

Subj: XOR Command meeting minutes 1/9/95
Date: 95-01-12 12:17:40 EST
From: Jay_Elrod@notes.seagate.com
To: X3T10 Membership

X3T10/95-114r0

XOR Command Study Group meeting minutes - Document number X3T10/95-114r0
Meeting place: Harrah's, Lake Tahoe
Date: January 9, 1995

Attendees:

Gerry Houlder Seagate
Jay Elrod Seagate
Bill Hutchison Hewlett Packard
Doug Prins Q-Logic
Larry Lamers Adaptec
Paul Hodges IBM
Keith Staub Quantum
Paul Peterson Conner
James McGrath Quantum
John Baudrexl Fujitsu/Intellistor
George Penokie IBM
Bob Snively Sun Microsystems
Paul Boulay Hitachi Computer Products
Thai Nguyen StorageTek
Duncan Penman IIX
Giles Frazier IBM Austin
Jeanne Martin LLNL
Lansing Sloan LLNL
Gene Milligan Seagate

Gerry Houlder acted as chairman for the meeting. The issues discussed are summarized below.

- 1) Newly added 10 byte XDWRITE command - the NDisc bit will be moved from bit 5 to bit 2 (of byte 1) to more closely resemble the 16 byte XDWRITE command.
- 2) RAID Group Address mode page - the concept of multiple RAID groups, LBA ranges, etc. will be removed. (The whole mode page may be removed for the time being, until a need is demonstrated for a more-than-one-byte physical secondary address.)
- 3) Opcodes for XOR commands were presented, with no objection. The six new opcodes are:

80h = XDWRITE (16)
81h = REBUILD
82h = REGENERATE
50h = XDWRITE (10)
51h = XPWRITE
52h = XDREAD

- 4) The NDisc bit will be put back into the 16 byte XDWRITE command for those who wish to use the XDWRITE command (rather than REGENERATE) during a controller supervised rebuild or regenerate operation. (A "controller supervised" rebuild or regenerate is one where the devices do not have peer to peer capability, and the controller must send XOR commands to all devices involved.) The NDisc bit will occupy bit 2 of byte 1.
- 5) Once again, the issue of whether to consolidate the MAX XDWRITE SIZE and MAX XPWRITE SIZE fields in the RAID Control mode page was discussed. The two fields will be consolidated into one field which will occupy bytes 4 - 7 of the RAID Control mode page. Bytes 8 - 11 will be reserved.
- 6) It was recommended that there be information available in the RAID Control mode page regarding total buffer space available for XOR type operations. No final conclusion was reached as to how this would best be done. It was agreed, however, that there should be a "buffer full" status byte to indicate when there is no more room available in the buffer. This buffer full condition results (for example) from XOR data sitting in the buffer waiting for an XDREAD command. A "buffer full" status byte will be proposed to the appropriate committee.
- 7) There was discussion regarding how long XOR data should remain in the buffer of an XOR device following an XOR command which would be expected to be followed by an XDREAD command (10 byte XDWRITE, for example). This issue will be covered in the next release of the XOR document, stating the events which can cause the XOR data to be cleared.
- 8) It was recommended that an implementor's note be added to any commands which require a follow up XDREAD command, stating that linking of the XDREAD command to the associated prior command is recommended for data integrity.
- 9) It was recommended that the models at the front of the XOR document be updated to include more examples of possible RAID configurations, along with recommended procedures for executing RAID functions in those configurations.
- 10) For reasons of SPC compatibility it was recommended that the 16 byte XDWRITE command be restructured. As a result of this discussion, the Secondary Address field will be moved from bytes 12-14 to byte 14 (reducing it from 3 bytes down to 1 byte), and the Transfer Length field will be expanded from bytes 10-11 to bytes 10-13.
- 11) For reasons of compatibility with the REBUILD and REGENERATE commands, the Secondary Control field of the 16 byte XDWRITE command will be moved to bits 0-1 of byte 1.
- 12) It was recommended that an "XOR Disable" bit be added to the RAID Control mode page to allow the controller to disable any XOR functionality within the device. The thinking is that in some designs this may free up any buffer space which may have otherwise been

designated for use in XOR operations. Bit 1 of byte 2 of the RAID Control mode page will be used for the XOR Disable bit.

13) It was recommended that the Mirror bit be removed from the 16 byte XDWRITE command since the operation of the command is radically altered by this bit to the point that no XOR functions are performed. This bit will be removed in the next revision.

14) It was recommended that an XOR bit be added to Inquiry data, indicating functionality in the device. This will be proposed to the appropriate committee.

15) It was recommended that the statement on page 22, regarding how far an XPWRITE command shall have progressed before sending good ending status when write caching is enabled, be removed. This will be done.