

TO: Members of X3T9.2 and interested parties  
FROM: Paul Boulay, (408) 986-9770 x 205  
SUBJECT: Media Changer Device Capabilities Page

Revision 1 Note: This has been updated to reflect my understanding of the consensus of the March 19-20 Working Group meeting.

In section 16.3.3.1, the description of the EXCHANGE MEDIUM command capabilities is somewhat ambiguous and conflicts with a reference to these fields in the EXCHANGE MEDIUM command description.

(For the page see Table 16-26. The fields are in bytes 12 - 15 -- I refer to this set of fields collectively as the table. The description is the last paragraph on page 16-37, Rev 10c. For the command see section 16.2.1, 3rd paragraph on page 16-8, Rev 10c.)

The ambiguity arises because some reasonable medium changer implementations may be able to do some of the operations implied by a bit in the table but not others. A one bit in the table asserts that a class of operations are supported. So how should an implementation set this bit if it supports some of these operations but not others. I believe that we should provide some sort of guidance here.

In the reference to this table in the EXCHANGE MEDIUM command, the reference describes these bits as providing information about particular element combinations rather than element type combinations. I believe that this needs to be changed to agree with the description of the page.

To provide better understanding of the problem, consider a changer that supports most exchange operations but because of the physical design of the device does not implement an exchange where the source element and the destination 2 element happen to be the identical element. (This is the case given by the current text of the reference paragraph in the EXCHANGE MEDIUM command.) In such a device, how should the bit in the table be set?

Actions:

1. In the command change the reference to make it clear that the page gives information about exchanges in terms of element types. This was agreeable to the working group. My suggested wording for the 3rd paragraph on page 16-8 is:

The device capabilities MODE SENSE page (see 16.3.3.) provides a matrix with the supported source element type / destination 1 element type combinations for EXCHANGE MEDIUM commands with source element type the same as destination 2 element type.

2. In the page description: These 3 alternative approaches were considered:

- A. Interpret a bit set to mean that ALL such operations are supported.
- or- B. Interpret a bit set to mean that SOME of the operations are supported but others in the class may not be.
- or- C. Add an ALL/SOME bit off to the side someplace that guides user in the interpretation of this data.

Both options (A) and (B) garnered some support; nobody liked (C). Support for (B) flagged when the group recognized that (A) really does give the device driver more information. With this interpretation, if a bit is not set a simple device driver will use a set of primitive operations that performs some desired function. A smarter device driver might well try it on the possibility that the particular combination is supported after all.

Given that option (A) won out, in the name of consistency, the previous paragraph ends up changing too. My suggested wording for the last two paragraphs on page

16-37 is:

An XX->YY bit value of one indicates that the medium changer device supports all MOVE MEDIUM commands where the source is element type XX, the destination is element type YY and these element addresses are otherwise valid. An XX->YY bit value of zero indicates that these MOVE MEDIUM commands may or may not be valid depending on the particular elements requested. Those which are not valid will be rejected with ILLEGAL REQUEST.

An XX<>YY bit value of one indicates that the medium changer device supports all EXCHANGE MEDIUM commands where the source is element type XX, destination 1 is element type YY, destination 2 is the same type as the source element type and these element addresses are otherwise valid. An XX<>YY bit value of zero indicates that these EXCHANGE MEDIUM commands may or may not be valid depending on the particular elements requested. Those which are not valid will be rejected with ILLEGAL REQUEST.

The discussion also brought out some options to consider for SCSI-3. We could separate each of the tables into 2 - the idea being to allow the target to report different capabilities (1) for exchanges with the same ELEMENT as both source and destination and (2) for all other exchanges within a class.

Another concept was advanced that could generalize the reporting of capabilities. The idea is to define groups with like capabilities. Then report capabilities on a group basis. The example given was a system where a particular transport element is able to serve only one drive. Other examples of the usefulness of this concept were given.

On a somewhat related subject, I was asked to draft an implementers note describing the way the transport geometry page could be used. Actually this was an early (and probably not very successful) attempt to report grouping information. Suggested text:

IMPLEMENTORS NOTE: This page reports information about the way transport elements are physically clustered in a system. The model for this is a medium changer device with more than one independent robotics subsystem, where each of these supports multiple transport elements. The elements that are supported by a particular robotics subsystem form a set. This sort of information is helpful for optimization and error recovery in such a large system. (Recall that in the model for this device type, section 16.1, the individual transport element is addressed not the robotics subsystem. An element is defined to be a place where a unit of media may be at any point in time.)