

**Proposed
Draft**

**Serial ATA
International Organization**

**Version 7
21 July 2014**

**Serial ATA Revision 3.2 Technical Proposal 061
Title : Define/Use Sequential NCQ command(s)**

This is an internal working document of the Serial ATA International Organization. As such, this is not a completed standard and has not been approved. The Serial ATA International Organization may modify the contents at any time. This document is made available for review and comment only.

Permission is granted to the Promoters, Contributors and Adopters of the Serial ATA International Organization to reproduce this document for the purposes of evolving the technical content for internal use only without further permission provided this notice is included. All other rights are reserved and may be covered by one or more Non Disclosure Agreements including the Serial ATA International Organization participant agreements. Any commercial or for-profit replication or republication is prohibited. Copyright © 2000 to 2014 Serial ATA International Organization. All rights reserved.

This Draft Specification is NOT the final version of the Specification and is subject to change without notice. A modified, final version of this Specification ("Final Specification") when approved by the Promoters will be made available for download at this Web Site: <http://www.sata-io.org>.

THIS DRAFT SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE. Except for the right to download for internal review, no license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted or intended hereunder.

THE PROMOTERS DISCLAIM ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO USE OF INFORMATION IN THIS DRAFT SPECIFICATION. THE PROMOTERS DO NOT WARRANT OR REPRESENT THAT SUCH USE WILL NOT INFRINGE SUCH RIGHTS.

THIS DOCUMENT IS AN INTERMEDIATE DRAFT FOR COMMENT ONLY AND IS SUBJECT TO CHANGE WITHOUT NOTICE.

* Other brands and names are the property of their respective owners.

Copyright © 2000 to 2014 Serial ATA International Organization. All rights reserved.

Author Information

Author Name	Company	Email address
Ralph Weber	Western Digital	Ralph.Weber@WDC.com

Workgroup Chair Information

Workgroup (Phy, Digital, etc...)	Chairperson Name	Email address
Digital	James Hatfield	James.C.Hatfield@seagate.com

Document History

Version	Date	Comments
0	23 June 2014	Initial release.
1	25 June 2014	Incorporated comment from 23 June Digital call.
2	26 June 2014	At the request of the SATA editor, an excess instance of the spelled out 'Native Command Queuing' was sent to the scrapheap.
3	7 July 2014 am	Changed definition of "sequential NCQ command" as requested by the 30 June Digital call. As requested by James Boyd and Jim Hatfield, added "supported as sequential" bits to the NCQ send and receive log. Also added a comparison of SATA vs SCSI command ordering to the introduction.
4	7 July 2014 pm	Changed 'all' to 'any' definition of "sequential NCQ command" as requested by the 7 July Digital call.
5	8 July 2014	Added processing beginning/completion requirements to the definition of "sequential NCQ command" based on review comments from WD firmware developers.
6	21 July 2014	Remove highlighting for changes made in v5 to produce a clean revision for forwarding to the board.
7	21 July 2014	Member review. Changed TPR D188 to TPR 061.

Introduction

During the June T13 Plenary the Digital group chair referenced the SATA revision 3.2 Gold definition for “immediate NCQ command” and outlined a means by which the concept of an “immediate NCQ command” might be *enhanced* to create a solution for NCQ handling of General Purpose Logging operations that write logs in a required order (e.g., writing the SCT Command log).

This proposal instantiates the wise advice provided by the Digital group chair by defining a “sequential NCQ command” in a manner based upon but not identical to the “immediate NCQ command” definition.

For some, the result might (incorrectly) be viewed as creating the following relationships between NCQ and SCSI command queuing.

immediate NCQ command ↔ head of queue
sequential NCQ command ↔ ordered

The use of different names in SATA is intentional. A careful examination shows that NCQ behavior differs substantially from the so-called SCSI equivalent.

- The NCQ behaviors apply only in a SATA-specific queued environment, which is not universal in SATA and does not exist at all in PATA. The SCSI behaviors are specified in the architecture model and are universal to all SCSI transports.
- The NCQ behaviors apply only to those commands to which the SATA specification attaches them. The SCSI behaviors are specified in the architecture model and are applicable to all SCSI commands.
- The SCSI behaviors affect only when a command is delivered to the device server for processing. Device server action is outside the scope of the SCSI architecture model. SATA has no such layering of requirements and nothing appear to be outside the scope of the SATA specification.
- Processing of an immediate NCQ command does not begin until the processing has been completed for all previously accepted commands that have begun a DMA transfer. A SCSI head of queue command is delivered immediately to the device server where processing at the “usual pace”.
- Processing of a sequential NCQ command is ordered only with respect to other sequential NCQ commands. SCSI ordered commands are delivered to the device server in sequence with respect to all other commands received by the SCSI target device.

The motivations for the SATA definitions are clear and valid in the single-host realm of SATA. However, SCSI to ATA translation layers will have a practically impossible time employing immediate/sequential NCQ commands to effect the behaviors specified for head of queue and ordered commands by the SCSI architecture model.

Once defined, the “sequential NCQ command” is applied to the NCQ encapsulation of WRITE LOG DMA EXT.

Additions are shown in blue underline and removals in ~~red-strikeout~~.

1 Technical Specification Changes

1.1 Copy of “immediate NCQ command” glossary entry (for reference)

4.1.1.73 immediate NCQ command

An immediate native command queuing (NCQ) command is a NCQ command that shall be processed:

- a) after any command previously accepted by the device that the device has transmitted a DMA Setup FIS and has not reached command completion; and
- b) before any NCQ command previously accepted by the device that the device has not transmitted a DMA Setup FIS.

1.2 New “sequential NCQ command” glossary entry

4.1.1.xyz sequential NCQ command

A sequential NCQ command is a NCQ command that shall be processed:

- a) after any sequential NCQ commands previously accepted by the device have completed processing; and
- b) before any sequential NCQ commands subsequently accepted by the device have begun processing.

1.3 Application of “sequential NCQ command” to READ LOG DMA EXT

13.6.7 RECEIVE FPDMA QUEUED

13.6.7.1 RECEIVE FPDMA QUEUED definition

The 512 Byte Block DMA IN subcommands make use of this transport command. The RECEIVE FPDMA QUEUED command supports LBA mode only and uses 48 bit addressing only. The format of the command is defined in Figure 367.

Some RECEIVE FPDMA QUEUED subcommands are processed as sequential NCQ commands (see 4.1.1.xyz).

...

13.6.7.6 READ LOG DMA EXT subcommand (01h)

13.6.7.6.1 READ LOG DMA EXT overview

The READ LOG DMA EXT subcommand is a sequential NCQ command (see 4.1.1.xyz).

The READ LOG DMA EXT subcommand functionality and behavior is dependent on all requirements of the READ LOG DMA EXT command and the IDENTIFY DEVICE command defined in ACS-3.

1.4 Application of “sequential NCQ command” to WRITE LOG DMA EXT

13.6.8 SEND FPDMA QUEUED

13.6.8.1 SEND FPDMA QUEUED definition

The 512 Byte Block DMA OUT subcommands make use of this transport command. The SEND FPDMA QUEUED command supports LBA mode only and uses 48 bit addressing only. The format of the command is defined in Figure 372.

[Some SEND FPDMA QUEUED subcommands are processed as sequential NCQ commands \(see 4.1.1.xyz\).](#)

...

13.6.8.8 WRITE LOG DMA EXT subcommand (02h)

13.6.8.8.1 WRITE LOG DMA EXT overview

[The WRITE LOG DMA EXT subcommand is a sequential NCQ command \(see 4.1.1.xyz\).](#) The WRITE LOG DMA EXT subcommand functionality and behavior is dependent on all requirements of the WRITE LOG DMA EXT subcommand and the IDENTIFY DEVICE command defined in ACS-3.

1.5 Definition of support bits

13.7.6 NCQ send and receive log (13h)

13.7.6.1 NCQ send and receive log overview

To determine the supported SEND FPDMA QUEUED and RECEIVE FPDMA QUEUED subcommands and their respective features, host software may read log 13h.

This log shall be supported if the SEND FPDMA QUEUED and RECEIVE FPDMA QUEUED command is supported (i.e., IDENTIFY DEVICE data Word 77 bit 6 is set to one.)

Dword	Bits	Description
0	Subcommands Supported	
	31..2	Reserved
	1	Supports HYBRID EVICT (see 13.7.6.2)
	0	Supports Data Set Management (see 13.7.6.3)
1	Data Set Management	
	31..1	Reserved
	0	Supports Trim (see 13.7.6.4)
2	Supports Read Log	
	31..2 +	Reserved
	1	Supports Sequential Read Log (see 13.7.6.a)
	0	Supports Read Log (see 13.7.6.5)
3	Supports Write Log	
	31..2 +	Reserved
	1	Supports Sequential Write Log (see 13.7.6.b)
	0	Supports Write Log (see 13.7.6.6)
4..127		Reserved

Figure 383 – NCQ SEND and RECEIVE log (13) data structure definition

...

13.7.6.a Supports Sequential Read Log

If the Supports Sequential Read Log bit is set to one, then the device supports the READ LOG DMA EXT subcommand (see 13.6.7.6) of the RECEIVE FPDMA QUEUED command as a sequential NCQ command. If the Supports Sequential Read Log bit is cleared to zero, then the device does not support the READ LOG DMA EXT subcommand of the RECEIVE FPDMA QUEUED command as a sequential NCQ command.

If the Supports Sequential Read Log bit is set to one, then the Supports Read Log bit (see 13.7.6.5) shall be set to one.

...

13.7.6.b Supports Sequential Write Log

If the Supports Sequential Write Log bit is set to one, then the device supports the WRITE LOG DMA EXT subcommand (see 13.6.7.6) of the SEND FPDMA QUEUED command as a sequential NCQ command. If the Supports Sequential Write Log bit is cleared to zero, then the device does not support the WRITE LOG DMA EXT subcommand of the SEND FPDMA QUEUED command as a sequential NCQ command.

If the Supports Sequential Write Log bit is set to one, then the Supports Write Log bit (see 13.7.6.6) shall be set to one.