

# Ordered SCSI Protocol

(T10/02-283r0)

Dave Peterson, Cisco Systems, Inc.

## Background

Today's SCSI protocol contains no support for the ordering of SCSI commands, even though there are ordering requirements. No ordering support works ok for parallel SCSI, but not for the serial SCSI transports (e.g., FCP, iSCSI). To resolve these ordering requirements the serial SCSI transports must implement ordering within the transport itself, thus adding additional complexity to the protocol. If ordering functionality was built into the SCSI protocol, transports could simply do what they are supposed to do, i.e., transport the data (reliably and guaranteed is a plus).

## FCP-2

To allow for ordering in FCP-2, the concept of a CRN was introduced into the FCP\_CMND IU as a single byte (to not require any change to the FCP\_CMND structure). A single byte will allow for a maximum of 255 outstanding commands, which is ok for today's (tape) applications, but is probably too small for the future. It should be noted for disk, if there are any ordering requirements, the application will simply issue the command and wait for the response. This behavior may introduce a significant performance penalty.

The original intent was to use CRN at the target/transport level, but it eventually was implemented at the lun/application client level. This was fine but there was no hook in place for the application client to supply a CRN, so the CRN was introduced into SAM-2. In addition, changes at the application client level are required to implement CRN.

## iSCSI

To allow for ordering in iSCSI, the concept of a CmdSN was introduced into the SCSI Command PDU as a 4-byte field. The CmdSN is a session-wide parameter across multiple connections and the target must deliver the command for processing in CmdSN order. As such, commands sent to one logical unit may stall waiting for a prior CmdSN to arrive for a different logical unit.

## Proposal

### Introduce command numbering into the SCSI protocol.

This can be implemented in the following manner:

A. Add a 4-byte (or 2-byte) command number to the CDB structure(s). The command reference number could also be placed ahead of the opcode field (i.e., my goal is to add ordering to the SCSI protocol).

Example CDB layout for an ERASE(16) command producing a 20-byte CDB:

**Table 1 — ERASE(20) command**

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (93h)							
1	Reserved			FCS		LCS	IMMED	LONG
2	Reserved							
3	PARTITION							
4	(MSB)							
5								
6								
7								
8	LOGICAL BLOCK ADDRESS							
9								
10								
11	(LSB)							
12	Reserved							
13	Reserved							
14	Reserved							
15	CONTROL							
16	(MSB)							
17								
18	COMMAND REFERENCE NUMBER							
19	(LSB)							

B. Use a bit in the control byte to indicate a 4-byte command number follows the CDB structure.

**Table 2 — ERASE(16) command**

Bit Byte	7	6	5	4	3	2	1	0	
0	OPERATION CODE (93h)								
1	Reserved			FCS	LCS	IMMED	LONG		
2	Reserved								
3	PARTITION								
4	(MSB)								
5									
6									
7									
8		LOGICAL BLOCK ADDRESS							
9									
10									
11								(LSB)	
12	Reserved								
13	Reserved								
14	Reserved								
15	Vendor Specific	Reserved			CRN	NACA	Obsolete	LINK	

Command Reference Number structure:

**Table 3 — Command Reference Number structure**

Bit Byte	7	6	5	4	3	2	1	0	
0	(MSB)								
1									
2		COMMAND REFERENCE NUMBER							
3								(LSB)	

C. Implement using the variable length CDB format.

**Table 4 — ERASE(16) service action**

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB)							
9	SERVICE ACTION (XXXXh)							(MSB)
10	(MSB)							
11	COMMAND REFERENCE NUMBER							
12								
13	(LSB)							
14	Reserved			FCS	LCS	IMMED	LONG	
15	Reserved							
16	PARTITION							
17	(MSB)							
18								
19								
20	LOGICAL BLOCK ADDRESS							
21								
22								
23								
24	(LSB)							
25	Reserved							
26	Reserved							
27	CONTROL							