

T10/05-167 revision 2

Date: June 07, 2005

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: SAS-2: New NOTIFY to indicate imminent power failure

Overview

Both parallel SCSI and FC provide a method to inform a SCSI device that a power failure may occur within a short time (e.g., 1 millisecond). On both of those interfaces the imminent power fail is indicated by electrical signals on the connector. SAS has no pins on it's connectors that could be used for this function. However, SAS provides the NOTIFY primitive which can be sent directly from an expander or initiator without any request from the application client at any time.

This proposal uses one of the reserved NOTIFY primitives to indicate to all the logical units in the SCSI target device that a power failure may occur within a time defined in a mode page. The effect of this NOTIFY would be as if all the logical units received a CLEAR TASK SET task management function.

After the NOTIFY is received all the phys on the target device would issue an OPEN_REJECT(RETRY) to any OPEN received after the NOTIFY until the power fail timer expires.

If the logical unit supports the retry delay timer it should return the remaining time as indicated by the power fail timer in the RETRY DELAY TIMER field of the RESPONSE information unit for each command that receives a BUSY status.

In addition to clearing the task set the logical unit would be required to stop processing any write commands on the next block boundary (i.e., all write activity would continue until a block boundary was reached then all writing would stop and the command aborted).

If the power fail timer expires (i.e., the power really did not fail) the logical units establishes an UNIT ATTENTION of all initiators of POWER FAIL EVENT COMMANDS CLEARED (2F01h).

SAS 2 changes

0.0.1 Primitive summary

Table 1 defines the primitives not specific to the type of connection.

Table 1 — Primitives not specific to type of connection (part 1 of 2)

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
AIP (NORMAL)	NoConn		E		I	E	T	Single
AIP (RESERVED 0)	NoConn				I	E	T	Single
AIP (RESERVED 1)	NoConn				I	E	T	Single
AIP (RESERVED 2)	NoConn				I	E	T	Single
AIP (RESERVED WAITING ON PARTIAL)	NoConn				I	E	T	Single
AIP (WAITING ON CONNECTION)	NoConn		E		I	E	T	Single
AIP (WAITING ON DEVICE)	NoConn		E		I	E	T	Single
AIP (WAITING ON PARTIAL)	NoConn		E		I	E	T	Single
ALIGN (0)	All	I	E	T	I	E	T	Single
ALIGN (1)	All	I	E	T	I	E	T	Single
ALIGN (2)	All	I	E	T	I	E	T	Single
ALIGN (3)	All	I	E	T	I	E	T	Single
BREAK	All	I	E	T	I	E	T	Redundant
BROADCAST (CHANGE)	NoConn	I	E		I	E	T	Redundant
BROADCAST (SES)	NoConn			T	I	E	T	Redundant
BROADCAST (RESERVED 1)	NoConn				I	E	T	Redundant
BROADCAST (RESERVED 2)	NoConn				I	E	T	Redundant
BROADCAST (RESERVED 3)	NoConn				I	E	T	Redundant
BROADCAST (RESERVED 4)	NoConn				I	E	T	Redundant
BROADCAST (RESERVED CHANGE 0)	NoConn				I	E	T	Redundant
BROADCAST (RESERVED CHANGE 1)	NoConn				I	E	T	Redundant
CLOSE (CLEAR AFFILIATION)	STP	I					T	Triple
CLOSE (NORMAL)	Conn	I		T	I		T	Triple
CLOSE (RESERVED 0)	Conn				I		T	Triple
CLOSE (RESERVED 1)	Conn				I		T	Triple
EOAF	NoConn	I	E	T	I	E	T	Single
ERROR	All		E		I	E	T	Single
HARD_RESET	NoConn	I	E		I	E	T	Redundant
NOTIFY (ENABLE SPINUP)	All	I	E				T	Single
NOTIFY (POWER FAILURE EXPECTED)	All	I	E				T	Single
NOTIFY (RESERVED 1)	All				I	E	T	Single
NOTIFY (RESERVED 2)	All				I	E	T	Single
OPEN_ACCEPT	NoConn	I		T	I		T	Single

Table 1 — Primitives not specific to type of connection (part 2 of 2)

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
OPEN_REJECT (BAD DESTINATION)	NoConn		E		I		T	Single
OPEN_REJECT (CONNECTION RATE NOT SUPPORTED)	NoConn	I	E	T	I		T	Single
OPEN_REJECT (NO DESTINATION)	NoConn		E		I		T	Single
OPEN_REJECT (PATHWAY BLOCKED)	NoConn		E		I		T	Single
OPEN_REJECT (PROTOCOL NOT SUPPORTED)	NoConn	I		T	I		T	Single
OPEN_REJECT (RESERVED ABANDON 0)	NoConn				I		T	Single
OPEN_REJECT (RESERVED ABANDON 1)	NoConn				I		T	Single
OPEN_REJECT (RESERVED ABANDON 2)	NoConn				I		T	Single
OPEN_REJECT (RESERVED ABANDON 3)	NoConn				I		T	Single
OPEN_REJECT (RESERVED CONTINUE 0)	NoConn				I		T	Single
OPEN_REJECT (RESERVED CONTINUE 1)	NoConn				I		T	Single
OPEN_REJECT (RESERVED INITIALIZE 0)	NoConn				I		T	Single
OPEN_REJECT (RESERVED INITIALIZE 1)	NoConn				I		T	Single
OPEN_REJECT (RESERVED STOP 0)	NoConn				I		T	Single
OPEN_REJECT (RESERVED STOP 1)	NoConn				I		T	Single
OPEN_REJECT (RETRY)	NoConn	I		T	I		T	Single
OPEN_REJECT (STP RESOURCES BUSY)	NoConn		E	T	I			Single
OPEN_REJECT (WRONG DESTINATION)	NoConn	I		T	I		T	Single
SOAF	NoConn	I	E	T	I	E	T	Single

^a The Use column indicates when the primitive is used:
a) NoConn: SAS physical links, outside connections;
b) Conn: SAS physical links, inside connections;
c) All: SAS physical links, both outside connections or inside any type of connection; or
d) STP: SAS physical links, inside STP connections.

^b The From and To columns indicate the type of ports that originate each primitive or are the intended destinations of each primitive:
a) I for SAS initiator ports;
b) E for expander ports; and
c) T for SAS target ports.
Expander ports are not considered originators of primitives that are passing through from expander port to expander port.

^c The Primitive sequence type columns indicate whether the primitive is sent as a single primitive sequence, a repeated primitive sequence, a continued primitive sequence, a triple primitive sequence, or a redundant primitive sequence (see 7.2.4).

0.0.2 Primitive encodings

Table 2 defines the primitive encoding for primitives not specific to type of connection.

Table 2 — Primitive encoding for primitives not specific to type of connection (part 1 of 2)

Primitive	Character			
	1 st	2 nd	3 rd	4 th (last)
AIP (NORMAL)	K28.5	D27.4	D27.4	D27.4
AIP (RESERVED 0)	K28.5	D27.4	D31.4	D16.7
AIP (RESERVED 1)	K28.5	D27.4	D16.7	D30.0
AIP (RESERVED 2)	K28.5	D27.4	D29.7	D01.4
AIP (RESERVED WAITING ON PARTIAL)	K28.5	D27.4	D01.4	D07.3
AIP (WAITING ON CONNECTION)	K28.5	D27.4	D07.3	D24.0
AIP (WAITING ON DEVICE)	K28.5	D27.4	D30.0	D29.7
AIP (WAITING ON PARTIAL)	K28.5	D27.4	D24.0	D04.7
ALIGN (0)	K28.5	D10.2	D10.2	D27.3
ALIGN (1)	K28.5	D07.0	D07.0	D07.0
ALIGN (2)	K28.5	D01.3	D01.3	D01.3
ALIGN (3)	K28.5	D27.3	D27.3	D27.3
BREAK	K28.5	D02.0	D24.0	D07.3
BROADCAST (CHANGE)	K28.5	D04.7	D02.0	D01.4
BROADCAST (SES)	K28.5	D04.7	D07.3	D29.7
BROADCAST (RESERVED 1)	K28.5	D04.7	D01.4	D24.0
BROADCAST (RESERVED 2)	K28.5	D04.7	D04.7	D04.7
BROADCAST (RESERVED 3)	K28.5	D04.7	D16.7	D02.0
BROADCAST (RESERVED 4)	K28.5	D04.7	D29.7	D30.0
BROADCAST (RESERVED CHANGE 0)	K28.5	D04.7	D24.0	D31.4
BROADCAST (RESERVED CHANGE 1)	K28.5	D04.7	D27.4	D07.3
CLOSE (CLEAR AFFILIATION)	K28.5	D02.0	D07.3	D04.7
CLOSE (NORMAL)	K28.5	D02.0	D30.0	D27.4
CLOSE (RESERVED 0)	K28.5	D02.0	D31.4	D30.0
CLOSE (RESERVED 1)	K28.5	D02.0	D04.7	D01.4
EOAF	K28.5	D24.0	D07.3	D31.4
ERROR	K28.5	D02.0	D01.4	D29.7
HARD_RESET	K28.5	D02.0	D02.0	D02.0

Table 2 — Primitive encoding for primitives not specific to type of connection (part 2 of 2)

Primitive	Character			
	1 st	2 nd	3 rd	4 th (last)
NOTIFY (ENABLE SPINUP)	K28.5	D31.3	D31.3	D31.3
NOTIFY (POWER FAILURE EXPECTED)	K28.5	D31.3	D07.0	D01.3
NOTIFY (RESERVED 1)	K28.5	D31.3	D01.3	D07.0
NOTIFY (RESERVED 2)	K28.5	D31.3	D10.2	D10.2
OPEN_ACCEPT	K28.5	D16.7	D16.7	D16.7
OPEN_REJECT (BAD DESTINATION)	K28.5	D31.4	D31.4	D31.4
OPEN_REJECT (CONNECTION RATE NOT SUPPORTED)	K28.5	D31.4	D04.7	D29.7
OPEN_REJECT (NO DESTINATION)	K28.5	D29.7	D29.7	D29.7
OPEN_REJECT (PATHWAY BLOCKED)	K28.5	D29.7	D16.7	D04.7
OPEN_REJECT (PROTOCOL NOT SUPPORTED)	K28.5	D31.4	D29.7	D07.3
OPEN_REJECT (RESERVED ABANDON 0)	K28.5	D31.4	D02.0	D27.4
OPEN_REJECT (RESERVED ABANDON 1)	K28.5	D31.4	D30.0	D16.7
OPEN_REJECT (RESERVED ABANDON 2)	K28.5	D31.4	D07.3	D02.0
OPEN_REJECT (RESERVED ABANDON 3)	K28.5	D31.4	D01.4	D30.0
OPEN_REJECT (RESERVED CONTINUE 0)	K28.5	D29.7	D02.0	D30.0
OPEN_REJECT (RESERVED CONTINUE 1)	K28.5	D29.7	D24.0	D01.4
OPEN_REJECT (RESERVED INITIALIZE 0)	K28.5	D29.7	D30.0	D31.4
OPEN_REJECT (RESERVED INITIALIZE 1)	K28.5	D29.7	D07.3	D16.7
OPEN_REJECT (RESERVED STOP 0)	K28.5	D29.7	D31.4	D07.3
OPEN_REJECT (RESERVED STOP 1)	K28.5	D29.7	D04.7	D27.4
OPEN_REJECT (RETRY)	K28.5	D29.7	D27.4	D24.0
OPEN_REJECT (STP RESOURCES BUSY)	K28.5	D31.4	D27.4	D01.4
OPEN_REJECT (WRONG DESTINATION)	K28.5	D31.4	D16.7	D24.0
SOAF	K28.5	D24.0	D30.0	D01.4

0.0.2.1 NOTIFY

NOTIFY may be transmitted in place of any ALIGN (see 7.2.5.2) being transmitted for clock skew management (see 7.3), rate matching (see 7.13), or STP initiator phy throttling (see 7.17.2). Substitution of a NOTIFY may or may not affect the ALIGN rotation (i.e., the NOTIFY may take the place of one of the ALIGNs in the rotation through ALIGN (0), ALIGN (1), ALIGN (2), or ALIGN (3) or it may delay the rotation). A specific NOTIFY shall not be transmitted a second time until at least three ALIGNs or different NOTIFYS have been transmitted.

NOTIFY shall not be forwarded through expander devices. Expander devices shall substitute an ALIGN for a NOTIFY if necessary.

SAS target devices are not required to detect every transmitted NOTIFY.

The versions of NOTIFY representing different reasons are defined in table 3.

Table 3 — NOTIFY primitives

Primitive	Description
NOTIFY (ENABLE SPINUP)	Specify to an SAS target device that it may temporarily consume additional power while transitioning into the active or idle power condition state.
NOTIFY (POWER FAILURE EXPECTED)	Specify to an SAS target device that power may be lost within a time specified by the power fail timeout timer (see 0.0.2.1.1).
NOTIFY (RESERVED 1)	Reserved.
NOTIFY (RESERVED 2)	Reserved.

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-3) and/or the START STOP UNIT command (see SBC-2), as described in 10.2.9.

SAS initiator devices and expander devices shall use NOTIFY (ENABLE SPINUP) while attached to SAS target devices (i.e., devices that report SSP target support in their IDENTIFY address frames). They shall transmit one NOTIFY (ENABLE SPINUP) after power on when the enclosure is ready for initial spin-up. After the initial NOTIFY (ENABLE SPINUP), they shall transmit NOTIFY (ENABLE SPINUP) periodically. Otherwise, the selection of when and how often to transmit NOTIFY (ENABLE SPINUP) is outside the scope of this standard.

NOTE 1 - The SAS initiator device or expander device uses NOTIFY (ENABLE SPINUP) to avoid exceeding enclosure power supply capabilities during spin-up of multiple SAS target devices. It may choose to rotate transmitting NOTIFY (ENABLE SPINUP) across all of its ports, distributing it to N ports at a time if the enclosure power supply is capable of powering N SAS target devices spinning up at a time. An expander device may allow this timing to be configured by a NVROM programming with enclosure-specific sequencing patterns, or may employ more complex, dynamic interaction with the enclosure power supply.

NOTE 2 - NOTIFY (ENABLE SPINUP) should be transmitted as frequently as possible to avoid incurring application layer timeouts.

I_T nexus loss, logical unit reset, and hard reset shall not cause a SAS target device to spin-up automatically on receipt of NOTIFY (ENABLE SPINUP).

SAS target devices with multiple SAS target ports shall honor NOTIFY (ENABLE SPINUP) from all SAS target ports equivalently (e.g., NOTIFY (ENABLE SPINUP) received on SAS target port A serves as a wakeup for a START STOP UNIT command received through SAS target port B).

[NOTIFY \(POWER FAILURE EXPECTED\) is transmitted by a SAS initiator port or expander port and is used to specify to a SAS target device that a power failure may occur within a time defined in the POWER FAILURE TIMEOUT field of the Protocol-Specific Logical Unit mode page \(see 0.0.2.1.1\). The NOTIFY \(POWER FAILURE EXPECTED\) shall be transmitted at least three times by the SAS initiator port or expander port.](#)

[If a target device supports NOTIFY \(POWER FAILURE EXPECTED\) and a NOTIFY \(POWER FAILURE EXPECTED\) is received, then the device server for each logical unit accessed by the SAS target port that received the NOTIFY \(POWER FAILURE EXPECTED\) shall:](#)

- a) [stop processing any write commands on a block boundary \(i.e., all write activity shall continue until a block boundary is reached then all writing shall stop\); and](#)
- b) [clear the task set \(i.e., perform a CLEAR TASK SET task management function\).](#)

If a target device supports NOTIFY (POWER FAILURE EXPECTED) and a NOTIFY (POWER FAILURE EXPECTED) is received, then each SAS phy within the target device shall:

- a) if there is a connection, then issue a BREAK on that connection; and
- b) issue an OPEN_REJECT (RETRY) to all OPEN requests after the NOTIFY (POWER FAILURE EXPECTED) is received until the power fail timeout timer expires (see 0.0.2.1.1).

If the power fail timeout timer expires, then each device server shall establish a unit attention condition for the initiator port associated with every I_T nexus with the additional sense code set to POWER FAIL EVENT COMMANDS CLEARED.

Editor's Note 1: Currently SAS does not provide for the NOTIFY in any state machines. This should be corrected. If directed the writer of this proposal will make those changes as part of this proposal.

~~NOTIFY (RESERVED 0)~~, NOTIFY (RESERVED 1), and NOTIFY (RESERVED 2) shall be ignored by all devices.

0.0.2.1.1 Protocol-Specific Logical Unit mode page - short format

The mode page policy (see SPC-3) for the Protocol-Specific Logical Unit mode page short format subpage shall be either shared or per target port. If a SAS target device has multiple SSP target ports, the mode page policy should be per target port. Parameters in this page shall affect all phys in the SSP target port if the mode page policy is per target port, and shall affect all SSP target ports in the SAS target device if the mode page policy is shared.

Table 146 defines the format of the page for SAS SSP.

Table 4 — Protocol-Specific Logical Unit mode page for SAS SSP - short format

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (18h)					
1	PAGE LENGTH (06h)							
2	Reserved			TRANSPORT LAYER RETRIES	PROTOCOL IDENTIFIER (6h)			
3	POWER FAILURE TIMEOUT							
4	Reserved							
7	Reserved							

The PARAMETERS SAVEABLE (PS) bit is defined in SPC-3.

The SPF field shall be set to zero for access to the short format mode page.

The PAGE CODE field shall be set to 18h.

The PAGE LENGTH field shall be set to 06h.

The PROTOCOL IDENTIFIER field shall be set to 6h indicating this is a SAS SSP specific mode page.

A TRANSPORT LAYER RETRIES bit set to one specifies that the target port shall support transport layer retries for XFER_RDY and DATA frames for the logical unit as described in 9.2.4. A TRANSPORT LAYER RETRIES bit set to zero specifies that transport layer retries shall not be used.

A POWER FAILURE TIMEOUT field specifies the maximum time, in one millisecond increments, that a logical unit shall respond to all commands received on all I T nexuses with a BUSY status after receiving a NOTIFY (POWER FAILURE EXPECTED). A 0000h value in this field is undefined by this standard.