T10/00−232 r9

To: John Lohmeyer, chairperson, T10
From: Ken Moe Email: kenneth.moe@sun.com
Date: April 25, 2001
Subject: Asymmetric SCSI behavior

1 Introduction

A significant number of SCSI storage subsystems have the property of asymmetric accessibility to logical units through various target ports. Typically, one target port may provide full performance access to a logical unit, while another target port, possibly on a different physical controller, may provide either lower performance access or may support a subset of the available SCSI commands to the same logical unit. In some cases, the logical unit may be modified to provide full performance access to the limited target port if the original full performance target port fails. This proposal proposes a set of SCSI tools necessary to properly support such asymmetric access and failure recovery.

This revision addresses the recommendations provided during the April 23, 2001 conference call. There are a number of editorial changes, and two technical changes, included in this revision. The two technical changes are;

- *The device Identification page description in section 8.4.4 was revised to match another recently approved proposal, T10/01−065r2.*
- *The alignment of the ASYMMETRIC ACCESS STATE field was changed in both the REPORT TARGET PORT GROUP and SET TARGET PORT GROUP descriptors (tables XX3 and YY3 respectively), and the alignment of the TARGET PORT GROUP field was also changed in the SET TARGET PORT GROUP descriptor. These changes were made to make the alignment consistent with the new alignment of TARGET PORT GROUP in INQUIRY (see 8.4.4).*

All changes to the recommended modifications to SPC−3 (beginning with section 3 in this proposal) are highlighted. Strikeouts are shown and additions are shaded.

Thanks are owed to many for their help in developing this proposal, but I feel a special thanks is owed to four individuals. Bob Snively for his excellent initial version of this proposal which got the proposal off to a fine start. Rob Elliott for his help and guidance on day to day matters has been invaluable. I regard Rob as a coeditor of this proposal. George Penokie and Ralph Weber for their comprehensive review of this proposal, and their absolutely critical recommendations and corrections. I also owe a thanks to these individuals for their patience with my probably violating every rule in the "unwritten" style guide, including the unforgivable usage of the word "execute".

1.1 Overview

Symmetric access to logical units is very desirable, since it provides for very rapid recovery from link failures and it provides the infrastructure that supports dynamic load balancing capabilities. Symmetric access is characteristic of almost all simple disk drives and JBOD's. Symmetric access may be managed by the simple SCSI mechanisms already defined in SAM−2, SPC−2 and other documents.

Asymmetric access is useful, since it may be implemented for very large storage subsystems with very simple and low−cost storage controller configurations. Asymmetric access requires additional SCSI mechanisms to indicate that target ports are fully accessible and that target ports have only partial access to a particular logical unit. Mechanisms are also required to allow controlled transfer of the full performance functionality from one target port to another target port for a particular logical unit. The proposals in this document provide those mechanisms.

1.2 Definition of target port asymmetric access state and of target port group

Logical units may be connected to the service delivery subsystem via multiple target ports (see SAM−2). A target port asymmetric access state defines the performance properties and allowable command set for a logical unit when accessed through the target port maintaining that state.

A target port group is defined as a set of target ports that are in the same target port asymmetric access state at all times.

A logical unit may be connected to multiple target port groups. Logical units support asymmetric logical unit access if different target port groups may be in different target port asymmetric access states.

2 Parameters and states for managing asymmetrical access to SCSI logical units

The proposal is designed to be included in SPC−3. The commands indicating that the logical unit implements asymmetric access are extracted from SPC−2. The commands selected by the working group for target port discovery will be carried into SPC−3, but were originally defined for SCC−2.

Target Port Asymmetric Access State

All target ports in a target port group that supports asymmetric access to logical units shall be in one of the following target port asymmetric access states with respect to the ability to access a particular logical unit:

Active/Optimized :

While in the active/optimized state the target port group should be capable of accessing the logical unit. All commands operate exactly as specified in the appropriate command set standards.

Active/Non−Optimized:

While in the active/non−optimized state the device server shall support all commands that the logical unit supports. These commands shall operate exactly as specified in the appropriate command set standards. The execution of certain commands, especially those involving data transfer or caching, may operate with lower performance than they would if the target port group were in the active/optimized state.

Standby:

While in the standby state all target ports in a target port group are capable of performing a limited set of commands. The standby state is intended to provide a state from which it should be possible to provide a higher level of accessibility, should this become necessary for any reason, to a logical unit by transitioning to either the active/optimized or active/non−optimized states.

Commands that operate in the standby state are those necessary for:

- *Diagnosing and testing the logical unit and its paths*
- *Identifying the path*
- *Identifying the logical unit*
- *Determining the operational state of the logical unit*
- *Determining the active/inactive state of the unit*
- *Manage or remove logical unit or element reservations*
- *Testing Service delivery subsystem*

The commands that shall operate normally in the standby state are listed in section 5.6.4.4.

Unavailable:

While in the unavailable state the device server shall accept only a limited set of commands specified in the appropriate command set standards. The unavailable state is intended for the situation when the target port accessibility to a logical unit may be severely restricted due to, for example, a hardware error and therefore it may not be possible to transition from this state to either the active/optimized, active/non−optimized or standby states. The unavailable state is also intended for minimizing any disruption when using the downloading microcode mode of the WRITE BUFFER command.

The commands that shall operate normally in the unavailable state are listed in section 5.6.4.5.

Management function

The following management operation may be performed for each SCSI target device:

Change access to logical unit:

The access of a target port group to a logical unit may be modified. The previously active/optimized target ports in a target port group may be changed to active/non−optimized, standby, or unavailable in a vendor specific manner. (see 4.7 in this document)

This function can be done explicitly or automatically. See 4.3 in this document to determine whether the SCSI target device requires explicit or implicit change.

Discovery requirements

The following information shall be discoverable by an appropriate mechanism

Identify asymmetric logical units access requirement:

A value is provided to indicate that only one target port group at a time is allowed normal access to a logical unit. (see 4.4, in this document)

Identify target port group to logical unit:

A value is provided identifying the target port and target port group through which a command is being passed (see 4.5 in this document).

Report target port asymmetric access state

A value is provided identifying the current target port asymmetric access state (see 4.7, in this document). This also represents the state of the target port group since all target ports in the target port group must be at the same state.

Report target port groupings:

A list of available targets port groups is provided for the logical unit (see 4.7 in this document). The software driver may be required to assist in the identification of target ports and target port groupings and their relationship to other identification parameters.

3 CHANGES REQUESTED IN OTHER DOCUMENTS ~~OR SPC ANNEX~~

This proposal does not require changes to documents other than SPC.

~~3.1 SCC COMMANDS ALLOWED IN THE PRESENCE OF VARIOUS RESERVATIONS~~

~~It is recommended that, until appropriate changes can be made in the subsequent version of the SCC−2 document, an informative section be added to Annex B which includes the following table for the SCC commands allowed in the presence of various reservations. The section will be similar to the text and tables in SPC−2 sections B.1 and B.2. This section is not essential to this proposal but is included as an accommodation to the editor.~~

~~B.x SCC commands~~

~~This subclause should be placed into the model clause of the next version of the SCC standard when, and if, a new version of that standard is published. It should replace all the individual command descriptions of how reservations work.~~

~~Reservation restrictions are placed on commands as a result of access qualifiers associated with the type of reservation. The details of which commands are allowed under what types of reservations are described in table B.x. For the reservation restrictions placed on commands for the reserve/release management method see table B.x column[A]. For the reservation restrictions placed on commands for the persistent reservations management, see columns under [B] in table B.x.~~

~~In table B.x the following key words are used:~~

~~allowed:~~ ~~Commands issued by initiators not holding the reservation or by initiators not registered when a registrants only persistent reservation is present should complete normally.~~

~~conflict:~~ ~~Commands issued by initiators not holding the reservation or by initiators not registered when a registrants only persistent reservation is present shall not be performed and the device server shall terminate the command with a RESERVATION CONFLICT status.~~

~~Commands from initiators holding a reservation should complete normally. The behavior of commands from registered initiators when a registrants only persistent reservation is present is specified in table B.x.~~

~~A command that does not explicitly write the medium shall be checked for reservation conflicts~~

before the command enters the current task state for the first time. Once the command has entered the current task state, it shall not be terminated with a RESERVATION CONFLICT due to a subsequent reservation.

A command that explicitly writes the medium shall be checked for reservation conflicts before the device server modifies the medium or cache as a result of the command. Once the command has modified the medium, it shall not be terminated with a RESERVATION CONFLICT due to a subsequent reservation.

For each command, this standard, SPC−2, or a related command standard defines the conditions that result in RESESRVATION CONFLICT. Depending on the particular command standard the conditions are defined in that standard's device model clause or in the subclauses that define the specific commands. An annex in SPC−2 contains the RESERVATION CONFLICT information for some of the command sets.

Table B.x −− SCC commands that are allowed in the presence of various reservations

Command	Addressed LU is reserved by another initiator [A]	Addressed LU has this type of persistent reservation held by another initiator [B]				
		From any initiator		From registered initiator (RO all types)	From initiator not registered	
		Write Excl	Excl Access		Write Excl −RO	Excl Access −RO
MAINTENANCE (IN)	Allowed	Allowed	Allowed	Allowed	Allowed	Allowed
MAINTENANCE (OUT)	Conflict	Conflict	Conflict	Allowed	Conflict	Conflict
REDUNDANCY GROUP (IN)	Allowed	Allowed	Allowed	Allowed	Allowed	Allowed
REDUNDANCY GROUP (OUT)	Conflict	Conflict	Conflict	Allowed	Conflict	Conflict
SPARE (IN)	Allowed	Allowed	Allowed	Allowed	Allowed	Allowed
SPARE (OUT)	Conflict	Conflict	Conflict	Allowed	Conflict	Conflict
VOLUME SET (IN)	Allowed	Allowed	Allowed	Allowed	Allowed	Allowed
VOLUME SET (OUT)	Conflict	Conflict	Conflict	Allowed	Conflict	Conflict
Key: LU=Logical Unit, Excl=Exclusive, RO=Registrants Only						

4 PROPOSAL FOR CHANGES TO SPC−3 DOCUMENT

4.1 Overview of Proposed SPC Documentation Changes

The proposed text for inclusion into SPC−3 is emphasized by non−italic *text.*

4.2 Section 3.1 Definitions

3.1.x Target Port Group: A set of target ports that are in the same target port asymmetric access state at all times.

3.1.y Target Port Asymmetric Access State: The characteristic that defines the properties of a target port and the allowable command set for a logical unit when commands and task management functions are routed through the target port maintaining that state.

3.1.z Target Port Group Asymmetric Access State: The target port asymmetric access state common to the set of target ports in a target port group.

4.3 Section 5 Changes

The following two rows are to be added to table 10 in section 5.

Table 10 −− SPC commands that are allowed in the presence of various reservations

Command	Addressed LU is reserved by another initiator [A]	Addressed LU has this type of persistent reservation held by another initiator [B]				
		From any initiator		From registered initiator (RO all types)	From initiator not registered	
		Write Excl	Excl Access		Write Excl −RO	Excl Access −RO
REPORT TARGET PORT GROUPS	Allowed	Allowed	Allowed	Allowed	Allowed	Allowed
SET TARGET PORT GROUPS	Conflict	Conflict	Conflict	Allowed	Conflict	Conflict
Key: LU=Logical Unit, Excl=Exclusive, RO=Registrants Only						

The following text is provided immediately after section 5.5 to define the behavioral model for asymmetric logical units access, pushing clause 5.6 and subsequent functions one number higher.

5.6 Asymmetric logical unit access

5.6.1 Introduction

Logical units may be connected to the service delivery subsystem via multiple target ports (see SAM−2). Asymmetric logical unit access occurs when the access characteristics of one port may differ from those of another port. SCSI devices with target ports implemented in separate physical units may need to designate differing levels of access for the target ports associated with each logical unit. While commands and task management functions may be routed to a logical unit through any target port, the performance may not be optimal, and the accepted command set may be less complete than when the same commands and task management functions and are routed through a different target port. When a failure on the path to ~~an active~~ one target port is detected, the SCSI target device may perform automatic internal reconfiguration to make a logical unit accessible from a different set of target ports or may be instructed by the initiator to make a logical unit accessible from a different set of target ports.

A target port characteristic called target port asymmetric access state (see 5.6.4) defines properties ~~of target ports~~ of a target port and the allowable command set for a logical unit when commands and task management functions are routed through the target port maintaining that state.

A target port group is defined as a set of target ports that are in the same target port asymmetric access state at all times. A target port group asymmetric access state is defined as the target port asymmetric access state common to the set of target ports in a target port group. The grouping of target ports is vendor specific.

A logical unit may have commands and task management functions routed through multiple target port groups. Logical units support asymmetric logical unit access if different target port groups may be in different target port group asymmetric access states.

An example of asymmetric logical unit access is a SCSI Controller device with two separated controllers where all target ports on one controller are in the same asymmetric access state with respect to a logical unit and are members of the same target port group. Target ports on the other controller are members of another target port group. The behavior of each target port group may be different with respect to a logical unit, but all members of a single target port group are always in the same target port asymmetric access state with respect to a logical unit.

Figure XX: Target port group example

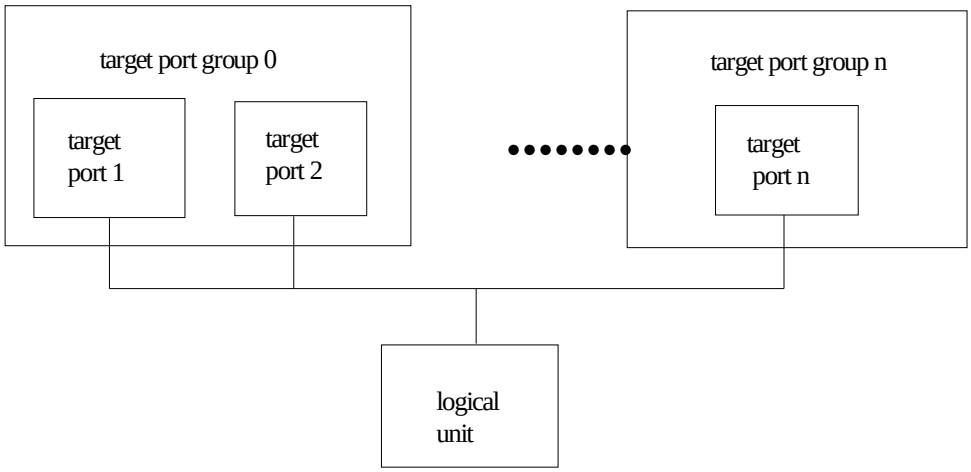

5.6.2 Explicit and implicit asymmetric logical unit access

Asymmetric logical unit access may be managed explicitly by an application client using the REPORT TARGET PORT GROUPS and SET TARGET PORT GROUPS commands.

Alternatively, asymmetric logical unit access may be managed implicitly by the SCSI target device based on the type of transactions being routed through each target port and the internal configuration capabilities of the target port group(s) through which the logical unit may be accessed. In general, the logical units attempt to maintain full performance across the target port groups that are busiest and that show the most reliable performance, allowing other target port groups to select a lower performance target port asymmetric access state.

If both explicit and implicit asymmetric logical unit access are implemented the precedence of one over the other is vendor specific.

5.6.3 Discovery of asymmetric logical unit access behavior

SCSI logical units with asymmetric logical unit access may be identified by inspecting the data returned by the INQUIRY command. The value in the asymmetric logical units access (ALUA) field (see 7.6.2) indicates whether or not the logical unit supports asymmetric logical unit access and if so whether implicit or explicit management is supported.

5.6.4 Target Port Asymmetric Access States

5.6.4.1 Overview

For all SCSI target devices that report in the INQUIRY data that they support asymmetric logical unit access, all of the target ports in a target port group shall be in the same target port asymmetric access state with respect to the ability to route information to a logical unit. The possible target port asymmetric access states are:

a) active/optimized;
b) active/non−optimized;
c) standby; and,
d) unavailable.

5.6.4.2 Active/Optimized

When commands and task management functions are being routed through a target port in the active/optimized target port asymmetric access state, the device server shall function as specified in the appropriate command set standards. All target ports within a target port group should be capable of immediately accessing the logical unit.

The SCSI target device shall participate in all task management functions as defined in SAM−2.

5.6.4.3 Active/Non−optimized

When commands and task management functions are being routed through a target port in the active/non−optimized target port asymmetric access state, the device server shall function as specified in the appropriate command set standards.

The processing of some commands, especially those involving data transfer or caching and task management function, may operate with lower performance than they would if the target port were in the active/optimized target port asymmetric access state.

The SCSI target device shall participate in all task management functions as defined in SAM−2.

5.6.4.4 Standby

When being accessed through a target port in the standby target port asymmetric access state, the device server shall support those of the following commands that it supports while in the active/optimized target port asymmetric access state:

a) INQUIRY;
b) LOG SELECT;
c) LOG SENSE;
d) MODE SELECT;
e) MODE SENSE;
f) REPORT LUNS (for logical unit number 0);
g) RECEIVE DIAGNOSTIC RESULTS;
h) SEND DIAGNOSTIC;
i) REPORT TARGET PORT GROUPS;
j) SET TARGET PORT GROUPS;
k) REQUEST SENSE;
l) PERSISTENT RESERVE IN;
m) PERSISTENT RESERVE OUT;
n) Echo buffer modes of READ BUFFER; and,
o) Echo buffer modes of WRITE BUFFER.

The device server may support other commands.

For those commands that are not supported, the device server shall return CHECK CONDITION status with the sense key set to NOT READY and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN STANDBY STATE.

The SCSI target device shall participate in all task management functions as defined in SAM−2.

5.6.4.5 Unavailable

When being accessed through a target port in the unavailable target port asymmetric access state, the device server shall accept only a limited set of commands. The unavailable target port asymmetric access state is intended for situations when the target port accessibility to a logical unit may be severely restricted due to, for example, a hardware error. Therefore it may not be possible to transition from this state to either the active/optimized, active/non−optimized or standby states. The unavailable target port asymmetric access state is also intended for minimizing any disruption when using the downloading microcode mode of the WRITE BUFFER command.

While in the unavailable target port asymmetric access state, the device server shall support those of the following commands that it supports while in the active/optimized state:

a) INQUIRY (the peripheral device qualifier shall be set to 001b (see 7.6.2));
b) REPORT LUNS (for logical unit number 0);
c) REPORT TARGET PORT GROUPS;
d) SET TARGET PORT GROUPS;
e) REQUEST SENSE;
f) Echo buffer modes of READ BUFFER;
g) Echo buffer modes of WRITE BUFFER; and,
h) Download microcode mode of WRITE BUFFER.

The device server may support other commands.

For those commands that are not supported, the device server shall return CHECK CONDITION status with the sense key set to NOT READY and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN UNAVAILABLE STATE.

The SCSI target device is not required to participate in all task management operations.

5.6.5 Transitions between target port asymmetric access states

The movement from one target port asymmetric access state to another is called a transition.

During a transition between target port asymmetric access states the device servers shall respond to a command in one of the following ways:

a) If during the transition the logical unit is inaccessible, then the transition is performed as a single indivisible event and the device server shall respond by either returning BUSY status or returning CHECK CONDITION status with the sense key set to NOT READY and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, TRANSITIONING BETWEEN ASYMMETRIC ACCESS STATES; or,

b) If during the transition the target ports in a target port group are able to access the requested logical unit, then the device server shall support those of the following commands that it supports while in the active/optimized asymmetric access state:

a) INQUIRY;
b) REPORT LUNS (for logical unit number 0);
c) REPORT TARGET PORT GROUPS;
d) SET TARGET PORT GROUPS;
e) REQUEST SENSE;
f) Echo Buffer modes of READ BUFFER; and,
g) Echo Buffer modes of WRITE BUFFER.

The device server may support other commands when those commands are routed though a target port that is transitioning between asymmetric access states.

For those commands that are not supported during a transition, the device server shall return CHECK CONDITION status with the sense key set to NOT READY and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, TRANSITIONING BETWEEN ASYMMETRIC ACCESS STATES.

The SCSI target device is not required to participate in all task management functions.

If the transition was explicit and it failed, then the device server shall return CHECK CONDITION status with the sense key set to HARDWARE ERROR and an additional sense code of SET TARGET PORT GROUPS COMMAND FAILED. The ~~failing~~ target port group that encountered the error should complete a transition to the unavailable target port asymmetric access state. If a target port group asymmetric access state change occurred as a result of the failed transition the device server shall establish a unit attention condition for all initiators other than the one ~~running~~ that sent the SET TARGET PORT GROUPS command with an additional sense code of ASYMMETRIC ACCESS STATE CHANGED ~~and an additional sense code qualifier of TBD~~.

If the transition was implicit and it failed, then the device server shall establish a unit attention condition for all initiators with an additional sense code of IMPLICIT ASYMMETRIC ACCESS STATE TRANSITION FAILED.

~~The target port asymmetric access state of a device server after a transition failure is vendor specific.~~ An implicit CLEAR TASK SET task management function may be performed following a transition failure.

Once a transition is completed, the new target port asymmetric access state may apply to some or all tasks entered into the task set before the completion of the transition. The new target port asymmetric access state shall apply to all tasks received by the device server after completion of a transition.

After an implicit target port asymmetric access state change, a device server shall establish a unit attention condition for all initiators with an additional sense code of ASYMMETRIC ACCESS STATE CHANGED ~~and an additional sense code qualifier of TBD~~.

After an explicit target port asymmetric access state change, a device server shall establish a unit attention condition with an additional sense code of ASYMMETRIC ACCESS STATE CHANGED ~~and an additional sense code qualifier of TBD~~ for all initiators other than the initiator that issued the SET TARGET GROUPS command.

5.6.6 Implicit asymmetric logical units access management

SCSI target devices with implicit asymmetric logical units access management are capable of setting the target port group asymmetric access state of each target port group using mechanisms other than the SET TARGET PORT GROUPS command.

All logical units that report in the INQUIRY data (see 7.6.2) that they support asymmetric logical units access and support implicit asymmetric logical unit access (ALUA field contains 01b or 11b) shall:

a) implement the INQUIRY command vital product data device identifier page identifier types 4h and 5h as described in 8.4.4. ~~This provides the necessary information to identify target ports reported by the REPORT TARGET PORT GROUPS command. The 05h identifier type defines the target port group~~; and,

b) support the REPORT TARGET PORT GROUPS command as described in 7.x. ~~This command provides a complete list of all the target port groups, the relative target port identifiers of the members of each target port group, and the current target port group asymmetric access state for each target port group.~~

5.6.7 Explicit asymmetric logical units access management

All logical units that report in the INQUIRY data (see 7.6.2) that they support explicit asymmetric logical units access (ALUA field contains 10b or 11b) shall:

a) implement the INQUIRY command vital product data device identifier page identifier types 4h and 5h as described in 8.4.4;

b) support the REPORT TARGET PORT GROUPS command as described in 7.x. ~~This command provides a complete list of all the target port groups, the relative target port identifiers of the members of each target port group, and the current target port group asymmetric access state for each target port group~~; and,

c) support the SET TARGET PORT GROUPS command as described in 7.y. ~~This command may be used to explicitly set the target port group asymmetric access state for one or more target port groups. The allowable combination of asymmetric access states for the target port groups is vendor specific. If the SET TARGET PORT GROUPS attempts to establish an invalid combination of target port asymmetric access states, the device server shall return CHECK CONDITION status with the~~

~~sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN PARAMETER LIST.~~

5.6.8 Behavior after power cycling or hard resets

For all SCSI target devices that report in the INQUIRY data (see 7.6.2) that they support only explicit asymmetric logical unit access (ALUA field contains 10b), the target port shall preserve the target port asymmetric access state across any reset and across any power cycle.

4.4 Section 7.6.2, Standard INQUIRY data

Two new bits will be placed in byte 5 of INQUIRY data to indicate that the logical unit has asymmetric targets and supports asymmetric access behavior. The text to be included in this section is defined below:

The content of the asymmetric logical unit access (ALUA) field is described in table ZZ0.

Table ZZ0 ALUA field contents

VALUE	*Description*
00b	The SCSI target device does not support asymmetric logical unit access or supports a form of asymmetric access that is vendor specific. Neither the REPORT TARGET GROUPS nor the SET TARGET GROUPS commands is supported.
01b	Only implicit asymmetric logical unit access (see 5.6.6) is supported. The SCSI target device is capable of changing target port asymmetric access states without a SET TARGET PORT GROUPS command. The REPORT TARGET PORT GROUPS command is supported and the SET TARGET PORT GROUPS command is not supported.
10b	Only explicit asymmetric logical unit access (see 5.6.7) is supported. The SCSI target device only changes target port asymmetric access states as requested with the SET TARGET PORT GROUPS command. Both the REPORT TARGET PORT GROUPS command and the SET TARGET PORT GROUPS command are supported.
11b	Both explicit and implicit asymmetric logical unit access (see 5.6.6 and 5.6.7) are supported. Both the REPORT TARGET PORT GROUPS command and the SET TARGET PORT GROUPS commands are supported.

4.5 Section 8.4.4 , Device Identification Page

The changes in this section are based on another, approved, proposal T10/01−065r2.

SPC presently defines a 4−byte "relative target port identifier" that, when included, with the association value of 1, indicates the target port (relative to some arbitrary internal ordering) through that the command is passed. The device identification page allows simultaneous presentation of identifiers, including the target port WWN, the logical units WWN, the relative target port identifier, and any other similar information.

The relative target identifier is used to locate the actual target ports that will be reported and controlled using the REPORT TARGET PORT GROUPS and SET TARGET PORT GROUPS commands.

An additional identifier entry is included for target port group asymmetric access discovery. The tables are updated and added as below.

~~***Table 181 − Identifier type***~~

~~***Value***~~	~~***Description***~~
~~5h~~	~~If the ASSOCIATION value is 1h, the IDENTIFIER value contains the target port group as defined in table ZZ1. The CODE SET field shall be set to 1h and the IDENTIFIER LENGTH field shall be set to 4. If the ASSOCIATION value is not 1h, use of this identifier type is reserved.~~
~~6h−Fh~~	~~reserved~~

Table xyz−Identifier type

Value	Identifier type
0h	Vendor specified
1h	T10 vendor ID
2h	EUI−64
3h	NAA
4h	Relative target port
5h	Target port group
6h−Fh	reserved

8.4.4.8 Target port group identifier format

If the identifier type is target port group (5h) and the ASSOCIATION value is 1h, the IDENTIFIER field identifies the target port group to which the target port is a member (see 5.6). Table yy defines the IDENTIFIER format. In this case, the CODE SET field shall be set to 1h and the IDENTIFIER LENGTH field shall be set to 4. If the ASSOCIATION value is not 1h, use of this identifier type is reserved.

Table yy−Target Port Group IDENTIFIER format

Bit **Byte**	7	6	5	4	3	2	1	0
0	reserved							
1								
2	TARGET PORT GROUP							
3								

Table ZZ1−Target Port Group Identifiers

Bit / Byte	7	6	5	4	3	2	1	0
0	reserved							
1	(MSB) TARGET PORT GROUP IDENTIFIER							
2	(LSB)							
3	reserved							

See 7.x for the definitions of the TARGET PORT GROUP IDENTIFIER field.

4.6 Section 7.23.6 , new ASC/ASCQ

The additional sense codes and qualifiers are defined in section 7.23.6 of the SPC−2 document. The ASC/ASCQ codes would include:

LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN STANDBY STATE (proposed value = 04/0B)
LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN UNAVAILABLE STATE (proposed value = 04/0C)
LOGICAL UNIT NOT ACCESSIBLE, TRANSITIONING BETWEEN ASYMMETRIC ACCESS STATES (proposed value = 04/0A)
SET TARGET PORT GROUPS COMMAND FAILED (proposed value = 67h/0Ah)

The following two ASC/ASCQ codes are for unit attention conditions.

IMPLICIT ASYMMETRIC ACCESS STATE TRANSITION FAILED (proposed value = ~~67h/0Bh~~ 2Ah/07h)

ASYMMETRIC ACCESS STATE CHANGED (proposed value = 2Ah/06h)

4.7 Target port group management

The committee has elected to place the management functions in the MAINTENANCE IN command (operation code A3) and the MAINTENANCE OUT command (operation code A4), primarily defined in SCC−2. Service action 0A has been assigned to Asymmetric Set Target Port Group and Asymmetric Report Target Port Group functions.

The text, placed at the same level as a standard command in section 7, will read:

7.x REPORT TARGET PORT GROUPS

The REPORT TARGET PORT GROUPS command (see table XX1) requests that the device server send target port group information to the application client. This command shall be supported for all SCSI target devices that report in the INQUIRY data (see 7.6.2) that they support asymmetric logical unit access (non zero ALUA field).

The REPORT TARGET PORT GROUPS command is a service action of the MAINTENANCE IN command. Additional MAINTENANCE IN service actions are defined in SCC−2 and in this standard. The MAINTENANCE IN service actions defined only in SCC−2 apply only to SCSI devices that return a device type of 0Ch or the sccs bit equal to one in their standard INQUIRY data.

Table XX1 − REPORT TARGET PORT GROUPS command

Byte \ Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (A3h)							
1	reserved			SERVICE ACTION (0Ah)				
2	reserved							
3								
4								
5								
6	(MSB)							
7	ALLOCATION LENGTH							
8								
9								(LSB)
10	reserved							
11	CONTROL							

The ALLOCATION LENGTH field indicates how much space has been allocated for the returned parameter data. If the length is not sufficient to contain all the parameter data, the first portion of the data shall be returned. This shall not be considered an error. The actual length of the parameter data is available in the REPORT FIELD LENGTH field in the parameter data. If the remainder of the parameter data is required, the application client should send a new REPORT TARGET PORT GROUPS command with an ALLOCATION LENGTH field large enough to contain all the data.

The format for the parameter data provided in response to a REPORT TARGET PORT GROUPS command is shown in table XX2.

Table XX2 − REPORT TARGET PORT GROUPS parameter data

Byte \ Bit	7	6	5	4	3	2	1	0
0	(MSB)							
1	REPORT FIELD LENGTH (n − 3)							
2								
3								(LSB)
4	REPORT TARGET PORT GROUP descriptor(s) (see Table XX3)							
n								

Returning REPORT TARGET PORT GROUPS parameter data may require the enabling of a nonvolatile memory. If the nonvolatile memory is not ready, the device server shall return CHECK CONDITION status, rather than wait for the SCSI target device to become ready. The sense key shall be set to NOT READY and the additional sense data shall be set as described in the TEST UNIT READY command (see Table 24).

The REPORT FIELD LENGTH field specifies the length in bytes of the list of target port groups. If the ALLOCATION LENGTH field in the CDB is too small to transfer all of the ~~identifier~~ descriptor, the length shall not be adjusted to reflect the truncation.

The format of the report target port group descriptor is defined in Table XX3. There shall be one REPORT TARGET PORT GROUP descriptor for each target port group.

Table XX3 − REPORT TARGET PORT GROUP descriptor

Bit / Byte	7	6	5	4	3	2	1	0
0	reserved				**ASYMMETRIC ACCESS STATE**			
1	reserved							
2	~~(MSB)~~ TARGET PORT GROUP ~~IDENTIFIER~~							
3								~~(LSB)~~
4	reserved				~~ASYMMETRIC ACCESS STATE~~ **reserved**			
5	STATUS CODE							
6	vendor unique							
7	TARGET PORT COUNT (x)							
8	(MSB)							
9								
10	RELATIVE TARGET PORT ~~IDENTIFIER~~ (first target port of target port group)							
11								(LSB)
...								
4x+4	(MSB)							
4x+5								
4x+6	RELATIVE TARGET PORT ~~IDENTIFIER~~ (last target port of first target port group)							
4x+7								(LSB)

The ASYMMETRIC ACCESS STATE field contains the target port group asymmetric access state (see ZZ2). Active/optimized, active/non−optimized, standby and unavailable indicate the current target port group asymmetric access state defined by 5.6.4.

The TARGET PORT GROUP ~~IDENTIFIER~~ field contains an identification of the target port group for which information shall be sent to the application client.

Table ZZ2 Asymmetric access state

Code	State
0h	Active/optimized
1h	Active/non−optimized
2h	Standby
3h	Unavailable
4h−Eh	Reserved
Fh	Transitioning between states

~~Active/optimized, active/non−optimized, standby and unavailable indicate the current target port group asymmetric access state defined by 5.6.4.~~

The STATUS CODE field indicates why a target port group may be in a specific target port groups asymmetric access state. It provides a mechanism to indicate error conditions (see table ZZ3).

Table ZZ3 Status Code field

Code	State
00h	No status available
01h	target port group asymmetric access state changed by SET TARGET PORT GROUPS command
02h	target port group asymmetric access state changed by implicit asymmetrical logical unit access behavior
03h−FFh	reserved

The TARGET PORT COUNT field specifies the number of target ports that are in that target port group. Every target port group shall contain at least one target port. The RELATIVE TARGET PORT ~~IDENTIFIER~~ field contains an identification of a target port. This is the same value returned by identifier type 4h in the Device Identification VPD page (see 8.4.4). A target port shall be listed in ~~one and only~~ exactly one target port group.

~~The STATUS CODE field indicates why a target port group may be in a specific target port groups asymmetric access state. It provides a mechanism to indicate error conditions (see table ZZ3).~~

~~**Table ZZ3 Status Code field**~~

~~Code~~	~~State~~
~~00h~~	~~No status available~~
~~01h~~	~~target port group asymmetric access state changed by SET TARGET PORT GROUPS command~~
~~02h~~	~~target port group asymmetric access state changed by implicit asymmetrical logical unit access behavior~~
~~03h−FFh~~	~~reserved~~

7.y SET TARGET PORT GROUPS

The SET TARGET PORT GROUPS command (see table YY1) requests the device server to set the asymmetric access state ~~(see 5.6 for details regarding the transition between target port group asymmetric access states)~~ of all of the target ports in the specified target port group(s). See 5.6 for details regarding the transition between target port group asymmetric access states. This command is mandatory for all SCSI target devices that report in the INQUIRY data (see 7.6.2) that they support explicit asymmetric logical units access (ALUA field contains either 10b or 11b).

The SET TARGET PORT GROUPS command is a service action of the MAINTENANCE OUT command. Additional MAINTENANCE OUT service actions are defined in SCC−2 and in this standard. The MAINTENANCE OUT service actions defined only in SCC−2 apply only to SCSI devices that return a device type of 0Ch or the sccs bit equal to one in their standard INQUIRY data.

Table YY1 − SET TARGET PORT GROUPS command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (A4h)							
1	reserved			SERVICE ACTION (0Ah)				
2	reserved							
3								
4								
5								
6	(MSB)							
7	PARAMETER LIST LENGTH							
8								
9								(LSB)
10	reserved							
11	CONTROL							

The PARAMETER LIST LENGTH field specifies the length in bytes of the target port group management parameters that shall be transferred from the application client to the device server. A parameter list length of zero indicates that no data shall be transferred, and that no change shall be made in the asymmetric access state of any target port groups. If the parameter list length violates the vendor specific length requirements, then the device server shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.

The allowable values to which target port asymmetric access states may be set is vendor specific. Target port groups that are not specified in a parameter list may change asymmetric access states as a result of the SET TARGET PORT GROUPS command. This shall not be considered an implicit target port group asymmetric access state change.

If the SET TARGET PORT GROUPS attempts to establish an invalid combination of target port asymmetric access states, the device server shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN PARAMETER LIST.

The SET TARGET PORT GROUPS parameter data format is shown in Table YY2.

Table YY2 − SET TARGET PORT GROUPS parameter data

Bit Byte	7	6	5	4	3	2	1	0
0	~~(MSB)~~							
1	reserved							
2								
3								~~(LSB)~~
4	SET TARGET PORT GROUP descriptor(s)							
n								

The format of the SET TARGET PORT GROUP descriptor is defined in Table YY3.

Table YY3 − SET TARGET PORT GROUP descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	reserved				ASYMMETRIC ACCESS STATE			
1	reserved							
2	TARGET PORT GROUP							
3								

~~**Table YY3 − SET TARGET PORT GROUP descriptor**~~

~~Bit Byte~~	~~7~~	~~6~~	~~5~~	~~4~~	~~3~~	~~2~~	~~1~~	~~0~~
~~0~~	~~reserved~~							
~~1~~	~~(MSB)~~ ~~TARGET PORT GROUP IDENTIFIER~~							
~~2~~								~~(LSB)~~
~~3~~	~~reserved~~				~~ASYMMETRIC ACCESS STATE~~			

The ASYMMETRIC ACCESS STATE field~~. This field~~ specifies the asymmetric access state to which all of the target ports in the target port group shall transition (see table ZZ4). See 5.6.4 for a description of target port group asymmetric access states.

The TARGET PORT GROUP ~~IDENTIFIER~~ field specifies a target port group for which the asymmetric access state shall be changed.

TABLE ZZ4 ASYMMETRIC ACCESS STATE

Code	State
0h	Active/Optimized
1h	Active/Non−optimized
2h	Standby
3h	Unavailable
4h−Eh	Reserved
Fh	Illegal Request

~~See 5.6.4 for a description of target port group asymmetric access states.~~

If the SET TARGET PORT GROUPS command has been performed the completion of the command depends upon which of the following conditions apply:

a) if ~~the target port asymmetric access state is in transition, and~~ the transition is treated as a single indivisible event (see 5.6.5), then the SET TARGET PORT GROUPS command shall not complete until the transition ~~into~~ to the requested state has completed; or,

b) if ~~the target port asymmetric access state is in transition, and~~ the transition is not treated as a single indivisible event, i.e., the device server supports other commands (see 5.6.5) when those commands are routed though a target port that is transitioning between asymmetric access states, then the SET TARGET PORT GROUPS command may complete before the transition into the requested state has completed.

If the SET TARGET PORT GROUPS command cannot be performed the completion of the command depends upon which of the following conditions apply:

a) if ~~the information needed to process the command is not available~~ the processing of a SET TARGET PORT GROUPS command ~~may~~ requires the enabling of a nonvolatile memory and ~~If~~ the nonvolatile memory is not ready, then the device server shall return CHECK CONDITION status, rather than wait for the logical unit to become ready. The sense key shall be set to NOT READY and the additional sense data shall be set as described in the TEST UNIT READY command (see 7.28 Table 124) ~~This information should allow the application client to determine the action required to cause the device server to become ready~~; or,

b) if a failure occurred before the transition was completed then the device server shall return CHECK CONDITION status. The sense key shall be set to HARDWARE ERROR and the additional sense data shall be set to LOGICAL UNIT NOT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND FAILED.

The target port group asymmetric access state change behavior if two SET TARGET PORT GROUPS commands are performed concurrently is vendor specific. A target should not process multiple SET TARGET PORT GROUPS concurrently.