To: INCITS Technical Committee T10
From: Fred Knight, Network Appliance
Email: knight@netapp.com
Date: January 14, 2008
Subject: SAT-2 – Additional Power Management support

1. ***Revision history***
   Revision 0 (5 November 2007) First revision
   Revision 1 (14 January 2008) Second revision
   1) The START STOP UNIT command is now used to map the existing SCSI power condition field to IDLE, STANDBY, and SET FEATURES commands for ATA power management support.
   2) The previously proposed TUR changes have been removed from the proposal.
   3) The Power Condition Mode Page section has been re-written to propose a new Power Condition Mode sub-page and is now used for setting the STANDBY timer values, and the APM setting. The new mode page extends the old page by adding an ATA specific sub-page.

2. ***Related documents***
   sat2r01a – SCSI / ATA Translation - 2
   spc4r11 – SCSI Primary Commands - 4
   sbc3r11 – SCSI Block Commands – 3
   08-050r0 - SPC-4 Power Condition Mode Page proposal

3. ***Overview***
   Not enough cases of ATA STANDBY power condition are reported to the SCSI host. This proposal enables reporting of more STANDBY power conditions and makes them more consistent, and adds definitions for enabling, disabling, and reporting APM mode using a SPC defined Power Condition mode page.

4. ***Expected Use Cases:***
   Manual method:
   The manual method would be used by initiators that perform their own timing of I/O. When a particular idle time (no I/O sent to the device) passed, the initiator would use the START/STOP command to manually cause the device to enter the desired power state. The initiator could step through the power states (from ACTIVE, to IDLE, to STANDBY, to STOPPED) as subsequent time intervals are passed. Once in IDLE or STANDBY modes, any I/O sent to the device would return the device to ACTIVE mode (after a short delay), and the I/O

would be serviced. When in the STOPPED state, just as it does today, the SATL would return the NOT READY-INITIALIZING COMMAND REQUIRED error, and the initiator would respond by sending the START COMMAND before reissuing the I/O request.

For the manual mode of operation, the power condition mode pages need not be used at all.

Automatic method:

The automatic method would use the power condition mode pages. The initiator would specify values in the idle timer, standby timer, or APM fields.

If the standby timer field was set, the device would automatically transition to STANDBY mode after no I/O had occurred for that time period (as defined in the ATA spec). Any new I/O sent to the device would return the device to ACTIVE mode (after a short delay), and the I/O would be serviced (again, as already defined in the ATA spec).

The SATL will ignore the IDLE timer field since ATA does not have an IDLE timer.

If the APM field was set, the device would use APM mode. In this mode, the device is in full control of the power utilization (within the bounds specified by the value in the APM field). This field is not truly a timer, but rather the ATA APM value. To activate this mode, the initiator would program the value in the mode page, and could issue the START/STOP command with the power condition field set to LU_CONTROL (7h).

To disable APM mode, the initiator would issue a MODE SELECT and set the APM bit to one and the APM VALUE to zero.

To determine the current state of the device, the MODE SENSE command would be used to examine the current state of the power condition settings (this would return the values set via the most recent MODE SELECT Command.

Author's comments/questions are in blue text, spec changes are in red.

The following co-dependant change is being proposed for the SPC-4 specification in the POWER CONDITION Mode Page (section 7.4.12 and table 309 in SPC4r11, see proposal 08-050r0).

<This is listed for reference only. Refer to the latest 08-050 proposal for exact details.>

Table 309B – ATA Power Condition Sub-Page

Byte \ Bit	7	6	5	4	3	2	1	0
0	PS	SPF (1b)	Page Code (1Ah)					
1	Sub Page Code (01h)							
2	(MSB)							
3	Page Length (2)							(LSB)
4	Reserved							APM
5	APM VALUE							

<Should it be longer with more reserved fields?>

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

The APM bit specifies if APM mode is active.

If the APM bit is set to zero, then the device server shall ignore the APM value.

If the APM bit is set to one, and the APM VALUE is non-zero, then APM mode is enabled with the APM VALUE.
If the APM bit is set to one, and the APM VALUE is zero, then APM mode is disabled.

<End of SPC change from 08-050r0 >

8.12.2 TEST UNIT READY command translation

The SATL processes the TEST UNIT READY command as follows:

1) If any condition exists that prevents the SATL from issuing commands to the ATA device, the SATL should terminate the TEST UNIT READY command with CHECK CONDITION status with the sense key set to NOT READY with the additional sense code of LOGICAL UNIT NOT READY, CAUSE NOT REPORTABLE;
2) If the device is in the stopped state as the result of processing a START STOP UNIT command (see 9.11), then the SATL shall terminate the TEST UNIT READY command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code of LOGICAL UNIT NOT READY, INITIALIZING COMMAND REQUIRED;
3) If the ATA device is performing a self-test in the foreground mode, the SATL shall terminate the command with CHECK CONDITION status, with the sense key set to NOT READY, and the additional sense code set to LOGICAL UNIT NOT READY, SELF-TEST IN PROGRESS;
4) If the SATL is processing a FORMAT UNIT command for the emulated device (see 9.2), then the SATL shall terminate the 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, FORMAT IN PROGRESS;
5) If the ATA device supports the Removable Media feature set (i.e., ATA IDENTIFY DEVICE data word 82 bit 2 is set to one), then the SATL shall issue an ATA GET MEDIA STATUS command to the ATA device. If the ATA device completes the command with the NM bit set to one in the Error register, then

the SATL shall terminate the TEST UNIT READY command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code of MEDIUM NOT PRESENT; and

6) If the ATA device completed the most recent ATA command with the DF bit set to one in the Status register, then the SATL shall terminate the TEST UNIT READY command with CHECK CONDITION status with the sense key set to HARDWARE ERROR and the additional sense code of LOGICAL UNIT FAILURE.

If none of the conditions defined in items 1 through 6 exist, then the SATL shall issue an ATA CHECK POWER MODE command to the ATA device, and:

a) If the ATA CHECK POWER MODE command completes with an error, then the SATL shall terminate the 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 DOES NOT RESPOND TO SELECTION; or

b) If the ATA CHECK POWER MODE command completes without error, then the SATL shall complete the TEST UNIT READY command with GOOD status.

~~c) If either APM mode or the STANDBY timer has been enabled, and the ATA CHECK POWER MODE command completes without error and indicates the device is in STANDBY mode, then the SATL shall terminate the 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, INITIALIZING COMMAND REQUIRED.~~

9.11.1 START STOP UNIT command overview

The START STOP UNIT command provides a method for controlling the power state of a logical unit.

If a SATL receives a command that requires medium access while the device is in the Stopped state (see SBC-2), then the SATL shall return 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.

Table 41 shows the translation for fields specified in the START STOP UNIT CDB.

The POWER CONDITION field is used to specify that the logical unit be placed into a specific power condition or to adjust a timer as defined in table 41. If the POWER CONDITION field contains a value other than 0h, then the SATL shall not consider the ATA device to be in the stopped state (see 8.12.2). If this field is not supported and is set to a value other than 0h, then the SATL shall terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

TABLE 41 – START/STOP UNIT CDB field translations

Field	Description or reference	
OPERATION CODE	Set to 1Bh. See 9.11.2 and 9.11.3	
IMMED	The SATL shall implement this field as defined in 9.11.2 and 9.11.3	
POWER CONDITION	**See SBC3 - START STOP Command definitions**	
	0 – Start valid	The SATL shall process the LOEJ and START fields as defined in section 9.11.3
	1 - Active	The SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA verify command (see 3.1.23) to the ATA

		device with the Sector Count set to one and the LBA set to a value between zero and the maximum LBA supported by the ATA device in its current configuration[a]; 3) If the ATA verify command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; 4) If the ATA verify command completes without error and the IMMED bit is set to zero, then return GOOD status (see 9.11.2).

	2 - Idle	The SATL shall:
		1) If the IMMED bit is set to one, then return GOOD status;
		2) Issue an ATA flush command (see 3.1.11) to the ATA device;
		3) If the ATA flush command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR;
		4) If the ATA flush command completes without error, then
		a. if the LOEJ bit = 0, then to enable IDLE mode immediately, issue an ATA IDLE IMMEDIATE command to the ATA device with the following register values:
		i. Feature = 0;
		ii. Sector Count = 0;
		iii. LBA = 0;
		b. if the LOEJ bit = 1, then to enable IDLE mode immediately with head unload/park, issue an ATA IDLE IMMEDIATE command to the ATA device with the following register values:
		i. Feature register = 44h;
		ii. Sector Count register = 0;
		iii. LBA = 554E4Ch;
		5) If the ATA IDLE IMMEDIATE command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; and
		6) If the ATA IDLE IMMEDIATE command completes without error and the IMMED bit is set to zero, then return GOOD status (see 9.11.2).

	3 - Standby	The SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA flush command (see 3.1.11) to the ATA device; 3) If the ATA flush command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; 4) If the ATA flush command completes without error, then the SATL shall issue an ATA STANDBY IMMEDIATE command to the ATA device; 5) If the ATA STANDBY IMMEDIATE command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to

		COMMAND SEQUENCE ERROR; and 6) If the ATA STANDBY IMMEDIATE command completes without error and the IMMED bit is set to zero, then return GOOD status (see 9.11.2).
	7 – LU_control	The SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA flush command (see 3.1.11) to the ATA device; 3) If the ATA flush command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; 4) If the ATA flush command completes without error, then to enable APM mode, the value from the ATA Power Condition Sub-Page APM VALUE shall be set in the count field and the SATL shall issue an ATA SET FEATURES – Enable/disable advanced power management (i.e., subcommand 05h) command to the ATA device. 5) If the ATA SET FEATURES command completes with any error then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; and 6) If the ATA SET FEATURES command completes without error and the IMMED bit is set to zero, then return GOOD status (see 9.11.2).
	0xA – Force I_0	The SATL shall complete the command without error and return GOOD status. <Do we need anything here? Should it error?>
	0xB – Force S_0	The SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA flush command (see 3.1.11) to the ATA device; 3) If the ATA flush command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; 4) If the ATA flush command completes without error, then the count field shall be set to zero, and the SATL shall issue an ATA STANDBY command to the ATA device; 5) If the ATA STANDBY command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; and 6) If the ATA STANDBY command completes without error and the IMMED bit is set to zero, then return GOOD status (see 9.11.2).
	All others	The SATL shall terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
LOEJ	If the POWER CONDITION field = 0, ~~T~~the SATL shall implement this field as defined in 9.11.3	

START	If the POWER CONDITION field = 0, ~~T~~the SATL shall implement this field as defined in 9.11.3
CONTROL	6.4
[a] An ATA device medium access occurs when an LBA is specified whose data is not contained in ATA device's cache memory. If a value in LBA is specified for an ATA verify command where the data is contained in ATA device's cache memory, then an ATA device may not be in the Active power mode (see ATA8-ACS) after completion of the ATA verify command.	

9.11.2 Processing ending status if an error occurs

If an error occurs during the processing of the START STOP UNIT command and the IMMED bit is set to zero, then the SATL shall terminate the START STOP UNIT command with CHECK CONDITION status with a sense key set to ABORTED COMMAND, and the additional sense code set to the value specified for the error being reported (see table 42).

If an error occurs during the processing of the START STOP UNIT command and the IMMED bit is set to one, then the SATL shall terminate the START STOP UNIT command and return CHECK CONDITION status as a deferred error (see SPC-3) with a sense key set to ABORTED COMMAND, and the additional sense code set to the value specified for the error being reported (see table 42).

9.11.3 START STOP UNIT START bit LOEJ bit combinations

If the POWER CONDITION field is set to zero, ~~T~~the SATL shall perform the actions shown in table 42 in response to a START STOP UNIT command.

See Below

Table 42 — Definition of START and LOEJ bits in the START STOP UNIT CDB

START	LOEJ	Definition
0	0	The SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA flush command (see 3.1.11) to the ATA device; 3) If the ATA flush command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; 4) If the ATA flush command completes without error, then issue an ATA STANDBY IMMEDIATE command to the ATA Sector Count set to zero; 5) If the ATA STANDBY IMMEDIATE command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to COMMAND SEQUENCE ERROR; and 6) If the ATA STANDBY IMMEDIATE command completes without error and the IMMED bit is set to zero, then return GOOD status (see 9.11.2) [a].
0	1	If the ATA device supports the Removable Media feature set (see ATA/ATAPI-7), then the SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA MEDIA EJECT command to the ATA device; 3) If the ATA MEDIA EJECT command completes with any error, then process ending status according to the IMMED bit (see 9.11.2) with the additional sense code set to MEDIA LOAD OR EJECT FAILED; and 4) If the MEDIA EJECT command completes without error and the IMMED bit is set to zero, then return GOOD status. If the ATA device does not support the Removable Media feature set, then the SATL shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
1	0	The SATL shall: 1) If the IMMED bit is set to one, then return GOOD status; 2) Issue an ATA verify command (see 3.1.23) to the ATA Sector Count set to one and the LBA set to a value between zero and the maximum LBA supported by the ATA device in its current configuration [b]; and 3) If the IMMED bit is set to one, then return GOOD status when command completion is received for the ATA verify command (see 3.1.23) [c].
1	1	The SATL shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST, with the additional sense code set to INVALID FIELD IN CDB.

[a] After returning GOOD status for a START STOP UNIT command with the START bit set to zero, the SATL shall consider the ATA device to be in the Stopped power state (see SBC-2).

[b] An ATA device medium access occurs when an LBA is specified whose data is not contained in ATA device's cache memory. If a value in LBA is specified for an ATA verify command where the data is contained in ATA device's cache memory, then an ATA device may not be in the Active power mode (see ATA8-ACS) after completion of the ATA verify command.

[c] After returning GOOD status for a START STOP UNIT command with the START bit set to one, the SATL shall consider the ATA device to be in the Active power state (see SBC-2).

and the power condition field set to zero,

<….>

Section 9.11.3 (table 42 <above>) – TYPO correction:

4) If the ATA flush command completes without error, then issue an ATA STANDBY IMMEDIATE command to the ATA device with the Sector Count set to zero;

2) Issue an ATA verify command (see 3.1.23) to the ATA device with the Sector Count set to one and the LBA set to a value between zero and the maximum LBA supported by the ATA device in its current configuration [b]; and

Section 10.1.9:

The SCSI START/STOP command is used to specifically change power state. The Power Condition mode page allows setting of the ATA APM mode setting, and the ATA STANDBY timer value. It also provides information about current power states.

Table X – Standby Timer Mappings

Power Condition Mode page STANDBY CONDITION TIMER value (uses 100 millisecond granularity)	ATA count register value (uses 5 second / 30 second granularity / …)
1-12000 (.001 second to 1200 seconds)	(INT(Value - 1) / 50) + 1
12001 – 12599	0xFC
12600 (21 minutes)	0xFC
12601 – 12749	0xFF
12750 (21 minutes 15 seconds)	0xFF
12751 – 17999	0xF1
18000 – 198000 (30 minutes – 330 minutes)	(INT(Value / 18000) + 240
198001 – 287999	0xFD
288000 – 432000 (8 hours – 12 hours)	0xFD
Others	0xFD
Note: Times are approximate.	

Table 65+ – Power Condition Control mode page fields

Field	Changeable	Description or reference
PS	n/a	Unspecified (see 3.4.2)
SPF	no	Shall be set to zero.
Page Code	no	Shall be set to 1Ah
Page Length	no	Shall be set to 1Ah
IDLE	no	When processing a MODE SENSE command, the IDLE bit shall be returned as zero. When processing a MODE SELECT command, if the IDLE bit is set to one, then the SATL shall terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.
STANDBY	yes	When processing a MODE SENSE command, if a standby timer is enabled, the STANDBY bit shall be returned as one. If the standby timer is not enabled, the STANDBY bit shall be returned as zero. When processing a MODE SELECT command, if the STANDBY bit is set to one, then the SATL shall issue the ATA STANDBY command to the ATA device, and the value in the STANDBY CONDITION TIMER field shall be translated as defined in tabled X and used to set the Timer period value (TPV) (aka count field).
IDLE CONDITION	no	When processing a MODE SENSE command, this field

TIMER		shall be returned as zero. When processing a MODE SELECT command, this field shall be ignored.
STANDBY CONDITION TIMER	yes	When processing a MODE SENSE command: If the standby timer is not enabled, then the STANDBY CONDITION TIMER shall return zero. If a standby timer is enabled, then it shall be translated (see table X) and returned in this field (if the translated value is in a range, then the lowest value in that range shall be returned). When processing a MODE SELECT command: If STANDBY is set to one, then the value in this field shall be translated as defined in table X and used to set the Timer period value (TPV) (aka count field).

Table 65+ ATA Power Condition Sub-Page fields

Field	Changeable	Description or reference
PS	n/a	Unspecified (see 3.4.2)
SPF	no	Shall be set to one
Page Code	no	Shall be set to 1Ah
Page Sub Code	no	Shall be set to 01h
APM	yes	When processing a MODE SENSE command, if APM mode is enabled, the APM bit shall be returned as one. If APM mode is not enabled, the APM bit shall be returned as zero. When processing a MODE SELECT command, if the APM bit is set to one, then the SATL shall alter APM mode by issuing an ATA SET FEATURES command. If the APM VALUE contains a non-zero value, the ATA SET FEATURES – Enable/disable advanced power management (i.e., subcommand 05h) command shall be used and APM VALUE shall be used to set the power management level (count field). If APM VALUE contains a zero, the ATA SET FEATURES – Disable advanced power management (i.e., subcommand 85h) command shall be used.
APM VALUE	yes	When processing a MODE SENSE command, if APM mode has been enabled, the last value set shall be returned. If APM mode is not enabled, APM VALUE shall return zero. When processing a MODE SELECT command: If APM is set to one, then the value in this field shall be used in the ATA SET FEATURES COMMAND to set the power management level (count field).