

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 13 March 2008
 Subject: 08-008r2 SPC-4 Download microcode I_T nexus usage

Revision history

Revision 0 (4 December 2007) First revision, offshoot of 07-459r1.

Revision 1 (10 March 2008) Incorporated comments from January CAP WG - add bit to Extended INQUIRY VPD page to report behavior, clarify that the save must be complete before returning GOOD status on the command delivering the last microcode (mode 0Eh not mode 0Fh for deferred activation).

Revision 2 (13 March 2008) Incorporated comments from March CAP WG.

Related documents

spc4r11 - SCSI Primary Commands - 4 (SPC-4) revision 11

07-459 - SAM-4 SPC-4 SBC-3 Unit attention condition queuing (Rob Elliott, HP)

06-282r4 - SAM-4 SPC-4 WRITE BUFFER clarifications (Rob Elliott, HP)

Overview

1. SPC-4 is (purposely) vague on whether a sequence of WRITE BUFFER commands used to download microcode with offsets (modes 06h, 07h, 0Eh, and 0Fh) may be sent on more than one I_T nexus.

A new field in the Extended INQUIRY VPD page is proposed to report whether or not the device server supports multiple I_T nexuses. If it does, rules are defined to block commands with non-zero buffer offsets yet allow for restart of the sequence if the I_T nexus doesn't complete the sequence.

2. Modes 04h and 05h should not be required to deliver microcode in one command, as is stated in the model (this was only recently changed). The BUFFER ID field, BUFFER OFFSET field, and PARAMETER LIST LENGTH field are vendor-specific, and it was plausible under old rules to deliver microcode in multiple commands using them. The rule in the model was based on these original WRITE BUFFER mode descriptions:

- a) mode 04h: When the microcode download has completed successfully the device server shall establish a unit attention condition ... MICROCODE HAS BEEN CHANGED.
- b) mode 05h: When the download microcode and save command has completed successfully the device server shall establish a unit attention condition ... MICROCODE HAS BEEN CHANGED.
- c) mode 06h: After the last command completes successfully the device server shall establish a unit attention condition ... MICROCODE HAS BEEN CHANGED.
- d) mode 07h: After the last command completes successfully the device server shall establish a unit attention condition ... MICROCODE HAS BEEN CHANGED.

The mode 05h wording was assumed to be correct by 06-282r4 and also applied to mode 04h. Instead, the less restrictive mode 04h wording should have been applied to mode 05h.

Suggested changes to SPC-4 (with changes from 07-459 shown in tan)

5.15 Downloading and activating microcode

SCSI target device implementations may use microcode (e.g., firmware) that is stored in nonvolatile storage. Microcode may be changeable by an application client using the WRITE BUFFER command (see 6.38). The WRITE BUFFER command provides multiple methods for downloading microcode to the SCSI target device and activating the microcode.

Downloading and activating microcode involves the following steps:

- 1) Download: The application client transfers microcode from the Data-Out buffer to the device server in one or more WRITE BUFFER commands;
- 2) Save: After receiving the complete microcode, if defined by the download microcode mode, the device server saves the microcode to nonvolatile storage; and
- 3) Activate: After receiving the complete microcode and after saving it to nonvolatile storage if defined by the download microcode mode, the SCSI target device begins using the new microcode for the first time after an event defined by the download microcode mode.

After power on or hard reset, the SCSI target device shall use the last microcode that was saved to nonvolatile storage.

Table 53 defines the WRITE BUFFER download microcode modes with respect to the steps described in this subclause.

Table 53 — WRITE BUFFER download microcode modes

Mode	Download ^a	Save ^b	Activate ^c
Download microcode and activate (i.e., 04h)	yes ^d	no	yes
Download microcode, save, and activate (i.e., 05h)	yes ^d	yes	optional
Download microcode with offsets and activate (i.e., 06h)	yes ^e	no	yes
Download microcode with offsets, save, and activate (i.e., 07h)	yes ^e	yes	optional
Download microcode with offsets, save, and defer activate (i.e., 0Eh) ^f	yes ^e	yes	no
Activate deferred microcode (i.e., 0Fh)	no	no	yes

^a Entries in the Download column are as follows. For modes labeled yes, the application client delivers microcode in the WRITE BUFFER command(s). For modes labeled no, the application client does not deliver microcode with the WRITE BUFFER command.

^b Entries in the Save column are as follows. For modes labeled yes, the device server shall save the microcode to nonvolatile storage for use after each subsequent power on or hard reset and shall not return GOOD status for the final command in the WRITE BUFFER sequence until the microcode has been saved. For modes labeled no, the device server shall discard the microcode on the next power on or hard reset.

^c Entries in the Activate column are as follows. For modes labeled yes, the device server shall activate the microcode after completion of the final command in the WRITE BUFFER command sequence (i.e., the set of WRITE BUFFER commands that downloads the microcode and activates it). For modes labeled optional, the device server may or may not activate the microcode image upon completion of the final command in the WRITE BUFFER command sequence. For modes labeled no, the device server shall not activate the microcode upon completion of the final WRITE BUFFER command.

^d The application client ~~shall deliver the entire~~ delivers microcode in ~~one~~ a vendor-specific number of WRITE BUFFER commands (i.e., a WRITE BUFFER sequence). The device server should require that the microcode be delivered in a single command. The device server shall perform any required verification of the microcode prior to returning GOOD status for the final WRITE BUFFER command in a sequence (i.e., the WRITE BUFFER command delivering the last part of the microcode).

^e The application client delivers microcode in ~~multiple~~ one or more WRITE BUFFER commands, specifying a buffer ID and buffer offset in each command. If the device server does not receive the necessary WRITE BUFFER commands required to deliver the complete microcode before a logical unit reset occurs, an I_T nexus loss occurs, or a WRITE BUFFER command specifying a different download microcode mode is processed, the device server shall discard the new microcode. If the device server determines that it is processing the final WRITE BUFFER command (i.e., the WRITE BUFFER command delivering the last part of the microcode), it shall perform any required verification of the microcode prior to returning GOOD status for the command.

^f Microcode downloaded with this mode is called deferred microcode.

When microcode is activated due to processing a WRITE BUFFER command with a mode that causes activation after processing (see table 53), the device server shall establish a unit attention condition (see SAM-4) for the initiator port associated with every I_T nexus except the I_T nexus on which the WRITE BUFFER command was received with the additional sense code set to MICROCODE HAS BEEN CHANGED.

NOTE 1 - The application client should respond to the GOOD status for the WRITE BUFFER command the same way that it responds to a unit attention condition with an additional sense code set to MICROCODE HAS BEEN CHANGED (e.g., assume a hard reset has occurred).

When microcode is activated due to processing a WRITE BUFFER command with a mode that optionally causes activation after processing (see table 53), the device server shall establish a unit attention condition (see SAM-4) for the initiator port associated with every I_T nexus with the additional sense code set to MICROCODE HAS BEEN CHANGED.

When microcode is activated due to power on or hard reset, the device server may establish a unit attention condition (see SAM-4) for the initiator port associated with every I_T nexus with the additional sense code set to MICROCODE HAS BEEN CHANGED in addition to the unit attention condition for the power on or hard reset.

When deferred microcode (see table 53) is activated due to a command defined by its command standard as causing deferred microcode to be activated (e.g., the FORMAT UNIT command and the START STOP UNIT command in SBC-3), the device server:

- a) shall establish a unit attention condition (see SAM-4) for the initiator port associated with every I_T nexus with the additional sense code set to MICROCODE HAS BEEN CHANGED; and
- b) may establish other unit attention condition(s) as defined for the command (e.g., CAPACITY DATA HAS CHANGED for the FORMAT UNIT command).

If new microcode is saved before deferred microcode is activated, the deferred microcode is not activated and the new saved microcode is considered deferred.

...

If the device server is unable to process a WRITE BUFFER command with a download microcode mode because of a vendor specific condition ([e.g., the device server requires the microcode be delivered in order and the BUFFER OFFSET field is not equal to the previous WRITE BUFFER command's BUFFER OFFSET field plus PARAMETER LIST LENGTH field](#)), it shall terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to COMMAND SEQUENCE ERROR.

Table 54 defines the DM_MULT field in the Extended INQUIRY VPD page (see 7.7.4), which indicates how the device server handles multiple I_T nexuses using WRITE BUFFER download microcode modes (see table 53).

Table 54 — Multiple I_T nexus handling for WRITE BUFFER download microcode modes

DM_MULT field ^a	Description
0h	Handling of multiple I_T nexuses is vendor-specific
1h	<p>For modes that download microcode (see table 53), the device server shall:</p> <p>a) for a WRITE BUFFER command with the BUFFER OFFSET field set to 000000h on any I_T nexus, process the command. This shall cause any microcode downloaded on another I_T nexus to be discarded and establishes the I_T nexus of the WRITE BUFFER sequence; and</p> <p>b) for a WRITE BUFFER command with the BUFFER OFFSET field set to a non-zero value on a different I_T nexus than that of the current WRITE BUFFER sequence, terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to COMMAND SEQUENCE ERROR.</p> <p>For mode 0Fh (i.e., activate deferred microcode), the device server shall terminate any WRITE BUFFER command processed on a different I_T nexus than that used for the WRITE BUFFER commands with mode 0Eh (i.e., download with offsets, save, and defer activate mode) with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to COMMAND SEQUENCE ERROR.</p>
2h	<p>For modes that download microcode (see table 53), the device server shall allow the sequence of WRITE BUFFER commands to be processed on more than one I_T nexus.</p> <p>For mode 0Fh (i.e., activate deferred microcode), the device server shall allow a WRITE BUFFER command to be processed on a different I_T nexus than that used for the WRITE BUFFER commands with mode 0Eh (i.e., download with offsets, save, and defer activate mode).</p>
3h	<p>For modes that download microcode (see table 53), the device server shall:</p> <p>a) for a WRITE BUFFER command with the BUFFER OFFSET field set to 000000h on any I_T nexus, process the command. This shall cause any microcode downloaded on another I_T nexus to be discarded and establishes the I_T nexus of the WRITE BUFFER sequence; and</p> <p>b) for a WRITE BUFFER command with the BUFFER OFFSET field set to a non-zero value on a different I_T nexus than that of the current WRITE BUFFER sequence, terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to COMMAND SEQUENCE ERROR.</p> <p>For mode 0Fh (i.e., activate deferred microcode), the device server shall allow a WRITE BUFFER command to be processed on a different I_T nexus than that used for the WRITE BUFFER commands with mode 0Eh (i.e., download with offsets, save, and defer activate mode).</p>
All others	Reserved
^a This field is in the Extended INQUIRY Data VPD page (see 7.7.4).	

For all WRITE BUFFER command modes that download microcode (see table 53), the COMMAND SPECIFIC field (see 6.24.4.2) located in the command timeouts descriptor of the parameter data returned by the REPORT SUPPORTED OPERATION CODES command (see 6.24) indicates the maximum time that access to the SCSI device is limited or not possible through any SCSI ports associated with a logical unit that processes a WRITE BUFFER command that activates microcode.

7.7.4 Extended INQUIRY Data VPD page

The Extended INQUIRY Data VPD page (see table 370) provides the application client with a means to obtain information about the logical unit.

Table 370 — Extended INQUIRY Data VPD page

Byte\Bit	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIFER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (86h)							
2	Reserved							
3	PAGE LENGTH (3Ch)							
4	Reserved		SPT			GRD_CHK	APP_CHK	REF_CHK
5	Reserved			GROUP_SUP	PRIOR_SUP	HEADSUP	ORDSUP	SIMPSUP
6	Reserved				WU_SUP	CD_SUP	NV_SUP	V_SUP
7	Reserved							LUICLR
8	Reserved				Reserved DM_MULT			
9	Reserved							
63	Reserved							

...

[The DM_MULT field is defined in table 54 \(see 5.15\).](#)