

July 24, 1998

T10/98-164, revision 1



To: NCITS Technical Committee T10

From: Bob Snively

Subject: Command behavior under reservations

This document reviews the presently defined behavior of commands encountering a legacy reservation and proposes the proper behavior for commands encountering the various classes of persistent reservation. The attached chart indicates those commands that may be executed in the presence of each type of reservation. If the entry in the chart is blank, a reservation conflict is the expected behavior. This work is based on a proposed chart by Ralph Weber, with additional inputs from the T11 SCSI working group meeting on May 6, 1998. Revision 1 adds the results of discussions at the July 15, 1998 working group meeting. That discussion led us to change most "always conflicts" commands to "shared access" commands, placing the responsibility for proper management of most control commands on a cooperative relationship among participating hosts. Revision 1 additionally assumes that those functions proposed for removal by George Penokie's document, 98-203, are in fact removed. It additionally removes those commands and types of reservations defined as obsolete in SCP-2. No change bars are used because the revision changes most lines in the document.

Please review note 10, which describes my assumptions about the SET LIMITS command, showing the assumptions that led to a different solution than that proposed by the working group.

For the persistent reservation definitions, logical unit and element reservation violations are treated identically and are expected to create a status of RESERVATION CONFLICT. This is possible because it is simple to identify what conflicting reservations may exist and which controlling host maintains the reservations.

Charts are provided for commands defined in SPC-2, SBC block devices, SSC sequential access devices, and SMC devices. The information in the left column of each chart is extracted from the documents, but extent reservation information has been dropped. To fit the necessary information in the chart, some command names have been abbreviated, although those familiar with the SCSI command set should easily be able to identify the commands.

The following charts are working documents intended to provide input for later reservation definitions. Sample wording is provided in the appropriate note. This

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

sample wording is expected to be added to SPC-2 and CSS, and as an erratum, supplement, or future revision to the SBC document. The sample wording takes into account the elimination of extent reservations.

Please review the SSC and SMC command sets, as they have not been published previously.

The following table lists the acronyms used in the reservation conflict charts.

Table 1: Key to acronyms

Acronym	Meaning
Reservation types and extents	
LU	Logical Unit
Excl	Exclusive
RO	Registrants Only
RS	Reads Shared
WP	Writes Prohibited
WS	Write Shared
WX	Writes Exclusive
RX	Reads Exclusive
Reg	The initiator performing the operation is a registrant
Not Reg	The initiator performing the operating is not a registrant
Indication of conditions that permit command to be executed in presence of reservations	
“blank”	Reservation conflict will occur
A	Allowed to execute, no reservation conflict will occur
Grouping of similar command behaviors	
R	Command is treated like READ
W	Command is treated like WRITE
AA	Command is always allowed, regardless of reservations.
AC	Always conflicting
SA	Shared Access Note that this has the same effect as the restrictions specified for write, but for architecturally different reasons.
?	Indicates anomalous behavior in legacy reserve. Proposal included below table.



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 2: SPC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
		LU	RS WX	RX WX	RS WX	RS WX	RX WX		
COMPARE		A		A	A	A		R	1
COPY				A		A		W	2
COPY & VERIFY				A		A		W	2
INQUIRY	A	A	A	A	A	A	A	AA	3
LOG SELECT				A		A		SA	4
LOG SENSE	A	A	A	A	A	A	A	AA	3
MODE SELECT(6)				A		A		SA	4
MODE SELECT(10)				A		A		SA	4
MODE SENSE(6)				A		A		SA	4
MODE SENSE(10)				A		A		SA	4
PERS RES IN		A	A	A	A	A	A	AA	6
PERS RES OUT (Reg, Release)		A	A	A	A	A	A	AA	6
PERS RES OUT (Res, Preempt other)				A		A		SA	6
PERS RES OUT (Preempt same)		A	A	A	A	A	A	AA	6
PREVNT/ALLOW (prevent equal 0)	A	A	A	A	A	A	A	AA	3, 7
PREVNT/ALLOW (prevent not equal 0)				A		A		SA	4, 7
READ BUFFER				A		A		SA	4
RECEIVE DIAG.				A		A		SA	4
RELEASE(10)								AC	8
RELEASE(6)								AC	8



Table 2: SPC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
LU	RS WX	RX WX	RS WX	RS WX	RX WX	RX WX			
REPORT LUNS	A	A	A	A	A	A	A	AA	3
REQUEST SENSE	A	A	A	A	A	A	A	AA	3
RESERVE(10)								AC	8
RESERVE(6)								AC	8
SEND DIAG				A		A		SA	4
TEST UNIT RDY				A		A		SA	4
WRITE BUFFER				A		A		SA	4

Table 3: SBC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
LU	RS WX	RX WX	RS WX	RS WX	RX WX	RX WX			
FORMAT UNIT				A		A		SA	4
LOCK/UNL CACHE				A		A		SA	4
PRE-FETCH		A		A	A	A		R	1
READ(6)		A		A	A	A		R	1
READ(10)		A		A	A	A		R	1
READ CAPACITY	A	A	A	A	A	A	A	AA	3



Table 3: SBC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
LU	RS WX	RX WX	RS WX	RS WX	RX WX	RX WX			
READ DEFCT DATA				A		A		SA	4
READ LONG				A		A		SA	4
REASSIGN BLKS				A		A		SA	4
REBUILD				A		A		SA	4, 9
REGENERATE				A		A		SA	4, 9
SEEK(10)				A		A		SA	4
SET LIMITS(10)	A	A	A	A	A	A	A	AA	3, 10
START/STOP UNIT START = 1 and power condition = 0	A	A	A	A	A	A	A	AA	3, 11
START/STOP UNIT START = 0 or power condition not 0				A		A		SA	4, 11
SYNCH CACHE				A		A		SA	4
VERIFY		A		A	A	A		R	1
WRITE(6)				A		A		W	2
WRITE(10)				A		A		W	2
WRITE & VERIFY				A		A		W	2
WRITE LONG				A		A		SA	4
WRITE SAME				A		A		SA	4
XDREAD		A		A	A	A		R	1
XDWRITE				A		A		W	2
XDWRITE EXT				A		A		W	2
XPWRITE				A		A		W	2



Table 4: SSC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
		LU	RS WX	RX WX	RS WX	RS WX	RX WX		
ERASE				A		A		W	2
FORMAT MEDIUM				A		A		W	2
LOAD UNLOAD				A		A		SA	4
LOCATE		A		A	A	A		R	1
READ		A		A	A	A		R	1
READ BLK LIMITS	A	A	A	A	A	A	A	AA	3, 12
READ POSITION		A		A	A	A		R	1
READ REVERSE		A		A	A	A		R	1
RECOVER BRFD DATA				A		A		SA	4
REPORT DENSITY SUPPORTED	A	A	A	A	A	A	A	AA	3, 12
REWIND				A		A		SA	4
SET CAPACITY				A		A		W	2
SPACE		A		A	A	A		R	1
VERIFY		A		A	A	A		R	1
WRITE				A		A		W	2
WRITE FILEMARKS				A		A		W	2



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 5: SMC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
		LU	RS WX	RX WX	RS WX	RS WX	RX WX		
EXCHANGE MEDIUM				A		A		SA	4
INITIALIZE ELEMENT				A		A		SA	4
MOVE MEDIUM				A		A		SA	4
MOVE MEDIUM ATTACHED				A		A		SA	4
POSITION TO ELEMENT				A		A		SA	4
READ ELEMENT STATUS curdata = 0				A		A		SA	4
READ ELEMENT STATUS curdata = 1	A	A	A	A	A	A	A	AA	3
READ ELEMENT STATUS ATTACHED curdata = 0				A		A		SA	4
READ ELEMENT STATUS ATTACHED curdata = 1	A	A	A	A	A	A	A	AA	3
RELEASE ELEMENT(6)	A								13



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 5: SMC commands that are allowed in the presence of various reservations

Command	Addressed LU has RESERVE logical unit held by another initiator	Addressed LU has this type of persistent reservation held by another initiator						Class of Action	Notes
		Write Excl	Excl Access	Reg	Not Reg	Reg	Not Reg		
				Write Excl RO	Write Excl RO	Excl Access RO	Excl Access RO		
LU	RS WX	RX WX	RS WX	RS WX	RX WX	RX WX			
READ RELEASE ELEMENT(10)	A								13
REQUEST VOLUME ELEMENT ADDR				A		A		SA	4
RESERVE ELEMENT (6)								AC	5
RESERVE ELEMENT(10)								AC	5
SEND VOLUME TAG				A		A		SA	4

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54



Notes:

- 1. This is the defining note for all READ type reservation commands. The following sample wording is suggested for all READ type commands.

If the logical unit is reserved by a RESERVE(6) or a RESERVE(10) command, a reservation conflict shall occur when a [insert read type command name here] command is received from an initiator other than the one holding a logical unit reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access, a reservation conflict shall occur when a [insert read type command name here] command is received from an initiator other than the one holding the reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only, a reservation conflict shall occur when a [insert read type command name here] command is received from an initiator that is not registered. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer shall take place, and no change in the storage media's state or content shall take place.

- 2. This is the defining note for all WRITE type reservation commands. The following sample wording is suggested for all WRITE type commands.

If the logical unit is reserved by a RESERVE(6) or a RESERVE(10) command, a reservation conflict shall occur when a [insert write type command name here] command is received from an initiator other than the one holding a logical unit reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Write Exclusive or Exclusive Access, a reservation conflict shall occur when a [insert write type command name here] command is received from an initiator other than the one holding the reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only or Write Exclusive Registrants Only, a reservation conflict shall occur when a [insert write type command name here] command is received from an initiator that is not registered. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer shall take place, and no change in the storage media's state or content shall take place.

- 3. This is the defining note for all "always allowed" type commands. The following sample wording is suggested for all such commands.

The [insert name of inquiry type command here] command shall not be affected by reservations or persistent reservations.

- 4. This is the defining note for all "shared access" type commands. The following sample wording is suggested for all shared access type commands.

If the logical unit is reserved by a RESERVE(6) or a RESERVE(10) command, a reservation conflict shall occur when a [shared access type command name here] command is received from an initiator other than the one holding a logical unit reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Write Exclusive or Exclusive Access, a reservation conflict shall occur when a [insert shared access type command name here] command is received from an initiator other than the one holding the



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only or Write Exclusive Registrants Only, a reservation conflict shall occur when a [insert shared access type command name here] command is received from an initiator that is not registered. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer shall take place, and no change in the storage media's state or content shall take place.

- 5. This is the defining note for all "always conflicting" type commands. The following sample wording is suggested for all always conflicting type commands.

If the logical unit is reserved by any RESERVE(6), RESERVE(10), or PERSISTENT RESERVE OUT command, a reservation conflict shall occur when a [always conflicting type command name here] command is received from an initiator other than the one holding the reservation. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer shall take place, and no change in the storage media's state or content shall take place.

- 6. Text for PERSISTENT RESERVE IN and PERSISTENT RESERVE OUT is already contained in SPC-2.

- 7. The PREVENT ALLOW MEDIA REMOVAL behaves in such a manner that any initiator can allow permission to remove media, but only the initiator holding a reservation can prevent the removal of media. The actual removal process requires all initiators to agree that media can be removed before media removal can take place. The sample wording is obviously the proper combination of note 4 and note 3.

The PREVENT ALLOW MEDIA REMOVAL command with the Prevent value set to zero shall not be affected by reservations or persistent reservations.

If the logical unit is reserved by a RESERVE(6) or a RESERVE(10) command, a reservation conflict shall occur when a PREVENT ALLOW MEDIA REMOVAL command with a non-zero Prevent value is received from an initiator other than the one holding a logical unit reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Write Exclusive or Exclusive Access, a reservation conflict shall occur when a PREVENT ALLOW MEDIA REMOVAL command with a non-zero Prevent value is received from an initiator other than the one holding the reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only or Write Exclusive Registrants Only, a reservation conflict shall occur when a PREVENT ALLOW MEDIA REMOVAL command with a non-zero Prevent value is received from an initiator that is not registered. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Note that an interesting additional mechanism for resetting the prevent state might be desirable, as shown in proposal 1.

- 8. Text for RESERVE and RELEASE is already contained in SPC-2.
- 9. The behavior of the REBUILD and REGENERATE commands with a RESERVE reservation present is not defined. Reservation conflicts should occur in that case, as specified in proposal 2.
- 10. The SET LIMITS command had previously been allowed if no extent reservations overlapped with the blocks to be reassigned. Since the SET LIMITS command provides protection in links of commands only, and since the command for which protection is to be forced is not defined, I have assumed that the SET LIMITS command can always be executed without conflict. If the subsequent commands in the link conflict with a reservation, reservation conflict will be presented to that subsequent command only. If the subsequent commands in the link conflict with the limits established by the SET LIMITS command, a CHECK CONDITION with a sense key of DATA PROTECT will be presented. SET LIMITS and command linking are optional and not commonly implemented. See proposal 3.
- 11. The START STOP UNIT command was reviewed by the committee. The committee proposed that any unit be able to start the device, but that only those units sharing a persistent reservation should be able to stop the device. I have followed this proposal, but added power condition commands to those restricted by reservations. The following text should replace the second paragraph of 6.1.14 in SBC the next time SBC is updated or an annex is published:
 If the logical unit is reserved by a RESERVE command, a reservation conflict shall occur when a START STOP UNIT command is received from an initiator other than the one holding a logical unit reservation. If a persistent reservation is present from any initiator and if a START STOP UNIT command with a POWER CONDITIONS field equal to zero and a START bit equal to one, indicating the device is to be started, no reservation conflict shall occur. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Read Shared, Write Exclusive, Read Exclusive, or Exclusive Access, a reservation conflict shall occur when a START STOP UNIT command with the POWER CONDITIONS field not equal to zero or the START bit equal to zero is received from an initiator other than the one holding the reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only or Write Exclusive Registrants Only, a reservation conflict shall occur when a START STOP UNIT command with the POWER CONDITIONS field not equal to zero or the START bit equal to zero is received from an initiator that is not registered. The command shall be rejected with RESERVATION CONFLICT status if a reservation conflict occurs and no power state change shall occur.
- 12. The READ BLOCK LIMITS and REPORT DENSITY SUPPORT commands are assumed to be simple informative commands that can be read by any initiator. This is



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

different from the present definition for tape drive logical units reserved by the RESERVE command.

- 13. The RESERVE ELEMENT command has special reservation conflict text defined. This text must be extended for the case of persistent reservations as shown below:

The RESERVE ELEMENT and RELEASE ELEMENT commands provide the basic mechanism for contention resolution in multiple initiator systems. A reservation may only be released by a RELEASE command from the initiator that made it. It is not an error for an application client to attempt to release a reservation that is not currently valid, or is held by another initiator. In this case, the device server shall return GOOD status without altering any other reservation. If a persistent reservation exists for either the element or the addressed logical unit, a reservation conflict shall occur. The command shall be rejected with RESERVATION CONFLICT status if a reservation conflict occurs.

Proposals:

- 1. PREVENT ALLOW MEDIA REMOVAL clearing mechanism

At present, there are only two mechanisms for allowing media removal once one initiator has set the machine to prevent media removal. Either all initiators can release the prevent media removal condition or a hard reset can be performed. In the presence of multiple initiators and persistent reservations, an initiator that is the target of a Preempt service action or a Preempt and Clear service action no longer be considered in the media removal prevention algorithm. One mechanism to perform this would be to insert the following sample paragraph after the second paragraph of section 7.14 in SPC-2.

The execution of a PERSISTENT RESERVE OUT command with a Preempt service action or a Preempt and Clear service action shall implicitly set the Prevent value to 00b for the initiator or initiators targeted by the service action.

- 2. Reservation conflict for REBUILD and REGENERATE

At present, the reservation behavior for the REBUILD and REGENERATE commands is not specified. With the additional text included for persistent reservation, the following text, based on the shared access type reservation, is proposed for a new paragraph for the REBUILD command immediately following table 23, section 6.1.10 of SBC:

If the logical unit is reserved by a RESERVE(6) or a RESERVE(10) command, a reservation conflict shall occur when a REBUILD command is received from an initiator other than the one holding a logical unit reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Write Exclusive or Exclusive Access, a reservation conflict shall occur when a REBUILD command is received from an initiator other than the one holding the reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only or Write Exclusive Registrants Only, a



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

reservation conflict shall occur when a REBUILD command is received from an initiator that is not registered. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer shall take place, and no change in the storage media's state or content shall take place.

With the additional text included for persistent reservation, the following text is proposed for a new paragraph for the REGENERATE command immediately following table 27, section 6.1.11 of SBC:

If the logical unit is reserved by a RESERVE(6) or a RESERVE(10) command, a reservation conflict shall occur when a REGENERATE command is received from an initiator other than the one holding a logical unit reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Write Exclusive or Exclusive Access, a reservation conflict shall occur when a REGENERATE] command is received from an initiator other than the one holding the reservation. If the logical unit is successfully reserved by a PERSISTENT RESERVE OUT service action of Exclusive Access Registrants Only or Write Exclusive Registrants Only, a reservation conflict shall occur when a REGENERATE command is received from an initiator that is not registered. If a reservation conflict occurs, the command shall be rejected with RESERVATION CONFLICT status, no data transfer shall take place, and no change in the storage media's state or content shall take place.

3. No reservation conflict for SET LIMITS

At present, the SET LIMITS command always receives RESERVATION CONFLICT status if a legacy reserve is present. I have proposed that the command never conflict for persistent reservations, since it is only the subsequent commands in the link that are meaningfully influenced by the SET LIMITS command or reservations. Assuming we do not take the wiser course of making the SET LIMITS command obsolete, the following wording is proposed to replace the second paragraph of section 6.1.13 of SBC.

The SET LIMITS command shall not be affected by reservations or persistent reservations.



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54