To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 6 March 2008
Subject: 08-126r0 SPC-4 SBC-3 SAS-2.1 Power condition enhancements

Revision history
Revision 0 (6 March 2008) First revision

Related documents
sbc3r13 - SCSI Block Commands - 3 (SBC-3) revision 13
sas2r14 - Serial Attached SCSI - 2 (SAS-2) revision 14
spc4r12 - SCSI Primary Commands - 4 (SPC-4) revision 12
mmc6r01 - Multimedia Commands - 6 (MMC-6) revision 1
rbc-r10a - Reduced Block Commands (RBC) revision 10a
T13/ata8-acs-r4c - ATA Command Set (ATA8-ACS) revision 4c
01-134r2 - SAM-2 SPC-3 SPI-4 SBC-2 WAKEUP and reset cleanup (Rob Elliott, Compaq)
07-485r2 - SAT-2 Additional power management methods (Frederick Knight, NetApp)

Overview
With renewed interest in power management, and plans to add interface power management in SAS-2.1, the device level power management features (e.g., active/idle/standby power conditions) in SCSI should be enhanced.

1. Some SCSI standards currently define or previously defined a **Sleep** power condition to send a target into an “off” mode, ready for power to be removed. This mode consumes the least amount of power possible, and requires a reset to wake up.

The SCSI parallel interface (SPI-5) was the last T10 transport protocol to support this. It defined a WAKEUP task management function to wakeup from Sleep, defined as as a “bus reset condition,” not a special signal. Among non-T10 SCSI protocols, ATAPI supports it.

Three command sets have defined the sleep power condition. Entry is via the START STOP UNIT command, (which is defined in individual command sets, not in SPC), not the Power Condition mode page (in SPC):

a) MMC-n (Multi-Media Commands);
b) RBC (Reduced Block Commands); and
c) SBC-n (SCSI Block Commands).

NOTE 1 - SBC-1 defined the POWER CONDITION field set to 5h as “Place device into Sleep” condition. SBC-2 revisions 0, 1, and 2 mistakenly relocated that to 6h. Revision 3 restored it to 5h. 02-464r3 changed that value to obsolete in SBC-2 revision 9.

ATA8-ACS also defines a sleep power condition, and a SCSI equivalent is needed for good SCSI to ATA translation.

2. Some idle mode enhancements have also been deployed in the industry :

a) low-rpm idle. The disk keeps spinning at a reduced rate, providing a shorter recovery time than the full spin-down required by the standby power condition.
b) park heads. Reduce friction on the spinning disk during the idle power condition.

Proposal

1. Restore the sleep power condition to the SBC-3 START STOP UNIT command. This will also let SAT-2 provide a mapping into the ATA sleep state (see ATA8-ACS), which is implemented by some SATA drives but is currently unreachable via a SATL.

Editor’s Note 1: automatic link reset after the drive goes down after SLEEP command cannot cause a wakeup, or else we defeat the purpose. Requiring a hard reset solves that, but requires software to send a hard reset. If the initiator/expander loses power but the drive doesn’t, the

normal power up sequence will find the drive unresponsive.

2. Add the sleep power condition to the SPC-4 Power Conditions mode page, since nothing in this mode is particular to the peripheral device type, so the.

3. Add a SLEEP mode for SAS expander phys (intended for the upstream link to a host, not the downstream link to a drive that is asleep).

4. Enhance the idle power condition for block devices to support two variations:

c) low-rpm idle
d) park heads

07-485r2 proposes a SAFETY bit be added to the START STOP UNIT command CDB to control parking heads:

> ... shall cause the device to enter a higher safety level (e.g., for those devices that have movable read/write heads, park/unload the drive heads as specified above (see POWER CONDITION field = IDLE)).

A more general one-byte field is proposed to allow for further expansion.

Background: MMC-6 excerpts

4.1.10 Power Management

Table 12 — POWER CONDITION Power Conditions

Mode	Description
Sleep	The lowest power consumption, with power applied, occurs in the Sleep condition. When in the Sleep condition a MM Drive requires a WAKEUP task management function to be activated
Standby	In the Standby condition a MM Drive is capable of accepting commands, but media is not immediately accessible (e.g., the spindle is stopped).
Idle	In the Idle condition a MM Drive is capable of responding quickly to media access requests. However, a MM Drive in the idle condition may take longer, than in the active condition, to complete the execution of a command because it may have to activate some circuitry.
Active	In the Active condition a MM Drive is capable of responding immediately to media access requests, and operations complete execution in the shortest time compared to the other power conditions.

6.39.2.4 Power conditions

Table 634 — POWER CONDITION field

Code	Description
5h	Place Drive into Sleep State. Immed has no meaning when sleep state is requested. Before entering the sleep state, all buffers shall be successfully flushed by the Drive. If the sleep command is successful, the Host should not issue new commands after receiving the successful completion status. The Device shall de-power and disable the interface only after all Drives have successfully completed the sleep operation.

6.39.3.2 Power Condition Changes

When the Drive enters the sleep state, any queued GET EVENT/STATUS NOTIFICATION commands shall be removed from the command queue without command completion.

If any commands other then GET EVENT STATUS NOTIFICATION are in the command queue when a the sleep requested, the command shall be terminated with CHECK CONDITION status and sense bytes SK/ASC/ASCQ shall be set to ILLEGAL REQUEST/COMMAND SEQUENCE ERROR.

If a request to go to a power state fails, the Drive shall remain in the current power state and shall generate power management class event with the Power Event Field set to PwrChg-Fail.

All power state change requests, except sleep, that complete successfully shall generate a power management class event with the Power Event field set to PwrChg-Successful.

Notification of power states shall occur upon entering a new power state.

B.3.5 Power Management and Device Reset in SCSI

When a SCSI Device is in the Power Managed Sleep state, a reset through the service delivery subsystem are used to wake the device.

H.1 Power Management States

... Sleep defines the “Off” state. ...

Table 635 — H.1 Power Management Model States

State	Power Consumption	Drive Context Retained	Restore time
Active	As needed for operation	All	None
Idle	Less than Active	All	The Drive should be restored to active state within 1 second on any request to enter active state, independent of the de-powering process.
Standby	Less than idle	All buffers are empty before entering Standby state.	Vendor specific: Greater than or equal to Idle to Active
Sleep	Less than Standby	None, Buffer & All of command queues are empty before entering Sleep state.	Greater than or equal to Standby to Active. Vendor Specific. May Need full initialization. The Initiator may remove Vcc.

H.2 Power State Transitions

Sleep State:

Maximum power saving state. Buffers and all command queues, including GET EVENT STATUS NOTIFICATION commands, should be emptied before entering into the Sleep state. When the Drive enters the sleep state, any GET EVENT STATUS NOTIFICATION commands present in the command queue, should be removed from the command queue, without command completion. In this Sleep state, all functions are stopped and no commands, except for reset may be received. The unit is consuming less power than when in the Standby state. The Drive context is invalid in the Sleep state.

The Initiator software should fully initialize the Drive after exiting Sleep state, as all context may be lost in the Sleep state. Most devices provide a manual eject mechanism for removing/inserting a disc independent of any lock/unlock mechanism employed. Given this possibility, when the Logical Unit is unable to determine if media has been changed while the Drive was in the sleep state, the Drive should report NEW MEDIA on the next GET EVENT STATUS NOTIFICATION (Media Status) command.

In the Sleep state, the Initiator may completely remove power from the device by turning off Vcc.

...

A power-on or hard reset always returns the Power State to the Standby State. A Device Reset does not alter the current Power State, unless the current Power State is Sleep. A Device Reset received while in sleep state returns the Power State to Standby.

The Sleep state is entered when the Drive has been commanded to go to Sleep but Vcc is still applied to the device. Removing Vcc always takes the device to the Power Off State. Removing Vcc is recommended only when all Drives on a given bus are in Sleep State.

Background: RBC revision 10a excerpts

4.3.1 START STOP UNIT command state restrictions

A removable medium device shall be in either PREVENT state 00b or 10b in order to successfully execute a START STOP UNIT command with thePOWER CONDITIONS field set to the Sleep state (5).

If a removable medium device, in either PREVENT state 01b or 11b, receives a START STOP UNIT command with the POWER CONDITIONS field set to the Sleep state (5), the device shall respond with status set to CHECK CONDITION (02h), the sense key to ILLEGAL REQUEST (05h), and the ASC/ASCQ to ILLEGAL POWER CONDITION REQUEST (2Ch/05h).

A removable medium device in the SLEEP State shall eject the media without causing the media to spin up in accordance with the PREVENT/ALLOW MEDIUM REMOVAL command requirements.

Refer to clause 0 for a description of the POWER CONDITIONS field values.

5.4.1 Power Conditions

Table 8 — Power Conditions

Code		Description
5	M	Place device in Sleep condition

Sleep (condition 5): Devices in the Sleep state are at a lower power consumption level than when in the Standby condition and have very little of the drive circuitry consuming power. A device reset may be required before access to the device is allowed. Prior to entering the Sleep state the device shall ensure that logical blocks in cache have their most recent data value recorded on the physical medium.

7.1.1 Power condition change notification

RBC devices shall notify an initiator of the intent to change power conditions via asynchronous event notification. The status value shall be set to CHECK CONDITION (02h), the sense key to UNIT ATTENTION (06h), and the ASC to POWER CONDITION CHANGE NOTIFICATION (5Eh). The ASCQ shall be set to the value of the new power condition plus 40h as shown in Table 22.

Table 22 — Power condition sense code and qualifier values

ASC	ASCQ	Description
5Eh	45	POWER CONDITION CHANGE TO SLEEP

Background: SBC-2 revision 7 excerpts

If the START STOP UNIT command is issued with the POWER CONDITION field set to 5h the device server shall:

a) suspend any Power Condition timers that are active on receipt of the START STOP UNIT command until a wakeup;
b) not respond to commands and task management functions until a wakeup.

On receipt of a wakeup any previously active power condition timers shall be restored to those values indicated by the saved Power Condition mode page parameters. Before returning a function complete response the target port shall place itself into a condition capable of receiving commands and task management functions and shall create a unit attention condition for all initiators. The sense key shall be set to UNIT ATTENTION with the additional sense code set to LOW POWER CONDITION ACTIVE.

...

In the sleep power condition the device server shall only respond to a wakeup. When a target port has multiple logical units attached it shall enter the sleep power condition only after all the logical units have been placed into a sleep power condition.

Background: ATA8-ACS revision 4c excerpts

4.18.2 Power management commands

...

The SLEEP command moves a device to Sleep mode. The device's interface becomes inactive after the device reports command completion for the SLEEP command. A device only transitions from Sleep mode after processing a hardware reset, a software reset, or a DEVICE RESET command.

4.18.4 Power modes

...

Transition PM0:PM3: When a SLEEP command is received, the device shall make a transition to the PM3:Sleep state.

Transition PM1:PM3: When a SLEEP command is received, the device shall make a transition to the PM3:Sleep state.

Transition PM2:PM3: When a SLEEP command is received, the device shall make a transition to the PM3:Sleep state.

PM3: Sleep: This state shall be entered when the device receives a SLEEP command.

A device transitions from Sleep mode only after processing a hardware reset, a software reset, or a DEVICE RESET command. The time to respond may be as long as 30 s. Sleep state provides the lowest power consumption of any state.

In Sleep state, see the applicable transport standard for a description of the device's interface behavior. Transition PM3:PM2:, A device shall transition to the PM2:Standby state after processing a hardware reset, software reset, or DEVICE RESET command.

7.52 SLEEP - E6h, Non-Data

7.52.1 Feature Set

This 28-bit command is mandatory for devices implementing the Power Management feature set.

7.52.2 Description

This command is the only way to cause the device to enter Sleep mode. The device shall exit Sleep (i.e., State PM3) only after processing a hardware reset, a software reset, or a DEVICE RESET command.

A device shall not power-on in Sleep state.

Suggested changes to SBC-3

4.16 START STOP UNIT and power conditions

4.16.1 START STOP UNIT and power conditions overview

The START STOP UNIT command (see 5.19) allows an application client to control the power condition of a logical unit. This method includes specifying that the logical unit transition to a power condition.

In addition to the START STOP UNIT command, the power condition of a logical unit may be controlled by the Power Condition mode page (see SPC-4). If both the START STOP UNIT command and the Power Condition mode page methods are being used to control the power condition of the same logical unit, then the power condition specified by any START STOP UNIT command shall override the Power Condition mode page's power control.

There shall be no notification to the application client that a logical unit has transitioned from one power condition to another. The REQUEST SENSE command (see SPC-4) indicates if a logical unit is in the idle

power condition or the standby power condition and may indicate if a logical unit is in the stopped power condition.

Editor's Note 2: assuming the logical unit responds to no commands while in the sleep power condition, no change is needed above

---

Table 23 defines the power conditions supported by this standard.

Table 23 — Power conditions

Effect	Power condition				
	Active	Idle	Standby	Stopped	Sleep
Process commands and task management functions	yes	yes	yes	yes	no
Process media access commands	yes	yes	no	no	no
Park heads, if any	no	yes or no [c]	yes	yes	yes
Rotating media spinning, if any	yes	yes [d]	no	no	on
Target port(s) enabled	yes	yes	yes	yes	no
Maximum recovery time [b]	See SPC-4	See SPC-4	See SPC-4	15 s	30 s
References	4.16.2.3 SPC-4	4.16.2.4 SPC-4	4.16.2.5 SPC-4	4.16.2.6	4.16.2.7

[a] Transport protocol-specific interface power management may used to temporarily turn off a target port, but recovery is automatic.
[b] Transport protocol-specific spinup control signals (e.g., NOTIFY (ENABLE SPINUP) in SAS) may increase the recovery time beyond this maximum.
[c] Depends on the POWER CONDITION MODIFIER field.
[d] May be at a low rpm if specified by the POWER CONDITION MODIFIER field.

Power consumption of a SCSI target device with a logical unit in each power condition shall follow this equation:

sleep <= stopped <= standby <= idle <= active

...

4.16.2 START STOP UNIT and power conditions state machine

4.16.2.1 START STOP UNIT and power conditions state machine overview

The SSU_PC (start stop unit power condition) state machine for logical units implementing the START STOP UNIT command describes the logical unit power states and transitions resulting from settings by the START STOP UNIT command and settings in the Power Condition mode page (see SPC-4).

The SSU_PC states are as follows:

a) SSU_PC0:Powered_on (see 4.16.2.2) (initial state);
b) SSU_PC1:Active (see 4.16.2.3);
c) SSU_PC2:Idle (see 4.16.2.4);
d) SSU_PC3:Standby (see 4.16.2.5); ~~and~~
e) SSU_PC4:Stopped (see 4.16.2.6)~~.~~; and

f) SSU_PC5:Sleep (see 4.16.2.7).

The SSU_PC state machine shall start in the SSU_PC0:Powered_on state after power on.

NOTE 2 - NOTE 6 - The SSU_PC state machine is an enhanced version of the Power Condition state machine described in SPC-4.

Figure 4 describes the SSU_PC state machine.

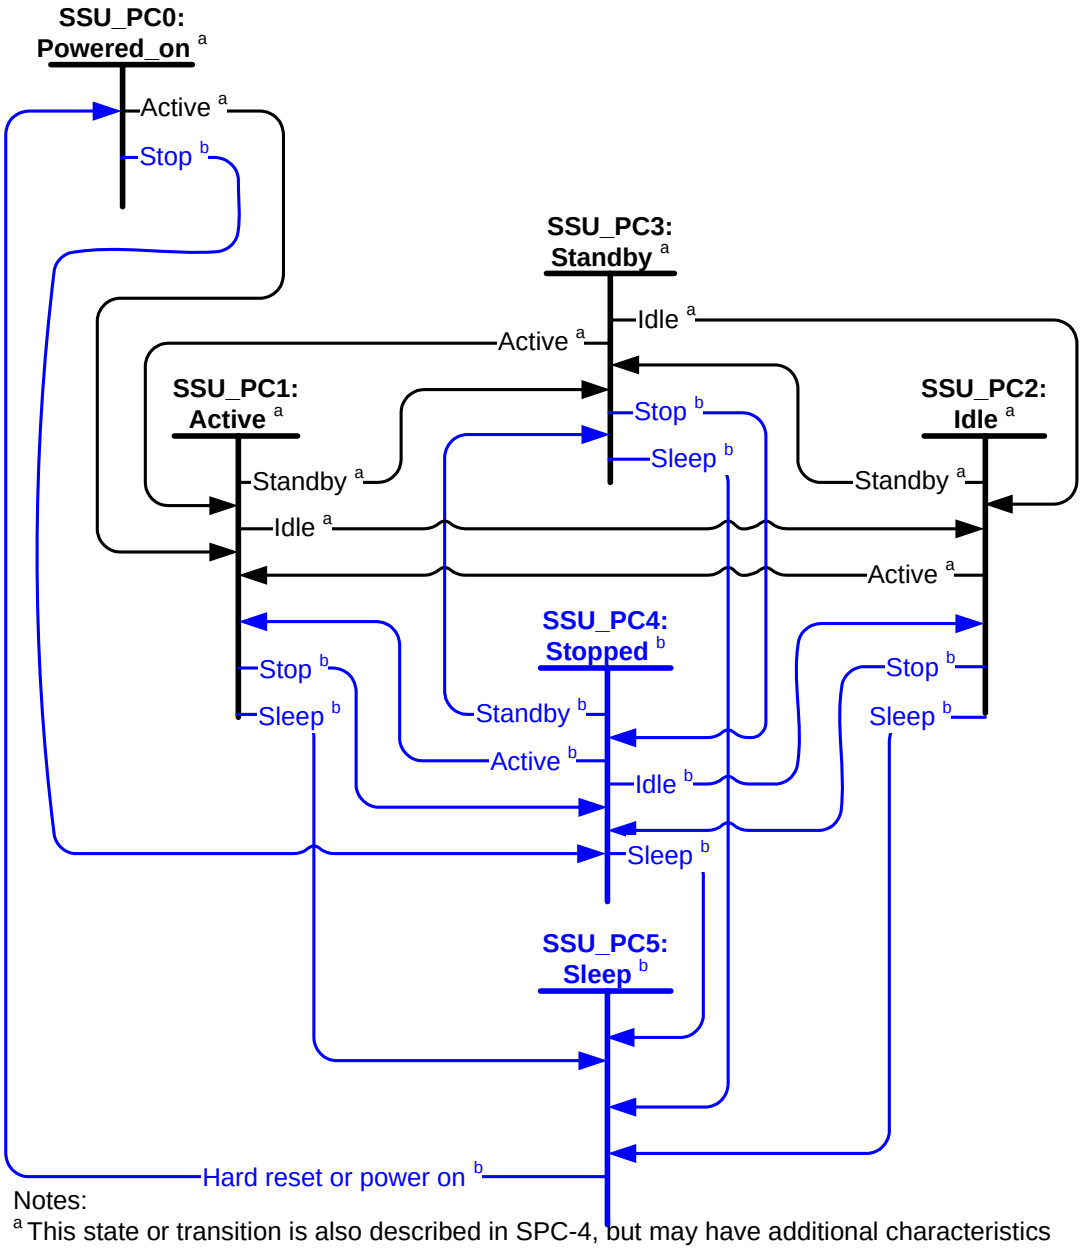

Notes:
[a] This state or transition is also described in SPC-4, but may have additional characteristics unique to this standard (e.g., a transition to or from a state described in this standard).
[b] This state or transition is described in this standard.

Figure 1 — Power condition state machine for logical units implementing the START STOP UNIT command

4.16.2.2 SSU_PC0:Powered_on state

4.16.2.2.1 SSU_PC0:Powered_on state description

The logical unit shall enter this state upon power on. This state consumes zero time.

4.16.2.2.2 Transition SSU_PC0:Powered_on to SSU_PC1:Active

This transition shall occur if:

a) the logical unit has been configured to transition to the SSU_PC1:Active state.

4.16.2.2.3 Transition SSU_PC0:Powered_on to SSU_PC4:Stopped

This transition shall occur if:

a) the logical unit has been configured to transition to the SSU_PC4:Stopped state.

4.16.2.3 SSU_PC1:Active state

4.16.2.3.1 SSU_PC1:Active state description

While in this state, if power on initialization is not complete, then the logical unit completes its power on initialization.

While in this state, after power on initialization is complete, then:

a) the logical unit is in the active power condition (see SPC-4);
b) if the idle condition timer is active (see SPC-4) and not disabled (see 5.19), then the idle condition timer is running; and
c) if the standby condition timer is active (see SPC-4) and not disabled (see 5.19), then the standby condition timer is running.

4.16.2.3.2 Transition SSU_PC1:Active to SSU_PC2:Idle

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the POWER CONDITION field set to IDLE;
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0; or
c) the idle condition timer is active (see SPC-4), enabled (see 5.19), and zero.

4.16.2.3.3 Transition SSU_PC1:Active to SSU_PC3:Standby

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the POWER CONDITION field set to STANDBY;
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0; or
c) the standby condition timer is active (see SPC-4), enabled (see 5.19), and zero.

4.16.2.3.4 Transition SSU_PC1:Active to SSU_PC4:Stopped

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero and the POWER CONDITION field set to START_VALID.

4.16.2.3.5 Transition SSU_PC1:Active to SSU_PC5:Sleep

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero and the POWER CONDITION field set to SLEEP.

Editor's Note 3: MMC-6 refuses to accept a SLEEP request if there any commands in the task set (it returns ILLEGAL REQUEST/COMMAND SEQUENCE ERROR). This way, entry into sleep doesn't need to abort commands. ATA8-ACS does not block its SLEEP command. Should SBC-3 do the same as MMC-6?

4.16.2.4 SSU_PC2:Idle state

4.16.2.4.1 SSU_PC2:Idle state description

While in this state:

a) the logical unit is in the idle power condition (see SPC-4);
b) the device server processes the REQUEST SENSE command as described in 4.16.1; and
c) if the standby condition timer is active (see SPC-4) and not disabled (see 5.19), then the standby condition timer is running.

4.16.2.4.2 Transition SSU_PC2:Idle to SSU_PC1:Active

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the START bit set to one;
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to ACTIVE; or
c) the device server processes a command that requires the logical unit to be in the SSU_PC1:Active state to process the command.

4.16.2.4.3 Transition SSU_PC2:Idle to SSU_PC3:Standby

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the POWER CONDITION field set to STANDBY;
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0; or
c) the standby condition timer is active (see SPC-4), enabled (see 5.19), and zero.

4.16.2.4.4 Transition SSU_PC2:Idle to SSU_PC4:Stopped

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero.

4.16.2.4.5 Transition SSU_PC2:Idle to SSU_PC5:Sleep

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero and the POWER CONDITION field set to SLEEP.

Editor's Note 4: MMC-6 refuses to accept a SLEEP request if there any commands in the task set (it returns ILLEGAL REQUEST/COMMAND SEQUENCE ERROR). This way, entry into sleep doesn't need to abort commands. ATA8-ACS does not block its SLEEP command. Should SBC-3 do the same as MMC-6?

4.16.2.5 SSU_PC3:Standby state

4.16.2.5.1 SSU_PC3:Standby state description

While in this state:

a) the logical unit is in the standby power condition (see SPC-4); and
b) the device server processes the REQUEST SENSE command as described in 4.16.1.

4.16.2.5.2 Transition SSU_PC3:Standby to SSU_PC1:Active

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the START bit set to one;
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to ACTIVE; or
c) the device server processes a command that requires the logical unit to be in the SSU_PC1:Active state to process the command.

4.16.2.5.3 Transition SSU_PC3:Standby to SSU_PC2:Idle

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the POWER CONDITION field set to IDLE;
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0; or
c) the device server processes a command that requires the logical unit to be in the SSU_PC2:Idle state to process the command.

4.16.2.5.4 Transition SSU_PC3:Standby to SSU_PC4:Stopped

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero.

4.16.2.5.5 Transition SSU_PC3:Standby to SSU_PC5:Sleep

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero and the POWER CONDITION field set to SLEEP.

Editor's Note 5: MMC-6 refuses to accept a SLEEP request if there any commands in the task set (it returns ILLEGAL REQUEST/COMMAND SEQUENCE ERROR). This way, entry into sleep doesn't need to abort commands. ATA8-ACS does not block its SLEEP command. Should SBC-3 do the same as MMC-6?

4.16.2.6 SSU_PC4:Stopped state

4.16.2.6.1 SSU_PC4:Stopped state description

While in this state:

a) the logical unit is in the stopped power condition;
b) the device server is not capable of processing medium access commands. The device server shallterminate each medium access command or TEST UNIT READY command processed while in this state with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, INITIALIZING COMMAND REQUIRED;
c) the device server processes the REQUEST SENSE command as described in 4.16.1; and
d) the power consumed by the SCSI target device should be less than or equal to that consumed than when the logical unit is in the SSU_PC1:Active, SSU_PC2:Idle, or SSU_PC3:Standby states.

4.16.2.6.2 Transition SSU_PC4:Stopped to SSU_PC1:Active

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the START bit set to one; or
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to ACTIVE.

The time from SSU_PC4:Stopped to SSU_PC1:Active shall be less than the maximum recovery time defined in table 23 (see 4.16.1).

4.16.2.6.3 Transition SSU_PC4:Stopped to SSU_PC2:Idle

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the POWER CONDITION field set to IDLE; or
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0.

4.16.2.6.4 Transition SSU_PC4:Stopped to SSU_PC3:Standby

This transition shall occur after:

a) the device server processes a START STOP UNIT command with the POWER CONDITION field set to STANDBY; or
b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0.

4.16.2.3.4 Transition SSU_PC4:Stopped to SSU_PC5:Sleep

This transition shall occur after the device server processes a START STOP UNIT command with the START bit set to zero and the POWER CONDITION field set to SLEEP.

Editor's Note 6: MMC-6 refuses to accept a SLEEP request if there any commands in the task set (it returns ILLEGAL REQUEST/COMMAND SEQUENCE ERROR). This way, entry into sleep doesn't need to abort commands. ATA8-ACS does not block its SLEEP command. Should SBC-3 do the same as MMC-6?

4.16.2.6 SSU_PC5:Sleep state

4.16.2.6.1 SSU_PC5:Sleep state description

This state shall only be supported by logical units accessible through SCSI target ports whose transport protocols all define reset events.

While in this state:

a) the logical unit is in the sleep power condition;
b) the device server is not capable of processing commands or task management functions. The device server shall terminate each command processed while in this state with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOW POWER CONDITION ACTIVE. The task manager shall terminate each task management function processed while in this state with a function response of FUNCTION REJECTED or SERVICE DELIVERY OR TARGET FAILURE;

Editor's Note 7: The target device may contain multiple target ports and multiple logical units. A target port is only turned off when all the logical units are in the sleep power condition, meaning that some could continue to receive commands while waiting for their siblings to go to sleep; thus the error reporting in b) (just letting the commands and TMFs timeout would be another viable approach). If there are multiple target ports, they are sent to sleep at the same time (when all the logical unit reach the sleep power condition).

---

c) the SCSI target device shall make inactive each SCSI target port through which the logical unit is accessible if all other logical units in the SCSI target device that are accessible through that SCSI target port are also in this state. It shall not make a SCSI target port inactive while any logical unit that is accessible through that SCSI target port is not in this state; and
d) the power consumed by the SCSI target device should be less than or equal to that consumed than when the logical unit is in the SSU_PC1:Active, SSU_PC2:Idle, SSU_PC3:Standby, or SSU_PC4:Stopped states.

4.16.2.6.2 Transition SSU_PC5:Sleep to SSU_PC0:Power_On

This transition shall occur after:

a) power on; or
b) hard reset.

The time from SSU_PC5:Sleep to SSU_PC1:Active via SSU_PC0:Power_On shall be less than the maximum recovery time defined in table 23 (see 4.16.1).

5.19 START STOP UNIT command

The START STOP UNIT command (see table 60) requests that the device server change the power condition of the logical unit (see 4.16) or load or eject the medium. This includes specifying that the device server enable or disable the direct-access block device for medium access operations by controlling power conditions and timers.

Logical units that contain cache shall write all cached logical blocks to the medium (e.g., as they would do in response to a SYNCHRONIZE CACHE command (see 5.20 and 5.21) with the SYNC_NV bit set to zero, the LOGICAL BLOCK ADDRESS field set to zero, and the NUMBER OF LOGICAL BLOCKS field set to zero) prior to entering into any power condition that prevents accessing the medium (e.g., before the rotating media spindle motor is stopped during transition to the stopped power condition).

If any deferred downloaded code has been received as a result of a WRITE BUFFER command (see SPC-4), then that deferred downloaded code shall replace the current operational code.

Table 60 — START STOP UNIT command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (1Bh)							
1	Reserved							IMMED
5	Reserved							
5	~~Reserved~~ POWER CONDITION MODIFIER							
4	POWER CONDITION				Reserved		LOEJ	START
5	CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value defined in table 60.

If the immediate (IMMED) bit is set to zero, then the device server shall return status after the operation is completed. If the IMMED bit set to one, then the device server shall return status as soon as the CDB has been validated.

The POWER CONDITION MODIFIER field specifies additional information about the requested power condition.

Table 61 — POWER CONDITION MODIFIER field

POWER CONDITION field	POWER CONDITION MODIFIER field	Description
All values that are not reserved	00h	Vendor-specific
2h (i.e., idle)	01h	Retract the spindle, if any.
	02h	Retract the spindle, if any, and operate at a low rpm.
All other combinations		Reserved

The POWER CONDITION field is used to specify that the logical unit be placed into a power condition or to adjust a timer as defined in table 61. If this field is supported and is set to a value other than 0h, then the START and LOEJ bits shall be ignored.

Table 61 — POWER CONDITION field

Code	Name	Description
0h	START_VALID	Process the START and LOEJ bits.
1h	ACTIVE	Place the ~~device~~logical unit into the active power condition.
2h	IDLE	Place the ~~device~~logical unit into the idle power condition.
3h	STANDBY	Place the ~~device~~logical unit into the standby power condition.
4h	Reserved	
5h	~~Obsolete~~ SLEEP	Place the logical unit into the sleep power condition.
6h	Reserved	
7h	LU_CONTROL	Transfer control of power conditions to the logical unit.
8h - 9h	Reserved	
Ah	FORCE_IDLE_0	Force the idle condition timer to zero.
Bh	FORCE_STANDBY_0	Force the standby condition timer to zero.
Ch - Fh	Reserved	

The contents of the CONTROL byte are defined in SAM-4.

If the START STOP UNIT command is processed with the POWER CONDITION field set to ACTIVE, IDLE, or STANDBY, then:

a) the logical unit shall transition to the specified power condition;
b) the logical unit shall change power conditions only after receipt of another START STOP UNIT command or a logical unit reset; and
c) the device server shall disable the idle condition timer if it is active (see SPC-4) and disable the standby condition timer if it is active (see SPC-4) until another START STOP UNIT command is processed that returns control of the power condition to the logical unit, or a logical unit reset occurs.

If the START STOP UNIT command is processed with the POWER CONDITION field set to SLEEP, then:

a) the logical unit shall transition to the sleep power condition; and
b) the logical unit shall change power conditions only after a power on or hard reset.

If the START STOP UNIT command is processed with the POWER CONDITION field set to LU_CONTROL, then the device server shall enable the idle condition timer if it is active (see SPC-4) and disable the standby condition timer if it is active (see SPC-4).

If the START STOP UNIT command is processed with the POWER CONDITION field set to FORCE_IDLE_0 or FORCE_STANDBY_0, then the device server shall:

a) force the specified timer to zero, cause the logical unit to transition to the specified power condition, and return control of the power condition to the device server; or
b) terminate a START STOP UNIT command that selects a timer that is not supported by the device server or a timer that is not active. The command shall be terminated with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

It is not an error to specify that the logical unit transition to its current power condition.

If the load eject (LOEJ) bit is set to zero, then the logical unit shall take no action regarding loading or ejecting the medium. If the LOEJ bit is set to one, then the logical unit shall unload the medium if the START bit is set to zero. If the LOEJ bit is set to one, then the logical unit shall load the medium if the START bit is set to one. If the START bit is set to zero, then the logical unit shall transition to the stopped power condition, disable the idle condition timer if it is active (see SPC-4), and disable the standby condition timer if it is active (see SPC-4). If the START bit set to one, then the logical unit shall transition to the active power condition, enable the idle condition timer if it is active, and enable the standby condition timer if it is active.

6.4.3 Block Device Characteristics VPD page

The Block Device Characteristics VPD page contains parameters indicating characteristics of the logical unit.

Table 132 defines the Block Device Characteristics VPD page.

Table 132 — Block Device Characteristics VPD page

Byte\Bit	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (B1h)							
2	Reserved							
3	PAGE LENGTH (3Ch)							
4	MEDIUM ROTATION RATE							
5								
6	Reserved							
7	Reserved				NOMINAL FORM FACTOR			
8	Reserved							~~Reserved~~ STOPPED_SUP
9	~~Reserved~~ MAXIMUM STOPPED RECOVERY TIME							
10	Reserved							
63								

...

A STOPPED_SUP bit set to one indicates that the device server supports the stopped power condition (see 4.6). A STOPPED_SUP bit set to zero indicates that device server does not support the stopped power condition.

The MAXIMUM STOPPED RECOVERY TIME field indicates the maximum time in 100 ms increments that the SCSI target device takes to return from the stopped power condition to the active power condition.

Editor's Note 8: Since power up in stopped mode exists, "to the active power condition" is not complete.

Suggested changes to SPC-4

Editor's Note 9: If the sleep power condition is added to SPC, then MMC-6 might also need modifications to refer to it.

3.1.5 active power condition: When a device server is capable of responding to all of its supported commands, including media access requests, without delay. See 5.9.

3.1.51 idle power condition: When a device server is capable of responding to all of its supported commands, including media access requests, but commands may take longer to complete than when in the active power condition. See 5.9.

3.1.150 sleep power condition: When a logical unit is not capable of accepting commands or task management functions. See 5.9.

3.1.150 standby power condition: When a device server is capable of ~~accepting commands, but not capable of processing media access commands~~responding to all of its supported commands, including media access requests, but commands may take longer to complete than when in the idle power condition. See 5.9.

5.9 Power conditions

5.9.1 Power conditions overview

The optional Power Condition mode page (see 7.4.12) allows an application client to control the power condition of a logical unit in a manner that may reduce power consumption of the SCSI target device. This control is invoked by enabling and setting the idle condition timer ~~and/or~~, the standby condition timer, and/or the sleep condition timer using the mode page. A change in the power condition of any logical unit in a SCSI target device may result in a change in the SCSI target device's power consumption.

In addition to the Power Condition mode page, the power condition of a logical unit may be controlled by the START STOP UNIT command (see SBC-3, MMC-6, or RBC). If both the Power Condition mode page and the START STOP UNIT command methods are being used to control the power condition of the same logical unit, then any START STOP UNIT command's power condition specification shall override the Power Condition mode page's power control and may disable the idle condition timer, ~~and~~ the standby condition timer~~s~~, and the sleep condition timer.

There shall be no notification to the application client that a logical unit has transitioned from one power condition to another. The REQUEST SENSE command (see 6.28) indicates if a logical unit is in the idle power condition or the standby power condition.

Command standards (see 3.1.18) may define for their peripheral device types additional power conditions (e.g., the stopped power condition defined by SBC-3 for direct-access block devices) and extensions to the REQUEST SENSE command for reporting power conditions.

No power condition shall affect the supply of any power required for proper operation of a service delivery subsystem.

Logical units that contain cache memory shall write all cached data to the medium for the logical unit (e.g., as a logical unit would do in response to a SYNCHRONIZE CACHE command as described in SBC-2) prior to entering into any power condition that prevents accessing the media (e.g., before a hard drive stops its spindle motor during transition to the standby power condition).

The power conditions are described in table 42.

Table 133 — Power Conditions

Power Condition	Description
active	While in the active power condition (see 3.1.5): a) A device server is capable of responding to all of its supported commands including media access requests; b) A logical unit completes processing of operations in the shortest time when compared to the time required for completion while in the idle or standby power conditions; and c) The SCSI target device may consume more power than when the logical unit is in the idle power condition (e.g., a disk drive's spindle motor may be active).
idle	While in the idle power condition (see 3.1.51): a) A device server is capable of responding to all of its supported commands including media access requests; b) A logical unit may take longer to complete processing a command than it would while in the active power condition (e.g., the device may have to activate some circuitry before processing a command); and c) The power consumed by the SCSI target device should be less than or equal to the power consumed when the logical unit is in the active power condition and may be greater than the power consumed when the logical unit is in the standby power condition.
standby	While in the standby power condition (see 3.1.150): ~~a) A device server is not capable of processing media access commands; and~~ a) A device server is capable of responding to all of its supported commands including media access requests; b) A logical unit may take longer to complete processing a command than it would while in the idle power condition (e.g., the device may have to activate a spindle motor); and c) The power consumed by the SCSI target device should be less than or equal to the power consumed when the logical unit is in the idle power condition (e.g., a disk drive's spindle motor is stopped).
sleep	While in the sleep power condition (see 3.1.xxx): a) the logical unit is not capable of processing commands and task management functions; and b) The power consumed by the SCSI target device should be less than or equal to the power consumed when the logical unit is in the standby power condition.

Power consumption of a SCSI target device with a logical unit in each power condition shall follow this equation:

sleep <= standby <= idle <= active

Editor's Note 10: Need to decide the best way to state the power consumption rules. Embed in table 133, separate <= list, or ?

...

5.9.2 Power condition state machine

5.9.2.1 Power condition state machine overview

The PC (power condition) state machine describes the logical unit power states and transitions resulting from Power Condition mode page settings.

The PC states are as follows:

a) PC0:Powered_on (see 5.9.2.2) (initial state);
b) PC1:Active (see 5.9.2.3);
c) PC2:Idle (see 5.9.2.4); ~~and~~
d) PC3:Standby (see 5.9.2.5)~~.~~; and
e) PC4:Sleep (see 5.9.2.6).

The PC state machine stall start in the PC0:Powered_on state after power on.

Figure 5 describes the PC state machine.

Figure 5 — Pow er condition state machine [add PC4:Sleep state]

5.9.2.2 PC0:Powered_on state

5.9.2.2.1 PC0:Powered_on state description

The logical unit shall enter this state upon power on. This state consumes zero time.

5.9.2.2.2 Transition PC0:Powered_on to PC1:Active

This transition shall occur after the logical unit is ready to begin its power on initialization.

Editor's Note 11: ATA devices can power up in standby mode; should that be allowed here, so SAT mapping is legal?

5.9.2.3 PC1:Active state

5.9.2.3.1 PC1:Active state description

While in this state, if power on initialization is not complete, then the logical unit shall complete its power on initialization.

While in this state, if power on initialization is complete, then:

a) The logical unit is in the active power condition (see table 42);
b) If the idle condition timer is active, then the idle condition timer is running; and
c) If the standby condition timer is active, then the standby condition timer is running.

5.9.2.3.2 Transition PC1:Active to PC2:Idle

This transition shall occur after:

a) The idle condition timer is active; and
b) The idle condition timer is zero.

5.9.2.3.3 Transition PC1:Active to PC3:Standby

This transition shall occur after:

a) The standby condition timer is active; and
b) The standby condition timer is zero.

5.9.2.3.4 Transition PC1:Active to PC4:Sleep

This transition shall occur after:

a) The sleep condition timer is active; and
b) The sleep condition timer is zero.

5.9.2.4 PC2:Idle state

5.9.2.4.1 PC2:Idle state description

While in this state:

a) The logical unit is in the idle power condition (see table 42);
b) The device server processes the REQUEST SENSE command as described in 6.28; and
c) If the standby condition timer is active, then the standby condition timer is running.

5.9.2.4.2 Transition PC2:Idle to PC1:Active

This transition shall occur after the device server processes a command that requires the logical unit to be in the PC1:Active state to process the command.

5.9.2.4.3 Transition PC2:Idle to PC3:Standby

This transition shall occur after:

a) The standby condition timer is active; and
b) The standby condition timer is zero.

5.9.2.4.4 Transition PC2:Idle to PC4:Sleep

This transition shall occur after:

a) The sleep condition timer is active; and
b) The sleep condition timer is zero.

5.9.2.5 PC3:Standby state

5.9.2.5.1 PC3:Standby state description

While in this state:

a) The logical unit is in the standby power condition (see table 42); and
b) The device server processes the REQUEST SENSE command as described in 6.28.

5.9.2.5.2 Transition PC3:Standby to PC1:Active

This transition shall occur after the device server processes a command that requires the logical unit to be in the PC1:Active state to process the command.

5.9.2.5.3 5.9.2.5.3 PC3:Standby to PC2:Idle

This transition shall occur after the device server processes a command that requires the logical unit to be in the PC2:Idle state to process the command.

5.9.2.5.4 Transition PC3:Standby to PC4:Sleep

This transition shall occur after:

a) The sleep condition timer is active; and
b) The sleep condition timer is zero.

5.9.2.6 PC4:Sleep state

5.9.2.6.1 PC4:Sleep state description

While in this state:

a) The logical unit is in the sleep power condition (see table 42); and
b) The device server processes commands and task management functions as described in table 42.

5.9.2.6.2 Transition PC4:Sleep to PC0:Powered_on

This transition shall occur after:

a) power on; or
b) hard reset.

7.4.12 Power Condition mode page

The Power Condition mode page provides an application client with methods to control the power condition of a logical unit (see 5.9). These methods include:

a) Specifying that the logical unit transition to a power condition without delay; and
b) Activating and setting of idle condition and standby condition timers to specify that the logical unit wait for a period of inactivity before transitioning to a specified power condition.

The mode page policy (see 6.9) for this mode page shall be shared.

When a device server receives a command while in a power condition based on a setting in the Power Condition mode page, the logical unit shall transition to the power condition that allows the command to be processed. If either the idle condition timer or the standby condition timer has been set, then they shall be reset on receipt of the command. On completion of the command, the timer(s) shall be started.

Logical units that contain cache memory shall write all cached data to the medium for the logical unit (e.g.,as a logical unit does in response to a SYNCHRONIZE CACHE command as described in SBC-2) prior to entering into any power condition that prevents accessing the media (e.g., before a hard drive stops its spindle motor during transition to the standby power condition).

The logical unit shall use the values in the Power Condition mode page to control its power condition after a power on or a hard reset until a START STOP UNIT command setting a power condition is received.

Table 319 defines the Power Condition mode page.

Table 134 — Power Condition mode page

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (1Ah)					
1	PAGE LENGTH (~~0Ah~~0Eh)							
2	Reserved							
3	Reserved					~~Reserved~~ SLEEP	IDLE	STANDBY
4	(MSB)	IDLE CONDITION TIMER						
7								(LSB)
8	(MSB)	STANDBY CONDITION TIMER						
11								(LSB)
12	(MSB)	SLEEP CONDITION TIMER						
15								(LSB)

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

The IDLE and STANDBY bits specify which timers are active.

If the IDLE bit is set to one and the STANDBY bit is set to zero, then the idle condition timer is active and the device server shall transition to the idle power condition when the idle condition timer is zero.

If the IDLE bit is set to zero, then the device server shall ignore the idle condition timer.

If the STANDBY bit is set to one and the IDLE bit is set to zero, then the standby condition timer is active and the device server shall transition to the standby power condition when the standby condition timer is zero.

If the STANDBY bit is set to zero, then the device server shall ignore the standby condition timer.

If both the IDLE and STANDBY bits are set to one, then both timers are active and run concurrently. When the idle condition timer is zero the device server shall transition to the idle power condition. When the standby condition timer is zero the device server shall transition to the standby power condition. If the standby condition timer is zero before the idle condition timer is zero, then the logical unit shall transition to the standby power condition.

The value in the IDLE CONDITION TIMER field specifies the inactivity time in 100 millisecond increments that the logical unit shall wait before transitioning to the idle power condition when the IDLE bit is set to one. The idle condition timer is expired when:

a) The IDLE CONDITION TIMER field is set to zero; or
b) The number of milliseconds specified by the value in the IDLE CONDITION TIMER field times 100 milliseconds has elapsed since the last activity (e.g., processing a command that requires the active power condition or performing a self test).

NOTE 3 - If the device server implements the idle power condition but does not implement an idle power condition timer, it reports that the IDLE bit is changeable, and that the IDLE CONDITION TIMER field has default value of 00000000h and is not changeable.

The value in the STANDBY CONDITION TIMER field specifies the inactivity time in 100 millisecond increments that the logical unit shall wait before transitioning to the standby power condition when the STANDBY bit is set to one. The standby condition timer is expired when:

a) The STANDBY CONDITION TIMER field is set to zero; or
b) The number of milliseconds specified by the value in the STANDBY CONDITION TIMER field times 100 milliseconds has elapsed since the last activity (e.g., processing any command or performing a self test).

NOTE 4 - If the device server implements the standby power condition but does not implement a standby power condition timer, it reports that the STANDBY bit is changeable, and that the STANDBY CONDITION TIMER field has default value of 00000000h and is not changeable.

The value in the SLEEP CONDITION TIMER field specifies the inactivity time in 100 millisecond increments that the logical unit shall wait before transitioning to the sleep power condition when the SLEEP bit is set to one. The sleep condition timer is expired when:

a) The SLEEP CONDITION TIMER field is set to zero; or
b) The number of milliseconds specified by the value in the SLEEP CONDITION TIMER field times 100 milliseconds has elapsed since the last activity (e.g., processing any command or performing a self test).

NOTE 5 - If the device server implements the sleep power condition but does not implement a sleep power condition timer, it reports that the SLEEP bit is changeable, and that the SLEEP CONDITION TIMER field has default value of 00000000h and is not changeable.

Editor's Note 12: Add controls for the stopped power condition here too, or is that OK as block device type specific?

7.7.4 Extended INQUIRY Data VPD page

The Extended INQUIRY Data VPD page (see table 423) provides the application client with a means to obtain information about the logical unit.

Table 423 — Extended INQUIRY Data VPD page

Byte\Bit	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (86h)							
2	Reserved							
3	PAGE LENGTH (3Ch)							
4	...							
7								
8	Reserved					~~Reserved~~ SLEEP_SUP	~~Reserved~~ STANDBY_SUP	~~Reserved~~ IDLE_SUP
9	~~Reserved~~ MAXIMUM IDLE RECOVERY TIME							
10	~~Reserved~~ MAXIMUM STANDBY RECOVERY TIME							
11	~~Reserved~~ MAXIMUM SLEEP RECOVERY TIME							
12	Reserved							
63								

...

A SLEEP_SUP bit set to one indicates that the device server supports the sleep power condition (see 5.9). A SLEEP_SUP bit set to zero indicates that device server does not support the sleep power condition.

A STANDBY_SUP bit set to one indicates that the device server supports the standby power condition (see 5.9). A STANDBY_SUP bit set to zero indicates that device server does not support the standby power condition.

An IDLE_SUP bit set to one indicates that the device server supports the idle power condition (see 5.9). An IDLE_SUP bit set to zero indicates that device server does not support the idle power condition.

The MAXIMUM IDLE RECOVERY TIME field indicates the maximum time in 100 ms increments that the SCSI target device takes to return from the idle power condition to the active power condition.

The MAXIMUM STANDBY RECOVERY TIME field indicates the maximum time in 100 ms increments that the SCSI target device takes to return from the standby power condition to the active power condition.

The MAXIMUM SLEEP RECOVERY TIME field indicates the maximum time in 100 ms increments that the SCSI target device takes to return from the sleep power condition to the active power condition.

Suggested changes to SAS-2.1

2.2 Approved references

At the time of publication, the following referenced standards or technical reports were approved:

ISO/IEC 14776-326, *Reduced Block Commands (RBC)* (ANSI INCITS 330-2000 with ANSI INCITS 330-2003/AM1)

ANSI INCITS TR-35-2004, *Methodologies for Jitter and Signal Quality Specification (MJSQ)*. When MJSQ is referenced from this standard, the FC Port terminology used within MJSQ should be substituted with SAS phy terminology.

2.3 References under development

At the time of publication, the following referenced standards were still under development. For information on the current status of the document, or regarding availability, contact the relevant standards body or other organization as indicated.

ISO/IEC xxxxx-xxx, *ATA Attachment-8 ATA/ATAPI Architecture Model (ATA8-AAM)* (T13/1700-D)

ISO/IEC xxxxx-xxx, *ATA Attachment-8 ATA/ATAPI Command Set (ATA8-ACS)* (T13/1699-D)

ISO/IEC 14776-414, *SCSI Architecture Model-4 (SAM-4)* (T10/1683-D)

ISO/IEC 14776-454, *SCSI Primary Commands-4 (SPC-4)* (T10/1731-D)

ISO/IEC 14776-323, *SCSI Block Commands-3 (SBC-3)* (T10/1799-D)

ISO/IEC 14776-372, *SCSI Enclosure Services-2 (SES-2)* (T10/1559-D)

ISO/IEC 14776-366, *Multi-Media Commands-6 (MMC-6)* (T10/1836-D)

NOTE 6 - For more information on the current status of these documents, contact the INCITS Secretariat at 202-737-8888 (phone), 202-638-4922 (fax) or via Email at incits@itic.org. To obtain copies of these documents, contact Global Engineering at 15 Inverness Way, East Englewood, CO 80112-5704 at 303-792-2181 (phone), 800-854-7179 (phone), or 303-792-2192 (fax) or see http://www.incits.org.

2.4 Symbols and abbreviations

See 2.1 for abbreviations of standards bodies (e.g., ISO). Units and abbreviations used in this standard:

Abbreviation	Meaning
ATA8-AAM	AT Attachment - 8 ATA/ATAPI Architecture Model standard (see 2.3)
ATA8-ACS	AT Attachment - 8 ATA/ATAPI Command Set standard (see 2.3)
MMC-6	Multi-Media Commands standard (see 2.3)
RBC	Reduced Block Commands standard (see 2.2)
SAM-4	SCSI Architecture Model - 4 standard (see 2.3)
SBC-3	SCSI Block Commands - 3 standard (see 2.3)
SPC-4	SCSI Primary Commands - 4 standard (see 2.3)

7.2.5.3.2 NOTIFY (ENABLE SPINUP)

NOTIFY (ENABLE SPINUP) is transmitted by a SAS initiator port or expander port and is used to specify to an SAS target device that it may temporarily consume additional power (e.g., while spinning-up rotating media) while transitioning into the active or idle power condition state. The length of time the SAS target device consumes additional power and the amount of additional power is vendor specific. NOTIFY (ENABLE SPINUP) shall interact with the device's power condition state transitions, controlled by the Power Conditions mode page (see SPC-4) and/or the START STOP UNIT command (see SBC-3, MMC-6, or RBC), as described in 10.2.10.

7.10 Power management

SATA interface power management is not supported in STP.

STP initiator ports shall not generate SATA_PMREQ_P, SATA_PMREQ_S, or SATA_PMACK. If an STP initiator port receives SATA_PMREQ_P or SATA_PMREQ_S, it shall reply with SATA_PMNAK.

If an expander device receives SATA_PMREQ_P or SATA_PMREQ_S from a SATA device while an STP connection is not open, it shall not forward it to any STP initiator port and shall reply with SATA_PMNAK. If

one of these primitives arrives while an STP connection is open, it may forward the primitive to the STP initiator port.

SCSI idle ~~and~~, standby, and sleep power conditions, implemented with the START STOP UNIT command (see SBC-3, MMC-6, or RBC) and the Power Condition mode page (see SPC-4), may be supported by SSP initiator ports and SSP target ports as described in 10.2.10.

ATA idle ~~and~~, standby, and sleep power modes, implemented with the IDLE, IDLE IMMEDIATE, STANDBY, STANDBY IMMEDIATE, SLEEP, and CHECK POWER MODE commands (see ATA8-ACS), may be supported by STP initiator ports. ~~The ATA sleep power mode, implemented with the SLEEP command, shall not be used.~~

10.2.6.5 START STOP UNIT command

The power condition states controlled by the START STOP UNIT command (see SBC-3, MMC-6, or RBC) for a SAS device are described in 10.2.10.

10.2.10 SCSI power conditions

10.2.10.1 SCSI power conditions overview

The logical unit power condition states from the Power Condition mode page (see SPC-4) and START STOP UNIT command (see SBC-3, MMC-6, or RBC), if implemented, shall interact with the NOTIFY (ENABLE SPINUP) primitive (see 7.2.5.3) to control temporary consumption of additional power (e.g., spin-up of rotating media) as described in this subclause.

The logical unit uses NOTIFY (ENABLE SPINUP) to:

a) initiate spin-up after power on; and
b) delay spin-ups requested by START STOP UNIT commands.

Table 23 defines the power conditions supported by this standard.

Table 424 — Power conditions

Power condition	Target port(s)	Maximum recovery time [b]	References
Active	on [a]	See SPC-4	10.2.10.2.3 10.2.10.2.8 SPC-4
Idle	on [a]	See SPC-4	10.2.10.2.4 10.2.10.2.9 SPC-4
Standby	on [a]	See SPC-4	10.2.10.2.5 SPC-4
Stopped	on [a]	15 s	10.2.10.2.6 SBC-3, MMC-6, RBC
Sleep	off	30 s	10.2.10.2.7 SBC-3, MMC-6, RBC

[a] Interface power management (see TBD) may used to temporarily turn off a target port, but recovery is automatic.
[b] NOTIFY (ENABLE SPINUP) may increase the recovery time beyond this maximum.

Editor's Note 13: Not sure what contents belong in this standard. Don't want to double-specify anything.

10.2.10.2 SA_PC (SCSI application layer power condition) state machine

10.2.10.2.1 SA_PC state machine overview

The SA_PC (SCSI application layer power condition) state machine describes how the SAS target device processes logical unit power condition state change requests and NOTIFY (ENABLE SPINUP) if it is a SCSI target device.

NOTE 7 - This state machine is an enhanced version of the logical unit power condition state machines described in SPC-4 and SBC-3.

This state machine consists of the following states:

a) SA_PC_0:Powered_On (see 10.2.10.2.2)(initial state);
b) SA_PC_1:Active (see 10.2.10.2.3);
c) SA_PC_2:Idle (see 10.2.10.2.4);
d) SA_PC_3:Standby (see 10.2.10.2.5);
e) SA_PC_4:Stopped (see 10.2.10.2.6)(specific to ~~SBC-3~~ logical units with certain peripheral device types);
f) SA_PC_5:Sleep (see 10.2.10.2.6)(specific to logical units with certain peripheral device types);
g) SA_PC_~~5~~6:Active_Wait (see 10.2.10.2.8)(specific to SAS devices); and
h) SA_PC_~~6~~7:Idle_Wait (see 10.2.10.2.9)(specific to SAS devices).

This state machine shall start in the SA_PC_0:Powered_On state after power on.

If the device server processes a START STOP UNIT command (see SBC-3, MMC-6, or RBC) with the IMMED bit set to one, it may complete the command before completing the transition, if any, specified by the POWER CONDITION field and the START bit.

Figure 220 describes the SA_PC state machine.

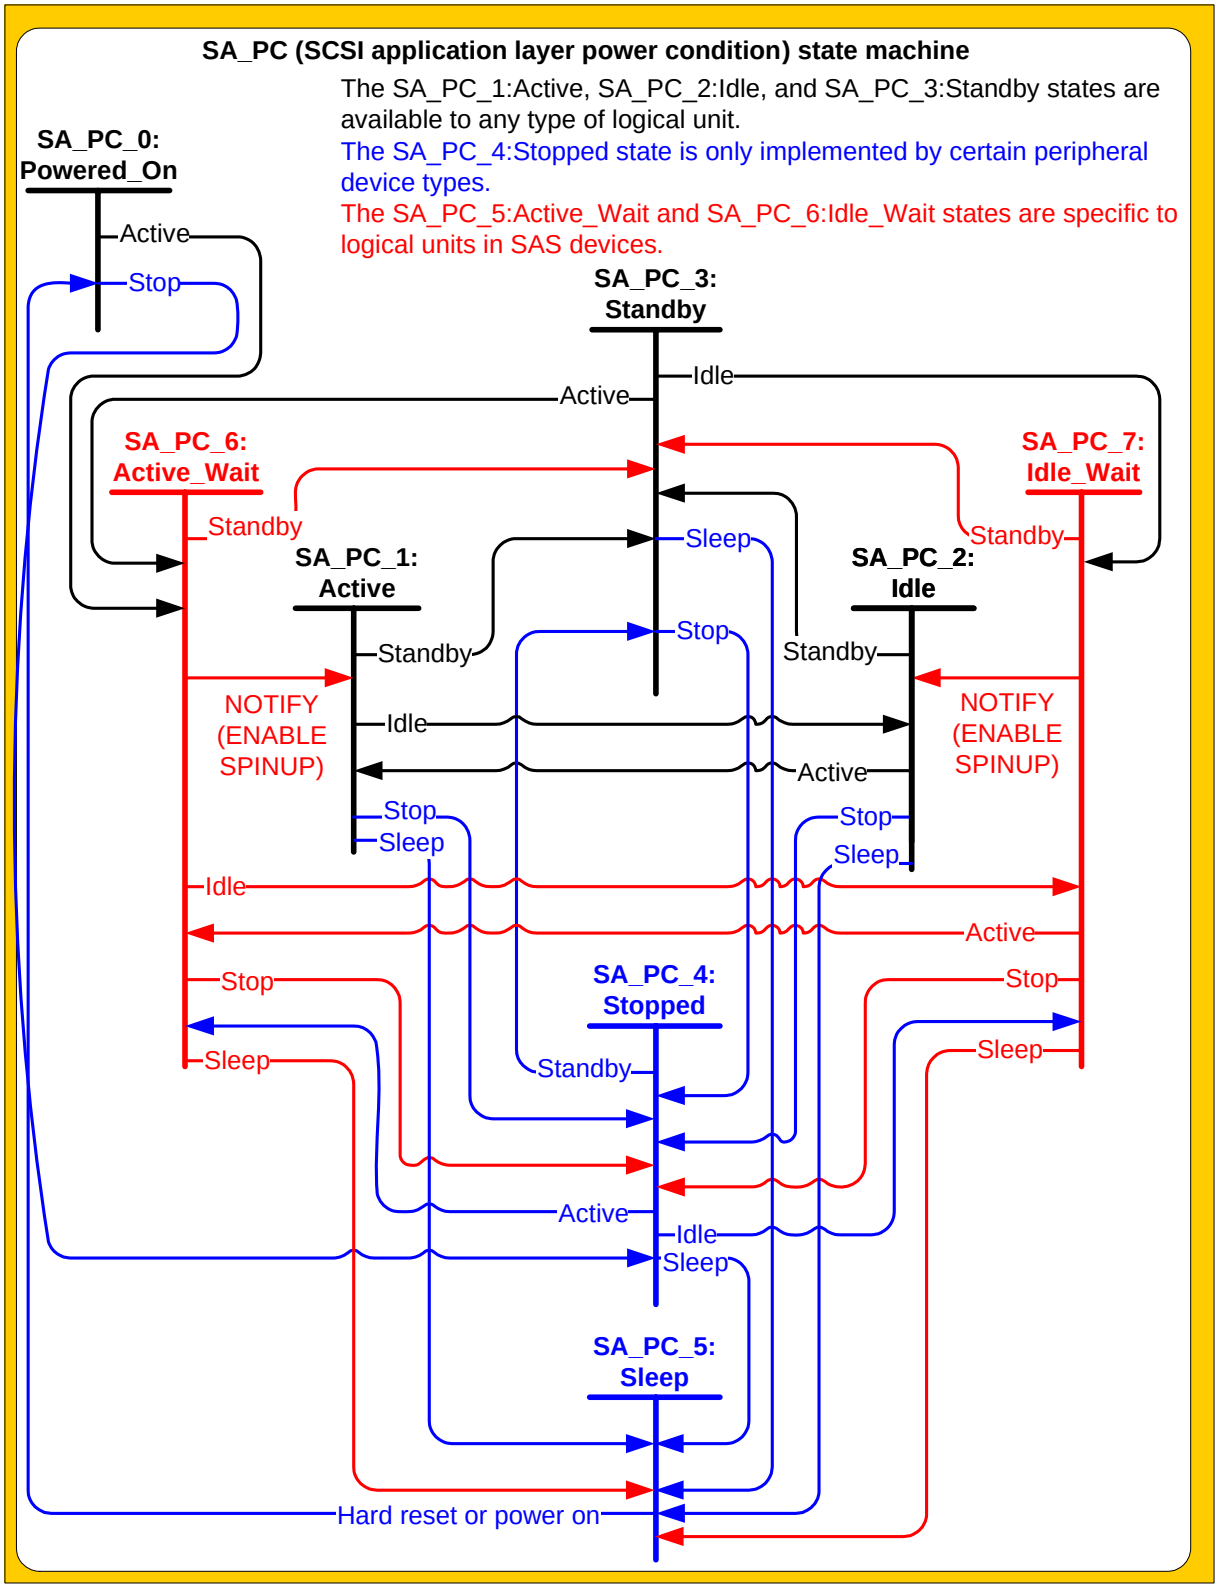

Figure 220 — SA_PC (SCSI application layer power condition) state machine for SAS [updated]

10.2.10.2.2 SA_PC_0:Powered_On state

10.2.10.2.2.1 State description

This state shall be entered upon power on. This state consumes zero time.

10.2.10.2.2.2 Transition SA_PC_0:Powered_On to SA_PC_4:Stopped

This transition shall occur if the SAS device has been configured to start in the SA_PC_4:Stopped state.

10.2.10.2.2.3 Transition SA_PC_0:Powered_On to SA_PC_~~5~~6:Active_Wait

This transition shall occur if the SAS device has been configured to start in the SA_PC_~~5~~6:Active_Wait state.

10.2.10.2.3 SA_PC_1:Active state

10.2.10.2.3.1 State description

While in this state, rotating media in block devices shall be active (i.e., rotating or spinning).

See SPC-4 for more details about this state.

10.2.10.2.3.2 Transition SA_PC_1:Active to SA_PC_2:Idle

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to IDLE is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0 is processed; or
c) the Power Condition mode page idle condition timer expires.

10.2.10.2.3.3 Transition SA_PC_1:Active to SA_PC_3:Standby

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to STANDBY is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0 is processed; or
c) the Power Condition mode page standby condition timer expires.

10.2.10.2.3.4 Transition SA_PC_1:Active to SA_PC_4:Stopped

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to zero is processed.

10.2.10.2.3.5 Transition SA_PC_1:Active to SA_PC_5:Sleep

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to SLEEP is processed.

10.2.10.2.4 SA_PC_2:Idle state

10.2.10.2.4.1 State description

While in this state, rotating media in block devices shall be active (i.e., rotating or spinning).

See SPC-4 for more details about this state.

10.2.10.2.4.2 Transition SA_PC_2:Idle to SA_PC_1:Active

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to one is processed;

b) a START STOP UNIT command with the POWER CONDITION field set to ACTIVE is processed; or
c) a command that requires the active power condition is processed.

10.2.10.2.4.3 Transition SA_PC_2:Idle to SA_PC_3:Standby

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to STANDBY is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0 is processed; or
c) the Power Condition mode page standby condition timer expires.

10.2.10.2.4.4 Transition SA_PC_2:Idle to SA_PC_4:Stopped

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to zero is processed.

10.2.10.2.4.5 Transition SA_PC_2:Idle to SA_PC_5:Sleep

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to SLEEP is processed.

10.2.10.2.5 SA_PC_3:Standby state

10.2.10.2.5.1 State description

While in this state, rotating media in block devices shall be stopped.

See SPC-4 for more details about this state.

10.2.10.2.5.2 Transition SA_PC_3:Standby to SA_PC_4:Stopped

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to zero is processed.

10.2.10.2.5.3 Transition SA_PC_3:Standby to SA_PC_5:Sleep

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to SLEEP is processed.

10.2.10.2.5.4 Transition SA_PC_3:Standby to SA_PC_~~5~~6:Active_Wait

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to one is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to ACTIVE is processed; or
c) a command that requires the active power condition is processed.

If the transition is based on a START STOP UNIT command with the IMMED bit set to zero, the device server shall not complete the command until this state machine reaches the SA_PC_1:Active state.

10.2.10.2.5.5 Transition SA_PC_3:Standby to SA_PC_~~6~~7:Idle_Wait

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to IDLE is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0 is processed; or
c) a command that requires the idle power condition is processed.

If the transition is based on a START STOP UNIT command with the IMMED bit set to zero, the device server shall not complete the command until this state machine reaches the SA_PC_2:Idle state.

10.2.10.2.6 SA_PC_4:Stopped state

10.2.10.2.6.1 State description

This state is only implemented in ~~block devices~~certain peripheral device types (see SBC-3, MMC-6, or RBC).

While in this state, rotating media shall be stopped.

See ~~SBC-3~~the command standard for the peripheral device type for more details about this state.

10.2.10.2.6.2 Transition SA_PC_4:Stopped to SA_PC_3:Standby

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to STANDBY is processed; or
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0 is processed.

10.2.10.2.6.3 Transition SA_PC_3:Stopped to SA_PC_5:Sleep

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to SLEEP is processed.

10.2.10.2.6.4 Transition SA_PC_4:Stopped to SA_PC_~~5~~6:Active_Wait

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to one is processed; or
b) a START STOP UNIT command with the POWER CONDITION field set to ACTIVE is processed.

If the transition is based on a START STOP UNIT command with the IMMED bit set to zero, the device server shall not complete the command until this state machine reaches the SA_PC_1:Active state.

10.2.10.2.6.5 Transition SA_PC_4:Stopped to SA_PC_~~6~~7:Idle_Wait

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to IDLE is processed; or
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0 is processed.

If the transition is based on a START STOP UNIT command with the IMMED bit set to zero, the device server shall not complete the command until this state machine reaches the SA_PC_2:Idle state.

10.2.10.2.7 SA_PC_5:Sleep state

10.2.10.2.7.1 State description

This state is only implemented in certain peripheral device types (see SBC-3, MMC-6, or RBC).

When a phy is inactive because the logical unit is in this state:

a) the phy transmitter shall transmit D.C. idle; and
b) the phy receiver shall be capable of detecting COMINIT.

See the command standard for the peripheral device type for more details about this state.

Editor's Note 14: Is D.C. idle necessary or is completely tristate better?

4.16.2.6.2 Transition SA_PC_5:Sleep to SA_PC_0:Powered_On

This transition shall occur after:

a) hard reset or power on.

10.2.10.2.8 SA_PC_~~5~~6:Active_Wait state

10.2.10.2.8.1 State description

This state shall only be implemented in SAS devices.

While in this state, rotating media in block devices shall be stopped. The device server shall be capable of processing commands and shall terminate each media access command or TEST UNIT READY command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, NOTIFY (ENABLE SPINUP) REQUIRED.

In response to a REQUEST SENSE command processed in this state, the device server shall return parameter data containing sense data with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, NOTIFY (ENABLE SPINUP) REQUIRED and return GOOD status for the command.

10.2.10.2.8.2 Transition SA_PC_~~5~~6:Active_Wait to SA_PC_1:Active

This transition shall occur if:

a) a NOTIFY (ENABLE SPINUP) is detected; or
b) the SAS device does not consume additional power as a result of the transition to SA_PC_1:Active.

10.2.10.2.8.3 Transition SA_PC_~~5~~6:Active_Wait to SA_PC_3:Standby

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to STANDBY is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0 is processed; or
c) the Power Condition mode page standby condition timer expires.

10.2.10.2.8.4 Transition SA_PC_~~5~~6:Active_Wait to SA_PC_4:Stopped

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to zero is processed.

10.2.10.2.8.5 Transition SA_PC_6:Active_Wait to SA_PC_5:Sleep

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to SLEEP is processed.

10.2.10.2.8.6 Transition SA_PC_~~5~~6:Active_Wait to SA_PC_~~6~~7:Idle_Wait

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to IDLE is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_IDLE_0 is processed; or
c) the Power Condition mode page idle condition timer expires.

If the transition is based on a START STOP UNIT command with the IMMED bit set to zero, the device server shall not complete the command until this state machine reaches the SA_PC_2:Idle state.

10.2.10.2.9 SA_PC_~~6~~7:Idle_Wait state

10.2.10.2.9.1 State description

This state shall only be implemented in SAS devices.

While in this state, rotating media in block devices shall be stopped. The device server shall be capable of processing commands and shall terminate each media access command or TEST UNIT READY command

with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, NOTIFY (ENABLE SPINUP) REQUIRED.

In response to a REQUEST SENSE command processed in this state, the device server shall return parameter data containing sense data with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, NOTIFY (ENABLE SPINUP) REQUIRED and return GOOD status for the command.

10.2.10.2.9.2 Transition SA_PC_~~6~~7:Idle_Wait to SA_PC_2:Idle

This transition shall occur if:

a) a NOTIFY (ENABLE SPINUP) is detected; or
b) the SAS device does not consume additional power as a result of the transition to SA_PC_2:Idle.

10.2.10.2.9.3 Transition SA_PC_~~6~~7:Idle_Wait to SA_PC_3:Standby

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to STANDBY is processed;
b) a START STOP UNIT command with the POWER CONDITION field set to FORCE_STANDBY_0 is processed; or
c) the Power Condition mode page standby condition timer expires.

10.2.10.2.9.4 Transition SA_PC_~~6~~7:Idle_Wait to SA_PC_4:Stopped

This transition shall occur if:

a) a START STOP UNIT command with the START bit set to zero is processed.

10.2.10.2.9.5 Transition SA_PC_7:Idle_Wait to SA_PC_5:Sleep

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to SLEEP is processed.

10.2.10.2.9.6 Transition SA_PC_~~6~~7:Idle_Wait to SA_PC_~~5~~6:Active_Wait

This transition shall occur if:

a) a START STOP UNIT command with the POWER CONDITION field set to ACTIVE is processed; or
b) a command that requires the active power condition is processed.

If the transition is based on a START STOP UNIT command with the IMMED bit set to zero, the device server shall not complete the command until this state machine reaches the SA_PC_1:Active state.

10.4.3.28 PHY CONTROL function

The PHY CONTROL function requests actions by the specified phy. This SMP function may be implemented by any management device server. In zoning expander devices, if zoning is enabled then this function shall only be processed from SMP initiator ports that have access to zone group 2 or the zone group of the specified phy (see 4.9.3.2).

Table 425 defines the request format.

Table 425 — PHY CONTROL request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (91h)							
2	ALLOCATED RESPONSE LENGTH							
3	REQUEST LENGTH (00h or 09h)							
4	(MSB) EXPECTED EXPANDER CHANGE COUNT							
5								(LSB)
6	Reserved							
8								
9	PHY IDENTIFIER							
10	PHY OPERATION							
11	Reserved							UPDATE PARTIAL PATHWAY TIMEOUT VALUE
12	Reserved							
23								
24	ATTACHED DEVICE NAME							
31								
32	PROGRAMMED MINIMUM PHYSICAL LINK RATE				Reserved			
33	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				Reserved			
34	Reserved							
35								
36	Reserved				PARTIAL PATHWAY TIMEOUT VALUE			
37	Reserved							
39								
40	(MSB) CRC							
43								(LSB)

The SMP FRAME TYPE field is defined in 10.4.3.2.2 and shall be set to the value defined in table 425.

The FUNCTION field is defined in 10.4.3.2.3 and shall be set to the value defined in table 425.

The ALLOCATED RESPONSE LENGTH field is defined in 10.4.3.2.4.

If the ALLOCATED RESPONSE LENGTH field is set to 00h, then the management device server shall:

a) set the RESPONSE LENGTH field to 00h in the response frame; and
b) return the first 4 bytes defined in table 428 plus the CRC field as the response frame.

If the ALLOCATED RESPONSE LENGTH field is not set to 00h, then the management device server shall:

a) set the RESPONSE LENGTH field in the response frame to the value defined in table 428 (i.e., 00h); and
b) return the response frame as specified by the ALLOCATED RESPONSE LENGTH field.

NOTE 8 - Future versions of this standard may change the value defined in table 428.

The REQUEST LENGTH field is defined in 10.4.3.2.5 and shall be set to one of the values defined in table 425 based on the LONG RESPONSE bit in the REPORT GENERAL response (see 10.4.3.4). A REQUEST LENGTH field set to 00h specifies that there are 9 dwords before the CRC field.

The EXPECTED EXPANDER CHANGE COUNT field is defined in the SMP CONFIGURE GENERAL request (see 10.4.3.18).

The PHY IDENTIFIER field specifies the phy (see 4.2.8) to which the SMP PHY CONTROL request applies.

Table 426 defines the PHY OPERATION field.

Table 426 — PHY OPERATION **field** (part 1 of 3)

Code	Operation	Description
00h	NOP	No operation.

[a] Phys compliant with previous versions of this standard did not reject this phy operation due to affiliations.
[b] Phys compliant with previous versions of this standard returned SMP FUNCTION REJECTED.

Table 426 — PHY OPERATION **field** (part 2 of 3)

Code	Operation	Description
01h	LINK RESET	If: a) a SAS phy is attached; b) a SATA phy is attached and there is no affiliation; or c) a SATA phy is attached and an affiliation exists for the STP initiator port with the same SAS address as the SMP initiator port that opened this SMP connection, then: a) if the specified phy is a physical phy, perform a link reset sequence (see 4.4) on the specified phy and enable the specified phy; and b) if the specified phy is a virtual phy, perform an internal reset and enable the specified phy. If a SATA phy is attached and an affiliation does not exist for the STP initiator port with the same SAS address as the SMP initiator port that opened this SMP connection, then the management device server shall return a function result of AFFILIATION VIOLATION in the response frame (see table 245 in 10.4.3.3). [a] See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device. Any affiliation (see 7.17.4) shall continue to be present. The phy shall bypass the SATA spinup hold state, if implemented (see 6.8.3.9). The management device server shall return the PHY CONTROL response without waiting for the LINK RESET phy operation to complete. While the LINK RESET phy operation is in progress, the management device server sets the NEGOTIATED PHYSICAL LINK RATE field and the NEGOTIATED PHYSICAL LINK RATE field to RESET_IN_PROGRESS in the SMP DISCOVER response (see 10.4.3.10).
02h	HARD RESET	If the specified phy is a physical phy, perform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the attached phy is a SAS phy or an expander phy, the link reset sequence shall include a hard reset sequence (see 4.4.2). If the attached phy is a SATA phy, the phy shall bypass the SATA spinup hold state. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device. If the specified phy is a virtual phy, perform an internal reset and enable the specified phy. Any affiliation (see 7.17.4) shall be cleared. The management device server shall return the PHY CONTROL response without waiting for the HARD RESET phy operation to complete. While the HARD RESET phy operation is in progress, the management device server sets the NEGOTIATED PHYSICAL LINK RATE field and the NEGOTIATED PHYSICAL LINK RATE field to RESET_IN_PROGRESS in the SMP DISCOVER response (see 10.4.3.10).

[a] Phys compliant with previous versions of this standard did not reject this phy operation due to affiliations.

[b] Phys compliant with previous versions of this standard returned SMP FUNCTION REJECTED.

Table 426 — PHY OPERATION **field** (part 3 of 3)

Code	Operation	Description
03h	DISABLE	Disable the specified phy (i.e., stop transmitting valid dwords and receiving dwords on the specified phy). The LINK RESET and HARD RESET operations may be used to enable the phy. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device.
04h	~~Reserved~~ SLEEP	Place the specified phy in sleep mode, where it does not transmit anything until its receiver detects COMINIT. The LINK RESET and HARD RESET operations may be used to enable the phy. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device.
05h	CLEAR ERROR LOG	Clear the error log counters reported in the REPORT PHY ERROR LOG function (see 10.4.3.11) for the specified phy.
06h	CLEAR AFFILIATION	Clear an affiliation (see 7.17.4) from the STP initiator port with the same SAS address as the SMP initiator port that opened this SMP connection. If there is no such affiliation, the management device server shall return a function result of AFFILIATION VIOLATION [b] in the response frame (see table 245 in 10.4.3.3).
07h	TRANSMIT SATA PORT SELECTION SIGNAL	This function shall only be supported by phys in an expander device. If the expander phy incorporates an STP/SATA bridge and supports SATA port selectors, the phy shall transmit the SATA port selection signal (see 6.6) which causes the SATA port selector to select the attached phy as the active host phy and make its other host phy inactive. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device. Any affiliation (see 7.17.4) shall be cleared. If the expander phy does not support SATA port selectors, then the management device server shall return a function result of PHY DOES NOT SUPPORT SATA. If the expander phy supports SATA port selectors but is attached to a SAS phy or an expander phy, the management device server shall return a function result of SMP FUNCTION FAILED in the response frame (see table 245 in 10.4.3.3).
08h	CLEAR STP I_T NEXUS LOSS	The STP I_T NEXUS LOSS OCCURRED bit in the REPORT PHY SATA function (see 10.4.3.12) shall be set to zero.
09h	SET ATTACHED DEVICE NAME	If the expander phy is attached to a SATA phy, set the ATTACHED DEVICE NAME field reported in the DISCOVER response (see 10.4.3.10) to the value of the ATTACHED DEVICE NAME field in the PHY CONTROL request.
All others	Reserved	

[a] Phys compliant with previous versions of this standard did not reject this phy operation due to affiliations.
[b] Phys compliant with previous versions of this standard returned SMP FUNCTION REJECTED.

Editor's Note 15: Interface power management will be invisible to the SCSI layers, but a device going to sleep should be visible (since there is no automatic wakeup). Add 04 SLEEP to 7.11 and

elsewhere.

If the operation specified by the PHY OPERATION field is unknown, the management device sever shall return a function result of SMP FUNCTION FAILED in the response frame (see table 245 in 10.4.3.3) and not process any other fields in the request.

If the PHY IDENTIFIER field specifies the phy which is being used for the SMP connection and a phy operation of LINK RESET, HARD RESET, or DISABLE is requested, the management device server shall not perform the requested operation and shall return a function result of SMP FUNCTION FAILED in the response frame (see table 245 in 10.4.3.3).

An UPDATE PARTIAL PATHWAY TIMEOUT VALUE bit set to one specifies that the PARTIAL PATHWAY TIMEOUT VALUE field shall be honored. An UPDATE PARTIAL PATHWAY TIMEOUT VALUE bit set to zero specifies that the PARTIAL PATHWAY TIMEOUT VALUE field shall be ignored.

The ATTACHED DEVICE NAME field is used by the SET ATTACHED DEVICE NAME phy operation and is reserved for all other phy operations. If a management application client detects the ATTACHED DEVICE NAME field set to zero in the DISCOVER response when a SATA device is attached, it shall set the ATTACHED DEVICE NAME field based on the IDENTIFY (PACKET) DEVICE data retrieved by an ATA application client in the same SAS initiator device as follows:

a) if IDENTIFY (PACKET) DEVICE data word 255 (i.e., the Integrity word) is correct and words 108-111 (i.e., the World Wide Name field) are not set to zero, set this field to the world wide name indicated by words 108-111 according to table 12 in 4.2.5;
b) if IDENTIFY (PACKET) DEVICE data word 255 (i.e., the Integrity word) is correct and words 108-111 (i.e., the World Wide Name) are set to zero, set this field to 00000000 00000000h; or
c) if IDENTIFY (PACKET) DEVICE data word 255 (i.e., the Integrity word) is not correct, set this field to 00000000 00000000h.

The PROGRAMMED MINIMUM PHYSICAL LINK RATE field specifies the minimum physical link rate the phy shall support during a link reset sequence (see 4.4.1). Table 427 defines the values for this field. This value is reported in the DISCOVER response (see 10.4.3.10). If this field is changed along with a phy operation of LINK RESET or HARD RESET, that phy operation shall utilize the new value for this field.

The PROGRAMMED MAXIMUM PHYSICAL LINK RATE field specifies the maximum physical link rates the phy shall support during a link reset sequence (see 4.4.1).Table 427 defines the values for this field. This value is reported in the DISCOVER response (see 10.4.3.10). If this field is changed along with a phy operation of LINK RESET or HARD RESET, that phy operation shall utilize the new value for this field.

Table 427 — PROGRAMMED MINIMUM PHYSICAL LINK RATE **and** PROGRAMMED MAXIMUM PHYSICAL LINK RATE **fields**

Code	Description
0h	Do not change current value
1h - 7h	Reserved
8h	1.5 Gbps
9h	3 Gbps
Ah	6 Gbps
Bh - Fh	Reserved for future physical link rates

If the PROGRAMMED MINIMUM PHYSICAL LINK RATE field or the PROGRAMMED MAXIMUM PHYSICAL LINK RATE field is set to an unsupported or reserved value, or the PROGRAMMED MINIMUM PHYSICAL LINK RATE field and PROGRAMMED MAXIMUM PHYSICAL LINK RATE field are set to an invalid combination of values (e.g., the minimum is greater than the maximum), the management device server shall not change either of their values and may return a function result of SMP FUNCTION FAILED in the response frame (see table 245 in 10.4.3.3). If it returns a function result of SMP FUNCTION FAILED, it shall not perform the requested phy operation.

The PARTIAL PATHWAY TIMEOUT VALUE field specifies the amount of time in microseconds the expander phy shall wait after receiving an Arbitrating (Blocked On Partial) confirmation from the ECM before requesting that the ECM resolve pathway blockage (see 7.12.4.5). A PARTIAL PATHWAY TIMEOUT VALUE field value of zero (i.e., 0 µs) specifies that partial pathway resolution shall be requested by the expander phy immediately upon reception of an Arbitrating (Blocked On Partial) confirmation from the ECM. This value is reported in the DISCOVER response (see 10.4.3.10). The PARTIAL PATHWAY TIMEOUT VALUE field is only honored when the UPDATE PARTIAL PATHWAY TIMEOUT VALUE bit is set to one.

The CRC field is defined in 10.4.3.2.8.

Table 428 defines the response format.

Table 428 — PHY CONTROL response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (91h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (00h)							
4	(MSB)			CRC				
7								(LSB)

The SMP FRAME TYPE field is defined in 10.4.3.3.2 and shall be set to the value defined in table 428.

The FUNCTION field is defined in 10.4.3.3.3 and shall be set to the value defined in table 428.

The FUNCTION RESULT field is defined in 10.4.3.3.4.

The RESPONSE LENGTH field is defined in 10.4.3.3.5 and shall be set to the value defined in table 428.

The CRC field is defined in 10.4.3.3.8.