



T10/00-232 r7

To: John Lohmeyer, chairperson, T10
From: Ken Moe Email: kenneth.moe@sun.com
Date: March 21, 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 February 26, 2001 conference call and from the Mar 7, 2001 CAP meeting. The most significant changes include the following:

- Merged options provided by "changing" asymmetric access state into transition behavior and dropped "changing" state.*
- Eliminated ordering restrictions on target port parameter list returned by REPORT TARGET PORT GROUPS command.*
- SET TARGET PORT GROUPS command modified to allow changing the asymmetric access state for a subset of the existing target port groups.*

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.

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 5.6.3

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 5.6.3.

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

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 3 be added to Annex B with a table for the SCC–2 commands allowed in the presence of various reservations. The table would be similar to the text and tables in SPC–2 sections B.1 and B.2. Included in this table would be the SET TARGET PORT GROUPS and REPORT TARGET PORT GROUPS commands. the row entries for the SET TARGET PORT GROUPS command will all be "Conflict" with the exception of the from registered initiator column that will be "Allowed". The row entries for the REPORT TARGET PORT GROUPS command will all be "Allowed".

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.x 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 those commands are routed through the target port maintaining that state.

3.1.x 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 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. ~~afforded via one port may differ from those afforded by another port.~~ 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 access to the logical unit may be possible through any target port, the performance may not be optimal, and the accepted command set may be less complete than when those commands are routed through a different target port. When a failure on the path to an active 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 and the allowable command set for a logical unit when accessed those commands 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 ports is vendor specific.

A logical unit may have commands and messages routed through be serviced by multiple target port groups. Logical units may receive information routed from 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 the logical unit, but all members of a single target port group are always in the same target port asymmetric access state with respect to accessing 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 performed routed through each target port and the internal configuration capabilities of the set of target ports through which the logical unit may be accessed. In general, the logical units attempt to maintain full performance across the links target port groups that are busiest and that show the most reliable performance, allowing links connections to other target port groups to select a assume one of the 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 behavior 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 behavior, 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 access route information to a logical unit. The possible target port asymmetric access states are:

- a) active/optimized;
- b) active/non-optimized;
- c) standby;
- d) unavailable; and,
- e) changing.

~~The moving from one target port asymmetric access state to another is called a transition. Any transition shall be performed as a single indivisible event. The time at which a transition is completed with respect to other functions being performed by the device server is vendor specific. Once a transition is completed, the new target port asymmetric access state may apply to some or all tasks queued 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.~~

~~Multiple SET TARGET PORT GROUPS commands may be queued concurrently. The order of processing of such commands is defined by the tagged queuing restrictions, if any, but each shall be processed as a single indivisible command without any interleaving of actions that may be required by other SET TARGET PORT GROUPS commands.~~

~~The device server shall respond to a command in one of the following ways;~~

- ~~a) the device server processes a command using the target port asymmetric access state from which the target port is changing;~~
- ~~b) the device server processes a command using the target port asymmetric access state to which the target port is changing;~~
- ~~c) if the target port asymmetric access state is transitioning, and the logical unit is inaccessible, the device server returns BUSY status;~~
- ~~d) if the target port is in the changing target port asymmetric access state, and the logical unit is inaccessible, the device server returns CHECK CONDITION status with the sense key set to NOT READY and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN CHANGING STATE;~~
- ~~e) if the explicit transition from one state to another failed, then the device server returns CHECK CONDITION status with the sense key set to HARDWARE ERROR and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND FAILED; or~~
- ~~f) if the implicit transition from one or state to another failed, then the device server returns CHECK CONDITION status with sense key set to UNIT ATTENTION and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, IMPLICIT ASYMMETRIC LOGICAL UNIT ACCESS TRANSITION FAILED.~~

~~The state of a device server after a transition failure is vendor specific. An implicit CLEAR TASK SET task management function may be performed. Queued tasks may be aborted if a transition failure occurs.~~

5.6.4.2 Active/Optimized

~~While in~~ **When being accessed 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 operations as defined in SAM-2.

5.6.4.3 Active/Non-optimized

~~While in~~ **When being accessed 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, 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 operations as defined in SAM-2.

5.6.4.4 Standby

~~While in~~ **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. ~~when accessed through a target port in the standby state.~~

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 operations as defined in SAM–2.

5.6.4.5 Unavailable

~~While in~~ **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 the commands specified in the appropriate command set standards. The unavailable **target port asymmetric access** 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 **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. ~~when accessed through a target port in the unavailable state.~~

For those commands that do not operate, 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. ~~when accessed through a target port in the unavailable state.~~

5.6.4.6 Changing

~~A device server may use the changing target port asymmetric access state to break a transition between two of the other target port asymmetric access states into two transitions, the first from the previous target port asymmetric access state to the changing state and the second from the changing state to the final target port asymmetric access state. In the changing state, the target ports in a target port group may not be able to access the requested logical unit.~~

~~While in the changing 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;~~
- ~~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 that the logical unit supports when accessed through a target port in the changing state.~~

~~For those commands that do not operate, 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 CHANGING STATE.~~

~~The SCSI target device is not required to participate in all task management operations when accessed through a target port in the changing state.~~

5.6.5 Transition 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 two 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 through a target port that is transitioning between asymmetric access states.

For those commands that do not operate 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 operations.

If the transition fails, the device server shall respond to a command in one of these ways:

- a) if the transition was explicit and it failed, then the device server returns CHECK CONDITION status with the sense key set to HARDWARE ERROR and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND FAILED; or,
- b) if the transition was implicit and it failed, then the device server returns CHECK CONDITION status with sense key set to UNIT ATTENTION and an additional sense code of LOGICAL UNIT NOT ACCESSIBLE, 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.

Once a transition is completed, the new target port asymmetric access state may apply to some or all tasks queued 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.

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 the present asymmetric access state of the target ports in the target port group. The status of the target port group may change at any time as required by implicit asymmetric logical units access behavior; 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 logical units access behavior state for each target port group.

5.6.7 Explicit asymmetric logical units access ~~behavior~~ 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 all 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 ~~or 11b~~), the target port shall preserve the target port asymmetric access state across any reset and across any power off period.

4.4 Section 7.6.2, *Standard INQUIRY data*

The text to be included in this section is defined below:

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

- a) ~~an ALUA field content of 00b indicates that the SCSI target device does not support asymmetric logical unit access or supports a form of asymmetric access behavior that is not specified in the SCSI family of standards.~~
- b) ~~an ALUA field content of 10b indicates that the SCSI target device only supports explicit asymmetric logical unit access as described in 5.6.6.~~
- c) ~~an ALUA field content of 01b indicates that the SCSI target device only supports implicit asymmetric logical unit access as described in 5.6.5.~~
- d) ~~an ALUA field content of 11b indicates that the SCSI target device supports both explicit and implicit asymmetric logical unit access behaviors.~~

Table ZZ0 ALUA field contents

<i>VALUE</i>	<i>Description</i>
00b	The SCSI target device does not support asymmetric logical unit access or supports a form of asymmetric access that is not specified in the SCSI family of standards. 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

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 state discovery and for asymmetric access target port group discovery. The tables are updated and added as below.

Table 181 – Identifier type

<i>Value</i>	<i>Description</i>
5h	If the ASSOCIATION value is 1h, the IDENTIFIER value contains the target port group and target port group asymmetric access state as defined in table ZZ1. For 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.
6h–Fh	reserved

Table ZZ1–Target Port Group and Target Port Group Asymmetric Access State Identifiers

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

See 7.x for the definitions of the TARGET PORT GROUP IDENTIFIER field and the ASYMMETRIC ACCESS STATE field.

4.6 Section 7.23.6 , new ASC/ASCQ

Proper error indications are defined in this section 7.23.6 of the SPC–2 document. The error indications 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, ~~TARGET PORT IN CHANGING STATE~~

TRANSITIONING BETWEEN ASYMMETRIC ACCESS STATES (proposed value = 04/0A)

LOGICAL UNIT NOT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND

FAILED (proposed value = 67h/0Ah)

LOGICAL UNIT NOT ACCESSIBLE, IMPLICIT ASYMMETRIC ~~LOGICAL UNIT~~

*ACCESS **STATE** TRANSITION FAILED* (proposed value = **67h/0Bh**)

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). ~~This service action shall be rejected by all other SCSI target devices that support the MAINTENANCE IN command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense data shall be set to INVALID FIELD IN CDB.~~

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. When a SCSI device returns a device type of 0Ch or the sees bit equal to one in its standard INQUIRY data, the implementation requirements for the SCC-2 MAINTENANCE IN service action are as specified in SCC-2. Otherwise the MAINTENANCE IN service action definitions and implementation requirements stated in this standard shall apply.

Table XX1 – REPORT TARGET PORT GROUPS command

Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (A3h)							
1	reserved			SERVICE ACTION (0Ah)				
2	reserved							
3	reserved							
4	reserved restricted							
5								
6	(MSB) ALLOCATION LENGTH (LSB)							
7								
8								
9								
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 with the MAINTENANCE IN service action is shown in table XX2.

Table XX2 – REPORT TARGET PORT GROUPS parameters

Bit	7	6	5	4	3	2	1	0
0	(MSB)							
1	REPORT FIELD LENGTH (n – 3)							
2								
3								
4	REPORT TARGET PORT GROUP Descriptor(s)							
n	(see Table XX3)							

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

Table XX3 – REPORT TARGET PORT GROUP parameters Descriptor

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

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, the length shall not be adjusted to reflect the truncation.

The TARGET PORT GROUP IDENTIFIER field specifies a two-byte identification of the target port group. A particular logical unit may be served by up to 65536 target port groups.

The ASYMMETRIC ACCESS STATE field is an encoded value (see ZZ2). This field specifies the target port group asymmetric access state in the SCSI target device that reports in the INQUIRY DATA that asymmetric logical unit access behavior is supported.

Table ZZ2 Asymmetric access state

<i>Codes</i>	<i>States</i>
0h	Active/optimized
1h	Active/non-optimized
2h	Standby
3h	Unavailable
4h	Changing Transitioning between states
5h-Fh	Reserved

Active/optimized, active/non-optimized, standby and unavailable and changing define the current target port group asymmetric access state defined by 5.6.4.

The TARGET PORT COUNT field specifies the number of target ports that have access to the logical unit **are** in that target port group. ~~Not all target ports may be connected or available, but all are always listed in the list and counted by this field. The minimum value is one representing the case when there is only one target port in the target port group.~~ Every target port group shall contain at least one target port.

The RELATIVE TARGET PORT IDENTIFIER field specifies a four-byte identification of a target port. ~~This is the same value provided for relative target port identifier in the vital product page 83h, the Device Identification Page.~~ This is the same value returned by identifier type 4h in the Device Identification VPD page (see x.x.x). Every target port shall be listed in one and only target port group.

~~The execution of a Returning REPORT TARGET PORT GROUPS parameter data command may require the enabling of a nonvolatile memory within the logical unit. 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 7.28 see Table 24). This information allows the application client to determine the action required to cause the device server to become ready.~~

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

Table ZZ3 Status Code field

<i>Code</i>	<i>State</i>
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 VENDOR UNIQUE STATUS field contains vendor specific details that accompany the status code.~~

7.y SET TARGET PORT GROUPS

The SET TARGET PORT GROUPS command (see table YY1) requests that the device server should set the **asymmetric access** state of all of the target ports in **the specified** target port group(s). 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). ~~This service action shall be rejected by all other devices that support the MAINTENANCE OUT command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense data shall be set to INVALID FIELD IN CDB.~~

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

Table YY2 – SET TARGET PORT GROUPS Parameter Data

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

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							
1	(MSB) TARGET PORT GROUP IDENTIFIER							
2	(LSB)							
3					ASYMMETRIC ACCESS STATE			

The SET FIELD LENGTH field specifies the length in bytes of the list of target port groups.

The TARGET PORT GROUP IDENTIFIER field specifies a two-byte identification of the target port group. A particular logical unit may be served by up to 65536 target port groups.

The TARGET PORT GROUP ASYMMETRIC ACCESS STATE field is an encoded value (see table ZZ4). This field specifies the asymmetric access state to which all of the target ports in the target port group shall transition.

TABLE ZZ4 TARGET PORT GROUP ASYMMETRIC ACCESS STATE

<i>Values</i>	<i>States</i>
0h	Active/Optimized
1h	Active/Non-optimized
2h	Standby
3h	Unavailable
4h–Fh	Reserved

Active/optimized, active/non-optimized, standby, and unavailable define the target port group asymmetric access state as defined by 5.6.4.

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

The completion of the SET TARGET PORT GROUPS command depends upon which of the following three conditions apply:

- ~~a. The command has been executed and the target port asymmetric access state change is in transition:~~

~~The SET TARGET PORT GROUPS command should complete immediately if the changing state is supported. If the changing state is not supported The SET TARGET PORT GROUPS command shall not complete until the transition has completed~~

- ~~b. The information needed to execute this command is not available:~~

~~The execution of a SET TARGET PORT GROUPS command may require the enabling of a nonvolatile memory within the logical unit. If the nonvolatile memory is not ready, 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). This information should allow the application client to determine the action required to cause the device server to become ready.~~

- ~~e. A failure occurred before the command was executed~~

~~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.~~

a) if the command has been performed, 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 the requested state has completed;

b) if the command has been performed, the target port asymmetric access state is in transition, and 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;

c) if the information needed to process the command is not available the processing of a SET TARGET PORT GROUPS command may require the enabling of a nonvolatile memory within the logical unit. 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,

d) 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.

Multiple SET TARGET PORT GROUPS commands may be queued concurrently. The order of processing of such commands is defined by the tagged queuing restrictions, if any, but each shall be processed as a single indivisible command without any interleaving of actions that may be required by other SET TARGET PORT GROUPS commands.