Date: Feb 22, 1991 X3T9.2/91-014 Rev 1

To: X3T9.2 Committee (SCSI)

From: George Penokie (IBM)

Subject: Control of SCSI Device Power Consumption

Proposal:

There is a large computer environment developing where reducing the amount of power consumed is critical. In this environment every milliwatt that can be saved is important. At this time most of the computers in this environment, that have hard drives attached, use the ATA, XT, or IDE type of interface. That interface has commands which allow the system to control the drives power conditions. But that interface has limitations and there is a desire to move to SCSI.

Before there can be a move to SCSI, however, something has to be defined within SCSI to allow an equivalent to the ATA low power conditions. In light of that I would like the following considered for SCSI-3.

7.x.x.x. Power States

The optional power states permit the initiator to modify the behavior of a target in a manner which reduces the power required to operate. There is no notification to the initiator that a target has entered into one of the low power states.

The lowest power consumption occurs in the Sleep state. When in the Sleep state a target needs a bus reset to be activated.

In the Shutdown state a target is capable of accepting commands, but media is not immediately accessible (eg spindle is stopped).

In the Standby state a target is capable of accepting commands, but media is not immediately accessible (eg spindle running but actuator is turned off).

In the Idle state a target is capable of responding quickly to media access requests. However, a target in the Idle state may take longer to complete the execution of a command because it may have to activate some circuitry.

In the Active state a target is capable of responding immediately to media access requests, and operations complete execution in the shortest possible time.

7.3.3.x. Power State Page

Table 7-xx: Power State Page

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	Reserved		Page Code (0Dh)				
1				Page Length (0Eh)				
2				Reserved				
3						Shutdown	Standby	Idle
4				Reserved				
5				Reserved				
6	(MSB)			Shutdown State Timer				
7								(LSB)
8				Reserved				
9				Reserved				
10	(MSB)			Standby State Timer				
11								(LSB)
12				Reserved				
13				Reserved				
14	(MSB)			Idle State Timer				
15								(LSB)

The power state page (Table 7-xx) provides the initiator the means to control the length of time a target will delay before lowering it's power requirements. There is no advance notification to the initiator that a target has entered into one of the low power states.

An shutdown bit of one indicates a target shall use the Shutdown State Timer to determine the length of inactivity time to wait before entering the Shutdown state. An shutdown bit of zero indicates a target shall not enter the Shutdown state.

A standby bit of one indicates a target shall use the Standby State Timer to determine the length of inactivity time to wait before entering the Standby state. A standby bit of zero indicates a target shall not enter the Standby state.

An idle bit of one indicates a target shall use the Idle State Timer to determine the length of inactivity time to wait before entering the Idle state. An idle bit of zero indicates a target shall not enter the Idle state.

The Shutdown State Timer field indicates the inactivity time in 100 millisecond increments that the target shall wait before entering the Shutdown State. This timer shall only count if the Idle State Timer and the Standby State Timer are both equal to zero.

The Standby State Timer field indicates the inactivity time in 100 millisecond increments that the target shall wait before entering the Standby State. This timer shall only count if the Idle State Timer is equal to zero.

The Idle State Timer field indicates the inactivity time in 100 millisecond increments that the target shall wait before entering the Idle State.

7.2.14. REQUEST SENSE Command

Table 7-34: REQUEST SENSE Command

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code (03h)							
1	Logical Unit Number			Reserved				
2	Reserved							
3	Reserved							
4	Allocation Length							
5	Control							

The REQUEST SENSE command (Table 7-34) requests that the target transfer sense data to the initiator.

The sense data:
(1) shall be available if the previous command to the specified I_T_x nexus terminated with CHECK CONDITION or COMMAND TERMINATED status
(2) shall be available if other information (e.g., medium position) is available in any field
(3) may be available if the previous command to the specified I_T_x nexus ended with an unexpected BUS FREE error (5.1.1).
(4) shall be available if the device is in one of the low power states.

If the target is in the Shutdown state, Standby state, or Idle state on receipt of the REQUEST SENSE command the target shall return a sense key of NO SENSE and additional sense codes of SHUTDOWN STATE ACTIVE, STANDBY STATE ACTIVE, or IDLE STATE ACTIVE. On completion of the command the target shall immediately return to the same power state which was active before the REQUEST SENSE command was received. Any active power state timer shall be suspended on receipt of the REQUEST SENSE command and resumed on completion of the command.

If the target has no sense data available to return, it shall return a sense key of NO SENSE and an additional sense code of NO ADDITIONAL SENSE INFORMATION.

The sense data shall be preserved by the target for the initiator until retrieved by the REQUEST SENSE command or until the receipt of any other command for the same I_T_x nexus (see 6.6). Sense data shall be cleared upon receipt of any subsequent command (including REQUEST SENSE) to the same I_T_x nexus.

IMPLEMENTORS NOTE: Some target implementations do not update sense data except on commands that return CHECK CONDITION or COMMAND TERMINATED status. Thus when polling for a logical unit to become ready, the initiator should issue TEST UNIT READY commands until GOOD status is returned. If desired, the initiator may issue REQUEST SENSE commands after the TEST UNIT READY commands that return CHECK CONDITION or COMMAND TERMINATED status to obtain the sense data.

The target shall return CHECK CONDITION status for a REQUEST SENSE command only to report errors specific to the command itself. For example:

(1) A non-zero reserved bit is detected in the command descriptor block.
(2) An unrecovered parity error is detected on the data bus.
(3) A target malfunction prevents return of the sense data.

If a recovered error occurs during the execution of the REQUEST SENSE command, the target shall return the sense data with GOOD status. If a target returns CHECK CONDITION status for a REQUEST SENSE command the sense data may be invalid.

IMPLEMENTORS NOTE: The sense data appropriate to the selection of an invalid logical unit is defined in 6.5.3.

Targets shall be capable of returning eighteen bytes of data in response to a REQUEST SENSE command. If the allocation length is eighteen or greater and a target returns less than eighteen bytes of data the initiator should assume that the bytes not transferred would have been zeros had the target returned those bytes. Initiators can determine how much sense data has been returned by examining the allocation length parameter in the command descriptor block and the additional sense length in the sense data. Targets shall not adjust the additional sense length to reflect truncation if the allocation length is less than the sense data available.

Table 7-41: ASC and ASCQ Assignments

ASC AND ASCQ ASSIGNMENTS

D = DIRECT ACCESS DEVICE
T = SEQUENTIAL ACCESS DEVICE
L = PRINTER DEVICE
P = PROCESSOR DEVICE
W = WRITE ONCE READ MULTIPLE DEVICE
R = READ ONLY (CD-ROM) DEVICE
S = SCANNER DEVICE
O = OPTICAL MEMORY DEVICE
M = MEDIA CHANGER DEVICE
C = COMMUNICATION DEVICE

BYTE 12	13	DTLPWRSOMC	DESCRIPTION
5E	00	DTLPWRSO C	LOW POWER STATE ACTIVE
5E	01	DTLPWRSO C	IDLE STATE ACTIVE
5E	02	DTLPWRSO C	STANDBY STATE ACTIVE
5E	03	DTLPWRSO C	SHUTDOWN STATE ACTIVE

8.2.17. START STOP UNIT Command

Table 8-33: START STOP UNIT Command

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code (1Bh)							
1	Logical Unit Number			Reserved				Immed
2	Reserved							
3	Reserved							
4	Power States						LoEj	Start
5	Control							

The START STOP UNIT command (Table 8-33) requests that the target enable or disable the logical unit for media access operations.

An immediate (Immed) bit of one indicates that status shall be returned as soon as the command descriptor block has been validated. An Immed bit of zero indicates that status shall be returned after the operation is completed.

The power states field requests the logical unit to be placed into the power state defined in Table x-xx. There shall be no indication from the logical unit to any device that it has entered the requested power state.

Table x-xx: Power States

Code	Description
0h	Make no change in power states
1h	Place device into Idle state
2h	Place device into Standby state
3h	Place device into Shutdown state
4h	Place device into Sleep state
5-6h	Reserved
7h	Place device into the Active state

In the Idle, Standby, and Shutdown states the logical unit shall respond to any operations in the most expedient fashion possible.

In the Sleep state the logical unit shall only respond to a bus reset condition.

A load eject (LoEj) bit of zero requests that no action be taken regarding loading or ejecting the medium. A LoEj bit of one requests that the medium shall be be unloaded if the start bit is zero. A LoEj bit of one requests that the medium is to be loaded if the start bit is one.

A start bit of one requests the logical unit be made ready for use. A start bit of zero requests that the logical unit be stopped (media cannot be accessed by the initiator).

Targets that contain cache memory shall implicitly perform a SYNCHRONIZE CACHE command for the entire medium prior to executing the STOP UNIT command or entering any low power state. [Is this right? GOP]