

TO: X3T9.2 Committee (SCSI)

FROM: Gerry Houlder

SUBJECT: SCSI-3 Changes for Dual Port Feature

My first dual port features document (X3T9.2/90-136R3) was accepted for inclusion in SCSI-3, but it has come to my attention that my second document (X3T9.2/91-149R0) that adds a few more dual port features was not accepted. The X3T9.2 editors asked me to create a new document that includes only the new features of 91-149R0 before a vote can be taken to include the latest features. This makes it easier to identify the new items.

The new issues are:

- a) Redefine effect of SCSI BUS RESET message. A SCSI BUS RESET message on either port will reset the entire device (clearing all I/O processes & reservations, restoring saved MODE SELECT values, setting Unit Attention for all initiators on both ports).
- b) Define effect on unit attention condition. Unit Attention due to SCSI RESET is not generated for the other port. Other Unit attention conditions lump the other port with the "other initiators".
- c) Add new command for dual port (PORT STATUS) that allows enabling or disabling the other port, overriding reservation held by the other port, and returning status (port enabled/disabled status and reservation status) about both ports.
- d) Revise RESERVE command description to take into account the PORT STATUS command and other dual port specific effects.
- e) *As a result of the March working group, alternative modes of operation have been added which are selected by bits in the Control Mode page.*

My editorial comments are in brackets []. Proposed text additions are underlined and proposed text deletions are ~~over-struck~~. Changes from Rev 0 are indicated by italics.

[SIP Document changes:]

6.3 BUS DEVICE RESET

The BUS DEVICE RESET message is sent from an initiator to direct a target to clear all I/O processes on that SCSI device. This message has the same effect as a hard reset condition to the selected SCSI device. The target shall go to the BUS FREE phase following successful receipt of this message. The target shall create unit attention condition for all initiators.

For dual port implementations, if the Byprtm bit equals zero, the hard reset condition shall apply to both ports. All I/O processes and device reservations shall be cleared, all operating modes restored to initial conditions, and the unit attention condition set for all initiators on both ports. If the Byprtm bit equals one, the hard reset condition shall apply only to the port from which the message was received. Any active I/O processes, queued I/O processes, device reservations, and operating modes for the other port are unaffected and unit attention condition is not set for initiators on the other port.

[No other changes is SIP section 6.3.]

7.4.4 Unit Attention Condition

The target shall generate a unit attention condition on each valid logical unit whenever the target has been reset by a BUS DEVICE RESET message, a hard reset condition, or by a power on reset. For dual port implementations, the unit attention condition for:

- (1) the hard reset condition,
- (2) the BUS DEVICE RESET message, and
- (3) the BUS DEVICE RESET OTHER PORT message

~~only~~ affects the initiators on one or both ports as described in sections 5.1.2.1, 6.4.3, and 6.4. The target shall also generate a unit attention condition on the affected logical unit(s) for each *affected* initiator whenever one of the following events occurs:

[No other changes in SIP section 7.4.4]

[This command description should go into SBC document.]
 [All of this text is new even though underline is not used.]

_____ PORT STATUS Command

Table _____: PORT STATUS Command

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (11h)							
1	Logical Unit Number			Reserved		DisOP	EnOP	PR
2	Reserved							
3	Reserved							
4	Allocation Length							
5	Control							

The PORT STATUS command (Table _____) provides a way for the initiator to *perform* functions concerning the other port on dual port targets and to return status concerning the ports. This command is mandatory if the dual port option is implemented.

The Disable Other Port (DisOP) and Enable Other Port (EnOP) bits are used to control ~~enabling of~~ the other port as follows:

DisOP	EnOP	
----	----	
0	0	No change in condition of other port
0	1	Enable the other port
1	0	Disable the other port
1	1	<i>Reserved</i>

When the other port is disabled, the target *shall* not respond to selection from any initiator on *the disabled* port. When it is enabled the target responds normally to selection attempts.

A Priority Reservation (PR) bit of one requires the target to supersede any existing reservation for any initiator on either port with a logical unit reservation for *the initiator issuing the Port Status Command*. This action precedes any other actions defined for *the Port Status Command*. A PR bit of zero indicates that the DisOP and EnOP functions are prevented if a conflicting reservation exists, but the port status data *shall* be returned and this command *shall return*

GOOD status.

The Allocation Length specifies the number of bytes that has been allocated for returned PORT STATUS data. An allocation length of zero indicates that no PORT STATUS data is transferred. This condition is not considered as an error. Any other value indicates the maximum number of bytes that shall be transferred. The target shall terminate the DATA IN phase when Allocation Length bytes have been transferred or when all available data has been transferred, whichever is less.

Table ____: PORT STATUS Data Format

Bit	7	6	5	4	3	2	1	0
0	Data Length							
1	PAR		Reserved				OPE	
2	RID							
3	Reserved							

The PORT STATUS data shall reflect the status of the target after the execution of any DisOP, EnOP, and PR options requested by this command. The Data Length indicates the number of additional PORT STATUS data bytes available for transfer during the DATA IN phase. The Data Length does not include itself and is therefore defined to be 3.

An Other Port Enabled (OPE) bit of zero indicates that the other port is not enabled. An OPE bit of one indicates that the other port is currently enabled.

The Port with Active Reservation (PAR) field identifies which port the initiator with an active reservation is reserved to.

Bit 7	Bit 6	
0	0	No reservation is active
0	1	Reservation is owned by initiator on this port
1	0	Reservation is owned by initiator on other port
1	1	Reserved

The Reservation ID (RID) field indicates the SCSI ID of the initiator that owns the active reservation. When the PAR field indicates that no reservation is active, this field shall be set to zero.

[Changes needed for SCC document. Since actual SCC text not available yet, SCSI-2 wording is used as a reference.]

_____ RESERVE Command

[CDB Table is unchanged]

The RESERVE and RELEASE commands provide the basic mechanism for contention resolution in multiple-initiator systems. The RESERVE command (Table _____) is used to reserve a logical unit or, if the extent reservation option is implemented, extents within a logical unit. The third party reservation allows logical units or extents to be reserved for another specified SCSI device.

IMPLEMENTORS NOTE: The reservation queuing option in X3.131-1986 has been removed from SCSI-2.

8.2.12.1 Logical Unit Reservation (Mandatory).

If the extent bit is zero, this command shall request that the entire logical unit be reserved for the exclusive use of the initiator until the reservation is:

- a) superseded by another valid RESERVE command from the initiator that made the reservation,
- b) released by a RELEASE command from the same initiator that made the reservation,
- c) released by a BUS DEVICE RESET message from any initiator (for dual port implementations, on either port if the *Byprtm* bit equals zero or on the same port if the *Byprtm* bit equals one),
- d) released by a BUS DEVICE RESET OTHER PORT message from any initiator on the other port from the initiator that made the reservation,
- e) released by a hard RESET condition from any initiator on the same port, if the *BybthH* bit equals zero, as the initiator that made the reservation, or on either port if the *BybthH* bit equals one,
- f) released by a power on cycle,
- g) superseded by a PORT STATUS command with priority reserve option.

~~superseded by another valid RESERVE command from the initiator that made the reservation or until released by a RELEASE command from the same initiator that made the reservation, by a BUS DEVICE RESET message from any initiator, by a hard RESET condition, or by a power on cycle.~~

A logical unit reservation shall not be granted if the logical unit or any extent is reserved by another initiator. It shall be permissible for an initiator to reserve a logical unit that is currently reserved by that initiator. If the extent bit is zero, the reservation identification and the extent list length shall be ignored.

If the logical unit, or any extent within the logical unit, is reserved for another initiator, the target shall return

RESERVATION CONFLICT status.

If, after honoring the reservation, any other initiator attempts to perform any command on the reserved logical unit other than an INQUIRY, REQUEST SENSE, PREVENT ALLOW MEDIUM REMOVAL (with a prevent bit of one), PORT STATUS, or RELEASE command then the command shall be rejected with RESERVATION CONFLICT status.

For dual port implementations, devices on the other port (i.e., the port that does not include the initiator to which a reservation has been granted) shall also be denied access to the logical unit as described in the preceding paragraph.

[No other changes needed in RESERVE command.]

7.3.3.1 Control Mode Page

Table 96: Control Mode Page

Bit	7	6	5	4	3	2	1	0
Byte								
0	PS	Rsrvd	Page Code (0Ah)					
1	Page Length (06h)							
2	Reserved						RLEC	
3	Queue Algorithm Modifier			Reserved		QErr	DQue	
4	EECA	RAC	<i>Byprtm</i>	<i>BybthH</i>	Rsrvd	RAENP	UAAENP	EAENP
5	Reserved							
6	(MSB)							
7	Ready AEN Holdoff Period						(LSB)	

The By Port Message (*Byprtm*) bit when set to one indicates that the Bus Device Reset message applies only to the port over which the message was received. When the *Byprtm* bit equals zero the Bus Device Reset Message applies to both ports.

The By Both hard RESET (*BybthH*) bit when set to one indicates that the hard Reset applies to both ports. When the *BybthH* bit equals zero the hard RESET applies only to the port over which it was received.