



NVM Express®

Subsystem Local Memory Command Set Specification

Revision 1.0

December 20th, 2023

Please send comments to info@nvmexpress.org

NVM Express® Subsystem Local Memory Command Set is available for download at <http://nvmexpress.org>. The NVM Express Subsystem Local Memory Command Set Specification Revision 1.0 incorporates TP4131.

SPECIFICATION DISCLAIMER

LEGAL NOTICE:

© Copyright 2008 to 2023 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express Subsystem Local Memory Command Set, Revision 1.0 is proprietary to the NVM Express, Inc. (also referred to as “Company”) and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express Subsystem Local Memory Command Set, Revision 1.0 subject, however, to the Member’s continued compliance with the Company’s Intellectual Property Policy and Bylaws and the Member’s Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: “© 2008 to 2023 NVM Express, Inc. ALL RIGHTS RESERVED.” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN “AS IS” BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

NVM Express Workgroup
c/o VTM, Inc.
3855 SW 153rd Drive
Beaverton, OR 97003 USA
info@nvmexpress.org

Table of Contents

1	INTRODUCTION.....	6
1.1	Overview	6
1.2	Scope.....	6
1.3	Conventions.....	7
1.4	Definitions	7
1.4.2	Definitions from the NVM Express Base Specifications	7
1.4.3	Definitions specific to the Subsystem Local Memory Command Set.....	7
1.5	References	7
2	SUBSYSTEM LOCAL MEMORY COMMAND SET MODEL.....	8
2.1	Theory of operation.....	8
2.1.1	Namespaces	8
2.1.2	Storage Entities.....	9
2.1.3	Flush Command.....	9
2.1.4	Command Ordering.....	9
2.1.5	Fused Operation	9
2.1.6	Atomic Operation.....	10
2.1.7	Metadata Region (MR)	10
2.2	I/O Controller Requirements.....	10
2.2.1	Command Support	10
2.2.2	Log Page Support	10
2.2.3	Features Support	11
3	I/O COMMANDS FOR THE SUBSYSTEM LOCAL MEMORY COMMAND SET.....	12
3.1	Submission Queue Entry and Completion Queue Entry	12
3.1.1	Common Command Format.....	12
3.1.2	SLM Command Set Specific Status Values	12
3.2	SLM Command Set Commands.....	12
3.2.1	Memory Copy command	13
3.2.2	Memory Fill command.....	19
3.2.3	Memory Read command.....	20
3.2.4	Memory Write command	21
4	ADMIN COMMANDS FOR THE SUBSYSTEM LOCAL MEMORY COMMAND SET.....	22
4.1	Admin Command behavior for the SLM Command Set	22
4.1.2	Asynchronous Event Request command	22

4.1.3	Format NVM command	22
4.1.4	Get Features & Set Features commands	22
4.1.5	Get Log Page command	22
4.1.6	Identify Command	22
4.1.7	Namespace Attachment command	24
4.1.8	Namespace Management command.....	24
4.1.9	Sanitize command.....	25
4.2	I/O Command Set Specific Admin commands	25
5	EXTENDED CAPABILITIES.....	26

Table of Figures

Figure 1: NVMe Family of Specifications	6
Figure 2: I/O Controller – SLM Command Set Support	10
Figure 3: I/O Controller – Log Page Support.....	11
Figure 4: Status Code – Command Specific Status Values, SLM Command Set	12
Figure 5: Opcodes for SLM Command Set Commands	13
Figure 6: Memory Copy – Data Pointer.....	13
Figure 7: Memory Copy – Command Dword 2 and Dword 3.....	14
Figure 8: Memory Copy – Command Dword 10 and Dword 11.....	14
Figure 9: Memory Copy – Command Dword 12.....	14
Figure 10: Memory Copy – Copy Descriptor Format types	15
Figure 11: Memory Copy – Source Range Entries for Copy Descriptor Format 4h.....	16
Figure 12: SLM NS Source and SLM NS Destination Relationship Example.....	18
Figure 13: LBA NS Source and SLM NS Destination Relationship Example	18
Figure 14: Memory Copy – Command Specific Status Values.....	19
Figure 15: Memory Fill – Data Pointer.....	19
Figure 16: Memory Fill – Command Dword 10, Command Dword 11	20
Figure 17: Memory Fill – Command Dword 12.....	20
Figure 18: Memory Read – Data Pointer.....	20
Figure 19: Memory Read – Command Dword 10, Command Dword 11	20
Figure 20: Memory Read – Command Dword 12.....	20
Figure 21: Memory Write – Data Pointer.....	21
Figure 22: Memory Write – Command Dword 10, Command Dword 11	21
Figure 23: Memory Write – Command Dword 12.....	21
Figure 24: Identify – CNS Values.....	22
Figure 25: I/O Command Set Specific Identify Namespace data structure for the SLM Command Set	23
Figure 26: Format Data Structure for the SLM Command Set	24
Figure 27: Identify Controller Data Structure, SLM Command Set Dependent - General	24
Figure 28: SLM Command Set Specification Version Descriptor Field Values	24

1 Introduction

1.1 Overview

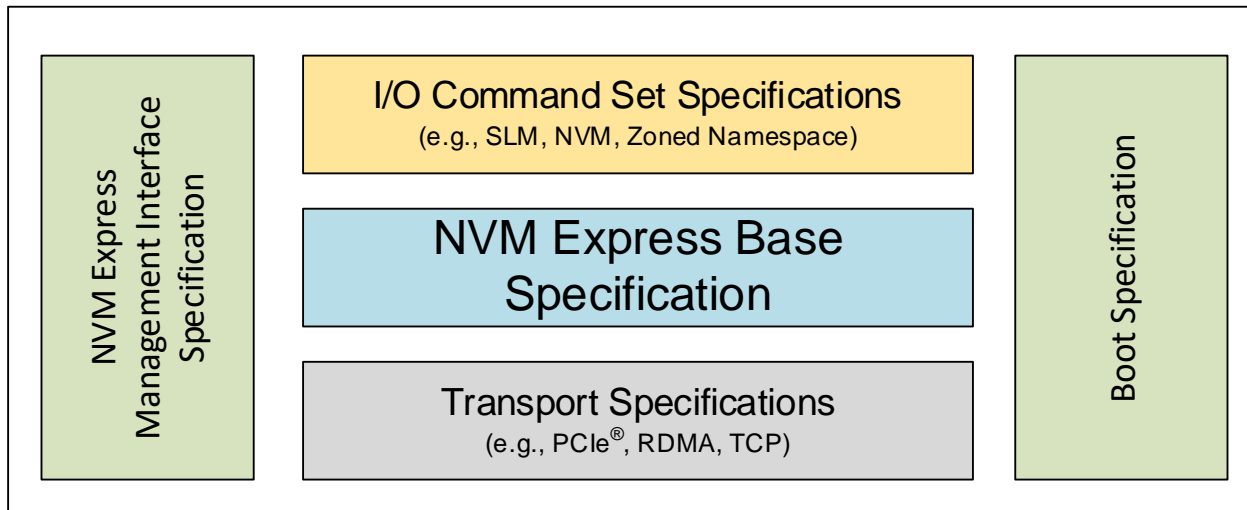
NVM Express® (NVMe®) Base Specification defines an interface for host software to communicate with a non-volatile memory subsystems over a variety of memory-based transports and message-based transports.

This document defines a specific NVMe I/O Command Set, the Subsystem Local Memory Command Set, which extends the NVM Express Base Specification.

1.2 Scope

Figure 1 shows the relationship of the Subsystem Local Memory Command Set Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



This specification supplements the NVM Express Base Specification. This specification defines additional data structures, Features, log pages, commands, and status values. This specification also defines extensions to existing data structures, features, log pages, commands, and status values. This specification defines requirements and behaviors that are specific to the Subsystem Local Memory Command Set. Functionality that is applicable generally to NVMe or that is applicable across multiple I/O Command Sets is defined in the NVM Express Base Specification.

If a conflict arises among requirements defined in different specifications, then a lower-numbered specification in the following list shall take precedence over a higher-numbered specification:

1. Non-NVMe specifications
2. NVM Express Base Specification
3. NVMe transport specifications
4. NVMe I/O command set specifications
5. NVMe-MI Specification

1.3 Conventions

This specification conforms to the Conventions section, Keywords section, and Byte, Word, and Dword Relationships section of the NVM Express Base Specification.

1.4 Definitions

1.4.2 Definitions from the NVM Express Base Specifications

This specification uses the definitions in the NVM Express Base Specification.

1.4.3 Definitions specific to the Subsystem Local Memory Command Set

This section defines terms that are specific to this specification.

1.4.3.1 Subsystem Local Memory (SLM)

Memory in an NVM subsystem that is accessible by commands defined in this specification and accessed by other I/O Command Sets (e.g., the Computational Programs Command Set).

1.4.3.2 memory namespace

A namespace that uses the Subsystem Local Memory Command Set.

1.4.3.3 compute namespace

A namespace that uses the Computational Programs Command Set.

1.5 References

NVM Express Base Specification, Revision 2.0, 02 June 2021. Available from <https://www.nvmexpress.org>.

2 Subsystem Local Memory Command Set Model

The NVM Express Base Specification defines a register level interface for host software to communicate with a non-volatile memory subsystem. This specification defines additional functionality for the Subsystem Local Memory (SLM) Command Set.

Each I/O Command Set is assigned a specific Command Set Identifier (CSI) value by the NVM Express Base Specification. The SLM Command Set is assigned a CSI value of 03h.

2.1 Theory of operation

An NVM subsystem may contain controllers that implement the SLM Command Set. The SLM Command Set is an I/O Command Set designed to provide host access to memory residing in an NVM subsystem. SLM storage is measured in bytes. This capability has multiple applications, including the movement of data used by Computational Programs.

A memory namespace contains volatile byte addressable memory and is associated with the SLM Command Set. The data in a memory namespace may be accessed by:

- Using commands in this command set;
- A compute namespace as described in the Computational Programs Command Set; or
- Any I/O Command Set that both specifies logical blocks and includes the Copy command specified by the NVM Command Set (e.g., the NVM Command Set, Zoned Namespace Command Set).

Unlike Controller Memory Buffer (CMB) and Persistent Memory Region (PMR) as defined by the NVM Express Base Specification, SLM is not mapped to a PCIe BAR and is only accessed by the methods described in this section.

Commands provided by the SLM Command Set are dword granular and dword aligned. This means that host accesses are required to be aligned on a dword boundary and the length is required to be some multiple of dwords. SLM is byte addressable by operations associated with the Computational Programs Command Set.

This command set has the following key attributes:

- Namespace accesses are dword addressable with dword granularity.
- Provides commands to transfer data between:
 - Host and SLM;
 - SLM and NVM; and
 - SLM and SLM.
- Provides a command to fill SLM.

2.1.1 Namespaces

A namespace is a set of resources that may be accessed by a host. In the case of SLM, the resources are memory accessed with memory semantics. Namespaces are as defined in the NVM Express Base Specification.

2.1.1.1 Reachability

This command set applies Reachability as defined in the NVM Express Base Specification to the operation of the Memory Copy command.

Note: Until the release of NVM Express Base specification following version 2.0, Reachability material is contained in TP4156.

2.1.1.2 Namespace Management

A Namespace Attachment command, if supported, shall support attaching namespaces associated with the Subsystem Local Memory command set (i.e., CSI 03h). Creation of namespaces associated with this I/O Command Set is outside the scope of this specification. A Namespace Management command, if supported, that specifies a namespace associated with the Subsystem Local Memory command set (i.e., CSI 03h) shall be aborted by the controller with a status code of Invalid I/O Command Set.

2.1.1.3 Interactions with Other Namespaces

In addition to access via this command set, data stored in a memory namespace may also be accessed by operations in another namespace (e.g., using a Memory Range Set in a compute namespace (refer to Computational Programs Command Set Specification), or using the Copy command in an NVM namespace (refer to NVM Command Set)).

A Memory Read command, Memory Write command, Memory Fill command, or Memory Copy command that access memory that is simultaneously being accessed by another namespace results in indeterminate behavior. It is expected that the host manages memory namespace accesses to prevent this condition.

Compute resources in a compute namespace may perform byte-addressable and byte-granular reads or writes to data in a memory namespace.

2.1.2 Storage Entities

Memory namespaces shall not be part of any Endurance Group or NVM Set as defined in the NVM Express Base Specification.

2.1.3 Flush Command

The Flush command (refer to the NVM Express Base Specification) shall have no effect on memory namespaces.

2.1.4 Command Ordering

Each command is processed as an independent entity without reference to other commands submitted to the same I/O Submission Queue or to commands submitted to other I/O Submission Queues. Specifically, the controller is not responsible for enforcing ordering between I/O commands. For example, if a Memory Write command to address x is submitted and then a Memory Read command for address x is submitted, there is no guarantee of the order of completion for those commands. If there are ordering requirements between these commands, the host enforces those requirements by submitting the Memory Read command after the Memory Write command completes.

2.1.5 Fused Operation

This command set does not support any Fused Operations.

2.1.6 Atomic Operation

This command set specifies a size in bytes of a write operation guaranteed to be written atomically under normal operation. No atomicity is guaranteed for power fail operation since memory namespaces contain only volatile memory.

- Accesses via I/O commands in this command set (e.g., accesses by overlapping write operations or between overlapping write and read operations) shall be atomic at a dword level.
- Concurrent accesses by overlapping I/O commands and compute operations by the Computational Programs Command Set shall be atomic at a byte level.
- Additional command set specific requirements are described in the appropriate I/O Command Set specification (e.g., Computational Programs Command Set Specification).

2.1.7 Metadata Region (MR)

Memory namespaces shall not support Metadata Regions.

2.2 I/O Controller Requirements

2.2.1 Command Support

This command set implements the command support requirements for I/O Controllers defined in the NVM Express Base Specification. Additionally, Figure 2 defines which commands are mandatory, optional, and prohibited for an I/O controller that supports this command set.

Figure 2: I/O Controller – SLM Command Set Support

Command	Command Support Requirements ¹
Memory Read	M
Memory Write	M
Memory Fill	O
Memory Copy	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

2.2.2 Log Page Support

This specification implements the log page support requirements for I/O Controllers defined in the NVM Express Base Specification. Figure 3 defines additional requirements for NVM Express Base Specification log pages that are mandatory, optional, and prohibited for an I/O controller that supports this command set.

Figure 3: I/O Controller – Log Page Support

Log Page Name	Log Page Support Requirements ¹
Reachability Groups ²	M
Reachability Associations ²	M
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. Defined in the NVM Express Base Specification Note: Until the release of NVM Express Base specification following version 2.0, Reachability material is contained in TP4156.	

2.2.3 Features Support

This specification implements the feature support requirements for I/O Controllers defined in the NVM Express Base Specification.

3 I/O Commands for the Subsystem Local Memory Command Set

This section specifies the I/O commands for this command set.

3.1 Submission Queue Entry and Completion Queue Entry

The Submission Queue Entry (SQE) data structure and the fields that are common to all NVMe I/O Command Sets are defined in the Submission Queue Entry – Command Format section in the NVM Express Base Specification. The Completion Queue Entry (CQE) data structure and the fields that are common to all NVMe I/O Command Sets are defined in the Completion Queue Entry section in the NVM Express Base Specification. The command specific fields in the SQE and CQE structures (i.e., Command Dword 2, Dword 3, and Dwords 10-15 of the SQE; and Dword 0, and Dword 1 of the CQE) for this command set are defined in this section.

3.1.1 Common Command Format

All commands in this command set follow the Common Command Format, as defined in the NVM Express Base Specification.

3.1.2 SLM Command Set Specific Status Values

Figure 4: Status Code – Command Specific Status Values, SLM Command Set

Value	Description	Commands Affected
88h	<p>Namespace Not Reachable: One or more of the Memory Copy Source Ranges specifies an NSID that is not a member of a Reachability Association (refer to the NVM Express Base Specification) with the destination namespace.</p> <p>Note: Until the release of NVM Express Base specification following version 2.0, Reachability material is contained in TP4156.</p>	Memory Copy

3.2 SLM Command Set Commands

This command set includes the commands listed in Figure 5. Commands are submitted as described in the NVM Express Base Specification.

Figure 5: Opcodes for SLM Command Set Commands

Opcode by Field			Combined Opcode ¹	Command ²	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer ³			
0b	000 00b	00b	00h	Memory Fill	3.2.2
0b	000 00b	01b	01h	Memory Copy	3.2.1
0b	000 00b	10b	02h	Memory Read	3.2.3
0b	000 01b	01b	05h	Memory Write	3.2.4
Notes: 1. Opcodes not listed are defined in the NVM Express Base Specification. 2. All SLM Command Set commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field. 3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.					

3.2.1 Memory Copy command

The Memory Copy command is used by the host to copy data from one or more source ranges in one or more source namespaces to a single consecutive address range in a destination memory namespace. Each source range may be in the same namespace or a different namespace with respect to any other source range and with respect to the destination memory namespace (i.e., the NSID in the Submission Queue Entry for the Memory Copy command) address range. All source ranges have the same Source Range Entries Descriptor type, and therefore all source ranges access data from the same type of namespace (e.g., all source namespaces are LBA based namespaces, or all source namespaces are memory namespaces).

If a Memory Copy command requests a copy between different NSIDs (i.e., the NSID in a Source Range is different than the NSID of the destination namespace (i.e., the NSID in the Submission Queue Entry for the Memory Copy command)) and there is no Reachability Association (refer to the NVM Express Base Specification) between those namespaces, then the command shall be aborted with the status code of Namespace Not Reachable. Note: Until the release of NVM Express Base specification following version 2.0, Reachability material is contained in TP4156.

This command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, and Command Dword 12 fields. If the command uses PRPs for the data transfer, then the PRP Entry 1 and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the SGL Entry 1 field is used. All other command specific fields are reserved.

Figure 6: Memory Copy – Data Pointer

Bits	Description
127:00	Data Pointer (DPTR): This field specifies the address of the data (i.e., Source Range Entries Descriptors, refer to Figure 11) to use for the command. Refer to the Common Command Format figure in the NVM Express Base Specification for the definition of this field.

Figure 7: Memory Copy – Command Dword 2 and Dword 3

Bits	Description
63:00	Length (LEN): This field specifies the number of bytes to transfer from the specified Source Range Entries (refer to Figure 11) to the destination. The value provided by this field shall be dword granular (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

Figure 8: Memory Copy – Command Dword 10 and Dword 11

Bits	Description
63:00	Starting Destination Address (SDADDR): This field specifies the byte address of the first location of the memory namespace to be written as part of the copy operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32. The value provided by this field shall be dword aligned (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

Figure 9: Memory Copy – Command Dword 12

Bits	Description														
31:12	Reserved														
11:08	Descriptor Format: Specifies the type of the Copy Descriptor Format that is used. The Copy Descriptor Format specifies the format of the Copy Descriptor which contains starting address, number of bytes, and parameters associated with the read portion of the operation.														
	<table border="1"> <thead> <tr> <th>Copy Descriptor Format type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Not supported¹</td> </tr> <tr> <td>1h</td> <td>Not supported¹</td> </tr> <tr> <td>2h</td> <td>Source Range Entries Copy Descriptor Format 2h is used. Refer to the NVM Command Set Specification for the description of this Source Range Entry.</td> </tr> <tr> <td>3h</td> <td>Source Range Entries Copy Descriptor Format 3h is used. Refer to the NVM Command Set Specification for the description of this Source Range Entry.</td> </tr> <tr> <td>4h²</td> <td>Source Range Entries Copy Descriptor Format 4h is used. Refer to Figure 10 for the description of this Source Range Entry.</td> </tr> <tr> <td>All Others</td> <td>Reserved</td> </tr> </tbody> </table>	Copy Descriptor Format type	Description	0h	Not supported ¹	1h	Not supported ¹	2h	Source Range Entries Copy Descriptor Format 2h is used. Refer to the NVM Command Set Specification for the description of this Source Range Entry.	3h	Source Range Entries Copy Descriptor Format 3h is used. Refer to the NVM Command Set Specification for the description of this Source Range Entry.	4h ²	Source Range Entries Copy Descriptor Format 4h is used. Refer to Figure 10 for the description of this Source Range Entry.	All Others	Reserved
	Copy Descriptor Format type	Description													
	0h	Not supported ¹													
	1h	Not supported ¹													
	2h	Source Range Entries Copy Descriptor Format 2h is used. Refer to the NVM Command Set Specification for the description of this Source Range Entry.													
	3h	Source Range Entries Copy Descriptor Format 3h is used. Refer to the NVM Command Set Specification for the description of this Source Range Entry.													
4h ²	Source Range Entries Copy Descriptor Format 4h is used. Refer to Figure 10 for the description of this Source Range Entry.														
All Others	Reserved														
07:00	Number of Ranges (NR): Specifies the number of Source Range Entries that are specified in the command. This is a 0's-based value.														
Notes:															
<ol style="list-style-type: none"> 1. This descriptor type is used for copy operations with LBA based destinations as defined in the NVM Command Set Specification. 2. If the Memory Copy command is supported, this format type shall be supported. 															

The controller shall indicate the Source Range Entries Copy Descriptor formats supported by the controller in the Copy Descriptor Formats Supported field in the Identify Controller data structure (refer to the NVM Express Base Specification).

Controller usage of Source Range Entries Copy Descriptor Formats 2h, 3h, and 4h is further qualified by whether the host has enabled these formats in the Host Behavior Support feature (refer to the NVM Express Base Specification). If the controller supports a Source Range Entries Copy Descriptor Format that has not been enabled, the controller shall process Copy commands as if that format is not supported (e.g., if that format is specified in the Descriptor Format field in Command Dword 12, the controller shall abort the command with a status code of Invalid Field in Command). A host that enables Source Range Entries Copy Descriptor Formats 2h, 3h, and/or 4h indicates to the controller that the host accepts the implications (e.g., for namespace access control) of the presence of a Source Namespace Identifier (SNSID) in these formats.

The data that the Memory Copy command provides (refer to the DPTR field in Figure 6) contains a list of Source Range Entries that describe the data to be copied to the destination range starting at the SDADDR. The Copy Descriptor Format type of the Source Range Entries is specified in the Descriptor Format field in Command Dword 12. If the Copy Descriptor Format specified in the Descriptor Format field is not supported by the specified namespace to which the Memory Copy command is addressed as defined by the NSID field, then the command shall be aborted with a status code of Invalid Field in Command.

Figure 10: Memory Copy – Copy Descriptor Format types

Copy Descriptor Format type	Description
0h	Not Supported
1h	
2h	Source data is described using LBAs (i.e., the source namespace supports an I/O Command Set that supports logical blocks) (refer to the NVM Command Set for details). The source namespace is an LBA based namespace.
3h	
4h	Source data is described using byte addressing (i.e., the source device supports the SLM Command Set). Protection Information is not supported, source namespace may differ from destination namespace.

The format for Source Range Entry type 2h and type 3h are described in the NVM Command Set Specification.

If the Copy Descriptor Format Type is set to 2h or to 3h, then the SNSID field in each Source Range Entries descriptor should not:

- match the value of the NSID field;
- be set to FFFFFFFFh; and
- be cleared to 0h.

If Copy Descriptor Format Type is set to 2h or to 3h and:

- an SNSID field matches the value of the NSID field, then the Memory Copy command shall be aborted with a status code set Incompatible Namespace or Format; and
- an SNSID field is set to FFFFFFFFh, or is cleared to 0h, then the Memory Copy command shall be aborted with a status code set to Invalid Field in Command.

Figure 11 shows the Copy Descriptor Format 4h descriptor and an example with 128 Source Range entries.

Figure 11: Memory Copy – Source Range Entries for Copy Descriptor Format 4h

Range	Bytes	Description						
Source Range 0	03:00	Source Namespace Identifier (SNSID): Specifies the source namespace for this Source Range entry.						
	07:04	Reserved						
	15:08	Starting Address (SADDR): The byte offset from the start of the namespace specified by the SNSID field to the address of the source data. The value provided in this field shall be dword aligned (i.e. bits 1:0 cleared to 00b). The command shall be aborted with a status code of Invalid Field in Command if bits 1:0 are not cleared to 00b.						
	19:16	Number of Bytes (NBYTE): The number of bytes that are to be transferred. The value provided in this field shall be dword granular (i.e. bits 1:0 cleared to 00b). The command shall be aborted with a status code of Invalid Field in Command if bits 1:0 are not cleared to 00b.						
	21:20	Reserved						
	23:22	Source Options:						
		<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15</td> <td>Fast Copy Only (FCO): If set to '1' the controller only performs fast copy operations (refer to the "Fast copy operations" section in the NVM Command Set Specification) for user data in this Source Range. If cleared to '0', this bit has no effect.</td> </tr> <tr> <td>14:00</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	15	Fast Copy Only (FCO): If set to '1' the controller only performs fast copy operations (refer to the "Fast copy operations" section in the NVM Command Set Specification) for user data in this Source Range. If cleared to '0', this bit has no effect.	14:00	Reserved
Bits		Description						
15	Fast Copy Only (FCO): If set to '1' the controller only performs fast copy operations (refer to the "Fast copy operations" section in the NVM Command Set Specification) for user data in this Source Range. If cleared to '0', this bit has no effect.							
14:00	Reserved							
	Reserved							
	31:24	Reserved						
Source Range 1	35:32	SNSID						
	39:36	Reserved						
	47:40	SADDR						
	51:48	NBYTE						
	55:54	Source Options						
	63:56	Reserved						
		...						
Source Range 127	4067:4064	SNSID						
	4071:4068	Reserved						
	4079:4072	SADDR						
	4083:4080	NBYTE						
	4087:4086	Source Options						
	4095:4088	Reserved						

If the number of Source Range entries (i.e., the value in the NR field) is greater than the value in the MSRC field (refer to Figure 25), then the Memory Copy command shall be aborted with a status code of Command Size Limit Exceeded.

If the value in the Length field (i.e., Command Dword 2 and Dword 3) is greater than the value in the MCMCL field (refer to Figure 25), then the Memory Copy command shall be aborted with a status code of Command Size Limit Exceeded.

For a Copy Descriptor Format Type set to 2h or to 3h:

- if a valid Source Range Entry specifies a Number of Logical Blocks field that represents a number of bytes (i.e., as determined by the LBA Data Size field (refer to the NVM Command Set Specification) of the LBA Format Data Structure associated with the namespace specified by the

SNID field) that is greater than the value in the MCMSSRL field (refer to Figure 25), then the Memory Copy command shall be aborted with a status code of Command Size Limit Exceeded;

- if the sum of all Number of Logical Blocks fields in all Source Range Entries represents a number of bytes (i.e., as determined by the LBA Data Size field associated with the namespace specified by the SNID field) that is greater than the value in the MCMCL field (refer to Figure 25), then the Memory Copy command shall be aborted with a status code of Command Size Limit Exceeded; and
- if the sum of all Number of Logical Blocks fields in all Source Range Entries represents a number of bytes to be written (i.e., as determined by the Data Size field (refer to Figure 26) associated with the namespace to which the command is sent) that does not match the value in the Length field, then the Memory Copy command shall be aborted with a status code of Invalid Field in Command.

For a Copy Descriptor Format Type set to 4h:

- If a valid Source Range Entry specifies a Number of Bytes field that is greater than the value in the MCMSSRL field (refer to Figure 25), then the Memory Copy command shall be aborted with a status code of Command Size Limit Exceeded;
- If the sum of all Number of Bytes fields in all Source Range entries is greater than the value in the MCMCL field (refer to Figure 25), then the Memory Copy command shall be aborted with a status code of Command Size Limit Exceeded;
- If the sum of all Number of Bytes fields in all Source Range entries does not match the value in the Length field, then the Memory Copy command shall be aborted with a status code of Invalid Field in Command; and
- If the NSID specified by the SNSID field in the Source Range Entry matches the namespace specified by the NSID field and addresses referenced by that Source Range Entry (i.e., SADDR field and NBYTE field) overlap with the address range specified for the destination (i.e., SDADDR field and LEN field), then the Memory Copy command shall be aborted with a status code of Overlapping I/O Range.

The number of Bytes requested to be written by the Memory Copy command is specified in Command Dword 2 and Dword 3. The data bytes specified by each Source Range Entry shall be copied to the destination byte range in the same order those bytes are listed in the Source Range entries (e.g., the bytes specified by Source Range entry 0 are copied to the lowest numbered byte address specified by the SDADDR field, the Bytes specified by Source Range entry 1 are copied to the next consecutive bytes specified by the SDADDR field). The read operations and write operations used to perform the copy may operate sequentially or in parallel. As a result, a Source Range Entries descriptor that contains an error may not be detected until after other Source Range Entries descriptors have been processed.

Figure 12 shows an example of the result of a Memory Copy command from memory namespaces to a memory namespace.

Figure 12: SLM NS Source and SLM NS Destination Relationship Example

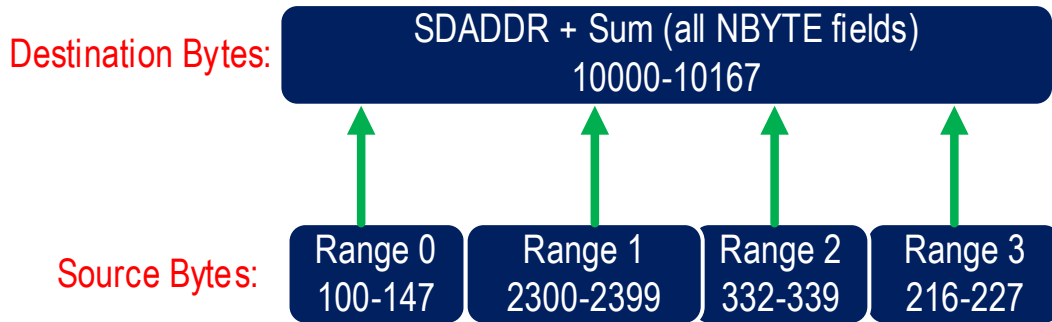



Figure 13 shows an example of the result of a Memory Copy command from LBA based namespaces formatted with 512 bytes per logical block to a memory namespace.

Figure 13: LBA NS Source and SLM NS Destination Relationship Example



For a Copy Descriptor Format Type set to 2h or to 3h:

If the source namespace for Source Range 0 is formatted with protection information (PI) and the PRINFOR.PRACT bit in Source Range 0 is set to '1', then for each Source Range, including Source Range 0:

- if the source namespace for that Source Range is formatted with PI, then the controller shall perform the data copying specified by the Copy command for that Source Range with PI stripped as specified in the NVM Command Set Specification Protection Information and Copying Across Different Namespaces section (note: until the release of the NVM Command Set Specification following version 2.0, this functionality is described in TP4130),

and

If the source namespace for Source Range 0 is formatted with protection information (PI), and the PRINFOR.PRACT bit in Source Range 0 is cleared to '0', then the controller shall abort the command with a status code of Incompatible Namespace or Format.

3.2.1.1 Command completion

Upon completion of the command, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

If the command completes with failure, then Dword 0 of the completion queue entry shall be set to the number of the lowest numbered Source Range entry that was not successfully copied (e.g., if Source Range 0, Source Range 1, Source Range 2, and Source Range 5 are copied successfully and Source Range 3 and Source Range 4 are not copied successfully, then Dword 0 is set to 3). If no data was written to the destination location, then Dword 0 of the completion queue entry shall be cleared to 0h.

Memory Copy command specific errors are defined in Figure 14.

Figure 14: Memory Copy – Command Specific Status Values

Value	Description
81h	<p>Invalid Protection Information: If the Copy Descriptor Format Type is set to 2h or to 3h and the protection information specified by the command is invalid due to:</p> <ul style="list-style-type: none"> The Protection Information Field Read (PRINFOR) field containing an invalid value for the Protection Information with which the source namespace was formatted (refer to the PI field in the Format NVM Command section in the NVM Express Base Specification and the DPS field in NVM Command Set Specification); or the EILBRT field in a Source Range Entry being invalid (refer to the NVM Command Set Specification).
83h	<p>Command Size Limit Size Exceeded: One or more of the Memory Copy command processing limits (i.e., non-zero value of the NR, MSSRL, and MCL fields in the Identify Namespace data structure (refer to Figure 25) was exceeded.</p>
85h	<p>Incompatible Namespace or Format: At least one source namespace and the destination namespace have incompatible formats (refer to the NVM Command Set Specification).</p>
86h	<p>Fast Copy Not Possible: The Fast Copy Only (FCO) bit was set to '1' in a Source Range entry and the controller was not able to use fast copy operations to copy the specified data (refer to the NVM Command Set Specification).</p>
87h	<p>Overlapping I/O Range: A source address range overlaps the destination address range.</p>
88h	<p>Namespace Not Reachable: One or more of the Memory Copy Source Ranges specifies an NSID that is not a member of Reachability Association (refer to the NVM Express Base Specification) with the destination namespace.</p>

3.2.2 Memory Fill command

The Memory Fill command fills a memory range of the specified namespace with zeros. The memory range is specified by the SB field and the FL field. Memory Fill commands are required to be dword aligned and of dword granularity.

The command uses Command Dword 10, Command Dword 11, and Command Dword 12.

Figure 15: Memory Fill – Data Pointer

Bits	Description
127:00	<p>Data Pointer (DPTR): This field is not used and shall be ignored by the controller.</p>

Figure 16: Memory Fill – Command Dword 10, Command Dword 11

Bits	Description
63:00	Starting Byte (SB): This field specifies the first byte of the memory range in the specified namespace to fill. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32. The value provided by this field shall be dword aligned (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

Figure 17: Memory Fill – Command Dword 12

Bits	Description
31:00	Fill Length (FL): This field specifies the number of bytes in the memory range to be filled. The value provided by this field shall be dword granular (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

3.2.2.1 Command completion

When the command is completed, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

3.2.3 Memory Read command

The Memory Read command reads data from the specified namespace. Memory Read commands are required to be dword aligned and of dword granularity.

The command uses Command Dword 10, Command Dword 11, and Command Dword 12.

Figure 18: Memory Read – Data Pointer

Bits	Description
127:00	Data Pointer (DPTR): This field specifies where data is transferred to. Refer to the Common Command Format figure in the NVM Express Base Specification for the definition of this field.

Figure 19: Memory Read – Command Dword 10, Command Dword 11

Bits	Description
63:00	Starting Byte (SB): This field specifies the first byte in the specified namespace to read. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32. The value provided by this field shall be dword aligned (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

Figure 20: Memory Read – Command Dword 12

Bits	Description
31:00	Read Length (RL): This field specifies the read length in bytes. The value provided by this field shall be dword granular (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

3.2.3.1 Command completion

When the command is completed, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

3.2.4 Memory Write command

The Memory Write command writes data to the specified namespace. Memory Write commands are required to be dword aligned and of dword granularity. The Namespace Optimal Write Granularity (NOWG) field in the I/O Command Set specific Identify Namespace data structure (CNS 05h, CSI 03h) may be used to discover optimal write granularity for a namespace.

The command uses Command Dword 10, Command Dword 11, and Command Dword 12.

Figure 21: Memory Write – Data Pointer

Bits	Description
127:00	Data Pointer (DPTR): This field specifies where data is transferred from. Refer to the Common Command Format figure in the NVM Express Base Specification for the definition of this field.

Figure 22: Memory Write – Command Dword 10, Command Dword 11

Bits	Description
63:00	Starting Byte (SB): This field specifies the first byte in the namespace to write. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32. The value provided by this field shall be dword aligned (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

Figure 23: Memory Write – Command Dword 12

Bits	Description
31:00	Write Length (WL): This field specifies the write length in bytes. The value provided by this field shall be dword granular (i.e., bits 1:0 cleared to 00b). If bits 1:0 are non-zero, then the controller shall abort the command with a status code of Invalid Field in Command.

3.2.4.1 Command completion

When the command is completed, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

4 Admin Commands for the Subsystem Local Memory Command Set

4.1 Admin Command behavior for the SLM Command Set

The Admin commands are as defined in the NVM Express Base Specification. The behavior for Admin commands defined in this command set is described in this section.

4.1.2 Asynchronous Event Request command

The Asynchronous Event Request command operates as defined in the NVM Express Base Specification.

4.1.3 Format NVM command

The Format NVM command operates as defined in the NVM Express Base Specification. The Format Index indicates a valid SLM Format from the SLM Format field in the SLM I/O Command Set specific Identify Namespace data structure (Figure 25).

4.1.4 Get Features & Set Features commands

This command set has no I/O Command Set specific features, and does not create any exceptions or changes from features defined in the NVM Express Base Specification.

4.1.5 Get Log Page command

This command set does not define any log pages.

4.1.6 Identify Command

This specification implements the Identify Command and associated Identify data structures defined in the NVM Express Base Specification. Additionally, this command set specifies the data structures defined in this section.

Each I/O Command Set is assigned a specific Command Set Identifier (CSI) value by the NVM Express Base Specification. This command set is assigned a CSI value of 03h.

Figure 24: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
05h	M ⁵	Identify I/O Command Set specific Namespace data structure for the specified NSID for the I/O Command Set specified in the CSI field.	Y	N	Y	4.1.6.1
06h	M	Identify I/O Command Set Specific Controller data structure for the controller processing the command.	Y	N	Y	4.1.6.2
Notes: 1. O/M definition: O = Optional, M = Mandatory. 2. The NSID field is used: Y = Yes, N = No. 3. The CDW10.CNTID field is used: Y = Yes, N = No. 4. The CDW11.CSI field is used: Y = Yes, N = No. 5. Mandatory for controllers that support the Namespace Management capability (refer to the NVM Express Base Specification).						

4.1.6.1 I/O Command Set specific Identify Namespace data structure (CNS 05h, CSI 03h)

Figure 25 defines the I/O Command Set specific Identify Namespace data structure for this command set. This data structure is used to provide information about a memory namespace based on its NSID.

Figure 25: I/O Command Set Specific Identify Namespace data structure for the SLM Command Set

Bytes	O/M ¹	Description	Capability Field
07:00	M	Namespace Size (NSZE): This field indicates the total size of the namespace in bytes. The value provided by this field shall be dword granular (i.e. bits 1:0 cleared to 00b).	No
08	M	Number of Formats (NF): Indicates the number of format descriptors supported by the namespace. Formats shall be allocated in order (starting with 0) and packed sequentially. This is a 0's based value. The maximum number of formats that may be indicated as supported is 32 (i.e., this field set to 31). Format Index fields with a value greater than the value set in this field are invalid and not supported. The supported SLM formats are indicated in bytes 512 to 1023 in this data structure. SLM Formats that are valid, but not currently available may be indicated by setting the Valid bit for that SLM Format to '0'. All memory namespaces shall clear this field to 0h (i.e., one format is supported).	Yes
12:09	O	Namespace Optimal Write Granularity (NOWG): This field indicates the size in dwords for optimal write performance for this namespace. The value is reported in terms of a power of two (2^n). Write operations are required to be aligned on a dword boundary and store multiples of dwords. The host should construct Memory Write commands that store multiples of NOWG dwords to achieve optimal performance.	No
Memory Copy Limits			
20:13	O	Memory Copy Maximum Copy Length (MCMCL): This field indicates the maximum number of bytes that may be specified in a Memory Copy command (i.e., the sum of the number of bytes specified in all Source Range entries). If the controller supports the Memory Copy command, then this field shall be set to a non-zero value.	No
24:21	O	Memory Copy Maximum Single Source Range Length (MCMSSRL): This field indicates the maximum number bytes that may be specified in the NBYTE field in each valid Source Range Entries Descriptor of a Memory Copy command (refer to section 3.2.4). If the controller supports the Memory Copy command, then this field shall be set to a non-zero value.	No
25	O	Memory Copy Maximum Source Range Count (MCMSRC): This field indicates the maximum number of Source Range entries that may be used to specify source data in a Memory Copy command. This is a 0's based value.	No
511:26		Reserved	
SLM Format List			
527:512	M	SLM Format 0 Support (F0): This field indicates the format 0 that is supported by the namespace. The format data structure is defined in Figure 26.	Yes
543:528	O	SLM Format 1 Support (F1): This field indicates the format 1 that is supported by the namespace. The format data structure is defined in Figure 26.	Yes
...			
1023:1008	O	SLM Format 31 Support (F31): This field indicates the format 31 that is supported by the namespace. The format data structure is defined in Figure 26.	Yes

Figure 25: I/O Command Set Specific Identify Namespace data structure for the SLM Command Set

Bytes	O/M ¹	Description	Capability Field
4095:1024		Reserved	
Notes:			
1. O/M definition: O = Optional, M = Mandatory.			

Figure 26: Format Data Structure for the SLM Command Set

Bits	Description
127	Valid (VAL): Indicates whether this data structure describes a valid format. If this bit is set to '1', this data structure describes a valid format. If this bit is cleared to '0', then the SLM format is not available. Formats with an index greater than 0 shall clear VAL to '0'. The SLM format with an index of 0 shall set VAL to '1'.
126:8	Reserved
7:0	Data Size (DS): is field indicates the granularity of the format of the SLM namespace. The value is reported in terms of a power of two (2 ⁿ) bytes. A value greater than 0 (i.e., 1 byte) is not supported.

4.1.6.2 Identify I/O Command Set specific Controller data structure (CNS 06h, CSI 03h)

Figure 27 defines the I/O Command Set Specific Identify Controller data structure for this command set.

Figure 27: Identify Controller Data Structure, SLM Command Set Dependent - General

Bytes	O/M	Description
03:00	M	Version (VER): This field contains a Specification Version Descriptor (refer to the NVM Express Base Specification) indicating the version of this specification supported by the controller, as defined in Figure 28.
4095:04		Reserved

Figure 28: SLM Command Set Specification Version Descriptor Field Values

Specification Versions ¹	MJR Field	MNR Field	TER Field
1.0	1h	0h	0h
Notes:			
1. The specification version listed includes lettered versions (e.g., 1.0 includes 1.0, 1.0a, 1.0b, etc.).			

4.1.7 Namespace Attachment command

The Namespace Attachment command operates as defined in the NVM Express Base Specification.

4.1.8 Namespace Management command

The Namespace Management command is not supported for memory namespaces. If a Namespace Management command specifies this I/O Command Set (i.e., CSI 03h), then the controller shall abort that command with a status code of Invalid Field in Command.

4.1.8.1 Host Software Specified Namespace Management Fields

Since the Namespace Management command is not supported, host specified namespace management fields are not applicable.

4.1.9 Sanitize command

Sanitize clears all bytes of memory namespaces to 0h. For memory namespaces, there is no difference in sanitize operation based on sanitize type (i.e., Block Erase, Crypto Erase, and Overwrite).

4.2 I/O Command Set Specific Admin commands

This command set does not define any Admin Commands.

5 Extended Capabilities

This command set does not add any Extended Capabilities beyond what are defined in the NVM Express Base Specification.