Date: May 03, 1991 X3T9.2/91-014 Rev 3

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 conditions

The optional power conditions 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 power conditions.

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

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

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

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

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

When the power conditions are being controlled by the Power Condition Page and a command is received while the device is in a low power condition the device may turn on just enough power to allow the command to execute. Only the timer(s) which match the level of power required to execute the command shall be reset and then restarted on completion of that command.

When the power conditions are being controlled by the Start/Stop Command the target shall change power conditions only on receipt of a Start/Stop

Command or a hard reset. Any Power Condition Page timers that are active on receipt of a Start/Stop Command which sets a power condition shall be suspended until a Start/Stop Command is received which returns control of the power condition to the target or a hard reset occurs. If a command is received which requires more power than allowed by the Start/Stop Command power condition the target shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOW POWER CONDITION ACTIVE.

7.3.3.x. Power Condition Page

Table 7-xx: Power Condition Page

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

The power condition page (Table 7-xx) provides the initiator the means to control the length of time a target will delay before changing its power requirements. There is no notification to the initiator that a target has entered into one of the power conditions.

On the receipt of a command only the timer(s) which match the level of power required to execute the command shall be reset and then restarted on completion of that command.

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

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

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

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

If the Idle bit is one, a value of zero in the Idle Condition Timer indicates the target shall immediately enter the Idle condition on completion of any command.

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

If the Standby bit is one and the Idle bit is zero, a value of zero in the Standby Condition Timer indicates the target shall immediately enter the standby condition on completion of any command.

If the Standby bit is one and the Idle bit is one, a value of zero in the Standby Condition Timer indicates the target shall immediately enter the Standby condition when the Idle Condition Timer equals zero.

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

If the Shutdown bit is one, the Standby bit is zero, and the Idle bit is zero, a value of zero in the Shutdown Condition Timer indicates the target shall immediately enter the shutdown condition on completion of any command.

If the Shutdown bit is one, the Standby bit is zero, and the Idle bit is one, a value of zero in the Shutdown Condition Timer indicates the target shall immediately enter the Shutdown condition when the Idle Condition Timer equals zero.

If the Shutdown bit is one and the Standby bit is one, a value of zero in the Shutdown Condition Timer indicates the target shall immediately enter the Shutdown condition when the Standby Condition Timer equals zero.

Figure 7-xx shows graphically the relationships between the different power conditions and their timers.

X3T9.2/91-014 Rev 3

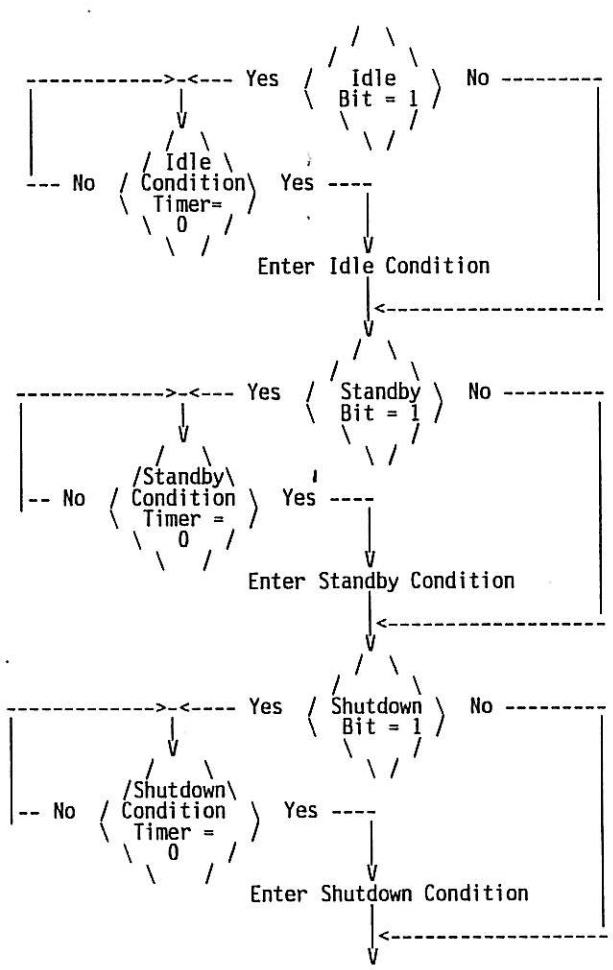

Figure 7-xx: Power conditions Flowchart

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 conditions.

If the target is in the Shutdown condition, Standby condition, or Idle condition on receipt of the REQUEST SENSE command the target shall return a sense key of NO SENSE and additional sense codes of SHUTDOWN condition ACTIVE, STANDBY condition ACTIVE, or IDLE condition ACTIVE. On completion of the command the target shall immediately return to the same power condition which was active before the REQUEST SENSE command was received. Any active power condition 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 condition ACTIVE
5E	01 DTLPWRSO C	IDLE condition ACTIVE
5E	02 DTLPWRSO C	STANDBY condition ACTIVE
5E	03 DTLPWRSO C	SHUTDOWN condition 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	Reserved			Power Conditions			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 conditions field requests the logical unit to be placed into the power condition defined in Table x-xx. The power conditions field is only valid if the Start bit is one. If the Start bit is zero the power conditions field shall be ignored by the logical unit.

There shall be no indication from the logical unit to any device that it has entered the requested power condition.

After a Start/Stop Command is used to set a power condition the logical unit shall change power conditions only on receipt of another Start/Stop Command or a hard reset.

Any Power Condition Page timers (see 7.x.x) that are active on receipt of a Start/Stop Command which sets a power condition shall be suspended until a Start/Stop Command is received which returns control of the power condition to the target or a hard reset occurs.

If a command is received which requires more power than allowed by the Start/Stop Command power condition the target shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOW POWER CONDITION ACTIVE.

It is not an error to request the logical unit be placed into a power condition which already exists.

Table x-xx: Power Conditions

Code	Description
0h	No change in power conditions or in which device is controlling power conditions
1h	Place device into the Active condition
2h	Place device into Idle condition
3h	Place device into Standby condition
4h	Place device into Shutdown condition
5h	Place device into Sleep condition
6h	Reserved
7h	Give control of power conditions to target

In the Sleep condition 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 condition.