

Date: May 04, 2004

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli), David Chambliss (IBM Almaden Research Center)

Subject: Command Classification

1 Overview

The following proposed wording represents changes to SCSI Block Commands - 2 (SBC-2) to enable the transmission of grouping or classification information on a per-command basis.

This proposal standardizes a function that allows application clients to group commands so that device servers can collect information about each group of command. A application client can then, in turn, read the information collected by the device server and use that information to do things like work load balancing. Generally the collection of this information is done using out-of-band methods (i.e., the information is not transmitted across using any SCSI protocols).

An example of how grouping or classification could be used would be if two applications were using a subsystem; one application primary streamed data and another primary does random access of data to a logical unit. If the streaming application groups all it's commands with one value (e.g., x) and the random application group all it's commands with another value (e.g., y) then a group x defined to hold performance metrics would collect all the performance metrics for the streamed commands together and a group y defined to also hold performance metrics would collect all the performance metrics for the random commands together. The result would be two sets of performance metrics (x and y). A management application could then read the performance metrics and determine if the performance of a specific group is acceptable.

This proposal does not attempt to define the groups. At this time it is not clear what the groups will be. Even after they have been defined it is not clear whether that is something that should be defined as an in-band, out-band, or set at the factory feature. If it is decided to be in-band a mode page can be defined at a later date with no effects that would relate to this proposal.

SBC-2 additions

1.1 Affected commands

It is proposed that a 5-bit field (GROUP NUMBER) be added to at least the following commands:

- a) READ (10)
- b) READ (12)
- c) READ (16)
- d) VERIFY (10)
- e) VERIFY (12)
- f) VERIFY (16)
- g) WRITE (10)
- h) WRITE (12)
- i) WRITE (16)
- j) WRITE AND VERIFY (10)
- k) WRITE AND VERIFY (12)
- l) WRITE AND VERIFY (16)

In addition the GROUP NUMBER field shall be added to the following commands:

[Editor's Note 1: make this one list of commands to add this to.](#)

- a) PRE-FETCH (10)
- b) PRE-FETCH (16)

- c) SYNCHRONIZE CACHE (10)
- d) SYNCHRONIZE CACHE (16)
- e) WRITE SAME (10)
- f) WRITE SAME (16)
- g) XDREAD (10)
- h) XDREAD (32)
- i) XDWRITE (10)
- j) XDWRITE (32)
- k) XDWRITEREAD (10)
- l) XDWRITEREAD (32)
- m) XPWRITE (10)
- n) XPWRITE (32)

The GROUP NUMBER field would be in bits 4-0 in byte 6 of 10-byte commands, byte 10 of 12-byte commands, byte 14 of 16-byte commands, byte 6 of 32-byte commands and byte 6 for variable length commands.

4.16 Grouping function

A grouping function is a function that collects information about attributes associated with commands (i.e., information about commands with the same group value are collected into the specified group). The definition of the attributes and the groups is outside the scope of this standard.

The collection of this information is outside the scope of this standard (i.e., the information may not be transmitted using any SCSI protocols).

NOTE 1 - An example of how grouping could be used would be if two applications use a subsystem: one application primarily streams data and another primarily accesses data randomly. If the streaming application groups all of its commands with one value (e.g., x), and the random application groups all of its commands with another value (e.g., y), then a group x defined to hold performance metrics collects all the performance metrics for the streamed commands together and a group y defined to also hold performance metrics collects all the performance metrics for the random commands together. The result is two sets of performance metrics (i.e., x and y). A management application then reads the performance metrics and determines if the performance of a specific group is acceptable.

5.9 READ (10) command

The READ (10) command (see table 1) requests that the device server transfer data to the application client. The most recent data value written in the addressed logical block shall be returned.

Table 1 — READ (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (28h)							
1	Reserved			DPO	FUA	Reserved		RELADR
2	(MSB) _____ LOGICAL BLOCK ADDRESS _____ (LSB)							
5								
6	Reserved			GROUP NUMBER				
7	(MSB) _____ TRANSFER LENGTH _____ (LSB)							
8								
9	CONTROL							

The GROUP NUMBER field specifies the group into which attributes associated with the command should be collected (see 4.16). A group number value of zero specifies any attributes associated with the command shall not be collected into any group.

