

Date: 5/24/2004

T10: T10 Technical Committee (SCSI)

From: Charles Binford, Sun Microsystems, Inc., (charles.binford@sun.com)

Subject: Report Target Port Groups Changes

Revision History

Revision 0 (April, 2004), first revision

Revision 1 (May 5, 2004) – editorial changes per CAP working group, as well as editorial changes received after the CAP working group.

Related Documents

SPC-3r17 – SCSI Primary Commands – 3, revision 17

Overview

After reviewing the Asymmetrical Logical Unit Access feature set in SPC-3 SUN proposes we add a new bit to the Report Target Port Groups parameter data. This new bit, PREF, allows a device server to optionally indicate a path preference for access to a logical unit. This new information allows a device server to indicate to a host multi-path driver what the preferred access path is regardless of the current state.

Examples:

- One of the arrays that SUN uses has hardware level performance gains if a LU is accessed via a specific path. While in Explicit failover mode, the new PREF bit will allow a host driver to detect and remedy the situation where the port group in the standby state is actually the preferred path without resorting to vendor specific commands.
- Another array doesn't have any performance gains based on path, and uses an implicit failover policy. In this case the preferred path indication is used to coordinate the path selection in the host multi-path drivers in a multi-initiator environment. To avoid thrashing on LU moves between controllers it is vital that the multiple host choose the same path for access to a given LU. Currently the path preference is indicated via vendor unique commands and data. This change would standardize the delivery of the path preference.

A second requested change is to add a clause that explains how a device that has symmetrical access can use these commands to communicate this to a host driver. Today some host drivers attempt to load balance IO across paths to a LU when they know the device supports symmetrical access. However, the determination of symmetrical access is vendor unique.

Details of Proposed Change

- Change name of Inquiry field from ALUA to TPGS, Target Port Group Support.
 - And all subsequent references to ALUA.
- Rename *5.8 Asymmetric logical unit access*, to *5.8 Target Port group access states*
- Change paragraph level of current *5.8 Asymmetric logical unit access*, and new *5.9 Symmetric logical unit access* to be sub-clauses of *5.8 Target Port group access states*.
- Delete the first sentence of the current 5.8.1.

New text for 5.8 Target Port group access states:

5.8 Target Port group access states

Logical units may be connected to the service delivery subsystem via multiple target ports (see SAM-2). The access to logical units through the multiple target ports may be symmetrical or asymmetrical.

SPC-3 r17 page 81, insert a new section 5.8.5 titled Preference Indicator

5.8.5 Preference Indicator

A device server may indicate one or more target port groups is a preferred target port group for accessing a logical unit by setting the PREF bit in the target port group descriptor parameter data (see table 152). The preference indication is independent of the asymmetric access state. An application client may use the PREF bit value in the target port group descriptor parameter data to influence its path selection to a logical unit (e.g. a target port group in the Standby state with the PREF bit set to one may be chosen over a target port group in the Active/Optimized state with the PREF bit set to zero).

The value of the PREF bit for a target port group may change whenever an asymmetric access state changes.

SPC-3 r17 page 204-205, 6.24 Report Target Port Groups command.

Add a PREF field to byte 0, bit 7 of table 152.

Table 152 — Target port group descriptor parameter data

Bit Byte	7	6	5	4	3	2	1	0
0	PREF	Reserved			ASYMMETRIC ACCESS STATE			
1	Reserved				U_SUP	S_SUP	AN_SUP	AO_SUP
2	(MSB)	TARGET PORT GROUP						(LSB)
3								
4	Reserved							
5	STATUS CODE							
6	Vendor unique							
7	TARGET PORT COUNT (x)							
	Relative target port(s)							
8	RELATIVE TARGET PORT 1							
11	⋮							
	⋮							
n-3	RELATIVE TARGET PORT x							
n								

Add description for PREF following table 152.

A PREF bit set to one indicates that the target port group is a preferred target port group for accessing the addressed logical unit (see 5.8.5 Preference Indicator). A PREF bit set to zero indicates the target port group is not a preferred target port group.

Symmetrical Access Indication

Suggesting a new section inserted at end of 5.8.

5.8.x Symmetric logical unit access

A device server that provides symmetrical access to a logical unit may use a subset of the asymmetrical logical access features to indicate this ability to an application client. Doing so allows an application client to use a common set of commands to determine how to manage port access to a logical unit.

Symmetrical logical unit access should be represented as follows:

- the TGPS field (see 6.4.2) in the standard INQUIRY data shall indicate implicit asymmetric access is supported,
- the REPORT TARGET PORT GROUPS command shall be supported, and
- the REPORT TARGET PORT GROUPS parameter data shall indicate the same state (e.g. all in the active/optimized state) for all target port groups with symmetrical access to the logical unit.