unsuccessful read operation or a successful write operation, while at beginning-of-partition, the device server shall:
   A) report a density code value as described for item B) if a previous MODE SELECT command has not established a density code for the currently mounted volume; or
   B) report a density code value as provided by the last successful MODE SELECT command for the currently mounted volume.

e) following a successful unload operation the device server shall report the most recent density code value as determined by items B) through D) above; or

f) following a ~~reset condition~~logical unit reset occurring while ready, the device server shall retain knowledge of the density code as determined by items B) through D) above.

8.3.4 Medium Partition mode page(1)

…

[Editor's note: this seems to be a shared mode page, so I_T nexus loss of all I_T nexuses clears it.]

A partition on format (POFM) bit of one specifies the MODE SELECT command shall not cause changes to the partition sizes or user data, either recorded or buffered. If POFM is set to one, actual media partitioning occurs when the device server receives a subsequent FORMAT MEDIUM command (see 7.2). When the FORMAT MEDIUM command partitions the media, it shall do so based on the contents of the mode data for Medium Partition mode pages (1-4). If POFM is set to one, field values specified by a MODE SELECT command for all Medium Partition mode pages (1-4) shall not be changed by the device server before the media is unloaded or until a logical unit reset or until I_T nexus loss of all I_T nexuses. Some field checking may be performed by the MODE SELECT command. However, there is no guarantee that any subsequent partitioning during a FORMAT MEDIUM command will complete with no errors.

Table 1 - TapeAlert log page parameter codes

…

1Eh Hardware A O C
Recommended application client message:
The tape drive has a hardware fault:
1. Eject the tape or magazine.
2. Reset the drive.
3. Restart the operation.
Probable cause:
The drive has a hardware fault that requires reset to recover.

Suggested changes to SRP

[Editor's note: Cris Simpson will define "I_T nexus loss" as a logout (SRP or IB level) - see 02-127]

Suggested changes to iSCSI

[Editor's note: Mallikarjun Chadalapaka will define "I_T nexus loss" as session logout]

Suggested changes to SPI-5

[Editor's note: George Penokie will define "I_T nexus loss" as when a reselection timeout occurs.]

Suggested changes to FCP-3

[Editor's note: Dave Peterson will define "I_T nexus loss" as logout]
[Editor's note: Dave will remove all the SCSI items from the Clearing effects table as suggested below.]

4.9 Clearing effects of task management, FCP, FC-FS, and FC-AL-2 actions

Tables 4 and 5 summarize the FCP target objects that are cleared as a result of Fibre Channel Link actions and SCSI operations, respectively. A 'Y' in the corresponding column of either table indicates the object is cleared to its default, saved, or initial value within the device upon successful completion of the specified action. The clearing actions are applicable only to Sequences and Exchanges associated with Fibre Channel protocol actions. Sequences and Exchanges associated with other actions follow rules specified in FC-FS or other relevant protocol standards. An 'N' in the corresponding column indicates the object is not affected by the specified action. A '-' in the column indicates that the action is not applicable. Rows indicating an effect for all initiator ports have the specified effect on all ports, regardless of the link that attaches the initiator port to the target.

[Editor's note: if CRN is made FCP-specific, then FCP-3 should continue to define its clearing effects. If it stays in SAM-2, FCP-3 should probably be silent.]

Table 4 - Clearing effects of link related functions

								ABTS
								TPRLO
							PRLI, PRLO	
						LOGO, PLOGI		
					Failed discovery after OLS			
				Failed discovery after LIP				
			Reset, LIP(y,x)					
		Target power cycle						
Target object								
PLOGI parameters For all logged-in initiator ports Only for initiator port associated with the action								
Open FCP Sequences Terminated For all initiator ports with open FCP Sequences Only for initiator port associated with the action Only for FCP Sequences associated with Aborted FCP								
Exchanges								
Login BB_Credit_CNT For all Logged-In L_Ports For transmitting L_Port only								
Hard Address Acquisition Attempted								
PRLI parameters cleared For all logged-in initiator ports Only for N_Port or L_Port associated with the action								
~~Open Tasks (FCP Exchanges) Aborted~~ ~~All tasks for all initiator ports with open tasks~~								

~~All tasks, only for initiator port associated with the action~~ ~~Only for specified task~~								
~~Target mode page parameters restored from saved pages~~ ~~(when saved pages are supported, or default mode page parameters when saved pages are not supported)~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~								
~~Pre-existing ACA, Unit Attention (7), and Deferred error conditions cleared~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~								
~~Device reservations~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~								
~~Persistent device reservations (10)~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~								
CRN (Command Reference Number) (set to one) For all initiator ports Only for initiator port associated with the action								
~~Prevent Allow Medium Removal state cleared to allow removal~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~								
~~Buffered data for XOR, EXTENDED COPY, COPY~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~								
~~Access controls data~~ ~~AccessID enrollment state to pending enrolled~~ ~~For all SCSI initiators in enrolled state~~ ~~Only for SCSI initiator port initiating action in enrolled state~~								

Table 5 - Clearing effects of initiator actions

	ABORT TASK SET	CLEAR TASK SET	LOGICAL UNIT RESET	TARGET RESET
Target object				
PLOGI parameters For all logged-in initiator ports Only for initiator port associated with the action				
Open FCP Sequences Terminated For all initiator ports with open FCP Sequences Only for initiator port associated with the action Only for FCP Sequences associated with Aborted FCP				
Exchanges				
Login BB_Credit_CNT For all Logged-In L_Ports For transmitting L_Port only				
Hard Address Acquisition Attempted				
PRLI parameters cleared For all logged-in initiator ports Only for N_Port or L_Port associated with the action				
~~Open Tasks (FCP Exchanges) Aborted~~ ~~All tasks for all initiator ports with open tasks~~ ~~All tasks, only for initiator port associated with the action~~				

~~Only for specified task~~				
~~Target mode page parameters restored from saved pages (when saved pages are supported, or default mode page parameters when saved pages are not supported)~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~Pre-existing ACA, Unit Attention (7), and Deferred error conditions cleared~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~Device reservations~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~Persistent device reservations (10)~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~CRN (Command Reference Number) (set to one)~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~Prevent Allow Medium Removal state cleared to allow removal~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~Buffered data for XOR, EXTENDED COPY, COPY~~ ~~For all initiator ports~~ ~~Only for initiator port associated with the action~~				
~~Access controls data~~ ~~AccessID enrollment state to pending enrolled~~ ~~For all SCSI initiators in enrolled state~~ ~~Only for SCSI initiator port initiating action in enrolled state~~				