

Accredited Standards Committee

NCITS, National Committee for Information Technology Standards

Doc: T10/97-242r1

Date: 10/31/97

Project: SPC-x

Ref Doc:

Reply to: Bob Snively

To: T10 Membership

From: Bob Snively

Subject: Additional persistent reservation type

During review of persistent reservation, a requirement for an additional type of persistent reservation was identified.

1.0 Overview of requirement

There are presently six types of reservation defined for persistent reservations. One of the types is defined to allow write access for only those initiators that have presently registered a reservation key with the device. Any initiator is allowed to read from the device and any initiator is allowed to create additional nonconflicting reservations. This type of reservation is intended to allow the managing initiator for a cluster of hosts to execute the Preempt and Clear or the Preempt service action to instantly remove the ability of a failing initiator to perform a write.

After careful review, it appears that an additional type of persistent reservation would be helpful in similar cluster environments. The reservation would allow read and write access for only those initiators that have presently registered a reservation key with the device. Unregistered initiators would not be allowed any media access. This type of reservation is intended to allow the managing initiator for a cluster of hosts to execute the Preempt and Clear or the Preempt service action to instantly remove the ability of a failing initiator to perform any media access. This provides the additional system benefit of providing a strong indication to the failing initiator that it shall not use data until it has reestablished itself among the cluster and performed a new registration. This is not intended to increase the security of data, but inadvertent misuse of data by initiators agreeing to the proper clustering software conventions will be prevented.

2.0 Proposed “Exclusive access, registrants only” type

At the earliest time allowed by the standards process, the following additions should be made to the SPC document’s present content.

Section 7.12.3.2 should have an additional entry placed immediately following the promised entry

for the “Write Exclusive, Registrants only” type of reservation. The new entry will be “Exclusive Access, Registrants Only” and will have type code 6. These two reservation types would then have an entry in table 41 as follows:

Table 41: Persistent Reservation Type Codes (continued)

Code	Name	Description
5h	Write Exclusive, Registrants Only	<p>Reads Shared: Any application client on any initiator may execute commands that perform transfers from the storage medium or cache of the logical unit to the initiator.</p> <p>Writes Exclusive: Any command that performs a transfer from an initiator that has not previously performed a Register service action with the device server, to the storage medium or cache of the logical unit, shall result in a reservation conflict.</p> <p>Additional Reservations Allowed: Any initiator may reserve the logical unit or extents or elements as long as the persistent reservations do not conflict with any reservations that are already known to the device server.</p>
6h	Exclusive access, Registrants Only	<p>Reads Exclusive: Any command that performs a transfer to an initiator that has not previously performed a Register service action with the device server, from the storage medium or cache of the logical unit, shall result in a reservation conflict.</p> <p>Writes Exclusive: Any command that performs a transfer from an initiator that has not previously performed a Register service action with the device server, to the storage medium or cache of the logical unit, shall result in a reservation conflict.</p> <p>Additional Reservations Allowed: Any initiator may reserve the logical unit or extents or elements as long as the persistent reservations do not conflict with any reservations that are already known to the device server.</p>
7-Fh	Reserved	

Table 42 would be modified as follows:

Table 42: New Persistent Reservation Conflicts with Existing

Persistent Reservation That is Being Attempted	Persistent Reservation That Is Held														
	Read Shared		Write Exclusive		Read Exclusive		Exclusive Access*		Shared Access*		Write Exclusive, Registrants Only		Exclusive Access, Registrants Only		
	LU	EX	LU	EX	LU	EX	LU	EX	LU	EX	LU	EX	LU	EX	
Read Shared	LU	N	N	Y	Y	Y	Y	Y	Y	N	N	Y	Y	Y	Y
	EX	N	N	Y	O	Y	O	Y	O	N	N	Y	O	Y	O
Write Exclusive	LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	EX	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O
Read Exclusive	LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	EX	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O
Exclusive Access	LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	EX	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O
Shared Access*	LU	N	N	Y	Y	Y	Y	Y	Y	N	N	Y	Y	Y	Y
	EX	N	N	Y	O	Y	O	Y	O	N	N	Y	O	Y	O
Write Exclusive, Registrants Only	LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N
	EX	Y	O	Y	O	Y	O	Y	O	Y	O	N	N	N	N
Exclusive Access, Registrants Only	LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N
	EX	Y	O	Y	O	Y	O	Y	O	Y	O	N	N	N	N

Key:

LU = Logical Unit scope	N = no conflict
EX = Extent or Element scope	Y = conflict
* = Conflicts with all reservation requests from other initiators	O = conflict occurs if extent or element reservation requests overlaps with existing extent or element reservation

Sincerely,

Robert Snively

Phone:(415) 786-6694

Email: bob.snively@sun.com