

To: T10 Technical Committee
From: Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
Date: 1 November 2000
Subject: Comments on 00-232r2 Asymmetric SCSI behavior

These comments apply to 00-232r2, the asymmetric SCSI behavior proposal. Comments are primarily from Dan Mazina and Bill Dallas at Compaq.

Add a Persist through Power Loss bit for Target Group States

A persistent bit will be defined in the Set/Report Target Group command parameters and the device identification page (VPD code page 83h) that indicates that the target group state is persistent for a particular logical unit (LU). Persistent target group state indicates that after a node is powercycled the target groups will attempt to regain their previous states with respect to the available logical units (LU). [Not correct: This bit only applies to a node powercycle condition and is not affected by port or device resets.]

This bit is writable for the Set Target Group command allowing the host to enable persistence for a logical unit (LU).

This bit readable for the Report Target Group and device identification page and indicates the current value of the persistence setting for a logical unit (LU).

If a target group does not support persistence, then this bit will be set to zero for all logical units. Attempts to set it to one will result in a check condition.

The persistent bit is defined to be bit 5 of the byte that contains the various target group states.

byte/bit	7	6	5	4	3	2	1	0
<i>n</i>			persist.	active/optim.	active/non-op	inactive	unavail	transit

Other modifications/clarifications:

- Read for Report and Inquiry
- Write - As a host, I don't want to try to set the bit to see if it can be set. Either it must always be settable or the data set is changed for Report Target Group that tells the OS if the Persistence can be set.

Add a Target Group Status Code and Detail

A status code and vendor unique status detail field will be defined in the Report Target Group command parameters that allows a target group to report any status conditions that may be affecting its target group state.

byte/bit	7	6	5	4	3	2	1	0
n	Reserved							
$n + 1$	TARGET GROUP IDENTIFIER							
$n + 2$			persist tent	active/ optim.	active/ non- op	inactive	unavail	transit
$n + 3$	STATUS CODE							
$n + 4$	VENDOR UNIQUE STATUS DETAIL							
$n + 5$	TARGET COUNT IN GROUP							

1.1.1. Status Code

The status code field indicates why a target may be in a specific state. It provides a mechanism to indicate error conditions.

- 0h No Status Available
- 1h Target Group Hardware Failure
- 2h Target Group Firmware Failure
- 3h Target Group Write Cache Failure
- 4h Target Group Read Cache Failure

1.1.2. Vendor Unique Status Detail

This field contains vendor specific details that accompany the status code field.

Other suggested modifications:

- Change Status code to a 2 byte field with the upper nibble defined to reflect who is reporting (e.g., application server, device server, etc.) or leave it a whole 2 byte field but define a field that reflects reporting entity.
- Change wording to reflect the status code reflects the posted state (active/optim, inactive, etc.). Log sense can be used to reflect history.
- Get rid of Vendor Unique Status Detail - put that in log sense information.
- Change Set Target Group parameters to include the setting of the above mentioned status code.
- 1.2.1 Status code value 5h get rid of SCSI to reflect just a backplane failure.
- 1.2.1 Status code value 6h get rid of Fibre to reflect just Port Failure.

Add Logical Unit Groups

This change is not required for 00-232 and will be handled as a separate proposal that is also used by Access Controls.

Asymmetrical SCSI behavior needs to include the concept of logical unit (LU) groups. Specific LU's may be dependent on each other and must have the same target group state value for all dependent LU's. These dependent LU's are defined as a LU group. Any change in target group

state for a single LU in a LU group will affect all LU's in the same group. LU groups may contain from one to n LU's. LU's are prohibited from being in two LU groups.

A new identifier type for the device identification page (VPD code page 83h) will be defined for LU groups. This will contain a value of 06h and is defined in the same way as identifier type 1h is (first 8 bytes is a Vendor ID).

1.1.3. LU Group Identification Descriptor

The LU group descriptor identifies a LU group that the addressed LU is a member of. If the LU is not a member of any LU groups, then no descriptor will exist. This descriptor is a World Wide Name (WWN) and is guaranteed to be unique.

1.1.3.1.1. LU Group Descriptor Format

byte/bit	7	6	5	4	3	2	1	0
0	Reserved				Code Set = 1 (binary)			
1	Reserved		Assoc. = 0		Type = 6h			
2	Reserved							
3	Identifier Length = 24							
4 – 27	Identifier							

State Exclusion Clarification

Due to the nature of asymmetrical behavior, a node with multiple target groups may exhibit exclusive behavior with respect to the states allowed for the target groups. In particular, a node may prevent two target groups from both claiming 'active' states to the same logical unit (LU) (or logical units within a logical unit (LU) group). This discussion considers the 'active/optimized' state and the 'active/non-optimized' state to be both 'active' states. Therefore, depending on the implementation of the asymmetrical behavior, a particular logical unit (LU) (or all logical units in a logical unit (LU) group) may have multiple target groups that have 'inactive' or 'unavailable' states with respect to the logical unit but only one target group will ever be 'active/*' at any one time. If the host uses the Set Target Group command to change a target group's state from 'inactive' to 'active/*', this can result in having the target group that previously was 'active/*' being implicitly moved into the 'inactive' state.

Editorial changes

Change "PATH STATE" in the new ASCQs to "TARGET GROUP" since that terminology is used

Make the proposed Status field writable by the application client, stored by the target group, so one cluster member can notify the others about why it took a target group offline (e.g. take it offline to download new microcode).

Make Download Microcode work in unavailable state rather than inactive state.

Ask T10 to reserve the new ASCQs and opcodes (function codes) if the proposal isn't finished and accepted this week.

The proposal targets asymmetrical behavior where this proposal can be applied to those controllers that exhibit symmetrical behavior. This can be very useful to application clients (hosts)

for symmetrical controllers (firmware upgrades etc).

One option: add a sentence clarifying that symmetric devices are a subset of asymmetric devices that just have one target group.

Another approach:

- Change wording throughout document that excludes symmetrical behavior in this proposal.

- Change Standard Inquiry bit name and definition from ATB (asymmetrical target behavior) to something like TGB (target group behavior).

- Change Standard Inquiry bit name and definition from IATB (implicit asymmetrical target behavior) to something like ITGB (implicit target group behavior).

- Change wording for IATB/ITGB that reflects that the controller or box that reports explicit target group behavior (host controlled) can also exhibit implicit target group behavior. Why restrict?

Page 9, second paragraph, third sentence change "device does not perform any" to "device may perform".

Change Report and Set Target Groups (7.x, 7.y) to reflect that when the command is rejected it is done with:

- o Check Condition
- o ILLEGAL REQUEST
- o then the ASC/ASQs

Expand the TARGET GROUP IDENTIFIER field from its current 1 byte to 2 bytes and leave a little room for future expansion.