

X3T10/95-260 R0

To: scsi @ symbios.com @ internet
cc:
From: Gerry Houlder
Date: 05/26/95 04:32:14 PM
Subject: Inadequacy of Number of Blocks in Mode Select/Sense

The Number of Blocks field in the MODE SELECT/SENSE block descriptor is defined as only 3 bytes. This is inadequate for some current products and will become more so in the future. The standard allows for setting 0xFFFFFFFF in this field when the actual capacity is greater than this, but this isn't useful for some current applications.

Some customers use this field on MODE SELECT to set the capacity that the drive will report. While this usage is not sanctioned by the SCSI-2 or -3 standards, there is enough customer demand to force vendor unique methods of doing this if the standard doesn't allow for it. I do not want to be forced into supporting several different vendor unique ways to do this; it is better to extend the standard to support it.

Another reason for a customer to want this is to split a drive into multiple LUNs. When capacity on one drive is too large to be handled as one LUN, a method is desired to split it into several LUNs. Using several block descriptors (one for each LUN) in MODE SELECT is one possible method of doing this. Does anyone have a better idea?

This could be fixed by defining a different block descriptor format for 10 byte MODE SELECT/SENSE commands that allocates 4 bytes for the Number of Blocks field. The block descriptor can remain 8 bytes long by defining it as follows:

Bytes 0-3: Number of Blocks Field
Byte 4: Density Code (same as byte 0 in present block descriptor)
Bytes 5-7: Block Length

I know that defining a new block descriptor for use with 10 byte MODE SELECT/SENSE commands has some backward compatibility problems, but few customers are actually using these commands yet and the added functionality will be worth the pain for the few that do. I am interested in knowing who has opinions on this.