

To: T10 Technical Committee
From: Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
Date: 18 July 2002
Subject: T10/02-149r1 SPC-3 Port-specific mode page clarifications

Revision History

Revision 0 (1 June 2002) first revision

Revision 1 (18 July 2002) incorporated feedback from the July CAP WG – allow target ports to share values in all cases, since that is how current dual port FC disk drives behave. Targets are advised they “should” maintain separate values per target port.

Related Documents

spc3r07 - SCSI Primary Commands - 3 revision 7

Overview

The Disconnect-Reconnect mode page, Protocol Specific LUN mode page, and Protocol Specific Port mode page are not clearly defined for target devices with multiple target ports (possibly using different protocols) and multiple logical units.

The Disconnect-Reconnect mode page and Protocol Specific Port mode page contain values that apply to the target port irregardless of the logical unit being accessed. SPC-3 (and the protocol standards) are not very clear about these scenarios:

1. If a target device has multiple logical units, and a field is changed via one logical unit, a unit attention shall be created on the other logical units indicating MODE PARAMETERS CHANGED.
2. If a target device has multiple target ports, it must implement a separate copy of each page for each target port.
3. If a target device has target ports using different protocols, the definition of the pages of each target port may be different (since different protocols have different requirements).
4. The settings for a port of one protocol may only be changed by accesses through that port.

The Protocol Specific LUN page description doesn't mention that if a target device has multiple target ports, the Protocol Specific LUN page also needs to be target-port specific. This is most evident if the target ports use different protocols.

Also, all three pages need to be shared by all initiator ports - no per-initiator values are supported. (SPC-3 defines three kinds of mode pages: shared, per-initiator port, and per-I_T.)

Suggested Changes

8.4.7 Disconnect-Reconnect mode page

The Disconnect-Reconnect mode page (see table 227) provides the application client the means to tune the performance of the service delivery subsystem. The name for this mode page, disconnect-reconnect, comes from the SCSI parallel bus interface. ~~A SCSI device based on any of the SCSI protocols may use appropriate parameters in the Disconnect-Reconnect mode page. The parameters appropriate to each SCSI protocol and their interpretation for that SCSI protocol may be specified in the individual SCSI protocol standards (see 3.1.57).~~

~~The device server communicates the parameter values in this mode page to the service delivery subsystem. Similarly the application client may also communicate parameter values to the service delivery subsystem. This communication is internal to the initiator or target device and is outside the scope of SCSI.~~

The Disconnect-Reconnect mode page controls parameters that affect one or more SCSI target ports. The parameters that may be implemented are specified in the SCSI protocol standard for

the SCSI target port. If a target device has multiple SCSI target ports, changes in the parameters for one SCSI target port should not affect other SCSI target ports.

The parameters for a SCSI target port affect its behavior regardless of which SCSI initiator port is forming an I_T nexus with the SCSI target port. The parameters may be accessed by MODE SENSE and MODE SELECT commands directed to any logical unit accessible through the SCSI target port. If a parameter value is changed, all logical units accessible through the SCSI target port shall establish a unit attention condition with an additional sense code of MODE PARAMETERS CHANGED for all SCSI initiator ports that form I_T_L nexuses through that SCSI target port other than the SCSI initiator port which made the change.

If a parameter that is not appropriate for the specific SCSI protocol implemented by the ~~SCSI device~~SCSI target port is non-zero, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

An interconnect tenancy is a period of time during which a ~~SCSI device~~SCSI port owns or may access the interconnect. For example, on arbitrated interconnects, a tenancy typically begins when a ~~SCSI device~~SCSI port successfully arbitrates for the interconnect and ends when the ~~SCSI device~~SCSI port releases the interconnect for use by other devices. Data and other information transfers take place during interconnect tenancies.

The BUFFER FULL RATIO field indicates to the ~~device server~~SCSI target port, during read operations, how full the buffer should be prior to requesting an interconnect tenancy. ~~Device server~~SCSI target ports that do not implement the requested ratio should round down to the nearest implemented ratio as defined in 5.3.

The BUFFER EMPTY RATIO field indicates to the ~~device server~~SCSI target port, during write operations, how empty the buffer should be prior to requesting an interconnect tenancy. ~~Device server~~SCSI target ports that do not implement the requested ratio should round down to the nearest implemented ratio as defined in 5.3.

The buffer full and buffer empty ratios are numerators of a fractional multiplier that has 256 as its denominator. A value of zero indicates that the ~~target~~SCSI target port determines when to request an interconnect tenancy consistent with the disconnect time limit parameter. These parameters are advisory to the ~~target~~SCSI target port.

NOTE 57 - As an example, consider a ~~device server~~SCSI target port with ten 512-byte buffers and a specified buffer full ratio of 3Fh. The formula is: INTEGER((ratio/256)*number of buffers). Therefore in this example $\text{INTEGER}((3\text{Fh}/256)*10) = 2$. During the read operations described in this example, the ~~device server~~SCSI target port should request an interconnect tenancy whenever two or more buffers are full.

The BUS INACTIVITY LIMIT field indicates the maximum time that the ~~target~~SCSI target port is permitted to maintain an interconnect tenancy without data or information transfer. If the bus inactivity limit is exceeded the ~~device server~~SCSI target port shall conclude the interconnect tenancy, within the restrictions placed on it by the applicable SCSI protocol. The contents of the DTDC field in this mode page also shall affect the duration of an interconnect tenancy. This value may be rounded as defined in 5.3. A value of zero indicates that there is no bus inactivity limit. Different SCSI protocols specify different units of measure for the bus inactivity limit.

The DISCONNECT TIME LIMIT field indicates the minimum time that the ~~target~~SCSI target port shall wait between interconnect tenancies. This value may be rounded as defined in 5.3. A value of zero indicates that there is no disconnect time limit. Different SCSI protocols specify different units of measure for the disconnect time limit.

The CONNECT TIME LIMIT field indicates the maximum duration of a single interconnect tenancy. If the connect time limit is exceeded the ~~device server~~SCSI target port shall conclude the interconnect tenancy, within the restrictions placed on it by the applicable SCSI protocol. The

contents of the DTDC field in this mode page also shall affect the duration of an interconnect tenancy. This value may be rounded as defined in 5.3. A value of zero indicates that there is no connect time limit. Different SCSI protocols specify different units of measure for the connect time limit.

The MAXIMUM BURST SIZE field indicates the maximum amount of data that the ~~device server SCSI target port~~ shall transfer during a single data transfer operation. This value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes, etc.). The relationship (if any) between data transfer operations and interconnect tenancies is specified in the individual SCSI protocol standards. A value of zero indicates there is no limit on the amount of data transferred per data transfer operation.

The enable modify data pointers (EMDP) bit indicates whether or not the ~~initiator allows the data transfer to be re-ordered by the target SCSI target port may transfer data out of order~~. If the EMDP bit is zero, the ~~target SCSI target port shall not re-order the data transfer transfer data out of order~~. If the EMDP bit is one, the ~~target SCSI target port~~ is allowed to ~~re-order the data transfer transfer data out of order~~.

The FAIR ARBITRATION field indicates whether the ~~target SCSI target port~~ should use fair or unfair arbitration when requesting an interconnect tenancy. The field may be used to indicate different fairness methods as specified in the individual SCSI protocol standards.

A disconnect immediate (DIMM) bit of zero indicates that the ~~target SCSI target port~~ may transfer data for a command during the same interconnect tenancy in which it receives the command. Whether or not the ~~target SCSI target port~~ does so may depend upon the ~~target SCSI target port's~~ internal algorithms, the rules of the applicable SCSI protocol, and settings of the other parameters in this mode page. A disconnect immediate (DIMM) bit of one indicates that the ~~target SCSI target port~~ shall not transfer data for a command during the same interconnect tenancy in which it receives the command.

The data transfer disconnect control (DTDC) field (see table 228) defines other restrictions on when multiple interconnect tenancies are permitted. A non-zero value in the DTDC field shall take precedence over other interconnect tenancy controls represented by other fields in this mode page.

Table 228 — Data transfer disconnect control

...

The FIRST BURST SIZE field indicates the maximum amount of data that may be transferred to the ~~target SCSI target port~~ for a command along with the command. This value is expressed in increments of 512 bytes; a value of one means 512 bytes, two means 1024 bytes, etc. A value of zero indicates that there is no first burst size limit.

[Editor's note: SPI-4 and FCP-2 called this the Logical Unit Control page.]

8.4.12 Protocol Specific LUN Logical Unit mode page

The Protocol Specific ~~LUN Logical Unit~~ mode page (see table 236) provides protocol specific controls that are associated with a logical unit. ~~See the SCSI protocol standard (see 3.1.61) for definition of the protocol specific mode parameters.~~

Table 236 — Protocol Specific LUN mode page

...

~~The PROTOCOL IDENTIFIER field (see 8.5.1) indicates the SCSI protocol to which the mode page applies. For MODE SENSE commands, the device server shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 (see 8.5.1) to indicate the SCSI protocol used by its service delivery subsystem. For MODE SELECT commands, the application client shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 indicating the SCSI protocol to which the protocol specific mode parameters apply. If a device server receives a mode page containing a protocol identifier value other than the one used by its service delivery subsystem, it~~

~~shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.~~

The Protocol Specific Logical Unit mode page controls parameters that affect both:

- a) one or more SCSI target ports; and
 - b) the logical unit
- during an I_T_L nexus.

The parameters that may be implemented are specified in the SCSI protocol standard for the SCSI target port. If a logical unit is accessible through multiple SCSI target ports, changes in the parameters for one SCSI target port should not affect other SCSI target ports.

The parameters for a SCSI target port and logical unit affect their behavior regardless of which SCSI initiator port is forming an I_T_L nexus with the SCSI target port and logical unit. If a parameter value is changed, the logical unit shall establish a unit attention condition with an additional sense code of MODE PARAMETERS CHANGED for all SCSI initiator ports that form I_T_L nexuses through that SCSI target port other than the SCSI initiator port which made the change.

The PROTOCOL IDENTIFIER field (see 8.5.1) indicates the SCSI protocol to which the mode page applies. For a MODE SENSE command, the device server shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 (see 8.5.1) to indicate the SCSI protocol used by the SCSI target port through which the MODE SENSE command is being processed. For a MODE SELECT command, the application client shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 indicating the SCSI protocol to which the protocol specific mode parameters apply. If a device server receives a mode page containing a PROTOCOL IDENTIFIER value indicating a protocol other than the one used by the SCSI target port through which the MODE SELECT command is being processed, it shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.

[Editor's note: SPI-4 and FCP-2 called this the Protocol-Specific Port Control page.]

8.4.13 Protocol Specific Port mode page

The Protocol Specific Port mode page provides protocol specific controls that are associated with a ~~port~~SCSI port. The page_0 format (see table 237) is used for subpage 00h and sub_page format (see table 238) is used for subpages 01h through FEh. See the SCSI protocol standard (see 3.1.61) for definition of the protocol specific mode parameters.

Table 237 — Page_0 format Protocol Specific Port mode page

...

Table 238 — Sub_page format Protocol Specific Port mode page

...

~~The PROTOCOL IDENTIFIER field (see 8.5.1) indicates the SCSI protocol to which the mode page applies. For MODE SENSE commands, the device server shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 (see 8.5.1) to indicate the SCSI protocol used by its service delivery subsystem. For MODE SELECT commands, the application client shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 indicating the SCSI protocol to which the protocol specific mode parameters apply. If a device server receives a mode page containing a protocol identifier value other than the one used by its service delivery subsystem, it shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.~~

The Protocol Specific Port mode page controls parameters that affect one or more SCSI target ports. The parameters that may be implemented are specified in the SCSI protocol standard for

the SCSI target port. If a target device has multiple SCSI target ports, changes in the parameters for one SCSI target port should not affect other SCSI target ports.

The parameters for a SCSI target port affect its behavior regardless of which SCSI initiator port is forming an I_T nexus with the SCSI target port. The parameters may be accessed by MODE SENSE and MODE SELECT commands directed to any logical unit accessible through the SCSI target port. If a parameter value is changed, all logical units accessible through the SCSI target port shall establish a unit attention condition with an additional sense code of MODE PARAMETERS CHANGED for all SCSI initiator ports that form I_T_L nexuses through that SCSI target port other than the SCSI initiator port which made the change.

The PROTOCOL IDENTIFIER field (see 8.5.1) indicates the SCSI protocol to which the mode page applies. For a MODE SENSE command, the device server shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 (see 8.5.1) to indicate the SCSI protocol used by the SCSI target port through which the MODE SENSE command is being processed. For a MODE SELECT command, the application client shall set the PROTOCOL IDENTIFIER field to one of the values shown in table 239 indicating the SCSI protocol to which the protocol specific mode parameters apply. If a device server receives a mode page containing a PROTOCOL IDENTIFIER value indicating a protocol other than the one used by the SCSI target port through which the MODE SELECT command is being processed, it shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.