



T10/00-232 r5

To: John Lohmeyer, chairperson, T10
From: Ken Moe Email: kenneth.moe@sun.com
Date: February 06, 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 by Mr. Ralph Weber during the January 15-19, 2001 T10 Plenary in Orlando, Florida. The most significant changes include the following:

- Changed "transition" state to "changing" state to distinguish the state from the process of transitioning from one state to another. The most significant impact of this change was to the Overview section 5.6.3.1 and to the "changing" state definition in section 5.6.3.6.*
- An explanation as to the purpose of the "standby" and "unavailable" states was added to 5.6.3.*
- Since the two bits returned as standard INQUIRY data to specify the Asymmetric access behavior have evolved to an encoded value, they were redefined as a field (reference section 7.6.2).*
- A number of changes were made to more closely conform with the accepted terminology and style conventions.*

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 which target ports are fully accessible and which 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 Behavior 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 behavior 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 behavior 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 behavior 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 Behavior State

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

Active/Optimized :

While in the active/non-optimized state the target port group should be capable of immediately 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 include 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*
- *Transferring sense data*
- *Manage or remove logical unit or element reservations*
- *Testing Service delivery subsystem*

The commands which shall operate normally in the standby state are listed in section 5.6.3

Unavailable:

While in the unavailable target port behavior state, the device server shall accept only a severely restricted set of the 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 which shall operate normally in the unavailable state are listed in section 5.6.3.

Changing:

While in the changing state the target port is in the process of changing from one state to another. The time at which the target port behavior state is changed with respect to other tasks being managed by the device server is vendor specific.

All target ports in a target port group are capable of performing a limited set of commands. The commands which shall operate normally in the changing 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 4.5 in this document)

This function can be done explicitly or automatically. See 4.2 in this document to determine whether the SCSI target device requires explicit or implicit change. [Alternatively, implicit or explicit behavior can be specified by adding a few bits to item 4.5 in this document.]

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 4.2, 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. This can be optionally associated with a target port world-wide name through either the SCSI command set or the Fibre Channel command set. (See 4.3, in this document)

Report target port behavior state

A value is provided identifying the current target port behavior state (See 4.3, 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. 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 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 which 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 following text is provided immediately after section 5.5 to define the behavioral model for asymmetric logical units access behavior, pushing clause 5.6 and subsequent functions one number higher. The proposed text for inclusion into SPC-3 is emphasized by non-italic text.

5.6 Asymmetric logical unit access behavior

5.6.1 Introduction

Logical units may be connected to the service delivery subsystem via multiple target ports (see SAM-2). A target port behavior state (see 5.6.3) 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 behavior state at all times. A target port group behavior state is defined as the target port behavior state common to the set of target ports in the target port group.

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 behavior states.

Note that logical units which support asymmetric logical unit access are not prevented from having multiple target port groups that are in the same target port behavior state at a given moment in time, and thereby exhibit symmetric behavior at that moment in time.

An example of such a device is a SCSI Controller device with two separated controllers where all target ports on one controller have the same behavior 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 at the same target port behavior state with respect to accessibility to the logical unit.

This clause defines the target port behavior states and provides a set of SCSI commands which may be used to manage them.

5.6.1.1 Symmetric and asymmetric logical unit access behavior

Symmetric access to logical units is often desirable, since it provides for very rapid access to a logical unit through alternate target ports when the connection to one target port is lost. Symmetrical access is characteristic of many simple SCSI devices such as disk drives and disk enclosures providing direct access to a set of disk drives.

Asymmetric access to logical units is useful for other reasons. 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 accessed 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 different set of target ports active or may be instructed by the initiator to make a different set of target ports active.

5.6.1.2 Explicit and implicit asymmetric logical unit access behavior

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

Alternatively, asymmetric logical unit access behavior may be managed implicitly by the SCSI target device based on the type of transactions being performed through each target port and the internal configuration capabilities of the set of target ports through which the logical unit can be accessed. In general, the logical units attempt to maintain full performance across the links that are busiest and which show the most reliable performance, allowing links to other target port groups to assume one of the lower performance states.

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

5.6.2 Discovery of asymmetric target behavior

SCSI logical units with asymmetric logical unit access behavior may be identified by inspecting the data returned by the performing an INQUIRY command to the logical unit. The values of the asymmetric logical units access behavior (ALUAB) field bit and implicit asymmetric logical unit access behavior bit (IAAB) (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. SCSI devices that do not support asymmetric logical unit access as described in this standard may implement similar vendor specific services. If the SCSI target device does not indicate support for asymmetric logical unit access behavior, the SCSI target device asymmetric or symmetric behavior is unspecified. SCSI devices that do not support asymmetric logical unit access as described in this standard may implement similar vendor specific services. When a SCSI target device supports asymmetric logical unit access behavior, the value of the IAAB bit ALUAB field (see 7.6.2) indicates whether or not the SCSI target device supports implicit management of the asymmetric logical units access behavior.

5.6.3 Target port behavior states

5.6.3.1 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 behavior state with respect to the ability to access a logical unit. A description of the possible target port behavior states follows:

- a) active/optimized
- b) active/non-optimized
- c) Inactive standby
- d) unavailable
- e) Transition changing

The moving from one target port behavior 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 behavior state may apply to some or all tasks queued before the completion of the transition. The new target port behavior state shall apply to all tasks received by the device server after successful completion of a transition. Multiple SET TARGET PORT GROUPS commands may be queued at the same time. 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 behavior state from which the target port is changing, or
- b) the device server processes a command using the target port behavior state to which the target port is changing, or
- c) if the target port behavior state is transitioning, and the logical unit is inaccessible, the device server returns BUSY status, or
- d) if the target port is in the changing target port behavior 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, or

- e) if the 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

5.6.3.2 Active/Optimized

~~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.~~

While in the active/optimized target port behavior 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.

5.6.3.3 Active/Non-optimized:

~~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.~~

While in the active/non-optimized target port behavior state, the device server shall function as specified in the appropriate command set standards.

The processing execution of certain 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 state.

The SCSI target device shall participate in all task management operations as defined in SAM-2.

5.6.3.4 Inactive Standby

While in the standby target port behavior state, the device server shall accept at least a more restricted set of the commands than are accepted in the active/optimized or active/non-optimized states.

~~All target ports in a target port group are capable of performing a limited set of commands.~~

~~The device server shall support all of the following commands that the logical unit supports when accessed through a target port in an active/optimized or active/non-optimized state:~~

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

- a) INQUIRY
- b) LOG SELECT
- c) LOG SENSE
- d) MODE SELECT (6/10)
- e) MODE SENSE (6/10)
- 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
- o) Echo buffer modes of WRITE BUFFER

The device server may support other commands when accessed through a target port in the ~~inactive~~ 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~~ INACTIVE STATE.

The SCSI target device shall participate in all task management operations as defined in SAM-2.

5.6.3.5 Unavailable

While in the unavailable target port behavior state, the device server shall accept at least a severely restricted set of the 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.

~~All target ports in a target port group cannot access the requested logical unit.~~

~~The device server shall support all of the following commands that the logical unit supports when accessed through a target port in an active/optimized or active/non-optimized state.~~

While in the unavailable target port behavior 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
- h) Download microcode mode of WRITE BUFFER

~~The device server may support other commands that the logical unit supports when accessed through a target port in an active/optimized or Active/Non-optimized state.~~ 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.

5.6.3.6 ~~Transition~~ Changing

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

The device server shall support all of the following commands that the logical unit supports when accessed through a target port in the active/optimized or active/non-optimized state:

While in the changing target port behavior 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
- 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 ~~Transition~~ **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 ~~TRANSITION~~ **CHANGING** STATE.

~~Commands that are supported shall operate according to one of the five following circumstances:~~

- ~~_____ a) the target port behavior state from which the target port is changing, or~~
- ~~_____ b) the target port behavior state to which the target port is changing, or~~
- ~~_____ c) the device server return BUSY status, or~~
- ~~_____ d) the device server 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~~
- ~~_____ e) The transition from one state to another failed and the device server shall returns CHECK CONDITION status with the sense key set to NOT HARDWARE ERROR and an additional sense code of LOGICAL UNIT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND FAILED.~~

~~The time at which a transition change is established with respect to other tasks being managed by the device server is vendor specific. Once the target port behavior state transition is completed, the new target port behavior state may apply to some or all tasks queued before the completion of the transition. The new target port behavior state shall apply to all tasks received by the device server after successful completion of the target port behavior state transition.~~

~~The execution of any target port behavior state transition shall be performed as a single indivisible event. Multiple SET TARGET PORT GROUPS commands may be queued at the the same time. The order of execution of such commands is defined by the tagged queuing restrictions, if any, but each is executed as a single indivisible command without any interleaving of actions that may be required by, for example, other SET TARGET PORT GROUPS commands.~~

The SCSI target device **shall** participate in all task management operations as defined in SAM-2.

5.6.3.7 Behavior after power cycling or hard resets

For all SCSI target devices that report in the INQUIRY data that they support only explicit asymmetric logical unit access behavior (ALUAB bit of one and IAAB bit of zero), the target port shall preserve the target port behavior state across any reset and across any power off period.

5.6.4 Implicit asymmetric logical units access behavior management

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

All logical units that report in the INQUIRY data that they support asymmetric logical units access behavior and support implicit asymmetric logical unit access behavior shall. (ALUAB field is 01b or 11b):

- a) implement the INQUIRY command vital product data device identifier page identifier types 4h and 5h as specified in 8.4.4. ~~The 04h identifier type defines the relative target identifier for the target port through which the INQUIRY command was passed.~~

This provides the necessary information to identify target ports defined by the REPORT TARGET PORT GROUPS command. The 05h identifier type defines the target port group and the present 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.

- b) support the REPORT TARGET PORT GROUPS command. 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 asymmetric logical units access behavior state for each target port group.

5.6.5 Explicit asymmetric logical units access behavior management

All logical units that report in the INQUIRY data that they support explicit asymmetric logical units access behavior shall. (ALUAB field is 10b or 11b):

- a) implement the INQUIRY command vital product data device identifier page identifier types 4h and 5h as specified in 8.4.4. ~~The 04h identifier type defines the relative target identifier for the target port through which the INQUIRY command was passed.~~
- b) support the REPORT TARGET PORT GROUPS command. 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 asymmetric logical units access behavior state for each target port group.
- c) support the SET TARGET PORT GROUPS command. This command may be used to explicitly set the target port group behavior state for all target port groups. The allowable combination of states for the target port groups is vendor specific. If the SET TARGET PORT GROUPS attempts to establish an invalid combination of 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.

5.6.5.1 Behavior after power cycling or hard resets

For all SCSI target devices that report in the INQUIRY data that they support only explicit asymmetric logical unit access behavior (ALUAB field is 10b or 11b—see 7.6.2), the target port shall preserve the target port behavior state across any reset and across any power off period.

4.2 Section 7.6.2, Standard INQUIRY data

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

Asymmetric logical unit access behavior (ALUAB field) :

- a) an ALUAB field value of 00b indicates that the SCSI target device does not support asymmetric logical unit behavior and therefore the asymmetric or symmetric behavior is unspecified.
- b) an ALUAB field value of 10b indicates that the SCSI target device only supports explicit asymmetric logical unit access behavior and supports both the REPORT TARGET PORT GROUPS and SET TARGET PORT GROUPS commands. Since only explicit logical unit behavior is supported when the ALUAB field value is 10b, the SCSI target device may only change target port states as requested with the SET TARGET PORT GROUPS command.
- c) an ALUAB field value of 01b indicates that the SCSI target device only supports implicit asymmetric logical unit access behavior and therefore is capable of changing target port behavior states without a SET TARGET PORT GROUPS command. When the ALUAB field is 01b, the REPORT TARGET GROUPS command is supported and the SET TARGET GROUPS command is not supported.
- d) an ALUAB field value of 11b indicates that the SCSI target device supports both explicit and implicit asymmetric logical unit access behavior and supports both the REPORT TARGET PORT GROUPS and SET TARGET PORT GROUPS commands.

- ~~a) An AAB bit of one indicates that the SCSI target device supports explicit asymmetric logical unit behavior and supports both the REPORT TARGET PORT GROUPS and SET TARGET PORT GROUPS commands.~~
- ~~b) An AAB bit of zero indicates the SCSI target device does not support explicit asymmetric logical unit behavior and does not support the SET TARGET PORT GROUPS command.~~

~~Implicit asymmetric logical unit access behavior (IAAB) bit:~~

- ~~a) An IAAB bit of one indicates the SCSI target device supports implicit asymmetric logical unit behavior and therefore is capable of changing target port behavior states without a SET TARGET PORT GROUPS command. An IAAB bit of one also indicates that the REPORT TARGET PORT GROUPS command is supported.~~
- ~~b) An IAAB bit of zero indicates that the SCSI target device does not support implicit asymmetric logical unit behavior and therefore may only change target port behavior states as requested with the SET TARGET GROUPS command, if supported.~~

The Table below ZZO describes the interaction between the ALUAB field, and IAAB bits.

Table ZZ0 ALUAB fields

ALUAB VALUE	<i>Description</i>
00b	No asymmetric behavior
01b	Implicit behavior only
10b	Explicit behavior only
11b	Implicit and explicit behavior

4.3 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 which 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 which will be reported and controlled using the REPORT TARGET PORT GROUPS and SET TARGET PORT GROUPS commands.

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

Table 181 – Identifier type

Value	<i>Description</i>
5h	If the ASSOCIATION value is 1h, the IDENTIFIER value contains the target port group and target port group behavior 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 Behavior State Identifiers

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

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

4.4 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

(proposed value = 04/0A)

LOGICAL UNIT NOT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND

FAILED (proposed value = ~~04/0D~~ 67h/0Ah)

4.5 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 that they support asymmetric logical unit access behavior (non zero ALUAB field – see 7.6.2). 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 the REPORT TARGET PORT GROUPS service action of the MAINTENANCE IN command defined in the SCC-2 standard. Additional MAINTENANCE IN service actions (that apply to SCC-2 devices and devices that set the SCCS bit in their Standard Inquiry data) are defined in SCC-2. Only those service actions of MAINTENANCE IN that are defined in this standard concern all SCSI devices. SCC-2 defines specific usage for bytes 4 and 5, and bit 1 in byte 10, however these fields are reserved for the REPORT TARGET PORT GROUPS command defined by this standard. Additional MAINTENANCE IN service actions are defined in SCC-2 and in this standard. (editors notes: it is up to the discretion of the editor as to whether or not the following sentences in this paragraph may be deleted) 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 sccs 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 (editors notes: it is left to the discretion of the editor if this would be improved by changing to "restricted for SCC-2")							
5								
6	(MSB) ALLOCATION LENGTH (LSB)							
7								
8								
9								
10	reserved						See SEC-2 restricted	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.

Table XX2 – REPORT TARGET PORT GROUPS parameters

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	REPORT FIELD LENGTH (= n)							
1									
2									
3	(LSB)								
4		reserved(first target port group)							
5									
6	(MSB)	TARGET PORT GROUP IDENTIFIER (first target port group)							
7		(LSB)							
8		reserved				TARGET PORT GROUP BEHAVIOR STATE			
9		STATUS CODE							
10		VENDOR UNIQUE STATUS							
11		TARGET PORT COUNT IN TARGET PORT GROUP (first target port group) (= x)							
12	(MSB)	RELATIVE TARGET PORT IDENTIFIER (first target port of first target port group)							
13									
14									
15	(LSB)								
...									
4x+8	(MSB)	RELATIVE TARGET PORT IDENTIFIER (last target port of first target port group)							
4x+9									
4x+10									
4x+11	(LSB)								
...									
n-4y-4		reserved (second target port group)							
n-4y-3									
n-4y-2	(MSB)	TARGET PORT GROUP IDENTIFIER (last target port group)							
n-4y-1		(LSB)							
n-4y		reserved				TARGET PORT GROUP BEHAVIOR STATE			
n-4y+1		STATUS CODE							
n-4y+2		VENDOR UNIQUE STATUS							
n-4y+3		TARGET PORT COUNT IN TARGET PORT GROUP (last target port group)(=y)							
...									
n	(MSB)	RELATIVE TARGET PORT IDENTIFIER (last target port of last target port group)							
n+1									
n+2									
n+3	(LSB)								

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 65,536 target port groups.

The TARGET PORT GROUP BEHAVIOR STATE field is an encoded value (see 7.6.2). This field specifies the target port group behavior state in the SCSI target device that reports in the INQUIRY DATA that asymmetric logical unit access behavior is supported.

Table ZZ2 Target Port Group Behavior State

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

Active/optimized, active/non-optimized, ~~Inactive~~, standby, unavailable and ~~Transition~~ Changing define the current target port group behavior state defined by 5.6.

The TARGET PORT COUNT IN TARGET PORT GROUP field specifies the number of target ports that have access to the logical unit 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.

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.

The execution of a REPORT 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 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). 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 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 Behavior State Changed By SET TARGET PORT GROUPS Command
02h	Target Port Group Behavior State Changed By Implicit Asymmetrical Logical Units 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 set the state of all of the target ports in a target port group for the specified logical unit. This command is mandatory for all SCSI target devices that report in the INQUIRY data that they support explicit asymmetric logical units access behavior (ALUAB field is either 10b or 11b – see 7.6.2). 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 the SET TARGET PORT GROUPS service action of the MAINTENANCE OUT OUT command defined in the SCC-2 standard. Additional MAINTENANCE OUT service actions (that apply to SCC-2 devices and devices that set the SCCS bit in their Standard Inquiry data) are defined in SCC-2. Only those service actions of MAINTENANCE OUT that are defined in this standard concern all SCSI devices. SCC-2 defines specific usage for bytes 4 and 5, and bit 1 in byte 10; however, these fields are reserved for the SET TARGET PORT GROUPS command defined by this standard.

Additional MAINTENANCE OUT service actions are defined in SCC-2 and in this standard. (editors notes: it is up to the discretion of the editor as to whether or not the following sentences in this paragraph may be deleted). 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. When a SCSI device returns a device type of 0Ch or the sccs bit equal to one in its standard INQUIRY data, the implementation requirements for the SCC-2 MAINTENANCE OUT service action are as specified in SCC-2. Otherwise the MAINTENANCE OUT service action definitions and implementation requirements stated in this standard shall apply.

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	reserved							
4	SCC-2 restricted (editors notes: it is left to the discretion of the editor if this would be improved by changing to "restricted for SCC-2")							
5								
6	(MSB) PARAMETER LIST LENGTH (LSB)							
7								
8								
9								
10	reserved						See SCC-2 restricted	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 state of any target port groups. The number of target port groups supported by a logical unit is vendor specific. The number of target port groups that shall be provided in the parameter list and the allowable values to which their states may be set is vendor specific. 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 SET TARGET PORT GROUPS parameter list (see table YY2) contains a four-byte field that contains the length in bytes of the parameter list.

Table YY2 – SET TARGET PORT GROUPS parameters

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) SET FIELD LENGTH (= n)							
1								
2								
3								
4	reserved							
5	(MSB) TARGET PORT GROUP IDENTIFIER (first group)							
6								
7	reserved				TARGET PORT GROUP BEHAVIOR STATE			
...								
n	reserved							
n+1	(MSB) TARGET PORT GROUP IDENTIFIER (last group)							
n+2								
n+3					TARGET PORT GROUP BEHAVIOR 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 65,536 target port groups.

The TARGET PORT GROUP BEHAVIOR STATE field is an encoded value. (see table ZZ4)
 This field specifies the state of **to which** all of the target ports in the target port group **shall transition**, for a SCSI target device that reports in the INQUIRY DATA that explicit asymmetric logical units access behavior is supported

TABLE ZZ4 TARGET PORT GROUP BEHAVIOR STATE

<i>Values</i>	<i>States</i>
0h	active/optimized
1h	active/non-optimized
2h	Inactive standby
3h	unavailable
4h–Fh	reserved

active/optimized, active/non-optimized, ~~Inactive~~ **standby**, and unavailable define the target port group behavior state as defined by 5.6.

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

- a. 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 logic 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.

- b. The command has been executed and the target port behavior state change is in transition:

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

- c. A failure occurred before the command was executed:

The device server shall return CHECK CONDITION status. The sense key shall be set to ~~NOT READY~~ HARDWARE ERROR and the additional sense data shall be set to LOGICAL UNIT NOT ACCESSIBLE, SET TARGET PORT GROUPS COMMAND FAILED.