To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
and Mallikarjun Chadalapaka (cbm@rose.hp.com)
Date: 18 July 2002
Subject: T10/02-232r2 SAM-3 SPC-3 SBC-2 SSC-2 Clearing effects of I_T nexus loss

Revision History

02-143r0 (16 April 2002) first revision
02-143r1 (19 April 2002) incorporated comments from Kevin Butt (IBM) on microcode clearing and comments from the clearing effects working group on SAM-2 changes.
02-143r2 (1 May 2002) incorporated comments from May CAP WG. The WG recommended the T10 plenary approve the changes to SPC-3 except for the mode page changes. Work will continue on the SAM-2, SBC-2, and SSC-2 changes.
02-232r0 (19 June 2002) removed accepted SPC-3 changes (incorporated spc3r07) and modified the SAM-2 changes based on internal discussion.
02-232r1 (14 July 2002) incorporated comments from the 6/21 Clearing effects conference call. The group recommended directing the changes to SAM-3 and removing WAKEUP.
02-232r2 (18 July 2002) incorporated comments from the July CAP WG – removed SSC-2 changes (Dave Peterson will drive as letter ballot comments), made mode pages impervious to I_T nexus loss, describe mode parameter header and block descriptors.

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
01-328r6 - Response to T10 letter ballot on SRP
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)
02-143r2 - Previous version of this proposal

Overview

Thanks to 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 Protocol standards are left with large tables listing behavior for commands; these tables are always incomplete and out of date.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-3.
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-3 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 protocols.
e) Cleans up relationships of these actions in SAM-3.

The affected items in SAM-3, SPC-3, SBC-2, and SSC-2 include:

- Background operations
- Deferred errors
- CRN value
- Sense data preservation
- Unit attention generation
- Tasks aborted
- Clear CA or ACA
- Reservations
- Persistent reservations (registration and reservation data, PRGENERATION value)
- Log pages (shared, per-initiator port)
- Mode pages (shared, per-initiator port, per-I_T)
- Asymmetric logical unit access state
- Power condition state
- Alias table associations
- INQUIRY data available without incurring media access delays
- Prevention of medium removal
- Extended Copy status data
- REPORT LUNS data available without incurring media access delays
- Downloaded microcode (effective time, discard partial)
- Mode parameter block descriptors number of blocks field
- Informational Exceptions Control timer/counter
- Access Controls (enabled, TransportID definition, state, proxy tokens, proxy LUN, data persistence)
- XOR data
- Tape capacity persistence
- Tape cleaning requirement persistence

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. FCP-2 entries marked "CONFLICT" are either complete reversals of the FCP-2 recommendation, or a clearing recommendation where FCP-2 was silent.

NOTE: SPC-3 entries accepted in 02-1345r2 have been removed - this is no longer a complete table.

Table 1. I_T nexus loss summary

What	Where	When (current phrase(s))	Proposed changes	FCP-2 clearing effects table
SAM-2:				
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	may be aborted by power on, hard resets, or logical unit resets, and shall not be aborted by I_T nexus loss	not mentioned
Deferred errors cleared	nowhere - SAM-2 5.5 should describe	not mentioned	may be cleared by power on, hard reset, logical unit resets and by I_T nexus losses of the initiator port that sent the command that spawned the background operation, but should not be cleared by I_T nexus loss.	cleared **[CONFLICT]**
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
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

What	Where	When (current phrase(s))	Proposed changes	FCP-2 clearing effects table
SAM-2:				
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
SPC-3:				
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	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	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	not mentioned **[CONFLICT]**
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]**
SBC-2:				
Mode parameters may also need initialization	SBC-2 4.2.1.5	[after logical unit resets]	after logical unit resets	[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

Suggested Changes to SAM-3

3.1.x background operation: Operation started by a command that operates after the task containing the command is no longer in the task set (see 5.5).

3.1.x deferred error: Errors generated by background operations (see 5.5).

3.1.39 hard reset: A ~~target~~ SCSI port action in response to a power on or a reset event in which the ~~target~~ SCSI port performs the operations described in ~~5.8.6~~N.1.3.

3.1.x I_T nexus loss: A SCSI port action in response to a hard reset or an I_T nexus loss notification event in which the SCSI port performs the operations described in N.1.5.

3.1.x I_T nexus loss notification event: A protocol-specific event that triggers an I_T nexus loss.

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~~N.1.4.

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~~. Logical unit reset events include processing the LOGICAL RESET task management function, processing the TARGET RESET task management function, and hard reset.

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

3.1.p power on: A SCSI device action in response to a power on event.

3.1.p power on event: Power being applied to a SCSI device, triggering a power on in the SCSI device.

3.1.84 reset event: A protocol specific event that triggers a hard reset ~~from a SCSI device as described in 5.8.6.~~.

~~**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 a SCSI target port as described in SPC-3.~~

4.14 The SCSI model for distributed communications

The SCSI model for communications between distributed objects is based on the technique of layering. In the layering technique, the initiator and target I/O systems are viewed as being logically composed of the ordered set of subsystems represented for convenience by the vertical sequence shown in figure 25.

..

[Editor's note: 6/21 call asked that something be added to figure 27 to cover the new notification services. However, these would operate backwards from the ULP/LLP described in this figure and require significant rewrite. We suggest this be deferred to a future proposal.]

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 nexus involving the SCSI port after the SCSI 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.
[Editor's note: if a CPQ (HP) letter ballot comment is accepted, CRN will be removed from SAM-2 and left to FCP to define, eliminating the need for this change.]
…

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);
    B) Any additional sense code whose ADDITIONAL SENSE CODE field contains 29h (e.g., POWER ON, RESET, OR BUS DEVICE RESET OCCURRED; POWER ON OCCURRED; SCSI BUS RESET OCCURRED; BUS DEVICE RESET FUNCTION OCCURRED; DEVICE INTERNAL RESET; TRANSCEIVER MODE CHANGED TO SINGLE-ENDED; ~~or~~ TRANSCEIVER MODE CHANGED TO LVD; or 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 directed to the logical unit processing the task.

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 specific 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

Write commands start background operations if a write-back cache is enabled, too.]

Some commands (e.g., commands with immediate bits like SEND DIAGNOSTIC, or write commands when a write cache is enabled) start background operations that operate after the task containing the command is no longer in the task set. These operations may be aborted by power on, hard resets, or logical unit resets. These operations shall not be aborted by I_T nexus loss.

Background operations may generate deferred errors, which are reported with autosense or in sense data for a subsequent completed command (see SPC-3). A deferred error may be cleared by power on, hard reset, logical unit reset, or by I_T nexus loss involving the initiator port that sent the command that spawned the background operation, if the device server has retained that information. Deferred errors should not be cleared by I_T nexus loss.

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:
a) it is transferred ~~by one of the methods listed below ;or until~~
b) another task from that initiator is entered into the task set;
c) a logical unit reset occurs; or
d) I_T nexus loss involving the initiator occurs.

The sense data may be transferred to the initiator through any of the following methods:
a) The REQUEST SENSE command (see SPC-2);
b) An asynchronous event report (see 5.8.4.2); or
c) Autosense delivery (see 5.8.4.3).
The following clauses describe the last two transfer methods.
…

5.8.5 Unit Attention condition

Each logical unit shall generate a unit attention condition whenever it receives a logical unit reset event or an I_T nexus loss notification event ~~whenever the logical unit has been reset~~(i.e., on a power on, hard reset, logical unit reset, or I_T nexus loss) ~~-~~as described in N.1.1~~5.8.7 or by a power-on reset~~. In addition, a logical unit shall generate a unit attention condition for each initiator port whenever one of the following events occurs:
a) A removable medium may have been changed;
b) The mode parameters in effect for this initiator port have been changed by another initiator port;
c) The version or level of microcode has been changed;
d) Tasks for this initiator port were cleared by another initiator port;
e) INQUIRY data has been changed;
f) The logical unit inventory has been changed;
g) The mode parameters in effect for the initiator port 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 port.

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~~unit attention condition with an additional sense code of POWER ON OCCURRED followed by one with an additional sense code of ~~a microcode change condition~~MICROCODE HAS BEEN CHANGED).
…

N Notification event services

The SCSI transport protocol layer uses these notification event services to notify the SCSI application layer of various protocol-specific events:

a) Nexus Lost(IN(I_T Nexus)); and
b) Transport Reset(IN(SCSI_Port)).

The services are described below.

Nexus Lost(IN(I_T Nexus))

Input arguments:
I_T Nexus: the specific nexus that is lost.

This notification event service is specified for both initiator devices and target devices. Each protocol defines the conditions, if any, that shall lead to generating this notification event to the application layer.

Transport Reset(IN(SCSI Port))

Input arguments:
SCSI_Port: the specific SCSI port in the SCSI device which is subject to hard reset.

This notification event service is specified for both initiator devices and target devices. Upon receiving this notification event, the SCSI application layer shall behave as if a hard reset has occurred (see N.1.3). Each protocol defines the conditions, if any, that would lead to generating this notification event to the application layer.

N.1 Power on, hard reset, logical unit reset, and I_T nexus loss

N.1.1 Overview

Figure 1 shows the relationships between these actions:

a) power on (see N.1.2);
b) hard reset (see N.1.3);
c) logical unit reset (see N.1.4); and
d) I_T nexus loss (see N.1.5),

and their associated events.

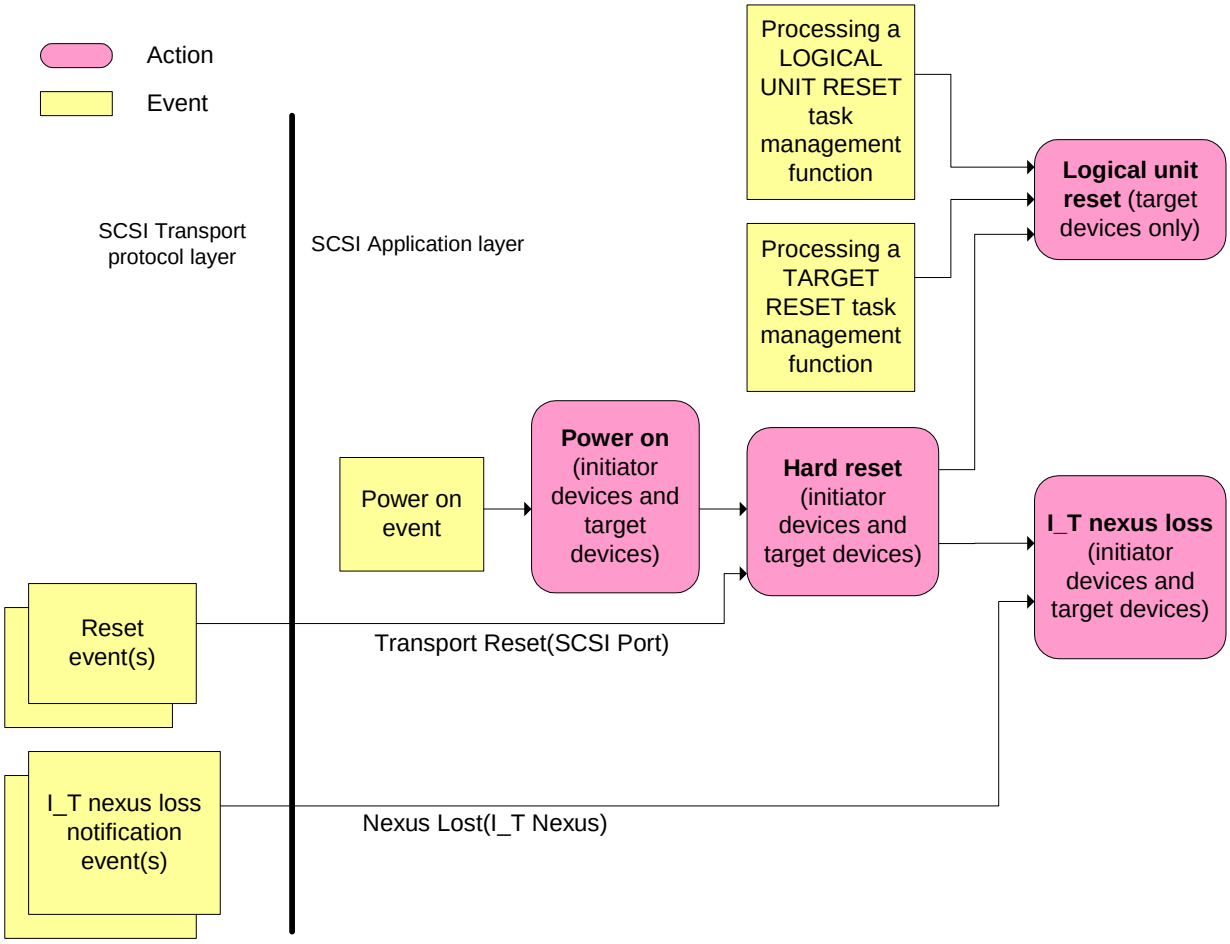

Figure 1. Power on, hard reset, logical unit, and I_T nexus loss relationships

Table xx shows the additional sense code that shall be used when a unit attention is reported for each of these actions by a logical unit. The protocol may define a more specific additional sense code for reset events than SCSI BUS RESET OCCURRED (e.g., TRANSCEIVER MODE CHANGED TO LVD). The actions are sorted from least-specific to most-specific. The most-specific action known to the logical unit shall be reported.

Table xx. Additional sense codes for power on, hard reset, logical unit, and I_T nexus loss

Action	Additional sense code
Logical unit cannot distinguish between the actions	POWER ON, RESET, OR BUS DEVICE RESET OCCURRED
Power on	POWER ON OCCURRED or DEVICE INTERNAL RESET
Hard reset	SCSI BUS RESET OCCURRED or protocol-specific (e.g., TRANSCEIVER MODE CHANGED TO SINGLE-ENDED or TRANSCEIVER MODE CHANGED TO LVD)
Logical unit reset	BUS DEVICE RESET FUNCTION OCCURRED
I_T nexus loss	I_T NEXUS LOSS OCCURRED

The logical unit may use the I_T NEXUS LOSS OCCURRED additional sense code if:

a) the initiator port retrieving the sense data is the initiator port that was involved in the I_T nexus loss, and the logical unit has maintained all per-initiator port state information since the I_T nexus loss; and
b) the I_T nexus being used to retrieve the sense data is the same I_T nexus that was lost, and the logical unit has maintained all per-I_T state information since the I_T nexus loss.

Otherwise, it shall use one of the less-specific additional sense codes (e.g., POWER ON OCCURRED).

N.1.2 Power on

[Editor's note: define this here so each protocol standard need not include it in its list of protocol-specific reset events. The additional sense code is defined in N.1.1.]
Power on is a SCSI device action in response to a power on event. When a SCSI device is powered on, it shall cause a hard reset.

Power on applies to both SCSI initiator devices and SCSI target devices.

~~5.8.6~~ N.1.3 Hard 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 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 protocol-specific actions in response to reset events. Protocols may include reset events that have no SCSI effects (e.g. Fibre Channel non-initializing LIP).

Hard reset applies to both SCSI initiator devices and SCSI target devices.

~~The~~ A ~~target~~ SCSI 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~~N.1.4.

~~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 performs a hard reset, it should terminate all its outstanding Execute Command remote procedure calls with SERVICE DELIVERY OR TARGET FAILURE.
[Editor's note: if it chooses not to, the commands will eventually time out (if it enforces any ULP timeouts). Thus the rule need not be a "shall."]

~~5.8.7~~ N.1.4 Logical unit 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~~N.1.3).

Logical unit reset applies only to SCSI target devices.

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

[Editor's note: Reservation and mode page effects should be left to the command standards to define. Mentioning SAM-3 effects like tasks being aborted *is* appropriate here. The "operating mode to the appropriate initial conditions" phrase might be worth keeping out of item d).]
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 N.1.1); ~~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. Move this into the list:]
~~In addition to the above, the logical unit shall p~~g) Perform any additional functions required by the applicable command set standards.

N.1.5 I_T nexus loss

An I_T nexus loss is a SCSI port action in response to a hard reset or an I_T nexus loss notification event. An I_T nexus loss notification event is an indication from the SCSI transport protocol that an I_T nexus no longer exists. Protocols may define I_T nexus loss notification events.

I_T nexus loss applies to both SCSI initiator devices and SCSI target devices.

When a target port detects an I_T nexus loss, each logical unit to which the I_T nexus has access shall process the I_T nexus loss by performing 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 N.1.1); and
d) Perform any additional functions required by the applicable command set standards.

If the logical unit retains state information for the I_T nexus that is lost, on subsequent I_T nexus re-establishment the logical unit should return a unit attention with an additional sense code of I_T NEXUS LOSS OCCURRED.

If the logical unit does not retain state information for the I_T nexus that is lost, it shall consider the subsequent I_T nexus re-establishment, if any, as the formation of a new I_T nexus for which there is no past history (e.g., return a unit attention with an additional sense code of POWER ON OCCURRED).

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.

Suggested changes to SPC-3

…

3.1.34. **hard reset**: A SCSI port action in response to a power on or a reset event in which the SCSI port performs the operations described in SCSI Architecture Model-2.

3.1.38 I_T nexus loss: A SCSI port action in response to a hard reset or an I_T nexus loss notification event in which the SCSI port performs the operations described in SCSI Architecture Model-2.

3.1.x I_T nexus loss notification event: ~~An~~ A protocol specific event that triggers I_T nexus loss as described in SAM-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 as described in SCSI Architecture Model-2.

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

~~**3.1.63 power on: Power being applied to a SCSI device.**~~

3.1.63 power on: A SCSI device action in response to a power on event.

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

3.1.72 **reset event**: ~~An~~ A protocol-specific event that triggers a hard reset as described in SCSI Architecture Model-2.

~~**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.~~

…

5.8 Power conditions

The optional power conditions permit the application client to modify the behavior of a target port and/or logical unit in a manner that may reduce power consumption. There is no notification to the application client that one of the power conditions has been entered. Power conditions may be controlled by the START STOP UNIT command for some device types (see SBC-2 or RBC) or the Power Condition mode page for all device types (see 8.4.11). If both methods are being used on the same logical unit then a power condition request made by a START STOP UNIT command shall override the Power Condition mode page's power condition.

No power condition shall affect the service delivery subsystem.

The logical unit power conditions are described in table 15.

Table 15 — Logical unit power conditions

Logical unit power condition	Description
active	The device server is capable of responding to all its supported commands including media access requests, and operations complete processing in the shortest time compared to the other power conditions.
idle	The device server is capable of responding to all its supported commands

	including media access requests. However, a device server in the idle condition may take longer than in the active logical unit power condition to complete processing a command because it may have to activate some circuitry.
standby	The device server is not capable of processing media access commands (e.g., the spindle is stopped).
~~sleep~~	~~The device server is not capable of accepting or processing commands.~~

~~The logical unit sleep power condition shall only be supported on logical units accessed through target ports using protocols that define wakeup events.~~

~~Table 16 — Target port power conditions~~

~~Target port power condition~~	**~~Description~~**
~~active~~	~~Logical units to which the target port has access that are not in the sleep logical unit power condition are capable of accepting commands routed from the target port.~~
~~idle~~	~~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.~~

~~The target port sleep power condition shall only be supported on target ports using protocols that define wakeup events.~~

Block devices that contain cache memory shall implicitly perform a SYNCHRONIZE CACHE command (see SBC-2 or RBC) for the entire medium prior to entering any power condition that prevents access the media (e.g., the spindle being stopped).

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.

~~The state diagram for the target port power conditions is shown in figure 5.~~

~~[]~~

~~Figure 5 — Target port power conditions state diagram~~

~~The target port power conditions state transitions are:~~

~~**TP0:TP1:** When all logical units in a target port reach the logical unit sleep power condition, the target port transitions to the target device sleep power condition.~~

~~**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.~~

The state diagram for the logical unit power conditions is shown in figure 6.

[]

Figure 6 — Logical Unit power conditions state diagram

[Editor's note: remove state LU8]

The logical unit power conditions state transitions (see table 17) are based on either automatic adjustments or application client controls.

State transition	Description	
	automatic adjustment	**application client control**
LU5:LU6	While in the active power condition, an idle time-out or a START STOP UNIT command with a power condition code of Ah causes a transition to idle.	While in the active or standby power condition, a START STOP UNIT command with a power condition code of 2h causes a transition to idle.
LU7:LU6	While in the standby power condition, any command that can be processed within the constraints of the idle power condition causes a transition to idle.	
LU6:LU7	While in the idle power condition, a~~n~~ standby time-out or a START STOP UNIT command with a power condition code of Bh causes a transition to idle.	While in the active or idle power condition, a START STOP UNIT command with a power condition code of 3h causes a transition to standby.
LU5:LU7	While in the active power condition, a standby time-out when the IDLE bit is set to zero causes a transition to standby.	
LU6:LU5 LU7:LU5	While in the idle or standby power conditions, any command that exceeds the power constraints of the idle power condition causes a transition to active.	While in the idle or standby power conditions, a START STOP UNIT command with a power condition code of 1h causes a transition to active.
~~LU6:LU8 LU5:LU8 LU7:LU8~~	~~While in the active, idle, or standby power conditions, a START STOP UNIT command with a power condition code of 5h causes a transition to the logical unit sleep power condition. When all logical units in a target port reach the logical unit sleep power condition, the target port transitions to the target device sleep power condition (see TP0:TP1 in figure 5).~~	
~~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.~~	

7.7 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 shall share mode parameter header and block parameter values across all initiator ports and I_T nexuses. I_T nexus loss shall not affect mode parameter header and block parameter values.

Logical units shall 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
shared	One copy of the mode page.
per-initiator port	Separate copies of the mode page for each initiator port.
per-I_T nexus	Separate copies of the mode page for each I_T nexus.

I_T nexus loss shall not affect mode page values.

On a logical unit reset, each mode parameter header, block parameter, and mode page shall revert to saved values if supported or default values if saved values are not supported.

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

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 (see 7.7); 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 (see 7.7).

…

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.

…

8.2 Log parameters

8.2.9 Self-Test Results log page

…

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,~~ or by issuing an exception command as defined in 5.4.3).

…

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

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 time 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 and the report counter across 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.

…

Suggested changes to SBC-2

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

~~3.1.30. **wakeup**: A target port returning from the sleep power condition to the active power condition (see SPC-3).~~

~~3.1.31. **wakeup event**: An event that triggers a wakeup from a target port as described in SPC-3.~~

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.

…

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 occurs:

a) a matching XDREAD command~~,~~;

b) logical unit reset~~,~~;

c) I_T nexus loss involving the initiator which sent the XDWRITE command

d) CLEAR TASK SET~~,~~;
d) ABORT TASK if the task matches the pending XDREAD~~,~~ or
e) 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.