COMMON READ LOG COMMAND
DRAFT DOCUMENT
LEVEL 3/18/87

Gregory G. Floryance

IBM Corporation
Storage Products Development
System Products Division
Hwy 52 & 37 Street NW
Rochester, Mn 55901
(507) 253-1869

TABLE OF CONTENTS

1.0 INTRODUCTION

1.1 BACKGROUND

The Mode Select page 2 parameters provide the capability to tune device performance on the SCSI bus. These parameters include Buffer Full/Empty Ratios, Bus Inactivity Limit and Disconnect/Connect Time Limits. In order to determine the proper parameter values, the initiator must have information available on device usage such as the number of read/write operations performed, average transfer size, etc. In a single initiator environment, it is possible for the initiator to log most of the necessary device usage information. However, in a multi-initiator environment, the ability to monitor bus activity caused by another initiator is very complex and impractical and inefficient. In addition, some critical device usage information is only known internally by the device and can not be deduced by the initiator through bus activity regardless of the number of initiators on the bus. Therefore, a means of reporting internal device usage to an initiator is needed if the potential of the page 2 Mode Select performance tuning parameters is to be fully realized.

Note: Bus performance tuning is not limited to development integration time. As the number and mix of devices are changed on a bus, the ability for a system to retune the operating characteristics of the attached devices (through occasional utilities or even in real time) becomes critically important. The ability for tuning a device is encouraged by the architected definition of the page 2 Mode Select parameters. However system performance tuning can not be fully utilized until an adequate reporting mechanism is architected and common parameter reporting formats are defined.

The Sequential access device READ LOG command provides the ability to report device internal information necessary for the Initiator to calculate the Mode Select parameter values. Currently, the command parameters are not defined which makes it's implementation vendor unique. In the case of DASD, the READ LOG command is not defined.

One piece of information critical to determining DASD performance characteristics is the device internal buffer overrun and underrun information. A buffer overrun on a read operation or a buffer underrun on a write operation will cause the transfer time for the command to be increased by a full revolution of the disk. Unlike read or write counters which can be maintained by a host, information on buffer over/under run counters is only available within the device itself. Therefore, a mechanism for reporting these values is necessary for DASD devices.

1.2 PROPOSAL SYNOPSIS

This proposal takes the existing Sequential Access device READ LOG command (1Fh) and makes it a common command for all device types. As currently specified, the contents of the data transferred is vendor unique with loose adherence to the CCS page concept. This proposal extends the page concept by defining page addressability within the CDB and also defines the structure of some pages which are needed to support the buffer over/under run counters in a common architected fashion.

2.0 READ LOG COMMAND

Peripheral Device Type: All
Operation Code Type: Optional

bit / byte	7	6	5	4	3	2	1	0
0	Op Code = (1Fh)							
1	LUN			PF	Reserved = 0			NLR
2	Page Code							
3 4	Allocation Length							
5	Vendor Uniq		Reserved = 0				Flag	Link

The READ LOG command is used to obtain statistical information maintained by the device about the device or attached medium.

Note: Changed the word Target to device.

The log data is available to the initiator at any time via the READ LOG command. Logging information can also be returned to the initiator on a REQUEST SENSE command in the REQUEST SENSE data.

Note: What does 'available .. at any time' mean. This should not be like the Inquiry command which executes through a Unit Atn.

Note: A request sense command can return anything at anytime following a CHECK CONDITION. Since there is no architected way to specifically ask for logging information in the Request Sense command, why bother with the 2nd sentence. This is as appropriate as adding a note that indicates physical cyl/hd/sector info is sometimes returned when a data error is reported. If you want, add a note to the Request Sense vendor unique field including both of these (and more).

The No log reset (NLR) bit, set to one, causes any log data to be transferred but not reset. This bit set to zero, causes the log data to be reset after the data has been transferred. The allocation length being zero and the NLR bit zero causes the log data to be reset without transferring any data to the initiator.

Note: Only reset all log data for PF=0 or PF=1 with Page Code 0. If Pf=1 and a page other than zero is specified, then only reset that pages log data.

The Allocation length indicates the number of bytes the initiator has set aside for the READ LOG data. The target shall terminate the data transfer when either the allocation length has been transferred or when all available log data has

been transferred to the initiator, whichever is less. An allocation length of zero is not considered an error.

Note: Reworded slightly

The Page Format (PF) bit of one indicates that the Read Log data conforms to page format. A value of zero indicates that the Read Log data is vendor unique and may or may not conform to page format.

Note: The PF bit may not be needed since this command was not defined in version 1. However, someone may have already implemented the command during the early version2 activity so leaving the bit in may be necessary. Also, note that this bit is shared by all paging commands and could be generally defined as the link and flag bits are.

The Page Code (PC) field specifies which page to return. The page codes are defined in Figure 1

Note: Page code zero is no longer vendor unique. This change may be OK if this command has not yet been implemented by someone. If it has, there should still be no problem since the PF bit had to be zero for the earlier version of SCSI-2. PF=0 could still be defined as vendor unique.

Page Code	Description
00h	Summary list of supported READ LOG pages
01h	Device Specific Log Data
02	Medium Specific Log Data
03	Buffer Over/Under run counters
04 - 7Fh	Rsvd
80 - FFh	Vendor Unique

The data returned by the device shall be in the same format as used for MODE SELECT/SENSE pages. (See 9.11 and 9.14). The buffer over/under run fields of page 3 should be at least 4 bytes in length.

131

2.1 SUMMARY LIST FORMAT (PF=1, PAGE CODE=0)

bit byte	7	6	5	4	3	2	1	0
0	Page Code = 0							
1	Parameter List Length							
2	1st supported page number							
.	.							
.	.							
1+n	Last supported page number							

This mode requests that a summary list of supported READ LOG pages for the device be sent to the Initiator. If any other pages are supported by the device, then support of page zero is mandatory.

Note: Typically, an initiator will first request page zero to determine the list of other supported pages within the device.

The Parameter List length field specifies the total number of supported pages. The count does not include bytes 0 or 1. The maximum value for this field is 254.

The supported page number list begins at byte 4 and contains a list of the supported pages for the device.

2.2 BUFFER OVER/UNDERRUN COUNTERS PAGE (PF=1, PAGE CODE=3)

bit / byte	7	6	5	4	3	2	1	0
0	Page Code = 3							
1	Parameter List Length							
2	Overrun Count							(MSB)
5								(LSB)
6	Underrun Count							(MSB)
9								(LSB)

This mode returns the device buffer overrun and underrun counters to the Initiator.

The OVERRUN and UNDERRUN COUNT fields contain the number of times a buffer overrun or underrun has occurred since the last time the counter was reset to zero. The count will be incremented for each occurrence of an underrun or overrun condition and can be incremented more than once for multiple occurrences during the execution of a single command.

Note: There is a big difference in an unsupported field of zero or a supported field of zero where no over/underruns have occurred. It may be necessary to either:

- provide a separate page for each counter
- provide sub-codes for each field
- require support for both or neither
- use a value of 'FFFFFFFF' to indicate invalid, counter overflow or not supported. (PREFERRED)

132

3.0 NOTES

Target or LUN addressability was considered and dropped for this command. It was first thought that Target/LUN addressability was needed for devices which may have buffers in both the the target and LUNs. However, at this time, the thrust of the command is driven by mechanical activities which would effect buffer over/underruns and not the management of other resources within the target or LUN. Further, if a target is designed as a bridge controller to separate buffered devices, it is the target design which should accommodate and massage information transferred through the READ LOG command so that performance tuning by the initiator should not have to consider an excessive number of variables in determining page 2 MODE SELECT values.

Other counter values to consider include:

- Number of Read operations performed
- Number of Write operations performed.
- Number of blocks read.
- Number of blocks written.
- Number of seeks > 2/3.
- Number of seeks > 1/3 < 2/3.
- Number of seeks > 1/6 < 1/3.
- Number of seeks > 1/12 and < 1/6.
- Number of seeks < 1/12.
- Number of zero length seeks.