



May 5, 1999

To: T10 Committee

From: Tom Coughlan
Compaq Computer Corporation
Mail Stop ZKO3-4/U14
110 Spitbrook Road
Nashua, New Hampshire
Telephone: 603-884-0933
E-mail: tom.coughlan@compaq.com

Subject: Still more on Persistent Reservation (99-182r1)

Change I. Clarification of the reservation behavior for Preempt and Preempt and Abort.

The current description of the Preempt service action (SPC-2 Rev. 9) fails to describe the reservation behavior for tasks that arrive after the Preempt is queued, but before the Preempt returns status. The proposed change is to treat these tasks in the same manner as tasks that were queued before the Preempt arrived, that is, in a vendor specific manner.

The current description of the Preempt and Abort service action is not clear about the reservation behavior for tasks that arrive after the device server has executed the Abort, but before GOOD status is returned for the Preempt and Abort. The proposed change is to require that these tasks be subject to the new reservation.

Change II. Clarification of Preempt and Abort behavior when there is no existing reservation.

The current revision of SPC-2 (Rev. 9) specifies the behavior of Preempt and Abort by providing a reference to the behavior of the Preempt service action. The description of Preempt describes two cases: a) preempting reservations, and b) preempting registrations.

The purpose of this change is to provide an explicit statement that the abort part of the Preempt and Abort shall occur in both of the two cases listed above. Although this fact can be inferred from a careful reading of the existing text, we believe that an explicit statement is needed. The reason is because it may be counter-intuitive

that simply preempting an initiator's registration also causes an abort of all the initiator's queued commands.

Detailed Changes.

The following changes are proposed for SPC-2.

5.3.2.5.2.2 Preempting reservations

Any registered initiator may preempt any persistent reservation with another persistent reservation by issuing a PERSISTENT RESERVE OUT command with a PREEMPT service action through a registered initiator with the following parameters:

- a) RESERVATION KEY set to the value of the initiator/logical unit pair's established reservation key;
- b) SERVICE ACTION RESERVATION KEY set to match the reservation key of the persistent reservation being preempted; and
- c) TYPE and SCOPE set to define a new persistent reservation. The scope and type of the persistent reservation created by preempting initiator may be different than the persistent reservation being preempted.

If the SERVICE ACTION RESERVATION KEY is associated with a reservation, the device server shall perform a preempt by doing the following as an uninterrupted series of actions:

- a) Remove the persistent reservation for the initiator identified by the SERVICE ACTION RESERVATION KEY specified in the PERSISTENT RESERVE OUT parameter list;
- b) Remove the registration for the initiator or initiators identified by the SERVICE ACTION RESERVATION KEY specified in the PERSISTENT RESERVE OUT parameter list (see 5.3.2.3);
- c) establish a persistent reservation for the preempting initiator;
- d) process tasks as defined in 5.3; and
- e) establish a unit attention condition for any initiator that lost its reservation and/or registration. The sense key shall be set to UNIT ATTENTION and the additional sense data shall be set to REGISTRATIONS PREEMPTED.

~~Tasks received after successful completion of the PERSISTENT RESERVE OUT command with the PREEMPT service action are subject to the persistent reservation restrictions established by the preempting initiator. Each task~~ After GOOD status has been returned for the PERSISTENT RESERVE OUT command, new tasks are subject to the persistent reservation restrictions established by the preempting initiator. Each task that:

- is received after the arrival, but before the completion of the PERSISTENT RESERVE OUT command with the PREEMPT service action, or
- is in the dormant, blocked, or enable state at the time the PERSISTENT RESERVE OUT command with the PREEMPT service action is received

shall be subject in a vendor specific manner to either the restrictions established by the persistent reservation being preempted or to the restrictions established by the preempting initiator. Completion status shall be returned for each task.

A PERSISTENT RESERVE OUT specifying a PREEMPT service action with the SERVICE ACTION RESERVATION KEY value equal to the reservation key is not an error. In that case the device server shall establish the new reservation.

5.3.2.5.3 Preempting an existing persistent reservation with the PREEMPT AND ABORT service action

The initiator's request for and the device server's responses to a PERSISTENT RESERVE OUT command PREEMPT AND ABORT service action are identical to the PREEMPT service action (see 5.3.2.5.2) except for the following additions. If no reservation conflict occurred, the device server shall do the following ~~as part of the~~ uninterrupted series of actions:

- a) Perform the uninterrupted series of actions described for the PREEMPT service action;
- a)b) Every task from all preempted initiators shall be terminated as if an ABORT TASK SET task management function had been performed by each of the preempted initiators. After the ABORT TASK SET function has completed, all ~~After GOOD status has been returned for the~~

- ~~PERSISTENT RESERVE OUT~~ command, new tasks are subject to the persistent reservation restrictions established by the preempting initiator;
- b)c) The device server shall clear any ACA or CA condition associated with an initiator being preempted and shall clear any tasks with an ACA attribute from that initiator. If TST=000b (see 8.3.4), then ACA or CA conditions for initiators, other than the initiator being preempted, shall prevent the execution of the PERSISTENT RESERVE OUT command that shall end with a status of ACA ACTIVE if NACA=1 (see SAM-2) or BUSY if NACA=0. If TST=001b, then ACA or CA conditions for initiators other than the initiator being preempted shall not prevent the execution of the PERSISTENT RESERVE OUT command; and
- e)d) For SCSI devices that implement the PREVENT ALLOW MEDIUM REMOVAL command, the device server shall perform an action equivalent to the execution of a PREVENT ALLOW MEDIUM REMOVAL command with the PREVENT field equal to zero for the initiator or initiators being preempted (see 7.13).

The actions described in the preceding list shall be performed for all initiators that are registered with the SERVICE ACTION RESERVATION KEY, without regard for whether the preempted initiator(s) hold the reservation.

Any asynchronous event reporting operations in progress are not affected by the PREEMPT AND ABORT service action.