

Date: July 18, 2003

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: SPC-3: Persistent All Registrants Fix Part 2

1 Overview Problem 1

As a result of the changes made in 03-230 some duplicate requirements relating to how unit attentions work when a RELEASE is received.

2 Solution

Remove the duplicate wording.

3 Proposal

Make the indicated changes:

3.0.0.0.0.1 Handling for released registrants only persistent reservations

When the persistent reservation holder (see 5.5.2.6) of a Write Exclusive – Registrants Only or Exclusive Access – Registrants Only type reservation becomes unregistered the persistent reservation shall be released. The device server shall establish a unit attention condition for every initiator port associated with a registered I_T nexus other than the initiator port associated with the I_T nexus that unregistered. The sense key shall be set to UNIT ATTENTION and the additional sense code shall be set to RESERVATIONS RELEASED.

The device server shall establish a unit attention for every initiator port associated with a registered I_T nexus whose reservation key was removed. The additional sense code shall be set as follows:

- a) If the service action was CLEAR, the additional sense code shall be set to RESERVATIONS PREEMPTED; or
- b) If the service action was PREEMPT or PREEMPT AND ABORT, the additional sense code shall be set to REGISTRATIONS PREEMPTED.

If the TYPE or SCOPE changed ~~or the reservation was released~~, the device server shall establish a unit attention for every initiator port associated with a registered I_T nexus whose reservation key was not removed except for the initiator port through which the command was issued. The additional sense code shall be set as follows:

- a) If the service action was PREEMPT or PREEMPT AND ABORT, the additional sense code shall be set to RESERVATIONS RELEASED;
- b) If the service action was REGISTER or REGISTER AND IGNORE with the SERVICE ACTION KEY field set to zero, the additional sense code shall be set to RESERVATIONS RELEASED; or
- ~~c) If the service action was RELEASE, the additional sense code shall be set to RESERVATIONS RELEASED.~~

[If a persistent reservation was released using a RELEASE service action see 5.5.2.7.2.](#)

3.0.0.0.0.2 Handling for released all registrants persistent reservations

A Write Exclusive – All Registrants or Exclusive Access – All Registrants type persistent reservation shall be released when the registration for the last registered I_T nexus is removed or when the TYPE or SCOPE is changed.

The device server shall establish a unit attention for every initiator port associated with a registered I_T nexus whose reservation key was removed. The additional sense code shall be set as follows:

- a) If the service action was CLEAR, the additional sense code shall be set to RESERVATIONS PREEMPTED; or
- b) If the service action was PREEMPT or PREEMPT AND ABORT, the additional sense code shall be set to REGISTRATIONS PREEMPTED.

If a persistent reservation was released using a RELEASE service action [see 5.5.2.7.2](#), ~~the device server shall establish a unit attention for every initiator port associated with a registered I_T nexus except for the initiator port through which the command was issued. The additional sense code shall be set to RESERVATIONS RELEASED.~~

4 Overview Problem 2

It is not clear what should happen if an All Registrants reservation is in effect, an ACA is in effect, and an initiator issues a Preempt and Abort service action to an initiator other than the faulted initiator.

The answer to that question is: The logical unit should honor the Preempt and Abort request.

That is not clear in the standard and wording will be proposed to fix that.

5 Solution

The logical unit should accept and execute the Preempt and Abort request and abort all commands.

6 Proposal

Change the following from section 5.5.2.7.5 Preempting and aborting as shown:

6.0.0.0.1 Preempting and aborting

The application client's request for and the device server's responses to a PERSISTENT RESERVE OUT command PREEMPT AND ABORT service action are identical to the responses to a PREEMPT service action (see 5.5.2.7.4) except for the following additions. If no reservation conflict occurred, the device server shall perform the following uninterrupted series of actions:

- a) [If the persistent reservation is not an all registrants type then:](#)
 - A) if the TST field is 000b (see 7.4.6) and ACA condition exists for initiator ports other than the initiator port associated with the [persistent reservation being preempted I_T nexus that is being preempted](#), the PERSISTENT RESERVE OUT command shall be terminated prior to processing with a status of ACA ACTIVE if the NACA bit equals one in the CDB CONTROL byte (see SAM-2) or BUSY if the NACA equals zero; [or](#)
 - B) if the TST field contains 001b, then ACA condition for initiator ports other than the initiator port associated with the [persistent reservations being preempted I_T nexuses that are being preempted](#) shall not prevent the processing of the PERSISTENT RESERVE OUT command;
- b) Perform the uninterrupted series of actions described for the PREEMPT service action (see 5.5.2.7.4);
- c) All tasks from the initiator ports associated with the [persistent reservations being preempted preempted I_T nexuses](#) (called preempted tasks) except the task containing the PERSISTENT RESERVE OUT command itself shall be terminated. Application client notification shall be provided, as specified by the TAS bit in the Control mode page (see 7.4.6) that applies to the initiator port associated with the [persistent reservation being preempted preempted I_T nexus](#) (called the preempted initiator port), as follows:
 - A) If the TAS bit is set to zero then all preempted tasks shall be terminated as if an ABORT TASK SET task management function had been performed by each preempted initiator port; or

- B) If the TAS bit is set to one then all preempted tasks from initiator ports other than the initiator port that sent the PREEMPT AND ABORT service action shall be terminated with a TASK ABORTED status (see SAM-2). Any preempted tasks from the initiator port that sent the PREEMPT AND ABORT service action shall be terminated as if an ABORT TASK SET task management function had been performed by that initiator port.

If a terminated task is a command that causes the device server to generate additional commands and data transfers (e.g., EXTENDED COPY), all commands and data transfers generated by the command shall be terminated before the ABORT TASK SET task management function is considered completed. After the ABORT TASK SET function has completed, all new tasks are subject to the persistent reservation restrictions established by the preempting initiator port;

- d) If the persistent reservation is not an all registrants type then the device server shall clear any ACA condition associated with an initiator port being preempted and shall clear any tasks with an ACA attribute from that initiator port;
- e) If the persistent reservation is an all registrants type then the device server shall clear any ACA condition and shall clear any tasks with an ACA attribute; and
- f) For logical units 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 port or initiator ports associated with the persistent reservation or reservations being preempted ~~I_T nexus or I_T nexuses being preempted~~ (see 6.13).

The actions described in the preceding list shall be performed for all I_T nexuses that are registered with the SERVICE ACTION RESERVATION KEY value, without regard for whether the preempted I_T nexuses hold the persistent reservation. If an all registrants persistent reservation is present ~~and the SERVICE ACTION RESERVATION KEY value is set to zero~~ the device server shall abort all tasks for all registered I_T nexuses.

7 Overview Problem 3

From Rob Elliott email dated 7/15/2003:

In the PREEMPT flowchart in spc3r05 and earlier, the "Existing persistent reservation? no" question lead to "Remove registration pointed to by the service action reservation key." When All Registrants was added in spc3r06, the arc was changed to go straight to "Done".

How should PREEMPT and PREEMPT AND ABORT be handled if there is no persistent reservation at all? Do nothing (per spc3r05), or go ahead and remove the specified registrations and (for PREEMPT AND ABORT) abort the related tasks (per spc3r06+)?

There doesn't seem to be much supporting text for the spc3r06+ figure except for the first 7 words in 5.5.2.7.4.4: "If there is a persistent reservation and if the SERVICE ACTION RESERVATION KEY field does not identify a persistent reservation holder the device server shall perform a preempt by doing the following in an uninterrupted series of actions:". This phrase was added in spc3r13 per 03-094r3 to better match the (possibly broken) figure 3.

I think the original behavior of "remove registrations (and abort tasks)" might make more sense.

See:

SPC-3 revision 5 section 5.5.3.7.3.1 figure 3 page 69

SPC-3 revision 14 section 5.5.2.7.4.1 figure 3 page 72

03-094r3 Yet more persistent reservations fixes

8 Response

It looks to me like the line in the flow chart is wrong. The text from 03-094r3 assumed (incorrectly it appears) that the flowchart was correct. Also, it appears a line of text from 03-094r3 was not placed into SPC-3. It stated (again incorrectly):

In addition a paragraph needs to be added to this section that states << If there is no persistent reservation there shall be no change to any registrations and a GOOD status shall be returned. >>

9 Solution

I believe the flowchart should be fixed (as shown below) and the paragraph that was not placed into SPC-3 not be placed into SPC-3. The paragraph in section 5.5.2.7.4.4 Removing registrations:

If there is a persistent reservation and if the SERVICE ACTION RESERVATION KEY field does not identify a persistent reservation holder the device server shall perform a preempt by doing the following in an uninterrupted series of actions:

Should be changed to:

If the SERVICE ACTION RESERVATION KEY field does not identify a persistent reservation holder, or there is no persistent reservation holder (i.e., there is no persistent reservation), the device server shall perform a preempt by doing the following in an uninterrupted series of actions:

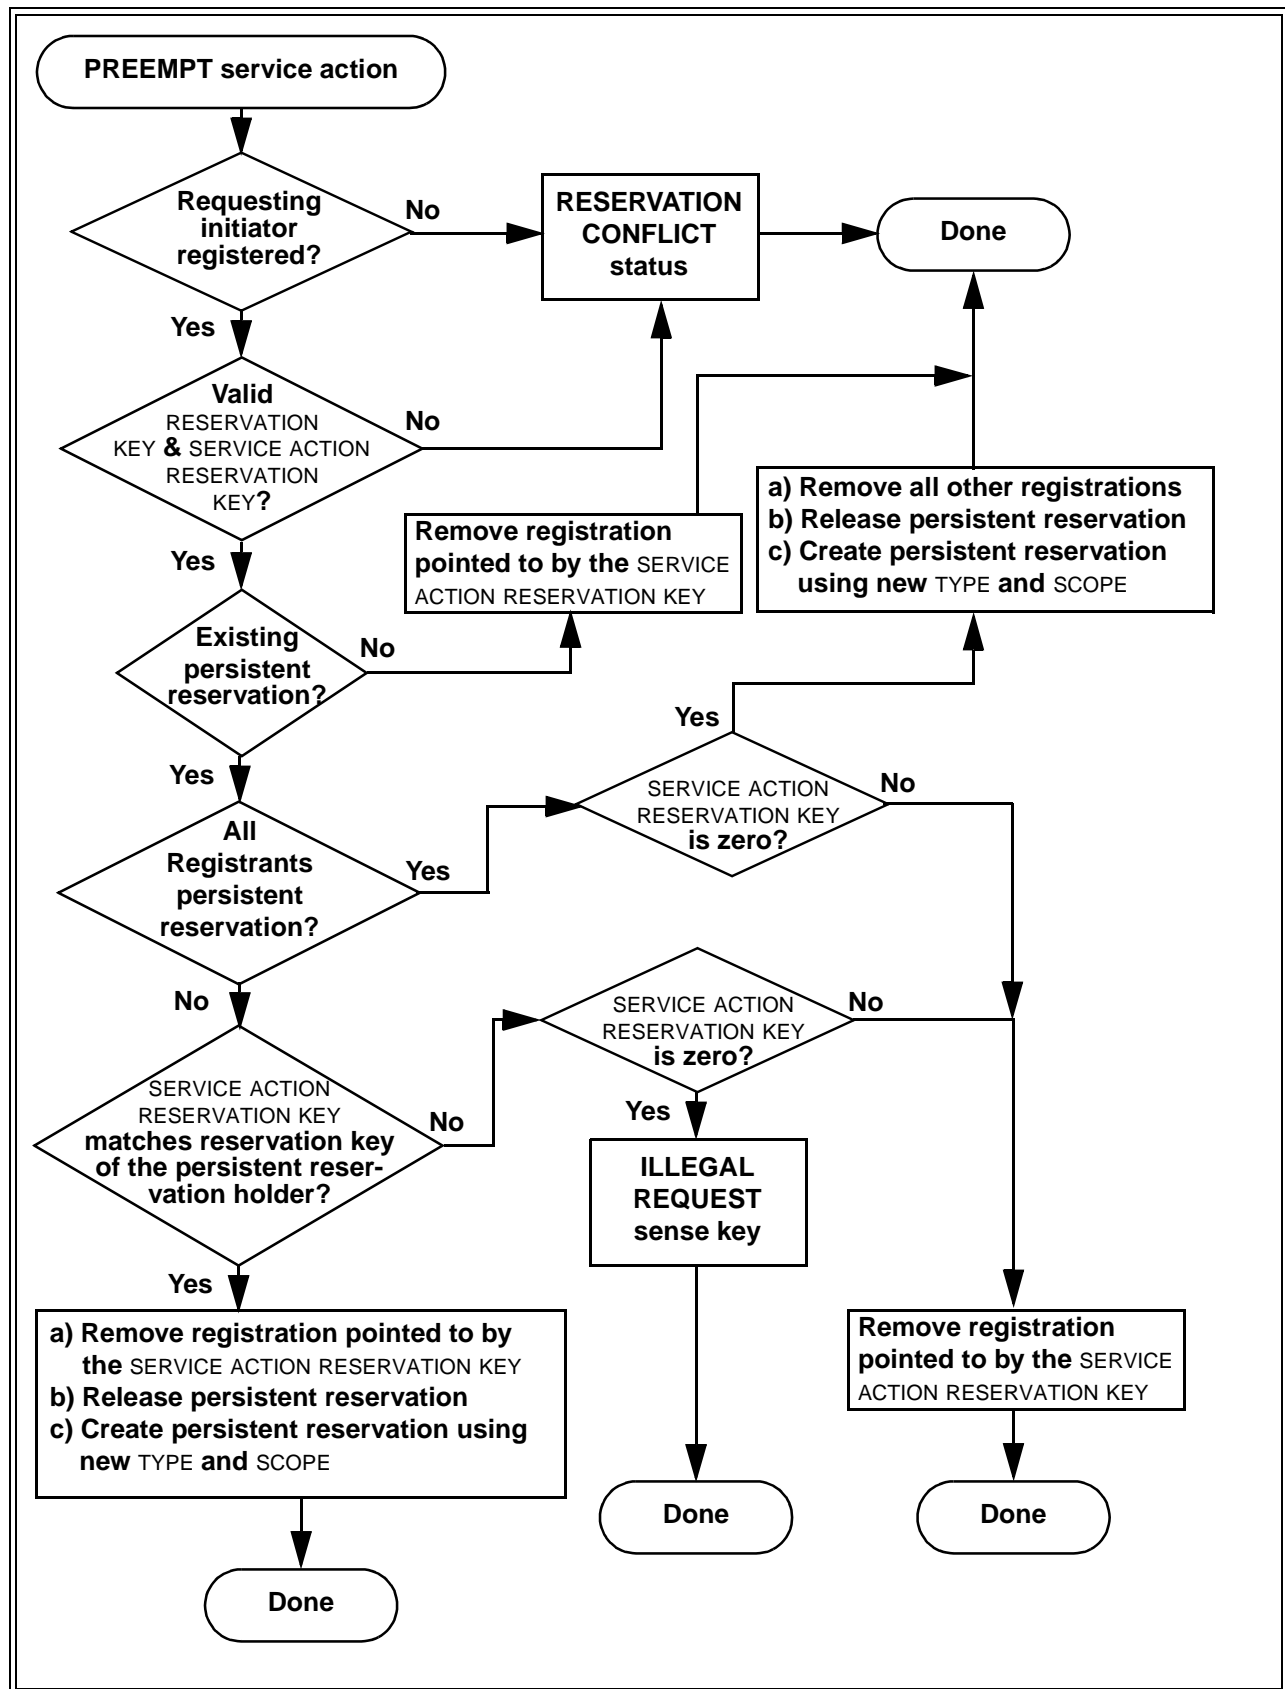


Figure 1 — Device server interpretation of PREEMPT service action