Multi-Port Changes for SAM-2 (T10/97-250R1)

Charles Monia

Digital Equipment Corporation

September 9, 1997

Based On:

- Multi-port proposals approved for SCSI-3.
  - X3T9.2/90-136R3, “Extensions for Dual-Port SCSI”.
  - X3T9.2/93-041R2, “SCSI-3 Changes for Dual-Port Feature”.
    - Amends X3T9.2/90-136R3.
- Other documents:
  - X3T9.2/91-149R0, ““Extensions for Dual-Port SCSI (again)”.

Review of Changes Approved for SCSI-3

- Change to parallel SCSI I/O Process Model
  - Nexus and I/O process now associated with a Port
    - P_I_T_L_Q implicitly replaces I_T_L_Q as nexus or I/O process identifier.
      - P = port I/D when process or nexus was created.
  - Nexus must always be restored on the port receiving the original command.

Dual-Port I/O Process Example

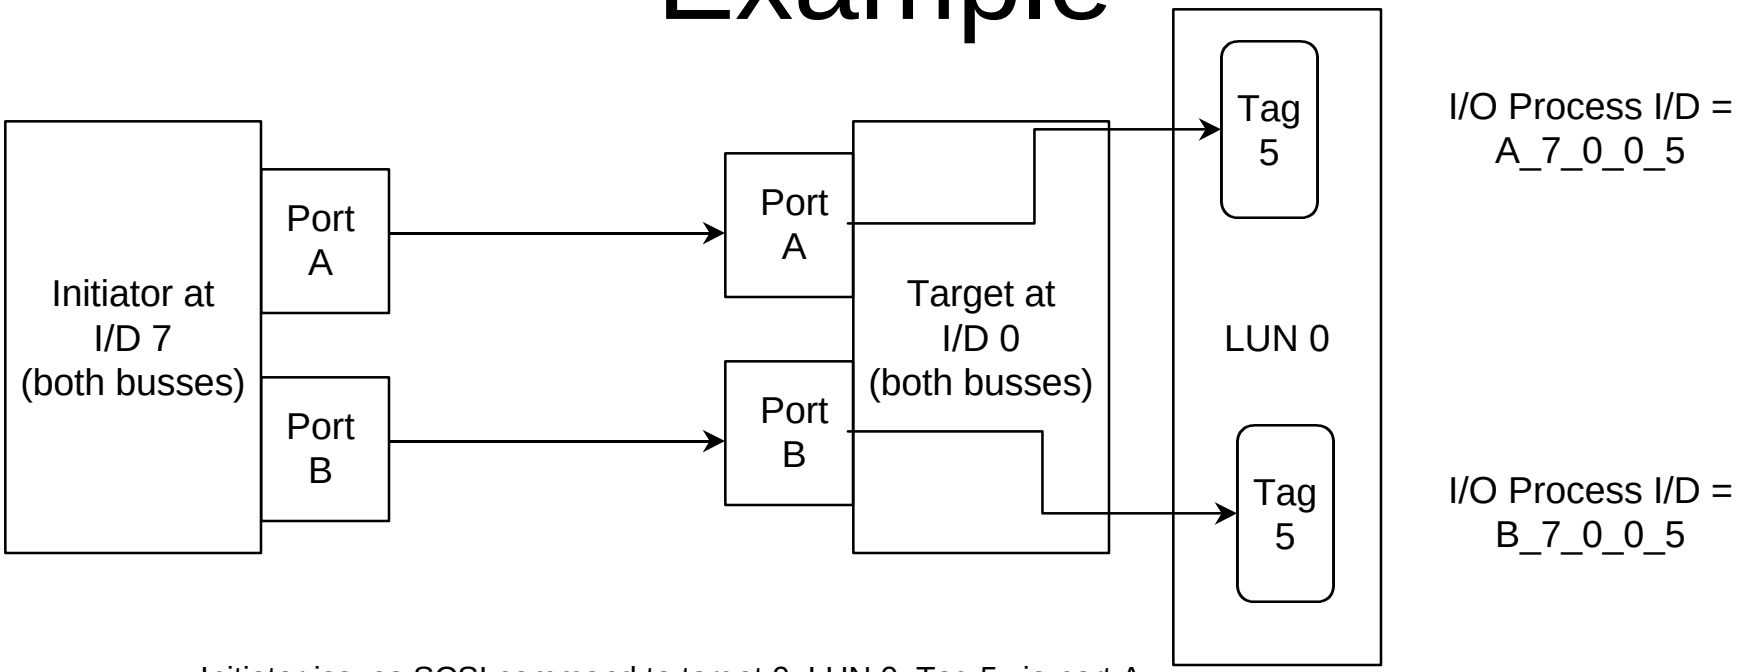

Initiator issues SCSI command to target 0, LUN 0, Tag 5 via port A.
Logical unit creates I/O process A_7_0_0_5.

Initiator issues SCSI command to target 0, LUN 0, Tag 5 via port B.
Logical unit creates I/O process B_7_0_0_5.

As seen by target, Initiator at A_7 is different than B_7.

Task Aborts

- Scope of ABORT TASK SET, ABORT TASK, CLEAR TASK SET limited to tasks whose port identifier is the same as the port receiving the task management request.
- Scope of hard reset and TARGET RESET optionally limited to tasks whose port identifier is the same as the port receiving the task management request.
- Controllable by initiator through Control Mode page.

Exceptions

- Contingent Allegiance -- Unchanged.
  - The faulted initiator is identified by Port I/D concatenated with bus I/D.
- Scope of hard resets can be controlled by the initiator via Control Mode page.
  - One or both ports.
  - Extent of Unit Attention due to hard reset corresponds to scope of the reset.

Device Reservations

- Ability to override a reservation.

New Command

- Port status and control command
  - Port control
    - Enable or disable other port
    - Supercede reservations on this or other port
  - Receive port status
    - Port Enabled or Disabled.
    - Reservations owned by an initiator on this/other port.

New Task Management Function

- TARGET RESET OTHER PORT
  - Directs a TARGET RESET to the “other port”.

Issues

- Proposed functions only support two ports.
  - Operations reference “other port”
- Ad hoc device reservation features
- Should multi-port behavior be hidden from host class driver?
  - Managed by transport layer?

Alternatives

- Dual-port limit
  - Do we want to extend to any number of ports?
  - Do we want port-specific hard reset and associated Unit Attention?
    - If eliminated, then multi-port becomes a special case of multi-initiator
    - Some issues could be addressed with Gene Milligan's multi-initiator proposal
  - Port control/statuscommand
    - Eliminate?
    - Extend to any number of ports?
    - Make protocol-specific and hide in protocol spec.?
- Reservation Management
  - Replace with Persistent Reserve?

Alternatives (cont)

- Concealing multi-port behavior from host class driver.
  - "Class driver" has no awareness of paths or ports.
  - All path and port management is done outside of the SCSI command set.
  - Port dependencies in task references are outside of SAM scope.
    - E.g., In parallel SCSI, the transport layer could automatically issue an ABORT TASK on the port associated with the task.
    - Path failures handled by transport layer.